ELSEVIER

Pattern Recognition Letters 17 (1996) 263-276

Pattern Recognition Letters

A Bayesian compatibility model for graph matching

Richard C. Wilson *, Edwin R. Hancock

Department of Computer Science, University of York, York, Y01 5DD United Kingdom

Received 10 May 1995; revised 2 October 1995

Abstract

This letter presents a new methodology for determining the compatibility coefficients required for performing graph matching by probabilistic relaxation. The adopted framework is Bayesian and commences by specifying the effects of segmentation errors in corrupting the connectivity structure or topology of the graphs under match. This model of relational constraint corruption leads to a pattern of compatibility coefficients that is completely determined by the global topological properties of the graphs under match. We illustrate the application of this new theory in two graph matching applications. The first of these is concerned with exploiting constraints provided by edges. Here the compatibility coefficient for consistent edges is equal to the inverse edge-density. Our second illustration extends the compatibility model to the level of graph faces; the required coefficients are again parameter-free. We provide experimental validation of our method in the matching of aerial images. Here we demonstrate that the theoretical values of our compatibility coefficients are close to their experimentally optimal values.

1. Introduction

Inexact graph matching is a critical process for effective intermediate level scene interpretation. In essence, the technique allows a corrupted relational description of noisy image entities to be matched against an idealised model graph. It is the process of eliciting stable relational descriptions of data and model that provides the main obstacle to effective matching (Sarkar and Boyer, 1993). Errors in the feature segmentation process disrupt the topology of the scene graph, corrupting the matching constraints provided by the model graph. Viewed from this perspective, the process is one of inexact graph matching aimed at locating the maximally consistent interpretation of a topologically corrupted scene-graph.

As a consequence of its pivotal role in practical vision applications, the graph matching problem has been the focus of sustained activity in a number of diverse methodological areas. Classically, inexact matching is accomplished by either graph-search (Horaud and Skordas, 1989; Shapiro and Haralick, 1985) or by relaxation labelling techniques (Bhanu and Faugeras, 1984; Cucka and Rosenfeld, 1992; Faugeras and Price, 1981; Kittler et al., 1993; Li, 1992; Price, 1985; Rangananth and Chipman, 1992; Wilson and Hancock, 1994). In graph-search departures from consistency are handled by locating subgraph isomorphisms (Shapiro and Haralick, 1985) or by identifying maximal cliques (Horaud and Skordas, 1989). As demonstrated by Hummel and Zucker (1983), relaxation approaches share the common goal of optimising a global criterion of match; in the case of discrete relaxation (Wilson and Hancock, 1994) the optimisation process is realised by symbolic label updating while in the case of probabilistic relaxation (Bhanu and Faugeras, 1984; Cucka and Rosenfeld, 1992; Kittler et al., 1993; Li, 1992; Price, 1985; Ran-

SSDI 0167-8655(95)00115-8

gananth and Chipman, 1992) match probabilities are updated in the light of local support. In both cases the main computational requirement is a means of gauging the quality and consistency of match (Hancock and Kittler, 1990; Kittler and Hancock, 1989; Wilson and Hancock, 1994).

In this letter it is the probabilistic relaxation method that concerns us (Bhanu and Faugeras, 1984; Cucka and Rosenfeld, 1992; Faugeras and Price, 1981; Kittler and Hancock, 1987; Kittler and Hancock, 1989; Kittler et al., 1993; Li, 1992; Price, 1985; Ranganath and Chipman, 1992). Here the consistency of match is gauged by a support function (Kittler and Hancock, 1989; Kittler et al., 1993; Li, 1992) which is defined over the nodes of the graphs under match. The support function combines evidence for putative matches by drawing on a model of labelling compatibility. Effective matching depends critically upon both the support function and the compatibility model. The support function should ideally provide an objective way of assessing the evidential support for putative matches. Compatibility coefficients on the other hand should model the way in which idealised relational constraints are violated in realistic image data. Despite significant recent advances in the theory underpinning the design of effective support functions (Kittler and Hancock, 1989; Kittler et al., 1993), the availability of an objective methodology for modelling the compatibility coefficients has proved to be more elusive.

Support function design has been an issue of central importance, aimed at overcoming the perceived limitations of the original non-linear Rosenfeld, Hummel and Zucker (1976) relaxation operator. In a nutshell, these efforts have concentrated on developing support functions that are both more objective in their quantification of the available evidence and which provide greater representational power (Kittler and Hancock, 1987; Kittler and Hancock, 1989; Kittler et al., 1993). For instance, Kittler and Hancock (1989) have cast the probabilistic relaxation process into a Bayesian framework in which the support function is specified in terms of probability distributions with precise meaning and well defined modelling roles. Moreover, this new theory allows support functions to be tailored to relational structures of specific neighbourhood topology (Kittler and Hancock, 1987). Renewed interest in the application of the probabilistic relaxation technique to matching problems has recently been stimulated by the work of Li (1992) who has presented an effective yet largely goal directed framework based around an attributed relational graph representation. The novelty of this work derives from the way in which the support function unifies measurement and symbolic information through the use of binary relations. Some of the heuristic elements of Li's (1992) support model have recently been overcome by Kittler, Christmas and Petrou (1993) who have succeeded in establishing the method within a Bayesian framework that significantly extends the earlier of work of Kittler and Hancock (1989).

Despite these efforts aimed at enhancing the evidence combining capabilities of probabilistic relaxation, the issue of compatibility modelling has attracted less attention. Most of the compatibility models reported in the literature are goal directed and require the tuning of extensive parameter sets. For instance, in an implementation aimed at 2D shape matching Bhanu and Faugeras (1984) have suggested and compared a number of alternative strategies for computing compatibility coefficients. Although different in specific detail, these strategies all share the common feature of drawing on a sigmoidal function of the weighted error between matched features. Faugeras and Price (1981) present a compatibility model which counts the relational co-occurrence of semantic labels. Price (1985) has developed this idea further using the compatibilities to grade relations according to a preferential hierarchy. These two basic modelling methods have recently attracted renewed interest. Revisiting the original Rosenfeld, Hummel and Zucker (1976) non-linear update formula under the guise of fuzzy relaxation, Ogawa (1994) utilises a sigmoidal compatibility model for shape matching while, Ranganath and Chipman (1992) explore the graph matching problem using relational co-occurrence compatibilities. The recent work of Kittler, Christmas and Petrou (1993) represents an important advance in trying the minimise the parametric complexity of the compatibility model. Their compatibility model is represented by well defined probability distributions for mixtures of measurements and symbols whose parameters may be determined using well established estimation methods. Despite this advance, the modelling of compatibility is aimed at capturing the variability of binary measurement relations rather than by directly modelling relational inexactness at the level of graph topology.

Our aim is this letter is to adopt a more direct approach that attempts to model the structural disruption of graphs caused by segmentation errors. The framework for this study is provided by the evidence combining support functions of Kittler and Hancock (1989). In contrast to the work of Kittler, Christmas and Petrou (1993), this framework demands a rigid dichotomy between measurements and symbols. From the modelling perspective, this dichotomy provides certain tangible advantages in terms of its capacity to represent the different sources of uncertainty in the matching process. We observe that the main limitation of existing methods derives from their failure to directly model the processes by which the topology of the scene-graph becomes corrupted; this frequently results in heuristic matching algorithms of unnecessary parametric complexity which are not easily controlled.

By contrast, our compatibility coefficients directly model constraints operating in the graphs under match and attempt to capture the effects of segmentation errors. According to our philosophy these compatibilities are purely symbolic in their modelling role since they represent the disrupted topology of scene structure. The novel aspect of the work reported here is to present a methodology which allows the compatibility coefficients to be determined using a topological constraint corruption process. The entire set of compatibility coefficients required to implement the relaxation scheme is completely determined by the numbers of nodes and edges in the scene graphs under match. This not only represents a significant advance in the existing methodology, it also represents an extremely desirable computational property since it obviates the need for an elaborate parameter estimation scheme.

The outline of this letter is as follows. Section 2 details the relaxation scheme used for graph matching. Section 3 develops the novel topological constraint corruption model required to compute the compatibility coefficients. Section 4 describes the experimental evaluation of the method for a matching application involving the fusion and registration of aerial infra-red images. Finally, Section 5 offers some conclusions.

2. Relaxation labelling

We are interested in matching tasks that can be abstracted in terms of relational graphs. The different images to be matched are given a graph representation in which the nodes are segmental entities and the edges signify the existence of a topological relationship between two nodes. We use the notation $G = (V, E, F)$ to denote such graphs, where V is the set of nodes, E is the set of edges and F is the set of first-order triangular faces. Edges are represented by Cartesian pairs of nodes, $(i, j) \in V \times V$. We use the term face to refer to Cartesian triples of the form $(i, j, k) \in V \times V \times V$. Each pair of nodes in the face is itself a consistent edge, i.e.

$$(i,j,k) \in F \Leftrightarrow (i,j) \in E \wedge (j,k) \in E \wedge (i,k) \in E.$$

In other words, the faces are first-oder cycles of the graphs. Our aim in matching is to associate nodes in two such graphs $G_D = (V_D, E_D, F_D)$ and $G_M = (V_M, E_M, F_M)$ using constraints provided by either edges or faces. Unmatchable entities are accommodated by augmenting the nodes in the model graph by a null label ϕ.

In order to perform the matching task, the relaxation scheme must iteratively update the probability of data node $J \in V_D$ matching model node $j \in V_M \cup \phi$. At iteration n of the relaxation scheme this probability is denoted by $P^{(n)}(J \to j)$. According to our viewpoint the modelling role of the initial probabilities, $P^{(0)}(J \to j)$, is to capture uncertainties in the raw measurement information. In other words, the initial probabilities represent matching affinities between nodes of the data and model graphs based on unary object attributes. It must be stressed that the modelling of these initial probabilities is not our prime concern in this letter. In our experimental evaluation we will make use of a conservative matching model based on line-length and relative orientation.

Commencing from a specification of the relaxation process in terms of the *a posteriori* label probabilities, Kittler and Hancock (1989) arrive at a probability update formula that bears of the hallmarks of the famous Rosenfeld, Hummel and Zucker (RHZ) (1976) non-linear relaxation operator. Although this Bayesian derivation is only strictly applicable during the first update process, the iterative RHZ formula can be regarded as performing probability updating providing that it is interpreted as an an implicit measurement filtering operation. The formula for updating probabilities between iterations n and $n + 1$ of the relaxation scheme is

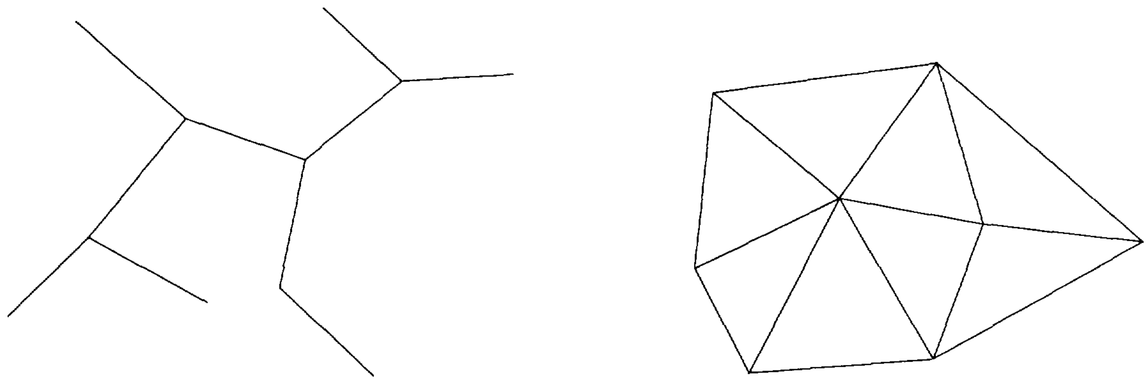

Fig. 1. Tree (left) and triangulation (right).

$$P^{(n+1)}(J \to j) = \frac{P^{(n)}(J \to j)Q(J \to j)}{\sum_{\lambda \in V_M \cup \phi} P^{(n)}(J \to \lambda)Q(J \to \lambda)}. \quad (1)$$

The crucial ingredient in this update formula is the support function $Q(J \to j)$. This combines evidence from the context conveying neighbourhood $\mathcal{G}_J$ of node J for the match $J \to j$. This neighbourhood consists of the set of nodes connected to node J by edges, i.e. $\mathcal{G}_J = \{I \in V_D \mid (I, J) \in E_D\}$; it incorporates information conveyed by the nodes which interact directly with J through sharing a common edge. There are many variants of the support function reported in the literature (see (Kittler and Hancock, 1989) for a review). Many of these are heuristic (Li, 1992; Rosenfeld et al., 1976) and suffer from internal inconsistencies in their specification. Others make recourse to unsatisfactory assumptions such as the weakness of contextual information.

Here we will use support functions designed using an internally consistent Bayesian evidence combining framework (Kittler and Hancock, 1989). This framework is capable of drawing on constraints applying to the consistent labellings of representational units of varying complexity. The simplest constraints are those that pertain to the consistent labellings of graph edges. As we will demonstrate later, edge-based constraints are most effectively exploited in the matching of tree-like graphs. These graphs contain few first-order cycles and have low edge densities. For graphs containing many triangular first order faces, the use of edge-based constraints alone may lead to internal inconsistencies in the specification of the support function. Under conditions where the neighbourhood structure of the graph is more conveniently expressed in terms of a triangulation, the support function may be formulated in terms of constraints applying at the level of first-order faces (see Fig. 1). In other words the choice of support function should closely reflect the neighbourhood topology of the graphs under match. Since the graphs used in our experimental study have relatively low edge densities, we will draw on triangulations only as an illustration of how our methodology extends to more complex relational structures. Suffice to say that we have conducted an exhaustive study concerned with the matching of triangulations which is reported in detail elsewhere (Finch et al., 1995).

Since our aim in this letter is principally to demonstrate a new methodology for determining compatibility coefficients for graph matching, we will not dwell on the nature of the support function. We therefore take as our principal demonstration vehicle the edge-based support function of Kittler and Hancock (1989). This support function has the form of a product-rule and is therefore not restricted to conditions of weak context, as is the case of the arithmetic average relaxation scheme applied by Kittler, Christmas and Petrou (1993). Details of its derivation are given in (Wilson and Hancock, 1992). The computation of support is achieved by summing over the potential label assignments to the nodes in the neighbourhood and taking the product over nodes

$$Q(J \to j) = \prod_{I \in \mathcal{G}_J} \sum_{i \in V_M \cup \phi} P^{(n)}(I \to i) R(I, J, i, j). \quad (2)$$

According to the Bayesian framework the compatibil-

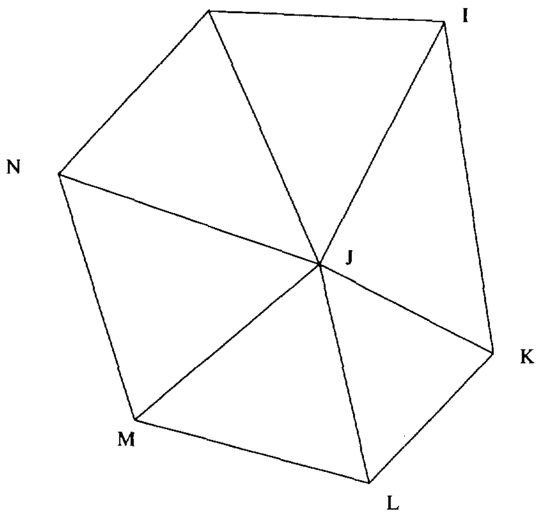

Fig. 2. Neighbourhood in a relational graph.

ity coefficient is specified by the mutual information measure

$$R(I,J,i,j) = \frac{P(I \to i, J \to j)}{P(I \to i)P(J \to j)}. \quad (3)$$

The Bayesian ingredients of this formula are the single node prior $P(J \to j)$ and the conditional prior $P(I \to i \mid J \to j)$. It is the conditional prior that measures the consistency between the symbolic matches $I \to i$ at the node I and $J \to j$ at node J. According to our Bayesian framework, the prior probabilities are responsible for representing a purely symbolic model of matching consistency. They do not, for instance, admit any form of attribute information. Unary attribute information is the modelled by the initial probabilities.

As we indicated earlier, detailed discussion of how support is computed at the face-level is beyond the scope of this letter. Details relating to the Bayesian derivation and efficient computation of a support function appropriate to the matching of triangulated graphs may be found in (Kittler and Hancock, 1987). However, for the purposes of our discussion it will suffice to note that the face-based counterpart of the Bayesian compatibility coefficient $R(I,J,i,j)$ defined in Eq. (3) is

$$S(I,J,K,i,j,k) = \frac{P(I \to i, J \to j, K \to k)}{P(I \to i)P(J \to j, K \to k)}. \quad (4)$$

The compatibility again takes the form of a mutual information measure, expressed additionally in terms of the prior probability $P(I \to i, J \to j, K \to k)$. This prior expresses the probability for for the triangular face $(I,J,K) \in F_D$ of the data graph to match onto model-graph triple (i,j,k); the model graph triples are potentially drawn from the space of combinations represented by the set $(V_M \cup \phi) \times (V_M \cup \phi) \times (V_M \cup \phi)$. In the next section we will present a model that can be exploited to compute both the edge and face compatibility coefficients. This model aims to describe the topological corruption of the data graph by segmentation errors.

3. Modelling constraint corruption

Application of the relaxation framework described in the previous section to the matching problem requires a probabilistic model of the graphs representing the different scenes. In constructing this model we would like to capture some of the uncertainties caused by segmentation error. These uncertainties include noise contamination, fragmentation due to over segmentation and merging due to under segmentation. They are manifest in the scene graphs as topological corruption, i.e. disruption of the local connectivity structure. Since the topology of the scenes is represented purely in terms of the labels assigned to nodes and the interconnectivity of the edges, this corruption process is purely symbolic. It therefore falls into the modelling domain of the compatibility coefficients in our relaxation scheme.

Our adopted modelling philosophy is that the relaxation operations are aimed at locating correspondences between the nodes in graph G_D and those in graph G_M. We take the view that the nodes in graph G_D represent data that must be matched in the light of constraints provided by graph G_M. These constraints are provided by the edges or faces appearing in the model graph G_M. Imperfect segmentation will corrupt the idealised pattern of constraints represented by the edge-set E_M or the face-set F_M. Segmental entities will invariably be lost through under-segmentation and extraneous entities will be introduced through over-segmentation effects. There will also be extraneous segments due to residual noise contamination. We would like to capture these effects in our probabilistic modelling of the label constraint process.

3.1. Edge compatibilities

We have recently reported a methodology which lends itself to this purpose (Hancock and Kittler, 1990). It is based on the idea of constraint corruption through the action of a label-error process. It can be viewed as providing a framework for softening the relational constraints represented by the edges of graph G_M. In order to admit the possibility of erroneous or extraneous nodes we have augmented the scene model with the label ϕ which facilitates null matches of the nodes in the graph G_M of the form $J \rightarrow \phi$. Because we admit the possibility of departures from consistency in the relaxation scheme we must potentially enumerate the computation of support over the complete combinatorial space of binary label configurations rather than over the edge-set alone. This implies that we must consider the possibility of match to any binary combination of augmented labels, i.e. $(j,i) \in (V_M \cup \phi) \times (V_M \cup \phi)$, with small yet finite probability. For this reason, we must assign the probability mass to the $|V_M \cup \phi \times V_M \cup \phi|$ terms entering the support function in Eq. (2) rather than to the $|E_M|$ configurations appearing in the edge-set. The specification of this distribution must capture departures from consistency.

The Bayesian basis for this constraint softening process is to compute the non-zero probabilities for each of the different combinatorial label configurations. Computation of these probabilities requires a model of the underlying constraint corruption process. In (Hancock, 1993; Wilson and Hancock, 1994) we modelled this process in terms of memoryless label corruption; the parameter of this process being the probability of label errors p. Underpinning our previous work was the objective of developing an error correcting matching process. Basic to this process was the idea of gauging local inconsistencies using Hamming distance. By contrast, our aim here is to use the concept of a label-error process to compute a global pattern of matching compatibilities. Rather than attempting to correct local matching errors by discrete relaxation operations (Wilson and Hancock, 1994), we aim the rectify them by imposing consistency in an evidence combining probabilistic relaxation process (Kittler and Hancock, 1989).

Following the methodology described in (Hancock, 1993) we adopt a binomial distribution of probability. Edges drawn from the graph G_M are uncorrupted and occur with total probability mass $(1-p)^2$. Edges with one node matched and one-node null-matched have total probability mass $2p(1-p)$. Edges in which both nodes are null-matched take the remaining probability mass, i.e. p^2. Matches involving non-null label pairs outside the edge-set V_M are completely forbidden and therefore account for zero total probability mass. In each of the three cases listed above the available mass of probability is distributed uniformly among the label configurations falling into the relevant constraint class. The resulting distribution of joint probability is specified by the following rule

$$P(I \rightarrow i, J \rightarrow j) = \begin{cases} \dfrac{(1-p)^2}{|E_M|} & \text{if } (j,i) \in E_M, \\ \dfrac{p(1-p)}{|V_M|} & \text{if } (j,i) \in (V_M \times \phi) \cup (\phi \times V_M), \\ p^2 & \text{if } (j,i) = (\phi, \phi), \\ 0 & \text{if } (j,i) \in V_M \times V_M - E_M. \end{cases} \tag{5}$$

The single-label priors required in the computation of compatibility coefficients are obtained by summing the joint probabilities in the axiomatic way with the following result

$$P(I \rightarrow i) = \begin{cases} \dfrac{1-p}{|V_M|} & \text{if } i \in V_M, \\ p & \text{if } i = \phi. \end{cases} \tag{6}$$

With the joint priors and the single-object priors to hand the compatibility coefficients required to implement the relaxation scheme are specified by the following rule

$$R(I,J,i,j) = \begin{cases} \dfrac{|V_M|^2}{|E_M|} & \text{if } (j,i) \in E_M, \\ 1 & \text{if } (j,i) \in (\phi \times V_M) \cup (V_M \times \phi), \\ 1 & \text{if } (j,i) \in \phi \times \phi, \\ 0 & \text{if } (j,i) \in V_M \times V_M - E_M. \end{cases} \tag{7}$$

This is a remarkable result: The graph-based constraint process is captured by a model which is entirely devoid of free parameters; it is specified purely in terms of the numbers of edge and nodes in the model-graph. Moreover, the strength of the constraints elicited from the edge-set for the model graph is gauged by the ratio $\rho_E = |V_M|^2/|E_M|$. It is the connectivity structure of the model graph that determines this ratio. When the model graph is tree-like and $|E_M| \simeq |V_M|$, then $\rho_E \simeq |V_M|$; it is under these conditions that the constraints are strongest. Under conditions in which the graph is fully connected, i.e. $|E_M| = |V_M|(|V_M|-1)/2$, in which case $\rho_E \simeq 2$, then the constraints are much weaker. It is interesting to note that even in this latter case, the compatibility coefficient is never small enough to merit using the weak-context approximation of the support function; the approximation is only valid under conditions where $\rho_E \simeq 1$.

Before proceeding it is worth stressing that the philosophy underpinning our compatibility model has been to average the effects of matching errors over the entire graph. This results in a global pattern of compatibilities that is both simple and economical in terms of parameter complexity. Localised compatibility models have a more complex structure. However, while drawing on more detailed local information concerning the structure of the graphs, they do so at the expense of greater demands in terms of parameter estimation or control. Specific examples include the binary attribute model of Christmas, Kittler and Petrou (1995) and the symbolic matching process of Wilson and Hancock (1994).

3.2. Face compatibilities

The weakening of pairwise constraints as the edge density increases eventually renders the support function given in Eq. (2) of little practical value for matching. Effective matching may, however, be facilitated by using relational constraints expressed at the face level. The model used to compute the matching compatibilities for edges can be extended to faces in a straightforward way. Application of the binomial error model leads to five distinct classes of face corruption:

- Faces of the data graph with all nodes matched to a face in the model graph are uncorrupted and occur with total probability mass $(1-p)^3$.
- Faces with two nodes matched to an edge in the model graph and one-node matched to the null label; these have total probability mass $p(1-p)^2$.
- Faces in which two nodes are null-matched have a probability of $p^2(1-p)$.
- Configurations in which all nodes in the triplet are null matches take the remaining probability mass, i.e. p^3.
- Matches involving non-null label triplets inconsistent with the above configurations are completely forbidden and therefore account for zero total probability mass.

In each of the cases listed above the available mass of probability is distributed uniformly among the label configurations falling into the relevant constraint class. Using this assumption we arrive at a rule which yields the joint probabilities for face and edge configurations:

$$P(I \to i, J \to j, K \to k) = \begin{cases} \frac{(1-p)^3}{|F_M|} & \text{if } (i,j,k) \in F_m, \\ \frac{p(1-p)^2}{|E_M|} & \text{if } (i,j) \in E_M \text{ and } k = \phi, \\ \frac{p^2(1-p)}{|V_M|} & \text{if } i \in V_M \text{ and } j = \phi \text{ and } k = \phi, \\ p^3 & \text{if } i = \phi \text{ and } j = \phi \text{ and } k = \phi, \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

The compatibility coefficients are derived by substituting the priors specified by Eqs. (8), (6) and (5) into Eq. (4). Their specification is given below in the following rule:

$$S(I,J,K,i,j,k) = \begin{cases} \frac{|E_M||V_M|}{|F_m|} & \text{if } (i,j,k) \in F_m, \\ \frac{|V_M|^2}{|E_M|} & \text{if } (i,k) \in E_M \text{ and } j = \phi \\ & \text{or if } (i,k) \in E_M \text{ and } j = \phi, \\ 1 & \text{if } j \in V_M \text{ and } i = k = \phi \\ & \text{or if } k \in V_M \text{ and } j = i = \phi \\ & \text{or if } i \in V_M \text{ and } j = k = \phi \\ & \text{or if } (i,j) \in E_M \text{ and } k = \phi \\ & \text{or if } j = k = i = \phi, \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

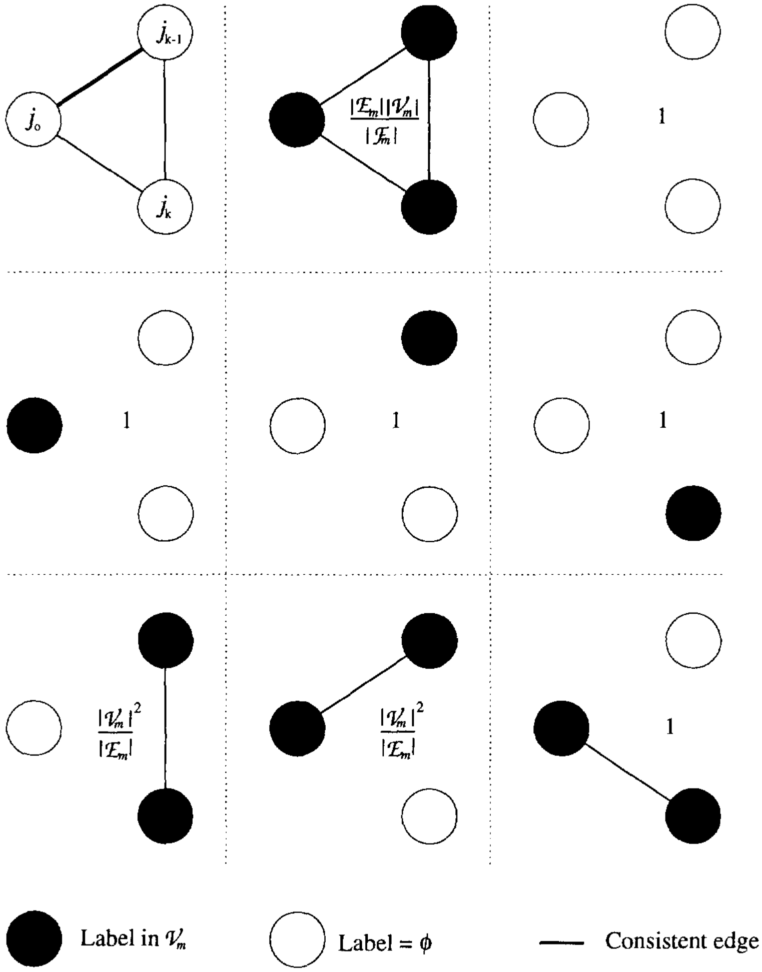

Fig. 3. The valid triplet labelling configurations.

Again the pattern of compatibilities is expressed in terms of the global topological properties of the model graph. The valid labelling configurations and their associated coefficients are illustrated in Fig. 3. The pattern of compatibilities clearly grades the different face-constraints according to their overall consistency. Fully consistent faces have higher compatibility values than partially matched faces which feature isolated edges or nodes. It is particularly interesting to note that the isolated edge compatibilities display several subtle features which discourage violation of the neighbourhood ordering relation. Partially matched faces containing a consistent trailing edge, i.e. $(k,j) \in E_M$, or an edge that preserves ordering in the contextual neighbourhood, i.e. $(i,k) \in E_M$, both receive a higher compatibility than the innovation of an isolated leading-edge that could potentially disrupt this ordering, i.e. $(i,j) \in E_M$. In other words, the compatibility pattern favours null matched nodes that are surrounded by a consistently ordered neighbourhood over matches that are connected to a plethora of incorrectly ordered yet individually consistent edges. It is this ability to impose ordering relations that enhances the internal consistency of our face-based relaxation scheme and represents the main advantage over purely edge-based compatibility models (Kittler et al., 1993; Wilson and Hancock, 1992) when triangulated graphs are being matched.

Finally, it is important to comment upon the relationship between the face compatibilities for trian-

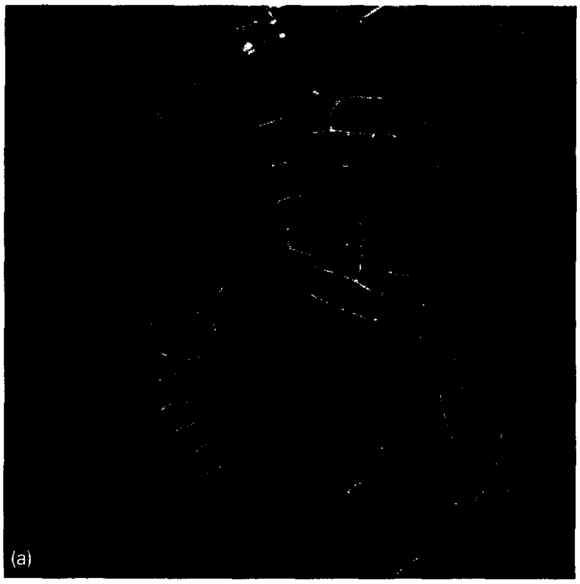

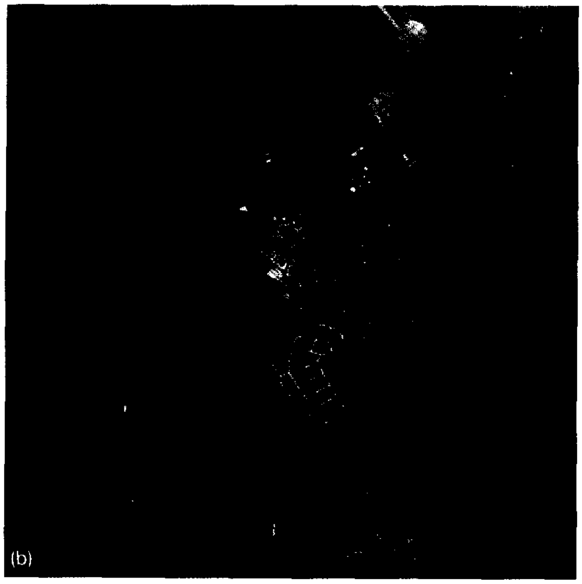

Fig. 4. Low- and high-altitude aerial images.

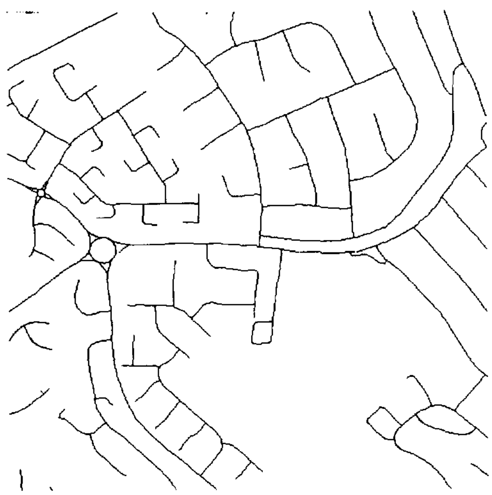

Fig. 5. Ordnance survey map data.

gulations and the edge compatibilities for trees. In the limit of large node-sets, the number of faces in a triangulation is approximately equal to the number of nodes, i.e. $|F_M| \simeq |V_M|$ while the number of edges is approximately twice the number of nodes, i.e. $|E_M| \simeq 2|V_M|$. As a result the face compatibility $\rho_F = (|E_M| \cdot |V_M|)/|F_M|$ is approximately given by $\rho_F \simeq 2|V_M|$ while the edge compatibility is approximated by $\rho_E \simeq |V_M|/2$. As viewed from the perspective of compatibility coefficients, the edges in the triangulation provide weaker constraints than their counterparts in a tree. Moreover, the face compatibility in the triangulation is twice as strong as the edge compatibility of the tree. In other words, our Bayesian model not only allows compatibilities to be computed in a parameter free way, it also provides a means of assessing the strength of the constraints provided by different types of relational structure.

4. Experimental evaluation

For the experimental aspects of this study we will be interested in matching road networks detected in aerial images. The data used in this study consists of infra-red line-scan data captured at varying altitudes together with ground truth map data. Example images are shown in Figs. 4(a) and 4(b). The corresponding map data is shown in Fig. 5. Inspection of the aerial images reveals that the road network is the dominant feature in the images which manifests itself as a web of connected intensity ridges in the raw data. We pose the matching problem as one of finding correspondences between junctions in the road network. In other words, the nodes of our graphs are T-junctions in the contour-web representing the road network. Edges represent the existence of connecting road structure. Economy of design dictates that transport networks should be of low edge-density. The network therefore contains very few first order faces and the use of the pairwise support function given in Eq. (2) is fully justified. Moreover, since the edges in the graph have trident topology at the nodes, there are are three times as many edges as nodes, i.e. $|E| = 3|V|$. In consequence, the edge compatibility is given by $\rho_E = |V|/3$.

The strategy that we adopt in eliciting segmental entities for matching is to first process the raw data to extract line contours representing the man-made road network. Extraction of the corresponding intensity ridge contours is achieved by applying orientational line detection kernels to the raw image and refining their output with a relaxation operator (Hancock, 1993). The refinement process enhances the connectivity of the extracted lines through the use of a model of local contour structure. Because the relaxation operator draws on an explicit dictionary model to represent this contour structure, the detected lines are rich in important junction features. Figs. 6(a) and 6(b) show the resulting line contours for the aerial images of Figs. 4(a) and 4(b). Although there are disparities between the results of segmentation and the cartographic information in the digital map, these are not severe. In other words, the detected junction features and their connecting line-structure appear to provide a viable basis for relational matching.

4.1. Initial match probabilities

The remaining model ingredient required to apply the relaxation formula to the feature matching process is a set of initial match probabilities between nodes. The role of the initial probabilities is to model transformational differences between the scenes under study, thereby complementing the relational constraint process which is modelled by the compatibility coefficients. It must be stressed that the modelling of these probabilities is not our prime concern in this paper. Rather, we aim to demonstrate the benefits to be derived from the compatibility model. Suffice to say that the modelling of the initial probabilities draws on geometric information derived from the properties of the lines forming junctions in the scenes under match. Essentially, we attempt to gauge the similarities between the trident line structures that form the T-junctions in the map and the image data. We do this by comparing both the angle differences within each trident and the ratio of the line-lengths. Uncertainties introduced by poor image segmentation or by geometric distortion must therefore be captured by imposing appropriate Gaussian probability distributions on the raw measurements. The parameters of this model are the mean orientation and scale differences between the two scenes under match, together with their corresponding variances. It should be stressed that there is no translational component to our model. Details of the model are outside the scope of this letter and may be found in (Kittler and Hancock, 1987).

4.2. Matching experiments

Our aims in this experimental section are twofold. In the first instance, we aim to demonstrate the effectiveness of our compatibility model in the matching of realistic imagery. Here we show that our very simple and purely symbolic compatibility model is capable of capturing realistic graph-errors. Our second aim is to show that our compatibility coefficients result in close to optimal performance.

In order to evaluate the performance of the relaxation scheme we have performed a number of experiments with the aim of posing tests of varying complexity to the matching algorithm. The data used in this study consists of aerial infra-red line scan images captured at various altitudes, together with ground-truth map data for the same region.

Our least demanding experiment involves matching a subregion of the map containing 13 junctions against a subregion of the low altitude image containing 15 junctions. Here we obtain 100% accuracy for T-junction matching. Our most demanding experiment involves matching the whole low altitude image against the whole high altitude image; there are 109 junctions in the low altitude image while the high altitude image contains 256 junctions. In this case 72% of T-junctions are matched correctly while 20% are matched to the null category. More details of the different experiments are summarised in Table 1. This gives a compilation of the results of matching either the entire datasets or selected subsets. In the column headed "Correct" we give the percentage of correct matches for the T-junctions. In the "Optimum" column we list the fraction of junctions for which a feasible match exists; these numbers represent the fraction of junctions in the first named image that are present in

(a) 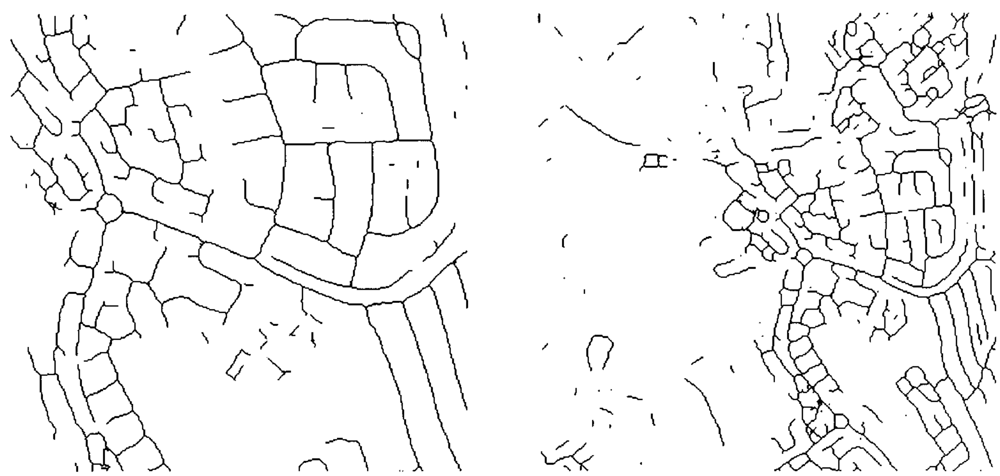 (b)

Fig. 6. Line contours.

Table 1
Comparison of experimental results

Matching (number of nodes)		ρ_E	Correct	Optimum	Corr/Opt	Null	Errors
Data graph	Model graph						
1. small map (13)	small low (15)	5	100%	100%	1.0	0%	0%
2. small low (15)	whole map (92)	30	87%	87%	1.0	7%	6%
3. small high (15)	whole map (92)	30	67%	80%	0.84	20%	13%
4. small low (15)	whole high (256)	85	67%	80%	0.84	20%	13%
5. whole low (109)	whole map (92)	30	44%	54%	0.81	36%	20%
6. whole map (92)	whole high (256)	85	45%	74%	0.61	36%	19%
7. whole low (109)	whole high (256)	85	72%	89%	0.81	28%	8%

the second named image. The "Null" column gives the fraction of junctions matched to the null-class. A perfect match would be represented by one in which the "Correct" and "Optimum" entries were identical and in which the "Optimum" and "Null" contents summed to unity. The fraction of residual matching errors is given in the final column of the table.

The table illustrates how the performance of the method degrades as the size of the datasets increases. It also illustrates the image-to-image matches are more reliable than map-to-image matches. On first consideration this may seem counterintuitive. However, inspection of Figs. 4 and 5 reveals differences at the segmentation level. This second point is demonstrated by the fraction of junctions available for matching in the "Optimum" column. However, even in the worst potential case of large image-to-map matches 61% of the feasible matches are correct.

The results of a typical matching experiment are shown in Fig. 7. Although the figure is confused by the multitude of matches, the lines denoting correct matches shown in the top half of the figure define a confocal envelope; those shown in the lower half of the figure are associated with incorrect matches and diverge from the envelope. It is worth noting that there is no pattern of organisation in the incorrect matches; they are distributed randomly across the scene. This is an encouraging observation since it means that the residual matching errors may be recoverable by a post-processing operation. One strategy would be to cluster the translational variables around the consistent matches provided by the relaxation process.

Our final experimental objective is to demonstrate that the pattern of compatibility coefficients predicted by our model give optimal performance when measured in terms of the experimental accuracy of match. To illustrate this point we have varied the value of the consistent edge compatibility about its theoretical value $\rho_E = |V_M|^2/|E_M$ and computed the fraction of correct matches. In other words, we aim to measure the

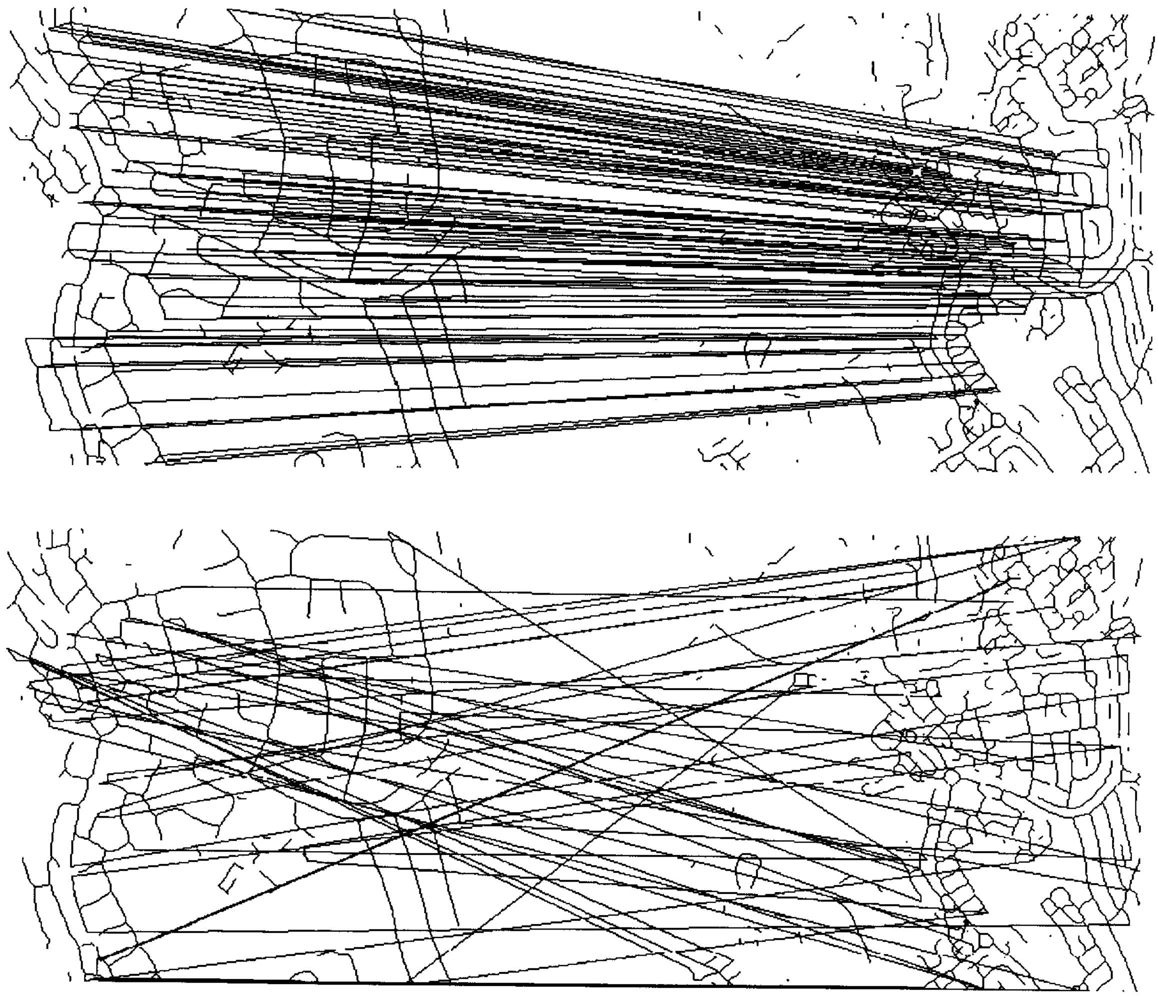

Fig. 7. Result of matching.

sensitivity of our matching method to changes in the compatibility coefficients. The results are displayed in Fig. 8. For the graphs under match $\rho_E = 84$. For the data under match, there are 97 junctions for which feasible matches exist. The matching accuracy shows a broad peak in the 40–90 compatibility range where it reaches its maximum value of 80%, when 78 junctions are matched correctly. Hand classification suggests that the maximum achievable accuracy is 89%.

5. Conclusions

In conclusion, we have reported an objective Bayesian methodology for computing the compatibility coefficients needed to implement relational matching as a probabilistic relaxation process. Drawing on a very simple model of memoryless matching errors, we derive patterns of compatibility coefficients that are entirely devoid of free parameters. The compatibility coefficients are completely specified by the global connectivity properties of the graphs under match. This is not only computationally attractive, it also provides a framework for assessing the relative strength of the constraints operating in different relational structures. Viewed from an experimental standpoint, the theoretical values of our compatibility coefficients are close to optimal in their performance.

Although very simple, our matching process ap-

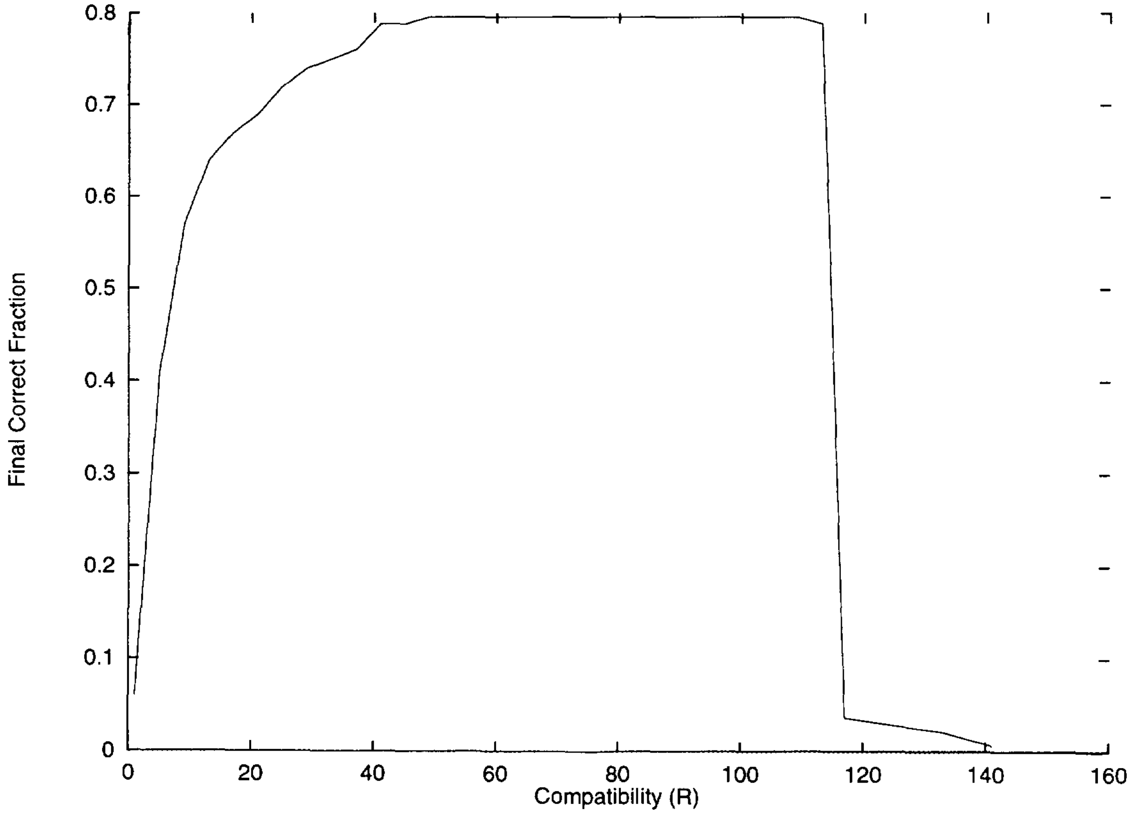

Fig. 8. Matching accuracy as a function of edge-compatibility.

pears to be an effective tool for image analysis. We acknowledge that we have used very little of the relational information available from the scenes under match. For instance, we have not tapped binary measurements of the sort exploited by Christmas, Kittler and Petrou (1995). Neither have we attempted to correct structural errors using graph-edit operations of the sort employed by Sanfeliu and Fu (1995). However, when weighed against these more sophisticated matching methods, we believe that one of the advantages of our probabilistic relaxation process lies in its economy of parameters and resulting ease of control. Suffice to say that we have extended the methodology described in this paper by both considering more complex relational entities (Finch et al., 1995) and by encompassing more elaborate control strategies such as graph editing (Wilson and Hancock, 1995).

References

Bhanu, B. and O.D. Faugeras (1984). Shape matching of 2-dimensional objects. *IEEE Pattern Anal. Mach. Intell.* 6, 137–156.

Christmas, W.J., J. Kittler and M. Petrou (1995). Structural matching in computer vision using probabilistic relaxation. *IEEE Pattern Anal. Mach. Intell.* 17, 749–764.

Cucka, P. and A. Rosenfeld (1992). Linear feature compatibility for pattern-matching relaxation. *Pattern Recognition* 25, 189–196.

Faugeras, O.D. and K.E. Price (1981). Semantic labelling of aerial images using stochastic relaxation. *IEEE Pattern Anal. Mach. Intell.* 3, 633–642.

Finch, A.M., R.C. Wilson and E.R.Hancock (1995). Matching Delaunay triangulations by relaxation labelling. In: *Computer Analysis of Images and Patterns*, Lecture Notes in Computer Science 970. Springer, Berlin, 350–358.

Hancock, E.R. (1993). Resolving edge-line ambiguities by relaxation labelling. *Proc. IEEE CVPR Conf.*, 300–306.

Hancock, E.R. and J. Kittler (1990). Discrete relaxation. *Pattern Recognition* 23, 711–733.

Horaud, R. and T. Skordas (1989). Stereo correspondence through feature grouping and maximal cliques. *IEEE Pattern Anal. Mach. Intell.* 11, 1168-1180.

Hummel, R.A. and S.W. Zucker (1983). On the foundations of relaxation labelling. *IEEE Pattern Anal. Mach. Intell.* 5, 267-287.

Kittler, J. and E.R. Hancock (1987). Contextual decision rule for region analysis. *Image and Vision Computing* 5, 156-165.

Kittler, J. and E.R. Hancock (1989). Combining evidence in probabilistic relaxation. *Internat. J. Pattern Recognition Artificial Intelligence* 3, 29-52.

Kittler, J., W.J. Christmas and M. Petrou (1993). Probabilistic relaxation for matching problems in machine vision. *Proc. Fourth Internat. Conf. on Computer Vision*, 666-673.

Li, S.Z. (1992). Matching invariant to translations, rotations and scale changes. *Pattern Recognition* 25, 583-594.

Ogawa, H. (1994). A fuzzy relaxation technique for partial shape matching. *Pattern Recognition Lett.* 15, 349-355.

Price, K.E. (1985). Relaxation matching techniques - a comparison. *IEEE Pattern Anal. Mach. Intell.* 7, 617-623.

Rangananth, S. and L.J. Chipman (1992). Fuzzy relaxation approach to inexact scene matching. *Image and Vision Computing* 10, 631-640.

Rosenfeld, A., R.A. Hummel and S.W. Zucker (1976). Scene labelling by relaxation operations. *IEEE Syst. Man Cybernet.* 6, 420-433.

Sanfeliu, A. and K.S. Fu (1995). A distance measure between attributed relational graphs. *IEEE Syst. Man Cybernet.* 13, 353-362.

Sarkar, S. and K.L. Boyer (1993). Perceptual organisation in computer vision; a review and proposal for a classificatory structure. *IEEE Syst. Man Cybernet.* 23, 382-399.

Shapiro, L.G. and R.M. Haralick (1985). A metric for comparing relational descriptions. *IEEE Pattern Anal. Mach. Intell.* 7, 90-94.

Wilson, R.C. and E.R. Hancock (1992). Relaxation matching of road networks in aerial images using topological constraints. *Proc. Internat. Society for Optical Engineering* 2059, 444-455.

Wilson, R.C. and E.R. Hancock (1994). Graph matching by discrete relaxation. In: E.S. Gelsema and L.N. Kanal, Eds., *Pattern Recognition in Practice IV: Multiple Paradigms, Comparative Studies and Hybrid Systems*. North-Holland, Amsterdam, 165-176.

Wilson, R.C. and E.R. Hancock (1995). Relational matching with dynamic graph structures. *Proc. Fifth Internat. Conf. on Computer Vision*, 450-456.