Active Graph Matching

Edwin R Hancock and Richard C Wilson
Department of Computer Science, University of York
York, Y01 5DD, UK

Abstract

This paper describes a novel approach to relational matching problems in machine vision. Rather than matching static scene descriptions, the approach adopts an active representation of the data to be matched. This representation is iteratively reconfigured to increase its degree of topological congruency with the model relational structure in a reconstructive matching process. The active reconfiguration of relational structures is controlled by a MAP update process. The final restored graph representation is optimal in the sense that it has *maximum a posteriori probability* with respect to the available attributes for the objects under match. The benefits of the technique are demonstrated experimentally on the matching of cluttered synthetic aperture radar data to a model in the form of a digital map. The operational limits of the method are established in a simulation study.

1 Introduction

Robust relational matching is a process of fundamental importance at intermediate levels in computer vision. It is a critical ingredient not only for effective stereopsis [5, 8, 15], but also for multi-sensor fusion [29], pose determination [22] and the matching of relational models [6, 7]. To be effective under realistic imaging conditions, the matching process must be robust to significant levels of noise, segmentation error and scene-clutter [31, 32]. It should additionally have stable performance when the available model is an over-representation of the data. Viewed from the perspective of structural matching, this means that scene elements may be missing or fragmented and that there may be significant numbers of extraneous elements. Under particularly severe conditions, extraneous entities my permeate the bona-fide elements, severely hindering the task of eliciting stable relational descriptions.

It is for these reasons that structural matching is invariably approached by inexact means [26, 27]. The matching process is frequently abstracted in terms of relational graphs [5, 15, 18, 21]; the critical ingredient being an efficient and robust way of searching a large space of matching possibilities when the data under study is corrupt or the model uncertain [5]. The available algorithms can be broadly divided into those that address the matching process in terms of attribute relations [5, 18] and those that attempt to operate with an effectively symbolic representation of the matching process [26, 27, 20, 31, 32]. Algorithms falling into the former category include the logarithmic conditional information of Boyer and Kak [5], and, the probabilistic relaxation scheme of Kittler, Christmas and Petrou [18]. Both algorithms draw on binary attribute relations and effectively aim to optimise a quadratic cost of inexact match. Higher-order attribute relations have recently been exploited by Yang and Kittler in a mean-field approach to matching [33]. Symbolic search-based approaches to the inexact matching process include the classical graph-distance metric of Shapiro and Haralick [27], and, the error-tolerant graph-edit idea of Messmer and Bunke [20]. Optimisation principals have found application in the simulated annealing ideas of Herault et al [16] and the configurational relaxation method of Wilson and Hancock [31, 32].

Underpinning these different approaches to the matching problem are a number of contrasting models of relational inexactness. In the case of attribute relations, observational inexactness is captured by Gaussian distributions over the available measurement space [5, 18]. Operating at a structural level, Shapiro and Haralick's [27] relational distance metric accommodates the effects of noise or segmentation error by inserting dummy nodes. Both the simulated annealing method of Herault [16] and the graph-edit approach of Messmer and Bunke [20], associate a cost function with the consistency of match. This idea of gauging the cost of symbolic consistency has recently been taken one stage further by Wilson and Hancock [31, 32]. Commencing from a model of the various types of noise and segmentation error anticipated to be present, they have developed a novel Bayesian criterion that can be used to gauge both the consistency of match and the degree of relational corruption. The criterion is compound exponential in character and measures consistency by the Hamming distance between the structural units of the graphs under match. Because the matching criterion is fine in its gauging of consistency, optimal matches may be located using straightforward gradient ascent methods.

Despite these contrasting views of relational inexactness, the techniques described above all share the common feature of exploiting a static representation of the structures under match.

*Richard Wilson is supported by EPSRC and DRA, Malvern under a CASE award; funding for this project comes from DRA, Malvern

Under conditions of extreme clutter and noise this can prove to be a severely limiting feature. If the genuine image entities are submersed in significant clutter, then little of the relational structure of the original model is likely to be preserved in the data. Although small levels of contaminating clutter can be accommodated by static means [5, 20, 27, 32], unless extreme care is taken, then there is a danger of diluting genuine relational structure to such an extent as to render the final result meaningless. Broadly speaking, there are two conventional strategies for controlling clutter in the matching process. The first of these attempts to associate a cost with clutter nodes and to tag or label them in an optimisation process. Algorithms falling into this category include the probabilistic relaxation technique of Kittler, Christmas and Petrou [18], and, the logarithmic conditional information of Boyer and Kak [5]. According to the second strategy, consistently matched subgraphs are identified in a search procedure. Barrow and Burstall provide a classical example in which the consistent matches are located by searching for maximal cliques of an association graph [4]. The methodology has recently been extended by Messmer and Bunke who use graph-edit [20] operations to locate consistent subgraph isomorphisms.

Each of the principal strategies for controlling clutter has an identifiable shortcoming. In the former case, unless the cost of null labels is carefully regulated, there is a danger of locating a globally null configuration. In the latter case, the persisting subgraphs may be so fragmented by the maximal clique finding process as to be of little practical value. Attempting to match a relational structure whose intrinsic topology is significantly corrupted by severe clutter can therefore only be anticipated to meet with a limited degree of success. It is the need to overcome the performance restrictions imposed by static relational descriptions abstracted away from the original image data that is central the work reported here. Our observation is that attempting to match static abstractions of cluttered or fragmented entities is only likely to meet with success if the level of contamination is relatively limited. Instead, we adopt the view that the matching task should be regarded as one of relational restoration or reconstruction in which the relevant abstraction is continually updated by reference to the available data.

In practice, this means that we dispense with static graph representations and operate instead with a dynamic relational model. By identifying and deleting extraneous entities in the data-graph we aim to directly rectify the cause of relational inexactness and restore the intrinsic topology of the graph structures under match. In other words, by rejecting relational noise, reconstructive matching permits constraints to be exploited more effectively. It must be stressed that there is a critical difference between our proposed reconstructive matching process and the maximal clique [13] or graph-edit ideas [20] conventionally applied to the identification of subgraph isomorphisms. Our proposed method is not simply concerned with modifying a relational abstraction that is isolated or decoupled from the data it represents. Instead a new relational description is constantly recomputed by removing extraneous entities from the raw data. It is this direct coupling between data and relational abstraction that is novel to our approach. Although the value of adaptive or dynamic representations has been recognised in the domain of deform-able models [30], they have hitherto not found widespread application in the context of relational modelling.

Our framework for realising these objectives is provided by *maximum a posteriori probability* estimation [8, 11]. Our MAP criterion allows us to assess the impact of each image entity on the relational structures representing both the data and the model under match. We can gauge the effect of adding or deleting nodes from the relational graphs in terms of a global consistency measure. This framework is sufficiently flexible to allow not only modification of the relational structure representing the data, but also to allow us to modify our relational model to rectify over-representation.

Conceptually, the process of reconfiguring the graph can be thought of as one of identifying relational outliers. As such, it is reminiscent of the use of robust statistics to reject outliers in data clustering [19, 24]. Instead of attempting to identify a set of inliers sharing a uniform co-ordinate transformation, our aim is to locate the entities that constitute a coherent structural cluster [13, 15] which satisfy relational constraints. Viewed from this perspective, our MAP criterion can be regarded as a relational clustering metric which can be used to reject a contaminating population of arbitrarily distributed outliers. Although the process of reconstructive matching has many aims in common with the constrained clustering technique of Rose, Gurewitz and Fox [23], it differs in one critical respect. Rather than employing a weighted Euclidean distance measure to gauge constraint violations, we attempt to exploit a more objective Bayesian framework in the modelling of relational consistency at the symbolic level. Following Wilson and Hancock [31, 32], we compute the probabilities of relational errors by compounding a series exponential functions of Hamming distance over the space of consistent mappings.

Evaluation of our method is based on the matching of Delaunay graphs representing Voronoi regions in 2D scenes [1, 2]. Here we demonstrate that the graph reconstruction process operates effectively provided that the fraction of clutter does not exceed 50% of the nodes in the corresponding uncorrupted data. It must be stressed that this high fidelity of reconstructive matching is obtained without any additional form of preprocessing. Construction of more elaborate relational descriptions, based for instance on perceptual groupings [14, 25], may be regarded as a form of structural pre-processing aimed at filtering away relational clutter. Unfortunately, the perceptual groupings that would prove most powerful in this respect are invariably extremely fragile and their robust identification requires a high degree of ingenuity to prove effective [25]. It must be stressed that our choice Delaunay graphs is dictated by convenience rather than by any intrinsic representational

limitation of our method. Nevertheless, it may be taken as indicative of the effectiveness of our technique that we can successfully match relying only on a very simple relational abstraction and without recourse to elaborate perceptual grouping operations.

The outline of this paper is follows. Section 2 introduces the formal ingredients of our method and describes our abstraction of scenes in terms of relational graphs. In Section 3 we describe our approach to modelling relational consistency using the concept of a label error process. Section 4 describes how this measure may be incorporated into the MAP estimation scheme as a model of the joint prior. The main theoretical contributions of the paper are detailed in Section 5, where we develop the formal criteria for node deletion and insertion. We present an experimental evaluation of our method in Section 6. This includes studies on both real and simulated data. The real world application is provided by the matching of hedge structures in SAR images against there cartographic representation in the form of a digital map. This provides a demanding test our of our methodology since the image data is fragmented and cluttered and the map an unreliable representation of ground truth. The simulation study complements the application, by allowing us to demonstrate the ability of our method to reject extraneous nodes under conditions of controlled clutter. Finally, Section 7 presents some conclusions and suggests directions for further investigation.

2 Relational Graphs

We abstract the matching process in terms of attributed relational graphs [5, 7, 18, 21, 27, 29]. According to this representation the nodes represent entities to be matched. The arcs represent relations operating between the nodes. In the experimental study reported in Section 6 we will be interested in matching hedge structures segmented from SAR images against their digital map representation. Here the nodes represent linear segments. Arcs denote the existence of a meaningful adjacency relation between pairs of line-segments. We establish the required relations on the basis of the adjacency of Voronoi regions generated from linear image segments. Here the relational structure used to abstract the image entities is a Delaunay graph. We have adopted this representation purely on the basis of convenience and it does not impose a limitation on the utility of our matching method. The framework presented in this paper is not only applicable to the plethora of computer vision tasks that are amenable to a Delaunay representation such as matching segmented regions or lines [5, 33], dot patterns [1, 2] and meshes [30]. It is also potentially amenable to a variety of alternative relational abstractions such as aspect graphs [6, 7], perceptual groupings [13, 15, 25] and feature relation graphs [29].

We use the triple $G = (V, E, \mathcal{A})$ to denote the graphs under match, where V is the set of nodes, E is the set of arcs and $\mathcal{A} = \{\underline{x}_i, \forall i \in V\}$ is a set of unary measurements associated with the nodes. Our aim in matching is to associate nodes in a graph $G_1 = (V_1, E_1, \mathcal{A}_1)$ representing data to be matched against those in a graph $G_2 = (V_2, E_2, \mathcal{A}_2)$ representing an available relational model. This matching process is facilitated using constraints provided by suitable relational subunits of the model graph G_2. Formally, the matching is represented by a function $f : V_1 \rightarrow V_2$ from the nodes in the data graph G_1 to those in the model graph G_2. The function f consists of a set of Cartesian pairs drawn from the space of possible matches between the two graphs, i.e. $f \subseteq V_1 \times V_2$; it provides a convenient device for indexing the nodes in the data graph G_1 against their matched counterparts in the model graph G_2. We use the notation $(u, v) \in f$ to denote the match of node $u \in V_1$ against node $v \in V_2$.

In performing the matches of the nodes in the data graph G_1 we will be interested in exploiting structural constraints provided by the model graph G_2. There are two issues at play in selecting structures appropriate to this task. If the structural units are too small then the matching process is impoverished in terms of the contextual information upon which it can draw in locating a consistent match. This limits the effectiveness of the matching scheme, rendering it susceptible to noise or error. If, on the other hand, the structural units are too large, then the matching process becomes excessively burdensome in terms of its computational requirements; the limitation stems from the need to explore the space of relational mappings between representational subunits. We will strike a compromise by using subgraphs that consist of neighbourhoods of nodes interconnected by arcs; for convenience we refer to these structural subunits or N-ary relations as supercliques.

The superclique of the node indexed j in the graph G_1 with arc-set E_1 is denoted by the set of nodes $C_j = j \cup \{i|(i, j) \in E_1\}$. We use the notation $R_j = (u_1, u_2, ..., u_{|C_j|})$ to denote the N-ary symbolic relation represented by the nodes of the superclique $C_j \subset V_1$ in the data graph G_1. The matched realisation of this superclique is therefore denoted by the relation $\Gamma_j = (f(u_1), f(u_2), ..., f(u_{|C_j|}))$. Our aim is to modify the match to optimise a measure of global consistency with the constraints provided by the model graph G_2. The constraints available to us are provided by the N-ary symbol relations on the supercliques of the model graph G_2. The critical ingredient in developing our matching scheme is therefore the set of feasible mappings between each superclique of graph G_1 and those of graph G_2. The set of feasible mappings, or dictionary, for the superclique C_j is denoted by $\Theta = \{S_i\}$ where $S_i = i \cup \{j|(i, j) \in E_2\}$. Each element S_i of Θ, is therefore a relation formed on the nodes of the model graph; we denote such consistent relations by $S_i = (v_1, v_2, ...)$. The dictionary of feasible mappings for the superclique C_j therefore consists of all the consistent relations that may be elicited from the graph G_2. In practice these relations are formed by performing cyclic permutation of the non-centre nodes for each superclique of the requisite size (an example of this mapping

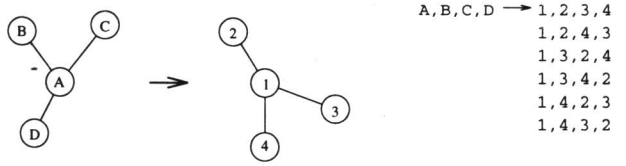

Figure 1: Example superclique mapping

process is shown in Figure 1); this process effectively preserves the adjacency structure of the model graph while leaving dictionary invariant to potential scene translations, scalings or rotations.

3 MAP Estimation

Our aim in performing active graph matching is to iteratively delete or reinstate the nodes of the data graph so as to optimise a global MAP criterion. At the computational level, this involves maintaining two separate representations of the data graph. The first of these represents the current relational description of the data, while the second is a recomputed graph from which a single node has been deleted. The means by which the graph is recomputed depends on the type of relational abstraction being employed. To provide a concrete example, in our experimental investigations we will be interested in matching Delaunay graphs representing Voronoi tessellation of the image plane [2]. Nodes of the Delaunay graph represent distinct Voronoi regions while the edges indicate region adjacency. Here the deletion of a node corresponds to removing a Voronoi region from the image plane. The Delaunay graph may be recomputed by growing those regions adjacent to the deleted node to fill the vacated space. In this way new relations are generated, i.e. new Delaunay edges are created between the surviving nodes.

To be more formal suppose $G_1 = (E_1, V_1, \mathcal{A}_1)$ is the current representation of the perceptual organisation of the data. Graph $G_1' = (E_1', V_1', \mathcal{A}_1')$, on the other hand, is the reconfigured relational structure formed by excluding the single perceptual entity represented by node u in G_1 and reassigning it to the null or outlier set Φ. The node-set $V_1' = V_1 - u$ and the attribute-set $\mathcal{A}_1' = \mathcal{A}_1 - \underline{x}_u^1$ are trivially recomputed by deleting the entry associated with node u. Determination of the edge-set E_1' is more complex and must be undertaken by recomputing the relational affinity of the nodes when the perceptual entity associated with node u is excluded from the raw image data. The match between the reconfigured graph G_1' and the model graph G_2 is represented by the function $f' : V_1' \rightarrow V_2$, where $f' = f - (u, f(u))$.

Our basic viewpoint is that modification of the graph structure G_1 is aimed at recovering a restored relational grouping G_1', under an assumed model of the data corruption process and a prior model of the structure to be recovered. From the standpoint of information theory, we therefore seek the grouped and matched configuration of nodes that has *maximum a posteriori probability* with respect to the available unary measurement information. From the computational perspective, the objective is to formulate iterative local decision schemes that are capable of both grouping and matching nodes. More formally, these iterative reassignment processes are aimed at partitioning the set of entities into those that are grouped to form a consistently matched graph f and those that are assigned to a set of relational outliers Φ. In other words, we seek the matched configuration of nodes that optimise the quantity

$$P(f, \Phi|\mathcal{A}_1, \mathcal{A}_2) = \frac{p(\mathcal{A}_1, \mathcal{A}_2|f, \Phi)P(f, \Phi)}{p(\mathcal{A}_1, \mathcal{A}_2)} \quad (1)$$

where $P(f, \Phi)$ is the joint prior for the current partition, $p(\mathcal{A}_1, \mathcal{A}_2|f, \Phi)$ and $p(\mathcal{A}_2, \mathcal{A}_2)$ are respectively the conditional measurement density and the probability density function for the sets of unary measurements.

We commence our development by making a number of conditional independence assumptions. Firstly we assume that the relational outliers are *a priori* independent of the match f, i.e.

$$P(f, \Phi) = P(f).P(\Phi) \quad (2)$$

This partitional model is extended to the corresponding conditional measurement densities which are also assumed independent of one-another

$$p(\mathcal{A}_1, \mathcal{A}_2|f, \Phi) = p(\underline{x}_u^1, \underline{x}_v^2, \forall(u, v) \in f|f)p(\underline{x}_u^1, \forall u \in \Phi|\Phi) \quad (3)$$

Furthermore, we assume that the different pairs of unary measurements associated by the matching f are conditionally independent of one another. Under these conditions, the joint conditional density appearing on the left-hand side in equation (1) may be factorised over the two subsets of nodes forming the partition. We therefore turn our attention separately to the set of Cartesian pairs that constitute the matching function and the nodes identified as relational outliers, i.e.

$$p(\mathcal{A}_1, \mathcal{A}_2|f, \Phi) = \prod_{(u,v)\in f} p(\underline{x}_u^1, \underline{x}_v^2|u, v). \prod_{u\in\Phi} p(\underline{x}_u^1|u) \quad (4)$$

Our basic assumption concerning the set of outliers is that it contains neither meaningful relational structure nor salient measurement information. As will become evident in Section 4, this is at complete odds with our model of $P(f)$ which is based on a relational consistency measure. This means that the outliers are conditionally independent of one-another and occur with a uniform probability P_ϕ. As a result of this assumption and applying the Bayes theorem, the above expression can be recast in terms of the *a posteriori* matching probabilities $P(u, v|\underline{x}_u^1, \underline{x}_v^2)$ and the matching priors $P(u, v)$

$$p(\mathcal{A}_1, \mathcal{A}_2|f, \Phi) = \prod_{(u,v)\in f} P(u, v|\underline{x}_u^1, \underline{x}_v^2)\frac{p(\underline{x}_u^1, \underline{x}_v^2)}{P(u, v)} \prod_{u\in\Phi} p(\underline{x}_u^1) \quad (5)$$

Turning our attention to the joint prior for the relational outliers and exploiting our uniformity assumption

$$P(\Phi) = P_{\phi}^{|\Phi|} \tag{6}$$

Graph reconstruction and matching proceed on the basis of node-by-node reassignment with the objective of locating the configuration of nodes that has the *maximum a posteriori probability*. Since this process involves modifying only a single node at a time, we confine our attention to the MAP ratio for the two realisations of the graph G_1 and G_1'. Moreover, because the unconditional measurement densities $p(\mathcal{A}_1, \mathcal{A}_2)$ and $p(\underline{x}_u^1, \underline{x}_v^2)$ are independent of the match labels and hence a static property of the data under consideration, the ratio in question depends only upon the *a posteriori* match probabilities and the matching priors. Furthermore, by virtue of the factorisation appearing in equation (4), the ratio involves only evaluating the effect reassigning the matched node u to the set of null entities or outliers

$$\frac{P(f, \Phi|\mathcal{A}_1, \mathcal{A}_2)}{P(f', \Phi|\mathcal{A}_1, \mathcal{A}_2)} = \frac{P(u, v|\underline{x}_u^1, \underline{x}_v^2)}{P(u, v).P_{\phi}} \frac{P(f)}{P(f')} \tag{7}$$

Essentially, we are concerned with developing an iterative reassignment process for locating an optimal partition between a set of outliers and a set of inliers that satisfy certain constraints upon relational consistency. In order to apply this global optimisation strategy we require three additional ingredients. These are models of the *a posteriori* matching probabilities and of the joint prior together with an effective algorithm control strategy. The most critical of these is a means of modelling the joint priors $P(f)$ so as to gauge the overall consistency of the surviving relational structure. It is the modelling of relational consistency that is the focus of our attention in the next section of this paper. The issues of control are discussed in Section 5. Details of the modelling of *a posteriori* matching probabilities are application dependent and are hence deferred until we address the issue of experimental validation of the technique in Section 6.

4 Global Consistency Criterion

In order to realise the graph reconfiguration process as one of iterative node reassignment, we require a metric of relational consistency. This model of the joint prior can be viewed as providing a means of imposing constraints on the relational cluster that constitutes the reconstructed graph. When viewed in this way, our reconstructive matching process has several conceptual similarities with the constrained clustering ideas of Rose, Gurewitz and Fox [23]. However, rather than using judiciously chosen Gibbs distributions to impose attribute constraints, we draw on an objective Bayesian model of relational corruption which is posed at the symbolic level. This results in a consistency metric which is a compound exponential function of relational distance (Hamming distance). Underlying this metric is a purely symbolic model of relational consistency, namely the dictionary of structure preserving mappings Θ.

Our strategy in modelling the joint prior is to construct an average consistency criterion [17] using the probabilities for the relational matches between subunits of the graphs G_1 and G_2. We effectively average the relational matching probabilities over the supercliques in graph G_1 using the supercliques in G_2 as a model. Fundamental to this approach, is an effective means of computing the probability relational matching between the supercliques the two graphs as specified by the function f. In other words, we must focus our attention on computing the probability of the matched relation Γ_j assigned to the superclique C_j of graph G_1.

As we noted in Section 2, the consistent labellings available for gauging the quality of match are represented by the set of relational mappings from C_j onto G_2, i.e. Θ. As demanded by the Bayes rule, we compute the probability of the required superclique matching by expanding over the basis configurations belonging to the dictionary Θ

$$P(\Gamma_j) = \sum_{S_i \in \Theta} P(\Gamma_j|S_i).P(S_i) \tag{8}$$

The development of a useful graph-mapping measure from this expression requires models of the processes at play in matching and of their roles in producing errors. These models are represented in terms of the conditional matching probabilities $P(\Gamma_j|S_i)$ and of the joint priors $P(S_i)$ for the consistent relations in the dictionary. In developing the required models we will limit our assumptions to the case of matching errors which are memoryless and occur with uniform probability distribution.

To commence our modelling of the conditional probabilities, we assume that the various types of matching error for nodes belonging to the same superclique are memoryless. In direct consequence of this assumption, we may factorise the required probability distribution over the symbolic constituents of the relational mapping under consideration. As a result the conditional probabilities may be expressed in terms of a product over label confusion probabilities

$$P(\Gamma_j|S_i) = \prod_{k=1}^{|S_i|} P(f(u_k)|v_k) \tag{9}$$

Our next step is to propose a model of the processes which give rise to erroneous matches. Since our matching process is structurally based, graphs with a relatively uniform composition give rise to a constant probability of error. We therefore assume that matching errors occur with a uniform probability distribution. If the probability of matching error is P_e, then the confusion probabilities appearing under the product of equation (9) may be assigned according to the following

distribution rule

$$P(f(u_k)|v_k) = \begin{cases} 1 - P_e & \text{if } f(u_k) = v_k \\ P_e & \text{if } f(u_k) \neq v_k \end{cases} \quad (10)$$

As a natural consequence of this distribution rule the joint conditional probability is a function of a single physically meaningful variable. This variable is the Hamming distance $H(\Gamma_j, S_i)$ between the assigned matching and the feasible relational mapping S_i. This quantity counts the number of conflicts between the current matching assignment Γ_j residing on the superclique C_j and those assignments demanded be the relational mapping S_i. With these ingredients, the resulting expression for the joint conditional probability acquires an exponential character

$$P(\Gamma_j|S_i) = K_{C_j} \exp[-k_e H(\Gamma_j, S_i)] \quad (11)$$

where $K_{C_j} = (1 - P_e)^{|C_j|}$. The exponential constant appearing in the above expression is related to the matching-error probability, i.e. $k_e = \ln \frac{(1-P_e)}{P_e}$. The expression may be regarded as providing a natural way of softening the hard relational constraints operating in the model graph. Having developed an exponential expression for the joint conditional matching probabilities, it only remains to specify the distribution of the prior probabilities for consistent relations in the dictionary. Here we adopt a uniform distribution of the available unit probability mass over the set of possibilities Θ, i.e. $P(S_i \in \Theta) = \frac{1}{|\Theta|}$. The final expression for the superclique matching probability is therefore

$$P(\Gamma_j) = \frac{K_{C_j}}{|\Theta|} \sum_{S_i \in \Theta} \exp[-k_e H(\Gamma_j, S_i)] \quad (12)$$

Before proceeding, it is important to comment on the structure of the above expression. The most striking and critical feature is that the consistency of match is gauged by a series of exponentials that are compounded over the dictionary of consistently mapped relations. It is this feature that distinguishes it from alternatives reported in the literature [5, 16, 18, 21]. Each relational mapping contributes a single exponential to the probability of match. In consequence, our method is able to operate in a robust manner when the space of relational mappings is large. Recent theoretical studies suggest that compound exponentials of the type defined above offer tangible benefits over linear or quadratic measures in terms of the number of relational mappings accommodated and the label-error probability of the resulting match [31]. Moreover, the importance of the different constraints is graded by a natural symbolic measure of relational affinity, namely Hamming distance; relational mappings of large Hamming distance contribute insignificantly while those of small Hamming distance dominate.

By gradually reducing P_e, the exponentials appearing in equation (12) approach their delta-function limits. This effectively corresponds to subjecting the softened relational constraints operating in the matching problem to a graded hardening. In the limit of vanishingly small error-probability the matching probabilities become binary in nature; their role is to effectively count the number of consistently matched relational units. Under these conditions our matching criterion becomes similar in function to the relational distance measure of Shapiro and Haralick [27]. It is also worth noting that had we adopted a confusion probability model based on multivariate Gaussian attribute relations rather than the purely symbolic specification of equation (10), then our criterion would be equivalent to that of Boyer and Kak [5] under conditions of small attribute deviations between the two graphs.

With the joint matching probabilities to hand, we can model the joint-prior for the global match by averaging over the complete set of mappings between the two graphs. In doing this we remain very close in spirit to the average consistency measure widely employed in relaxation schemes [11, 17, 18]. We therefore gauge the consistency for the global match $f : V_1 \rightarrow V_2$ by the quantity

$$P(f) = \frac{1}{|V_1|} \sum_{C_j \subset V_1} P(\Gamma_j) \quad (13)$$

This model of the joint-prior forms our basic measure of global relational consistency for the match between the graphs G_1 and G_2. As we indicated in the introduction to this paper our aim is to develop a dynamic matching process for reconfiguring and restoring corrupted graphs. This aim is realised by locating reconstructed graphs that are optimal in the MAP sense using constraints provided by a model graph. The role of the global consistency measure is to model the joint prior in such a scheme. Node reassignments between the graph and the outlier partition may then be sought to optimise the MAP ratio given in equation (7). Reconstructive matching places dual demands on the update process. In the first instance, nodes must be deleted or reinstated so as to bring the overall data-graph topology in line with that of the model. The second requirement is that the match residing on the surviving nodes of the data graph should be maintained in a state of optimal relational consistency.

In meeting these dual objectives it is important to stress that the model of the joint prior appearing in equation (13) can be effectively regarded as a relational clustering metric. Its role is to assess the impact of node insertions or deletions on the overall consistency between the data and model graphs at two different conceptual levels. In the next Section we will demonstrate that the metric not only fullfills the role of optimally partitioning the outlier nodes, but that it simultaneously imposes relational constraints on the nodes of the surviving graph. Moreover, we will demonstrate that the incremental change in the joint prior due to single node deletion is proportional to the Kullback-Leibler entropy.

5 Controlling the Matching Process

We now have all the formal ingredients necessary to realise the iterative graph reconstruction process. Our aim is to delete and reinstate nodes so as to monotonically increase the MAP ratio. The resulting partition between graph nodes and relational outliers is one that optimises our structural consistency measure. We also wish to maintain overall relational consistency among the remaining nodes of the reconfigured graph which have not been consigned to the set of outliers. Our requirements of the MAP criterion are twofold. In the first instance it must allow us to make decisions as to the admissibility of nodes as bona fide graph structure or relational clutter. Its second role is to allow us to reassign matches to the reconfigured graph so as to maintain consistency with the model graph.

5.1 Rejecting relational outliers

Since our model of structural consistency is averaged over the supercliques in the data-graph, in order to gauge the net effect of deleting a node we must examine those contributions that arise from modification of the supercliques containing the node in question. This set is constructed by identifying those nodes that form a superclique with node u in graph G_1, i.e. $C_u - \{u\}$, and determining the new superclique set for these nodes in the reconfigured graph G_1'. We let χ_u^+ denote the superclique set of object u in graph G_1 and χ_u^- denote the corresponding superclique set in the reconfigured graph G_1'. With this notation the change in the denominator of the MAP criterion caused by the deletion of the node u is proportional to

$$\Delta_u^- = P_\phi \sum_{j \in \chi_u^-} P(\Gamma_j) \tag{14}$$

By contrast, when considering the change in the numerator of the MAP criterion it is the superclique set χ_u^+ to which we turn our attention. The corresponding change to the MAP criterion is proportional to

$$\Delta_u^+ = \frac{P(u, f(u)|\underline{x}_u^1, \underline{x}_{f(u)}^2)}{P(u, f(u))} \sum_{j \in \chi_u^+} P(\Gamma_j) \tag{15}$$

With these two measures to hand, we can both delete and reinstate nodes in such a way as to monotonically increase the MAP ratio. We therefore delete node u provided $\Delta_u^+ < \Delta_u^-$ and reinstate the node if $\Delta_u^+ > \Delta_u^-$.

It should be noted that in assessing the change to the global MAP ratio we have confined our attention to the component of the average consistency criterion that is modified by node deletion. This effectively corresponds to ignoring the effect of the unmodified component which can be regarded as representing a constant pedestal consistency value. In order to justify this iterative reassignment approach, we will now illustrate that the modified consistency component is in fact proportional to the change in Kullback-Leibler entropy caused by node deletion. Consequently, the decision criteria specified in equations (14) and (15) not only locate the global MAP, they also ensure that the reconstructed graph is the structure that maximises the Kullback-Leibler entropy.

To proceed, we note that the canonical way of assessing the impact of the change in a probability distribution evaluated over a set of discrete entities is to compute the Kullback–Leibler divergence [3]. Our strategy for assessing the relevance of nodes in the data graph is therefore to compute the change in Kullback-Leibler entropy associated with reconfiguration of the graph and the consignment of nodes to the set of outliers Φ . In our application, we require a means of gauging the improvements associated with the deletion of nodes, which may be facilitated by comparing the matched relations from the original graph $\Gamma_u \in G_1$ with those in the modified graph $\Gamma_j' \in G_1'$. Adhering to our underlying philosophy of exploiting a relational model, we wish to gauge these differences in the light of the structure preserving relations residing in the dictionary, i.e. $S_i \in \Theta$. The Kullback-Leibler entropy which meets our requirements is

$$I_u(G_1, G_1') = -\sum_{j \in \chi_u^+} \sum_{S_i \in \Theta} P(\Gamma_j|S_i) \ln \frac{P(\Gamma_j|S_i)}{P(\Gamma_j'|S_i)} \tag{16}$$

Substituting for $P(\Gamma_j|S_i)$ from equation (11), and exploiting the exponential nature of the probability distribution

$$I_u(G_1, G_1') = -\sum_{j \in \chi_u^+} \sum_{S_i \in \Theta} \frac{k_e K_{C_j}}{|\Theta|} \times \Big(H(\Gamma_j, S_i) - H(\Gamma_j', S_i)\Big) \exp[-k_e H(\Gamma_j, S_i)] \tag{17}$$

The change in entropy associated with the iterative reassignment of the single node u is therefore proportional to the weighted change in Hamming distance over the set of modified supercliques. This is a result which has a pleasing physical intuition. Our maximum entropy criterion for deleting a node is that it minimises the average Hamming distance between the relations in G_1 and G_2.

Moreover, since the average effect of deleting a single node is to change one unit in a superclique, the improvement on deletion is at most one unit of Hamming distance, i.e. $H(\Gamma_j, S_v) - H(\Gamma_j', S_v) \simeq 1$. In consequence, the change in Kullback-Leibler entropy is proportional to the modified component of the average consistency criterion, i.e.

$$I_u(G_1, G_1') \simeq k_e \sum_{j \in \chi_u^+} P(\Gamma_j) \tag{18}$$

Configurations of optimum average consistency are therefore not only those of minimum average Hamming distance, they are also those that maximise the entropy of the relational model.

Viewed as a relational clustering process, the final partition between inliers, i.e. nodes belong to the set V_1, and outliers, i.e. nodes belonging to Φ, is one of maximum entropy.

5.2 Establishing and maintaining consistency

The goal of graph reconfiguration is to restore the topology of relational structure so that a more consistent match may be recovered and relational outliers rejected. In commencing the matching process, we have assumed that a maximally consistent match has already been established via the mapping f. In keeping with our MAP estimation philosophy, the optimal configuration is the one that maximises the quantity

$$\max_f P(f|\mathcal{A}_1, \mathcal{A}_2) \tag{19}$$

The assignment of nodes to optimise the global MAP may again be realised on a node-by-node basis. The match $f(u)$ is assigned as follows

$$f(u) = \arg\max_{v \in V_2} \frac{P(u, v|\underline{x}_u^1, \underline{x}_v^2)}{P(u, v)} \sum_{i \in \chi_u^+} P(\Gamma_i) \tag{20}$$

Consistency is maintained in the MAP sense provided that we update the matches for all nodes in the supercliques modified by the deletion or insertion processes.

Before proceeding, it is however pertinent to discus how we should exert control over the joint-prior in order to ensure consistency of the final match. This is most easily effected through the single parameter of the global criterion, namely the probability of matching errors. Since we expect an improvement in the accuracy of the labelling, one control strategy is to reduce the matching-error probability to a small terminal value according to some deterministic iteration dependent schedule. In our previously published [11, 31, 32] experimental investigations we have found the most effective strategy is to decay the value of P_e exponentially with iteration number. The error probability P_e may therefore be regarded as a control variable, much in the spirit of the temperature in an annealing schedule [8, 9].

6 Experiments

There are two main aspects to our experimental evaluation of the reconstructive graph matching process. The first of these may be regarded as an application vehicle and is concerned with matching hedge structures segmented from synthetic aperture radar (SAR) images against their cartographic representation in a digital Ordnance Survey map. This application study aims to demonstrate the effectiveness of the method when applied to highly cluttered imagery. To give some illustration of the difficulty of the task, the model consists of 23 linear segments while the data to be matched contains some 93 segments; the disparity between the size of the two data sets is not attributable solely to partial overlap, the scene data is in fact permeated by SAR clutter and unmapped segments.

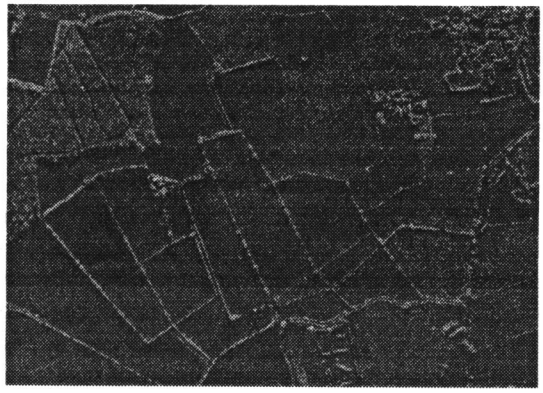

Figure 2: Raw SAR Image

Since the quantity of available SAR data is limited, the application is augmented by a simulation study aimed at evaluating some of the systematic limitations of the graph reconstruction method under conditions of controlled clutter and contamination. This simulation study compares the performance of the dynamic graph reconstruction technique with the use of static graph representations. Here we demonstrate that the dynamic approach consistently improves over the performance of a static method when the fraction of added clutter accounts for as much as 50% of the image entities. Even under these extreme conditions, it is capable of removing the vast majority of clutter while maintaining the overall consistency of the match.

6.1 Extracting line-segments from SAR data

For the application study we are interested in matching hedge structures detected in SAR images of rural scenes against their representation in the form of digital map data. An example image is shown in Figure 2 and the corresponding map data is shown in Figure 6. As we mentioned in the introduction, this is a complex matching task due to the segmentation problems posed by the raw data and the unreliable nature of the ground-truth map information. Hedge structures present themselves as intensity ridges of variable width in the SAR image. The raw intensity data is both noisy and exhibits anisotropies associated with the directionality of the radar used to sense the scene; the anticipated fidelity of feature detection may be expected to be subject to orientation dependent systematics. For instance, inspection of Figure 2 reveals shadowing associated with the directionality of the radar. At the cartographic level there are discrepancies between map and data. Many of the linear hedge structures recorded in the map data are absent from the SAR image while some genuine hedges are not mapped.

We commence the processing of the SAR data by applying a relaxation operator to the raw image data to extract intensity

Figure 3: Result of line detection

Figure 4: Contour grouping

ridges of good connectivity. The ridge extraction process relies on refining the output of a pair of directional line detection filters using a dictionary-based relaxation scheme [10]. There are two features of this process which enhance the performance of ridge detection and which merit special mention. In the first instance, the relaxation operator draws on an accurate noise model for the SAR data. For the data under study the noise process is multiplicative with the underlying intensity distribution. Rather than being modelled by a pure Rayleigh distribution, which is the case with additive Gaussian noise, when multiplicative noise is present the line-gradient distribution is specified by a product of Rayleigh and Bessel function components. The second noteworthy feature of the relaxation operator is its use of a dictionary model of contour structure to enhance the connectivity of the detected ridges. Example ridge contours are shown in Figure 3.

Although clean and noise free the detected line contours still exhibit gaps and are not in good agreement with their map representation. These gaps are due to constrictions and breaks in the SAR intensity profiles caused by genuine non-uniformities in the hedge structures. In order to overcome this limitation in the data, we have applied the contour grouping operator of Shashua and Ullman [28] to the ridge contours; results are shown in Figure 4. The contour grouping operator considerably enhances the coherence of the contours, enabling the extraction of meaningful linear features. Finally, Figure 5 shows the linear segments extracted from the linked contours; it is these entities that we will match against their cartographic counterparts.

Figure 5: Linear segments

6.2 Relational Modelling

The Ordnance survey map data which represents the model to be matched against the extracted linear segments is shown in Figure 6. As we mentioned in Section 2, we establish a graph representation by seeding a Voronoi tessellation of the image plane from the linear segments. Our relational abstraction of the scene is a Delaunay graph that represents spatial

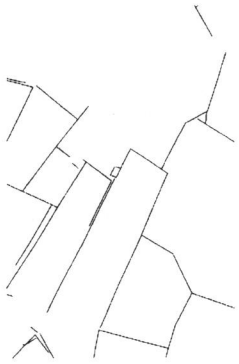

Figure 6: Digital map data

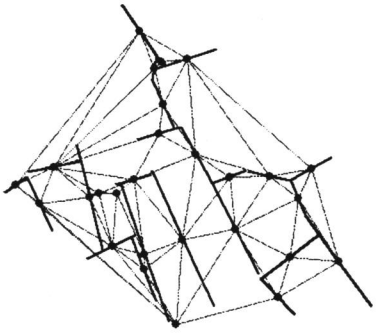

Figure 7: Model graph

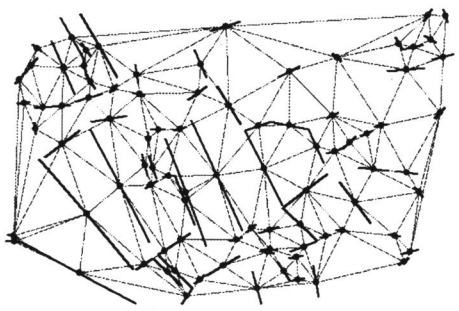

Figure 8: Data graph

adjacency of the Voronoi regions. Figures 7 and 8 show example graph structures generated by the Voronoi tessellations for the map model and SAR data respectively. Superimposed on the graphs are the relevant sets of line-segments. There are several features of these graphs that merit special mention. In the first instance, there is considerable variation in the sizes of the supercliques for the different nodes. The smallest contains only 4 nodes while the largest contains 9 nodes; this means that the lowest order of symbolic relation is 4 while the largest is 9. It is also clear that the data graph suffers from both relational drop-out and relational contamination; there are both spurious arcs and missing arcs. There are significant differences in topology between the two graphs to be matched.

The initial matches between the linear segments extracted from the SAR data and their map representation are established on the basis of the affinity between the vectors of unary node attributes in the two graphs. The matching probabilities are computed from exponential distributions of the Mahalanobis distance between attribute-vector pairs computed using an estimate of the variance-covariance matrix Σ, i.e.

$$P(u, f(u)|\underline{x}_u^1, \underline{x}_{f(u)}^2) = (1 - P_\phi) \times \frac{\exp[-\frac{1}{2}(\underline{x}_u^1 - \underline{x}_{f(u)}^2)^T \Sigma^{-1} (\underline{x}_u^1 - \underline{x}_{f(u)}^2)]}{\sum_{v \in V_2} \exp[-\frac{1}{2}(\underline{x}_u^1 - \underline{x}_v^2)^T \Sigma^{-1} (\underline{x}_u^1 - \underline{x}_v^2)]} \quad (21)$$

For the SAR application, we have used a single unary attribute, namely the orientation of the lines. Although there is clearly a considerable body of information remaining untapped by our matching affinity model, it must be stressed that this aspect of the matching process is not our prime concern in this paper. Rather, our aim is to demonstrate the benefits of the graph reconstruction method. In any case, since linear segments are invariably subject to fragmentation, the ability to recover an acceptable match based purely upon orientation information may provide certain operational advantages in terms of ease of control.

6.3 Matching experiments

The experimental matching study is based on 95 linear segments in the SAR data and 30 segments contained in the map. However only 23 of the SAR segments have feasible matches within the map representation. Figure 8 shows the initial Delaunay graph for the SAR data while Figure 7 shows the corresponding graph for the linear segments in the digital map.

Figure 9 shows the initial configuration of matches which has been used to seed the reconstructive matching algorithm; the black lines are correct matches while the grey lines are matching errors. It is important to stress that these initial matches have been optimised with respect to MAP process described in Section 5.2. In fact, they represent the best results obtainable by static means.

After application of the reconstructive matching technique, the results shown in Figure 10 are obtained. On comparison with Figure 9, it is clear that the main effect has been to consign the plethora of initially incorrect clutter fragments to the set of outliers. Some of these clutter elements form significant groupings of unmapped line structure that dominate whole sectors of the SAR scene. Moreover, some of the restorative benefits are clear if it is noted that the matching of certain relational groupings has improved as a result of clutter deletion; it is these clutter elements that significantly hindered the initial static match by disrupting key structure.

As a result of the reconstructive matching process, the majority of the clutter segments have been deleted from the SAR data graph. To give some idea of relative performance merit, in the case of the initial matching configuration 20 of the 23 matchable segments are correctly identified with 75 incorrect matches, while after application of the active matching method the final graph contains 19 correct matches and only 17 residual clutter; 59 nodes are consigned to the outlier set Φ.

6.4 Performance Evaluation

Although our SAR matching application has furnished a demanding test when noise and clutter are important factors, because of the limited quantities of available data it does not provide an ideal vehicle for demonstrating the performance characteristics of the graph reconstruction method. For this

Figure 9: Initial Matching Configuration

Figure 10: Reconstructive Match

reason, we conclude our experimental study by providing some results on synthetic data. Our aim here is to evaluate the ability of our method to operate under conditions of controlled clutter and to offer quantitative comparisons with conventional matching processes based on static graph representations.

In order to embark on this study, we have generated graphs consisting of a number of lines with random positions and orientations. The patterns have been given a relational abstraction by seeding a Voronoi tessellation from the line centre-points and computing the associated Delaunay graph. The nodes of the graph are therefore the random dots, while the arcs indicate that the associated Voronoi regions are adjacent to one-another. In order to simulate the effects of clutter and segmental dropout, we have both added lines at random locations and deleted random lines from the patterns used to seed the Voronoi tessellation. This has the effect of corrupting the topology of the associated Delaunay graph.

In order to demonstrate the disruptive effects of clutter upon the relational structures used in matching, Figure 11 shows the average error-probability as a function of the fractional degree of clutter. The most striking feature of this plot is that when measured in terms of fractional change in Hamming distance, the effect of clutter is amplified. For instance when the fraction of clutter nodes is 10%, approximately 20% of the nodes in any relation are corrupt. When a quarter of the nodes are clutter, then about one-half the edges in each relational unit are spurious.

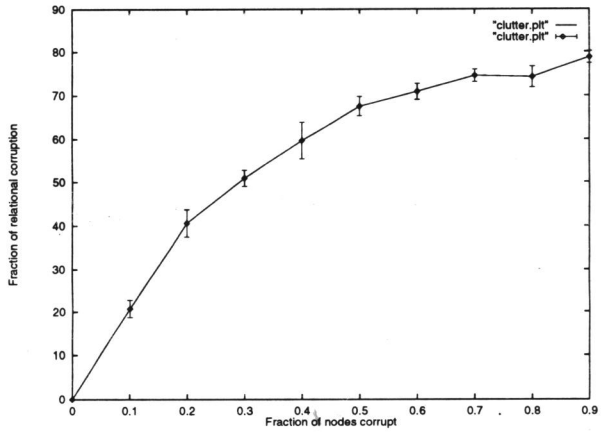

Figure 11: Effect of adding relational clutter

6.4.1 Iterative Reconstruction

We illustrate some of the dynamic qualities of our active matching process. The sequence of images in Figure 12 represents the iterative evolution and eventual reconstruction of the corrupted graph under the action of our active matching scheme. As nodes are assigned to the outlier partition, they become disjoint from the remaining graph. The level of both dropout and clutter in this scene is 50%. As the sequence evolves to stable convergence, the fraction of noise contamination falls from 60% to 20%. Additionally, there has been a 15% improvement in the consistency of match for the nodes retained in the reconfigured graph. In other words, there has been a significant increase in the consistency of match that is attributable to relational restoration. Moreover, very few consistently matched nodes are consigned in error to the outlier partition. After the process has reached stable convergence, 90% of the recoverable graph is restored and only 15% of the surviving nodes represent residual clutter. None of the genuine nodes in the reconstructed graph is matched in error.

6.4.2 Operational Limits

Finally, we have performed a series of simulation experiments with the aim of establishing the effective operating limits of our active matching technique and comparing these with the performance with some alternative methods. The two alternative strategies both employ static relational models. However, in order to provide meaningful comparison, all three algorithms draw on the consistency measure defined in equation (13). Constraint filtering applies the maximal clique algorithm [14] to remove residual inconsistencies from an association graph [4]. The association graph required as input to this algorithm is constructed from the set of candidate matches that optimise the relational consistency measure in equation (13). This is a hybrid algorithm which employs the classical maximal clique algorithm as a postprocessing step once optimisation ceases to find further iterative improvements [32]. Null-labelling, on

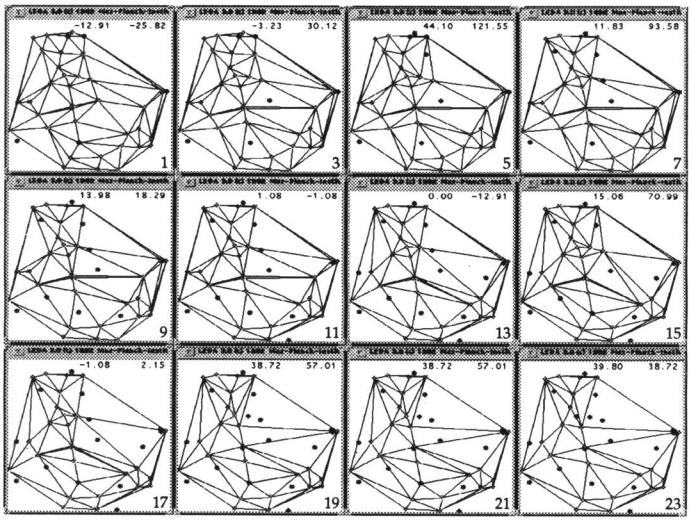

Figure 12: Reconfiguration sequence

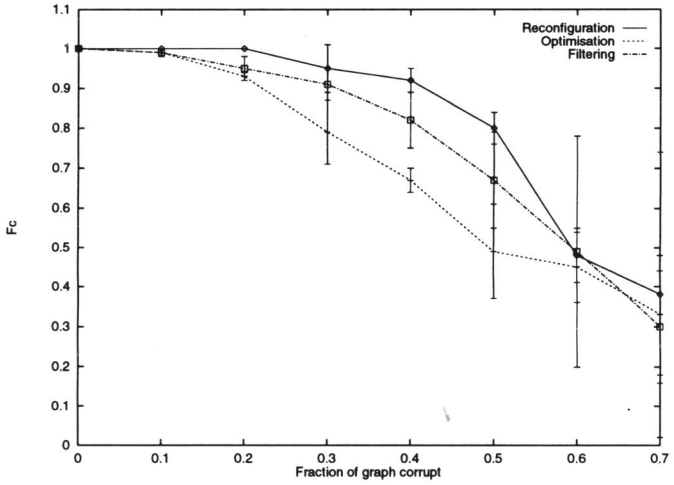

Figure 13: Effect of controlled levels of clutter

the other hand, involves identifying and tagging clutter nodes during the optimisation process. From the modelling perspective, this means that the probability distribution specified in equation (10) must be modified to include a probability mass P_ϕ associated with the null-category. This has the effect of modifying the configurational probability in equation (12) to take account of the number of null matched nodes $\Psi(\Gamma_j)$ on the superclique C_j in the following way [31]

$$P(\Gamma_j) = \frac{K'_{C_j}}{|\Theta|} \sum_{S_i \in \Theta} \exp[-(k_e H(\Gamma_j, S_i) + k_\phi \Psi(\Gamma_j)] \quad (22)$$

where $k_\phi = \ln \frac{(1-P_e)(1-P_\phi)}{P_\phi}$ and $K'_{C_j} = [(1-P_e)(1-P_\phi)]^{|C_j|}$.

Figure 13 shows the fraction of the original graph correctly matched as a function of the fraction of added clutter. The solid line represents the result of active matching. The dot-dashed line is the result of applying maximal clique finding to the association graph of the optimal matches. Finally, the dotted line is the result of null-labelling. It is clear that the active matching method consistently outperforms the two static methods. Errors do not appear until the fractional corruption exceeds 20%. Even when the clutter fraction is as high as 50%, then 80% matching accuracy is achievable. By contrast, the two static methods, are 10%-20% more susceptible to error. However, it is interesting to note that the maximal clique finding technique has a tangible performance edge over null-labelling. This is largely attributable to the difficulties encountered in controlling the parameter of the null match.

Although these simulation studies are far from exhaustive, they indicate that the matching method is robust to significant levels of graph corruption.

7 Conclusions

We have described a novel approach to relational matching. Rather than matching static representations of a scene, it dynamically reconfigures the graph so as to restore overall relational consistency. The scheme is formulated as *maximum a posteriori probability* restoration of the reconstructed scene graph. Commencing from a highly cluttered or over segmented representation of a scene, the data graph is incrementally reconfigured to optimise both its relational and topological similarity with a model graph. This process allows nodes of the graph to be iteratively deleted and reinserted subject to recomputation of the edge-set in the graph.

Our motivation in developing this new approach to matching has stemmed from a comparison of the different techniques available for controlling the matching of extraneous elements. This is conventionally approached by either explicitly labelling elements as null-matched or by filtering them from the final match using arc consistency tests. Under conditions of severe relational corruption both strategies fail. In the former case null matches may swamp the matching process, diluting the effect of relational constraints. In the latter method, the surviving relational structure may be so fragmented to be of little use as a meaningful image interpretation tool. By contrast, our proposed method, not only identifies extraneous nodes but by deleting them from the graph it also incrementally restores the consistency of the underlying relational structure by removing spurious edges. The tangible benefit of this new philosophy, is a graph matching technique that can operate effectively under very severe levels of clutter. This offers manifold potential benefits in the matching of very noisy or over-segmented image data against a parsimonious model.

References

[1] N Ahuja, "Dot Pattern Processing using Voronoi Neighbourhoods", *IEEE PAMI*, **4**, pp 336–343, 1982.

[2] N. Ahuja, B. An and B. Schachter, "Image Representation using Voronoi Tessellation", *CVGIP*, **29**, pp 286–295, 1985.

[3] D. Ackley, G. Hinton and T. Sejnowski, "A Learning Algorithm for Boltzmann Machines", *Cognitive Science*, **9**, pp 147–165, 1985.

[4] H.G. Barrow and R.M Burstall, "Subgraph Isomorphism, Matching Relational Structures and Maximal Cliques", *Information Processing Letters*, **4**, pp.83–84, 1976.

[5] K Boyer and A Kak, "Structural Stereopsis for 3D Vision", *IEEE PAMI*, 10, pp 144-166, 1988.

[6] S. Dickinson, A. Pentland and A Rosenfeld, "3D Shape Recovery using Distributed Aspect Matching", *IEEE PAMI*, **14**, pp 174–198, 1992.

[7] P.J Flynn and A.K.Jain, "CAD-Based Vision - from CAD Models to Relational Graphs", *IEEE PAMI*, **13**, pp 114-132, 1991.

[8] D Geman and S Geman, "Stochastic Relaxation, Gibbs Distributions and Bayesian Restoration of Images", *IEEE PAMI*, **6**, pp 721-741, 1984.

[9] D. Geiger and F. Girosi, "Parallel and Deterministic Algorithms from MRF's: Surface Reconstruction", *IEEE PAMI*, **13**, pp 401–412, 1991.

[10] E.R. Hancock, "Resolving Edge-Line Ambiguities by Relaxation Labelling", *Proceedings of IEEE CVPR Conference*, pp. 300–306, 1993.

[11] E.R. Hancock and J. Kittler, "Discrete Relaxation," *Pattern Recognition*, **23**, pp.711–733, 1990.

[12] E.R. Hancock and R.C Wilson, "A Bayesian Framework for Hierarchical Relaxation", *Proceedings 12th International Conference on Pattern Recognition*, pp. 7–12, 1994.

[13] T.C. Henderson, "Discrete Relaxation Techniques", *Oxford University Press*, 1990.

[14] R. Horaud, F.Veilon and T.Skordas, "Finding Geometric and Relational Structures in an Image", *Proceedings of the First European Conference on Computer Vision*, pp 374-384, 1990.

[15] R. Horaud and T. Skordas, "Stereo Correspondence through Feature Grouping and Maximal Cliques", *IEEE PAMI*, **11**, pp. 1168–1180, 1989.

[16] L. Herault, R. Horaud, F. Veillon and J-J. Niez, "Symbolic Image Matching by Simulated Annealing", *Proceedings of First British Machine Vision Conference*, pp. 319–324, 1990.

[17] R A Hummel and S W Zucker, "On the foundations of relaxation processes", *IEEE PAMI*, **5**, pp. 267-287, 1983.

[18] J. Kittler, W.J. Christmas and M.Petrou, "Probabilistic Relaxation for Matching Problems in Machine Vision", *Proceedings of the Fourth International Conference on Computer Vision*, pp. 666-674, 1993.

[19] P. Meer, D.Mintz and A Rosenfeld, "Robust Regression Methods for Computer Vision: A Review", *International Journal of Computer Vision*, **6**, 59–70, 1991.

[20] B.T. Messmer and H Bunke, " Efficient error-tolerant subgraph isomorphism detection", *Shape, Structure and Pattern Recognition, edited by D Dori and A Bruckstein*, 1994.

[21] S.Z. Li, "Matching Invariant to Translations, Rotations and Scale Changes", *Pattern Recognition*,**25**, pp. 583-594, 1992.

[22] A Pathak and O I Camps, "Bayesian View Class Determination", *Proceedings of IEEE CVPR Conference*, pp. 407–412, 1993.

[23] K. Rose, E.Gurewitz and G.C. Fox, "Constrained Clustering as an Optimisation Method", *IEEE PAMI*, **15**, pp 785–794, 1993.

[24] P.J. Rousseeuw, "Robust Regression and Outlier Detection", *Wiley, New York*, 1987.

[25] S. Sarker and K.L. Boyer, "Perceptual Organisation in Computer Vision: A Review and Proposal for a Classificatory Structure", *IEEE SMC*, **23**, pp 382–399, 1993.

[26] L Shapiro and R.M.Haralick, "Structural Description and Inexact Matching", *IEEE PAMI*, **3**, pp 504–519, 1981.

[27] L Shapiro and R.M.Haralick, "A Metric for Comparing Relational Descriptions", *IEEE PAMI*, **7**, pp 90-94, 1985.

[28] A Shashua and S Ullman, "Structural Saliency: the detection of globally salient structures using a locally connected network", *Proceedings of 2nd International Conference on Computer Vision*, pp. 321–327, 1988.

[29] Y.C. Tang and C.S.G. Lee, "A Geometric Feature Relation Graph Formalism for Consistent Sensor Fusion", *IEEE SMC*, **22**, pp 115–129, 1992.

[30] D Terzopoulos and M. Vasilescu, "Sampling and Reconstruction with Adaptive Meshes", *IEEE CVPR Conference* , pp 70–75, 1991.

[31] R.C. Wilson and E.R Hancock, "Graph Matching by Discrete Relaxation", *Pattern Recognition in Practice IV: Multiple Paradigms, Comparative Studies and Hybrid Systems*, **North Holland** pp. 165–177, 1994.

[32] R.C. Wilson and E.R Hancock, "Relational Matching by Discrete Relaxation", *Image and Vision Computing*, **13**, pp. 411–422, 1994.

[33] D. Yang and J. Kittler, "MFT-Based Discrete Relaxation for Matching High-Order Relational Structures", *Proceedings 12th International Conference on Pattern Recognition*, pp. 219–223, 1994.