

Matching Point-sets using Procrustes Alignment and the EM Algorithm

Bin Luo^{1,2}, E. R. Hancock¹

¹Department of Computer Science,
University of York, York YO1 5DD, UK.

²Anhui University, P.R. China
luo, erh@minster.cs.york.ac.uk

Abstract

This paper casts the problem of point-set alignment via Procrustes analysis into a maximum likelihood framework using the EM algorithm. The aim is to improve the robustness of the Procrustes alignment to noise and clutter. By constructing a Gaussian mixture model over the missing correspondences between individual points, we show how alignment can be realised by applying singular value decomposition to a weighted point correlation matrix. Moreover, by gauging the relational consistency of the assigned correspondence matches, we can edit the point sets to remove clutter. We illustrate the effectiveness of the method on matching stereograms. We also provide a sensitivity analysis to demonstrate the operational advantages of the method.

1 Introduction

The problem of point pattern matching has attracted sustained interest in both the vision and statistics communities for several decades. For instance, Kendall [6] has generalised the process to projective manifolds using the concept of Procrustes distance. Ullman [14] was one of the first to recognise the importance of exploiting rigidity constraints in the correspondence matching of point-sets. Recently, several authors have drawn inspiration from Ullman's ideas in developing general purpose correspondence matching algorithms using the Gaussian weighted proximity matrix. For instance Scott and Longuet-Higgins [10] locate correspondences by finding a singular value decomposition of the inter-image proximity matrix. Shapiro and Brady [11], on the other hand, match by comparing the modal eigenstructure of the intra-image proximity matrix. In fact these two ideas provide some of the basic groundwork on which the deformable shape models of Cootes *et al* [3] and Sclaroff and Pentland [9] build. This work on the co-ordinate proximity matrix is closely akin to that of Umeyama [15] who shows how point-sets abstracted in a structural manner using weighted adjacency graphs can be matched using an eigen-decomposition method. These ideas have been extended to accommodate parameterised transformations [16] which can be applied to the matching of articulated objects [17]. More recently, there have been several attempts at modelling the structural deformation of point-sets. For instance, Amit and Kong [2] have used a graph-based representation (graphical templates) to model deforming two-dimensional shapes in medical images. Lades *et al* [7] have used a dynamic mesh to model intensity-based appearance in images.

Broadly speaking the aim of point pattern matching is to recover the transformation between image and model co-ordinate systems. In order to estimate the transformation parameters a set of correspondence matches between features in the two co-ordinate systems is required. In other words, the feature points must be labelled. Posed in this way there is a basic chicken-and-egg problem. Before good correspondences can be estimated, there need to be reasonable bounds on the transformational geometry. Yet this geometry is, after all, the ultimate goal of computation. This problem is usually overcome by invoking constraints to bootstrap the estimation of feasible correspondence matches. If reliable correspondences are not available, then a robust fitting method must be employed [12, 11]. This involves removing rogue correspondences through outlier rejection.

The idea underpinning this paper is to provide a new framework for the maximum likelihood alignment of point-sets which allows linkage between alignment and correspondence. Specifically, we aim to realise iterative Procrustes alignment[6] using the EM algorithm. The possibility of hidden or missing data provides a natural way of representing the unknown correspondences between individual points. In the maximisation step of the algorithm, we align the points so that they minimise the weighted Procrustes distance. In the expectation step the positional residuals are used to estimate correspondence matching probabilities used in the weighting process. Editing to correct structural errors in the point-sets can be performed on the basis of the consistency of the pattern of matches. These three processes are interleaved and iterated to convergence.

Although the architecture of our method has much in common with the Procrustes soft-assign method of Rangarajan, Chui and Bookstein[8], it is couched in a statistical rather than an optimization framework.

2 Point-Sets

Our goal is to recover the parameters of a geometric transformation $\Phi^{(n)}$ that best maps a set of image feature points \mathbf{w} onto their counterparts in a model \mathbf{z} . In order to do this, we represent each point in the image data set by a position vector $\vec{w}_i = (x_i, y_i)^T$ where i is the point index. We will assume that all these points lie on a single plane in the image. In the interests of brevity we will denote the entire set of image points by $\mathbf{w} = \{\vec{w}_i, \forall i \in \mathcal{D}\}$ where \mathcal{D} is the point index-set. The corresponding fiducial points constituting the model are similarly represented by $\mathbf{z} = \{\vec{z}_j, \forall j \in \mathcal{M}\}$ where \mathcal{M} denotes the index-set for the model feature-points \vec{z}_j .

Later on we will show how the two point-sets can be aligned using singular value decomposition. In order to establish the required matrix representation of the alignment process, we construct two co-ordinate matrices from the point position vectors. The data-points are represented by the following matrix whose columns are the co-ordinate position vectors,

$$D = (\vec{w}_1 \quad \vec{w}_2 \quad \dots \quad \vec{w}_{|\mathcal{D}|}) \quad (1)$$

The corresponding point-position matrix for the model is

$$M = (\vec{z}_1 \quad \vec{z}_2 \quad \dots \quad \vec{z}_{|\mathcal{M}|}) \quad (2)$$

One of our goals in this paper is to exploit structural constraints to improve the recovery of transformation parameters from sets of feature points. We abstract the representation of correspondences using a bi-partite graph. Because of its well documented

robustness to noise and change of viewpoint, we adopt the Delaunay triangulation as our basic representation of image structure [13, 5]. We establish Delaunay triangulations on the data and the model, by seeding Voronoi tessellations from the feature-points [1].

The process of Delaunay triangulation generates relational graphs from the two sets of point-features. An example is shown in Figure 1. More formally, the point-sets are the nodes of a data graph $G_D = \{\mathcal{D}, E_D\}$ and a model graph $G_M = \{\mathcal{M}, E_M\}$. Here $E_D \subseteq \mathcal{D} \times \mathcal{D}$ and $E_M \subseteq \mathcal{M} \times \mathcal{M}$ are the edge-sets of the data and model graphs.

3 Dual Step EM Algorithm

The aim in this paper is to show how the Procrustes alignment of the two point-sets can be realised using the EM algorithm[4]. The ultimate goal of the alignment process is to identify point-to-point correspondences between the data and the model. Moreover, we are interested in the case where there are significant structural differences between the two-point sets due to the addition of noise or the occlusion and drop-out of certain feature-points.

The EM algorithm provides a natural framework for recovering the required correspondences. The method is concerned with finding maximum likelihood solutions to problems posed in terms of missing or hidden data. In the alignment problem it is the correspondences which are missing and the transformation parameters that need to be recovered. The utility measure underpinning the method is the expected log-likelihood function. Under the assumption that the positional errors between the aligned point-sets are Gaussian, then the maximum likelihood problem becomes one of minimising a weighted squared error measure. The weights used to control the different positional errors are in fact the *a posteriori* probabilities of the point correspondences. The EM algorithm iterates between two interleaved computational steps. In the expectation step the *a posteriori* correspondence probabilities are estimated from the current position errors by applying the Bayes formula to the Gaussian distribution functions. In the maximisation step the alignment parameters are estimated so as to maximise the expected log-likelihood function. This is equivalent to minimisation of the weighted error measure. Here we realise the maximisation step by adopting a matrix representation of the point-sets together with their putative correspondence probabilities and applying singular value decomposition to recover the alignment parameters. In practice we iterate this so-called Procrustes alignment on a weighted correspondence matrix.

3.1 Mixture Model

The idea underpinning the EM algorithm is to construct a mixture model over the hidden data to explain the distribution of the observed data. The ultimate goal is the set of maximum likelihood parameters Φ which explain the observed distribution of data. In our alignment problem, the observed data are the position vectors belonging to the set w . The parameters are the translation, rotation and scaling required by the Procrustes alignment of the point-sets.

The method commences by assuming that the different observations are independent of one-another. As a result we can factorise the joint conditional likelihood of the data over the individual point position vectors, i.e.

$$p(\mathbf{w}|\Phi) = \prod_{i \in \mathcal{D}} p(\vec{w}_i|\Phi) \quad (3)$$

The next step is to focus on the probability distribution $p(\vec{w}_i|\Phi)$. Here we assume that the observed data-point positions have arisen from the model-points via a measurement process. However, the original model point is hidden from us. We must therefore entertain the possibility that each data point may have originated via measurement from any of the model-points. This situation is expressed probabilistically by constructing a mixture model over the set of hidden model-data associations or correspondences. As a result, we write

$$p(\vec{w}_i|\Phi) = \sum_{j \in \mathcal{M}} p(\vec{w}_i|\vec{z}_j, \Phi) \pi_{i,j} \quad (4)$$

where $p(\vec{w}_i|\vec{z}_j, \Phi)$ is the probability distribution for the data-point position measurement or observation \vec{w}_i to have originated from the model-point \vec{z}_j under the set of alignment parameters Φ . The quantity $\pi_{i,j}$ is the mixing proportion required for the model-point \vec{z}_j in explaining the observation \vec{w}_i .

With these ingredients, the complete likelihood function that has to be maximised is

$$\mathcal{L} = \prod_{i \in \mathcal{D}} \sum_{j \in \mathcal{M}} p(\vec{w}_i|\vec{z}_j, \Phi) \pi_{i,j} \quad (5)$$

The idea underpinning the EM algorithm is to accommodate the hidden data, be re-couching the maximisation of the likelihood function in terms of the expected log-likelihood function. It was Dempster, Laird and Rubin [4] who originally showed that maximising the expected value of the log-likelihood function under hidden or missing data, was equivalent to maximising the following quantity

$$Q(\Phi^{(n+1)}|\Phi^{(n)}) = \sum_{i \in \mathcal{D}} \sum_{j \in \mathcal{M}} P(\vec{z}_j|\vec{w}_i, \Phi^{(n)}) \ln p(\vec{w}_i|\vec{z}_j, \Phi^{(n+1)}) \quad (6)$$

According to this viewpoint the *a posteriori* probabilities available at iteration n of the algorithm are used to compute the expectation value of the log-likelihoods of the missing data at iteration $n + 1$.

3.2 Maximisation

To develop a useful alignment algorithm we require a model for the measurement process. Here we assume that the observed position vectors, i.e. \vec{w}_i , are derived from the model points through a Gaussian error process. Suppose that the revised estimate of the position of the point \vec{w}_i under the set of alignment parameters $\Phi^{(n)}$ is $\vec{w}_i^{(n)}$. According to our Gaussian model of the alignment errors,

$$p(\vec{w}_i|\vec{z}_j, \Phi^{(n)}) = \frac{1}{2\pi\sqrt{|\Sigma|}} \exp\left[-\frac{1}{2}(\vec{z}_j - \vec{w}_i^{(n)})^T \Sigma^{-1} (\vec{z}_j - \vec{w}_i^{(n)})\right] \quad (7)$$

where Σ is the variance-covariance matrix for the point measurement errors. Here we assume that the position errors are isotropic, in other words the errors in the x and y directions are identical and uncorrelated. As a result we write $\Sigma = \sigma^2 I_2$ where I_2 is

the 2x2 identity matrix. With this model, maximisation of the expected log-likelihood function $Q(\Phi^{(n+1)}|\Phi^{(n)})$ reduces to minimising the weighted square error measure

$$\mathcal{E} = \sum_{i \in \mathcal{D}} \sum_{j \in \mathcal{M}} \xi_{i,j}^{(n+1)} (\vec{z}_j - \vec{w}_i^{(n)})^T (\vec{z}_j - \vec{w}_i^{(n)}) \quad (8)$$

where we have used the shorthand notation $\xi_{i,j}^{(n+1)}$ to denote the *a posteriori* correspondence probability $P(\vec{z}_j | \vec{w}_i, \Phi^{(n)})$. This is similar to the utility measure used by Rangarajan, Chui and Bookstein[8].

We would like to recover the maximum likelihood alignment parameters by applying Procrustes normalisation to the two point-sets. This involves performing singular value decomposition of a point-correspondence matrix. In order to develop the necessary formalism, we rewrite the weighted squared error criterion using a matrix representation. Suppose that $W^{(n)}$ is the data-responsibility matrix whose elements are the *a posteriori* correspondence probabilities $\xi_{i,j}^{(n)}$. With this notation the quantity \mathcal{E} can be expressed in the following matrix form

$$\mathcal{E} = Tr[M^T W^{(n)} M] - 2Tr[D^{(n+1)} W^{(n)} M^T] + Tr[D^{(n+1)T} W^{(n)} D^{(n+1)}] \quad (9)$$

Since the first and third terms of this expression do not depend on the alignment of the point-sets we can turn our attention to maximising the quantity

$$\mathcal{F} = Tr[D^{(n+1)} W^{(n)} M^T] \quad (10)$$

where $D^{(n+1)}$ is the revised matrix of point-positions which we aim to estimate via maximisation of \mathcal{E} . This quantity can be thought of as a weighted measure of overlap or correlation between the point-sets under the current alignment estimate. It is worth pausing to consider its relationship with measures exploited elsewhere in the literature on point pattern matching. The quantity $M^T D$ is simply the standard measure of overlap that is minimised in the work on least-squares alignment [16]. The matrix W , on the other hand, is just the inter-image proximity matrix used by Scott and Longuet-Higgins [10]. So, the utility measure delivered by the EM algorithm plays a synergistic role. The inter-image point proximity matrix weights the least-squares criterion.

The quantity \mathcal{F} can be maximised by performing a singular value decomposition. The procedure is as follows. The matrix $D^{(n+1)} W^{(n)} M^T$ is factorised into a product of three new matrices U , V and Δ , where Δ is a diagonal matrix whose elements are either zero or positive, and U and V are orthogonal matrices. The factorisation is as follows

$$D^{(n+1)} W^{(n)} M^T = U \Delta V^T \quad (11)$$

The matrices U and V define a rotation matrix Θ which aligns the principal component directions of the point-sets M and D . The rotation matrix is equal to

$$\Theta = V U^T \quad (12)$$

With the rotation matrix to hand we can find the Procrustes alignment which maximises the correlation of the two point sets. The procedure is to first bring the centroids of the two point-sets into correspondence. Next the data points are scaled to that they have the same

variance as those of the model. Finally, the scaled and translated point-sets are rotated so that their correlation is maximised.

To be more formal the centroids of the two point-sets are $\mu_D^{(n)} = E(\mathcal{D})$ and $\mu_M = E(\mathcal{M})$. The corresponding covariance matrices are $\Sigma_D^{(n)} = E((\mathcal{D}^{(n)} - \mu_D^{(n)})(\mathcal{D}^{(n)} - \mu_D^{(n)})^T)$ and $\Sigma_M = E((\mathcal{M} - \mu_M)(\mathcal{M} - \mu_M)^T)$.

With ingredients the update equation for re-aligning the data-points is

$$\vec{w}_i^{(n+1)} = \mu_M + \frac{Tr \Sigma_M}{Tr \Sigma_D} V U^T (\vec{w}_i^{(n)} - \mu_D^{(n)}) \quad (13)$$

3.3 Expectation

In the expectation step of the EM algorithm the *a posteriori* probabilities of the missing data (i.e. the model-graph measurement vectors, \vec{z}_j) are updated by substituting the revised parameter vector into the conditional measurement distribution. Using the Bayes rule, we can re-write the *a posteriori* measurement probabilities in terms of the components of the corresponding conditional measurement densities

$$P(\vec{z}_j | \vec{w}_i, \Phi^{(n+1)}) = \frac{\alpha_{i,j}^{(n)} p(\vec{w}_i, \vec{z}_j | \Phi^{(n)})}{\sum_{j' \in \mathcal{M}} \alpha_{i,j'}^{(n)} p(\vec{w}_i, \vec{z}_{j'} | \Phi^{(n)})} \quad (14)$$

The mixing proportions are computed by averaging the *a posteriori* probabilities over the set of data-points, i.e.

$$\alpha_{i,j}^{(n+1)} = \frac{1}{|\mathcal{D}|} \sum_{i \in \mathcal{D}} P(\vec{z}_j | \vec{w}_i, \Phi^{(n)}) \quad (15)$$

3.4 Structural Editing

The final step in the matching process is to edit the data point-set to remove unmatchable points which are noise or clutter. The aim here is to measure the consistency of the arrangement of correspondence matches on each neighbourhood of the Delaunay graph. We meet this goal by comparing the matched edges of the data-graph with those in the model graph.

In order to gauge the consistency of match, we represent the state of correspondence match between the nodes of the data-graph and those of the model graph using a function $f^{(n)} : \mathcal{D} \rightarrow \mathcal{M}$. The statement $f^{(n)}(i) = j$ means that the data-graph node i is matched to the model-graph node j at iteration n of the algorithm. The matches are assigned on the basis of the maximum *a posteriori* correspondence probabilities. In other words,

$$f^{(n)}(i) = \arg \max_{j' \in \mathcal{M}} P(\vec{z}_{j'} | \vec{w}_i, \Phi^{(n)}) \quad (16)$$

The pattern of the assigned matches is used to compute a probability of compatible correspondence arrangement. To compute this probability, we appeal to the model of structural pattern error recently reported by Wilson and Hancock [18] for graph-matching by discrete relaxation. Accordingly, we assume that the probability of erroneous edge insertion is P_e . By counting the number of consistently matched edges in the neighbourhood of each node we can the measure of consistency of correspondence match. The




Figure 1: An example of Delaunay graph editing by node deletion. The original graph is shown on the left, while the edited graph resulting from a node deletion is shown on the right.

number of consistently matched edges for the correspondence match $f^{(n)}(i) = j$ is

$$H_{i,j} = \sum_{(i,i') \in E_D} (1 - \epsilon_{i,i'}) \quad (17)$$

where the consistency of each edge is measured by the quantity

$$\epsilon_{i,i'} = \begin{cases} 1 & \text{if } (j, f(i')) \in E_M \\ 0 & \text{otherwise} \end{cases} \quad (18)$$

With this consistency measure to hand, the probability of compatible correspondence match is

$$\chi_{i,j}^{(n)} = \frac{\exp[-\mu H_{i,j}]}{\sum_{j' \in \mathcal{M}} \exp[-\mu H_{i,j'}]} \quad (19)$$

where $\mu = \ln \frac{1-P_e}{P_e}$. This probability is used to make decisions concerning the deletion of nodes from the data graph that fail to find a consistent correspondence match. The node i is deleted if $\chi_{i,j}^{(n)} < P_e$. Once a point has been deleted the remaining points are retriangulated. This process is illustrated in Figure 1.

4 Experiments

In this section we provide some experimental evaluation of the new alignment method. This is divided into two sections. First, we provide a simulation study to provide comparative sensitivity characteristics. Secondly, we provide some real world experiments on stereo images.

To evaluate the robustness of the novel approach, we furnish a sensitivity study. This compares the new iterative alignment method with the following three alternatives:

- The first method(Referred to as "Weight+SVD") is similar to that of Scott and Longuet-Higgins. This performs the singular value decomposition $W^{(0)} = U_S \Delta_S V_S^T$ on the initial inter-image weight matrix. Suppose that $\hat{\Delta}_S$ is the matrix obtained by setting the diagonal elements of Δ_S to unity, then the Scott and Longuet-Higgins

BMVC99




Figure 2: Sensitivity study

algorithm delivers an updated matrix of correspondence weights $\hat{W} = U_S \hat{\Delta}_S V_S^T$. The updated weight matrix can be used to align the point-sets using the method outlined earlier in this paper.

- The second algorithm (Referred to as "Single SVD") performs the singular value decomposition $DM^T = U\Delta V^T$ to find the rotation matrix $\Theta = VU^T$ that maximises the unweighted point correlation $Tr[DM^T]$. In other words, the algorithm is an unweighted and non-iterative or single-shot version of the EM algorithm presented in this paper.
- The third method (Referred to as "SVD") is based upon aligning and scaling in the principal component axes of the two point-sets.

The data used in our study is furnished by randomly generated point-sets. We have added two types of noise to the point-sets. Firstly, we have added Gaussian measurement errors to the positions of the points. The position errors are isotropic and of zero mean. The parameter of the noise process is the standard deviation of the measurement error. The second type of noise is structural error. Here we have added controlled fractions of clutter points at random locations. We have measured the accuracy of alignment by computing the root-mean-square (RMS) error between the final point positions and the corresponding ground truth model point.

Figure 2a shows the RMS error as a function of the standard deviation of the point position error. The main point to note from this plot is that for all four algorithms the RMS error increases linearly with the noise standard deviation. However, for the new algorithm (EM+Weight+SVD - shown as circle points), the rate of increase of the RMS error is much lower than the remaining three algorithms. In other words, the new algorithm gives more accurate alignments.

Figure 2b shows the fraction of points in correct correspondence as a function of the fraction of added clutter. The main point to note for this plot is that the new method (EM+Weight+SVD - shown as circles) is considerably more accurate in locating correspondences. Moreover, the two SVD-based methods perform only marginally better than the PCA alignment,

BMVC99



Figure 3: Correspondence after Procrustes alignment




Figure 4: Correspondence after graph editing

Finally, we investigate the effect of removing the graph-editing step from our algorithm. Figure 2c shows the comparison of the percentage of correspondence errors as a function of the percentage of added clutter. The dashed curve is the result obtained when editing is used, while solid curve is the result obtained when editing is omitted. The graph editing method always improves the performance by a good margin.

Finally, we provide some examples on real-world data. Here we use pairs of stereo images of an office scene to test the proposed algorithm. Figure 3 shows the correspondences using conventional Procrustes alignment. There exist significant correspondence errors. When the EM method is used, most of the false correspondences are removed (Figure 4).

5 Conclusions

In conclusion, we have shown how the process of Procrustes alignment can be formulated as a maximum likelihood estimation problem using the EM algorithm. This interpretation leads to a new point-set similarity measure in which point correspondence probabilities weight the standard least-squares point overlap distance. In other words, our new measure of point-set similarity combines the ideas already developed by Scott and Lonquet-Higgins, and Umeyama in a single statistical utility measure. Moreover, our new method both allows structural constraints to be imposed on Procrustes alignment and provides a framework for point-set editing to remove noise and clutter. The method leads to more accurate point-set alignment.

References

- [1] N. Ahuja. Dot Pattern Processing using Voronoi Neighbourhoods. *IEEE PAMI*, 4:336–343, 1982.
- [2] Y. Amit and A. Kong. Graphical Templates for Model Registration. *IEEE PAMI*, 18:225–236, 1996.
- [3] T.F. Cootes, C.J. Taylor, D.H. Cooper, and J. Graham. Active Shape Models - Their Training and Application. *CVIU*, 61:38–59, 1995.
- [4] A.P. Dempster, N.M. Laird, and D.B. Rubin. Maximum-likelihood from incomplete data via the EM algorithm. *J. Royal Statistical Soc. Ser. B (methodological)*, 39:1–38, 1977.
- [5] O.D. Faugeras, E. Le Bras-Mehlman, and J-D. Boissonnat. Representing Stereo Data with the Delaunay Triangulation. *Artificial Intelligence*, 44:41–87, 1990.
- [6] D.G. Kendall. Shape Manifolds: Procrustean metrics and complex projective spaces. *Bulletin of the London Mathematical Society*, 16:81–121, 1984.
- [7] M. Lades, J.C. Vorbruggen, J. Buhmann, J. Lange, C. von der Maalsburg, R.P. Wurtz, and W.Konen. Distortion-invariant object-recognition in a dynamic link architecture. *IEEE Transactions on Computers*, 42:300–311, 1993.
- [8] A. Rangarajan, H. Chui, and F. L. Bookstein. The softassign procrustes matching algorithm. In *Information Processing in Medical Imaging*, pages 29–42. Springer, 1997.
- [9] S. Sclaroff and A.P. Pentland. Modal Matching for Correspondence and Recognition. *IEEE PAMI*, 17:545–661, 1995.
- [10] G.L. Scott and H.C. Longuet-Higgins. An Algorithm for Associating the Features of 2 Images. *Proceedings of the Royal Society of London Series B-Biological*, 244:21–26, 1991.
- [11] L. S. Shapiro and J. M. Brady. Feature-based Correspondence - An Eigenvector Approach. *IVC*, 10:283–288, 1992.
- [12] L. S. Shapiro and J. M. Brady. Rejecting Outliers and Estimating Errors in an Orthogonal-regression Framework. *Phil. Trans. Roy. Soc. A*, 350:403–439, 1995.
- [13] M. Tuceryan and T Chorzempa. Relative Sensitivity of a Family of Closest Point Graphs in Computer Vision Applications. *Pattern Recognition*, 25:361–373, 1991.
- [14] S. Ullman. *The Interpretation of Visual Motion*. MIT Press, 1979.
- [15] S. Umeyama. An Eigen Decomposition Approach to Weighted Graph Matching Problems. *IEEE PAMI*, 10:695–703, 1988.
- [16] S. Umeyama. Least Squares Estimation of Transformation Parameters between Point sets. *IEEE PAMI*, 13:376–380, 1991.
- [17] S. Umeyama. Parameterised Point Pattern Matching and its Application to Recognition of Object Families. *IEEE PAMI*, 15:136–144, 1993.
- [18] R.C. Wilson and E.R. Hancock. Structural Matching by Discrete Relaxation. *IEEE PAMI*, 19:634–648, 1997.