

# Graph-Based Methods for Vision: A Yorkist Manifesto

Edwin Hancock and Richard C. Wilson

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

**Abstract.** This paper provides an overview of our joint work on graph-matching. We commence by reviewing the literature which has motivated this work. We then proceed to review our contributions under the headings of 1) the probabilistic framework, 2) search and optimisation, 3) matrix methods, 4) segmentation and grouping, 5) learning and 6) applications.

## 1 Introduction

For the past decade, we have been involved in a programme of activity aimed at developing a probabilistic framework for graph matching. This paper provides an overview of the main achievements of this programme of work. We commence, with a review of the literature to set our own work in context. The achievements are then summarised under the headings of 1) the probabilistic framework, 2) search and optimisation, 3) matrix methods, 4) segmentation and grouping, 5) learning and 6) applications. We conclude the paper by outlining our future plans.

The work summarised here has been undertaken with a number of present and former colleagues who have all made important contributions. These include Andrew Finch, Andrew Cross, Benoit Huet, Richard Myers, Mark Williams, Simon Moss, Bin Luo, Andrea Torsello, Marco Carcassoni and Antonio Robles-Kelly.

## 2 Literature Review

We set our work in context with a brief review of the related literature. Some of the pioneering work on graph matching was undertaken in the early 1970's by Barrow and Burstall [46] and by Fischler and Ensclager [55]. These two studies provided proof of concept for the use of relational structures in high-level pictorial object recognition. Over the intervening three decades, there has been a sustained research activity. Broadly speaking the work reported in the literature can be divided into three areas. The first of these is concerned with defining a measure of relational similarity. Much of the early work here was undertaken in the structural pattern recognition literature. For instance, Shapiro and Haralick [64] showed how inexact structural representations could be compared by

counting consistent subgraphs. This similarity measure was refined by Eshera and Fu [53] and by Sanfeliu and Fu [61] who showed how the concept of string edit distance could be extended to graphical structures. The formal basis of graph edit distance has recently been extended by Bunke and his coworkers [48,49] who have shown, among other things, that the edit distance is related to the size of the maximum common subgraph. More recently Tirthapura, Sharvit, Klein and Kimia have shown how the classical Levenshtein distance can be used to match shock-graphs representing 2D skeletal shapes [69].

Much of the work described above adopts a heuristic or goal directed approach to measuring graph similarity. The second issue addressed in our literature survey is that of how to develop more principled statistical measures of similarity. This endeavour involves the modelling of the processes of structural error present in the graph-matching problem. Wong and You [72] made one of the first contributions here by defining an entropy measure for structural graph-matching. Boyer and Kak [47] also adopted an information theoretic approach, but worked instead with attribute relations. Using a probabilistic relaxation framework Christmas, Kittler and Petrou [50] have developed a statistical model for pairwise attribute relations.

The third issue is that of optimisation. Here there have been several attempts to use both continuous and discrete optimisation methods to locate optimal graph matches. Turning our attention first to discrete optimisation methods, there have been several attempts to apply techniques such as simulated annealing [56], genetic search [52] and tabu search [71] to the graph matching problem. However, continuous optimisation methods provide attractive alternatives since their fixed points and convergence properties are usually better understood than their discrete counterparts. However, the main difficulty associated with mapping a discretely defined search problem onto a continuous optimisation method is that of embedding. There are several ways in which this embedding can be effected for the problem of graph matching. The most straightforward of these is to pose the graph-matching problem as that of recovering a permutation matrix which preserves edge or adjacency structure. For instance, Kosowsky and Yuille have cast the problem into a statistical physics setting and have recovered a continuous representation of the permutation matrix using mean-field update equations [73]. Gold and Rangarajan [54] have exploited the stochastic properties of Sinkhorn matrices to recover the matches using a soft-assign update algorithm. Umeyama [70] takes a more conventional least-squares approach and shows how an eigendecomposition method can be used to recover the permutation matrix. An alternative representation has recently been developed by Pelillo [60] which involves embedding the association graph. Matches are located by using the replicator equations of evolutionary game-theory to locate the maximal clique of the association graph, i.e. the maximum common subgraph, of the two graphs being matched. This method has subsequently also been applied to shock-graph matching [74].

Closely related to this work on recovering permutation structure by continuous embedding is the literature on spectral graph theory. This is a term applied

to a family of techniques that aim to characterise the global structural properties of graphs using the eigenvalues and eigenvectors of the adjacency matrix [51]. In the computer vision literature there have been a number of attempts to use spectral properties for graph-matching, object recognition and image segmentation. Umeyama has an eigendecomposition method that matches graphs of the same size [70]. Borrowing ideas from structural chemistry, Scott and Longuet-Higgins were among the first to use spectral methods for correspondence analysis [62]. They showed how to recover correspondences via singular value decomposition on the point association matrix between different images. In keeping more closely with the spirit of spectral graph theory, yet seemingly unaware of the related literature, Shapiro and Brady [65] developed an extension of the Scott and Longuet-Higgins method, in which point sets are matched by comparing the eigenvectors of the point proximity matrix. Here the proximity matrix is constructed by computing the Gaussian weighted distance between points. The eigen-vectors of the proximity matrices can be viewed as the basis vectors of an orthogonal transformation on the original point identities. In other words, the components of the eigenvectors represent mixing angles for the transformed points. Matching between different point-sets is effected by comparing the pattern of eigenvectors in different images. Shapiro and Brady's method can be viewed as operating in the attribute domain rather than the structural domain. Horaud and Sossa [57] have adopted a purely structural approach to the recognition of line-drawings. Their representation is based on the immanental polynomials for the Laplacian matrix of the line-connectivity graph. By comparing the coefficients of the polynomials, they are able to index into a large data-base of line-drawings. In another application involving indexing into large data-bases, Sengupta and Boyer [63] have used property matrix spectra to characterise line-patterns. Various attribute representations are suggested and compared. Shokoufandeh, Dickinson and Siddiqi [67] have shown how graphs can be encoded using local topological spectra for shape recognition from large data-bases.

### 3 Probability Distributions

We embarked on our programme of work motivated by the need for a probabilistic framework so that graph matching could be approached in a statistically robust manner. At the time (1992), there were three pieces of work that addressed this problem. Wong and You [72], in early work in the structural pattern recognition area had made one of the first attempts to cast the graph matching problem into an information theoretic setting by defining the relative entropy of two graphs. The second contribution came from Boyer and Kak [47], who cast the problem of structural stereopsis into a mutual information setting, and by choosing Gaussian distributions to measure the distribution of attributes, arrived at a least squares framework for relational matching. In a more ambitious piece of work, Christmas, Kittler and Petrou [75] showed how pairwise attribute relations could be modelled probabilistically and used to match graphs using an iterative probabilistic relaxation scheme. Taking a critical view of this existing

work, we identified two areas where there appeared to be scope for improvement and further work. First, the existing work in the area relied only on pairwise relations. In other words, there had been little attempt to model constraints on the graph matching problem beyond the edge-level. Second, there appeared to be a move away from the use of purely structural constraints in favour of attribute relations. Hence, we embarked on a programme of work aimed at modelling the distribution of structural errors and using the distribution to develop robust graph matching algorithms.

### 3.1 Edge and Face Compatibilities

Our initial study followed Christmas, Kittler and Petrou and pursued the graph-matching problem in the setting of probabilistic relaxation. Using the product-form support function developed by Kittler and Hancock [76], we set about the graph matching problem by attempting to model the compatibility co-efficients for graph matching [5,9]. In the probabilistic framework, these take the form of conditional probabilities. When the distribution of structural errors follows a uniform distribution, then we discovered that the compatibilities followed a particularly simple and intuitively appealing distribution rule. In the case of edges and triangular faces, the compatibilities were simply proportional to the edge and face densities in the graphs being studied. Moreover, the entire compatibility model was free of parameters.

### 3.2 Dictionaries

Encouraged by these findings, our next step was more ambitious. Here we aimed to cast the graph-matching problem into a discrete relaxation setting [2]. Rather than modelling structural corruption at the level of edge or face units, we aimed to extend the model to entire edge-connected neighbourhoods of nodes. To do this we adopted a dictionary model of structural errors [4,7]. This allowed for neighbourhoods to match to one-another under a cyclic permutation of nodes with possible dummy node insertions. Our structural model hence commenced from a dictionary of possible mappings between each data-graph neighbourhood and each permuted and padded neighbourhood of the model graph. With this structural model to hand, the next step was to develop probability distributions for the different types of error that can occur in graph matching by discrete relaxation. Discrete relaxation is an iterative process which relies on replacing symbolic label assignments so as to maximise a global consistency criterion [4]. Hence, at any iterative epoch there are two types of error that can occur. The first of these are label placement or assignment errors. The second are structural errors due to node and edge dropout or insertion. We commenced from a simple distribution model in which these two error processes were modelled by uniform and memoryless distributions. Based on these two assumptions, we arrived at a probability distribution for the label configurations contained in the structural dictionary. Again the model had a simple intuitive form. Each dictionary item had a probability that took an exponential form. The exponential was a function

of two variables. The first of these was the Hamming distance between the current configuration of label placements and those demanded by the permutations in the dictionary. The second variable was the size difference or number of padding nodes needed. The distribution was controlled by two parameters. These are the probability of label misplacement and the probability of dummy node insertion. The former parameter was gradually reduced to zero (annealed) as the discrete relaxation process iterated. The latter parameter was estimated from the size difference of the graphs being matched.

### 3.3 Graph Editing

In addition to allowing graph-matching, the probabilistic framework also allows us to rectify the effects of structural corruption. The aim here is to use maximum a posteriori probability estimation to restore noise corrupted graphs. This is done by developing statistical tests for the removal and re-insertion of nodes, together with their associated edges [11]. This method provides significant performance advantages over other methods for dealing with structural error such as null-labelling or locating maximal cliques of the association graph [4,7].

### 3.4 Edit Distance

One of the shortcomings of the method described above is that it requires the enumeration of a dictionary of structures so that it may be applied to the graph-matching problem. The number of dictionary items has a potentially exponential complexity with the degree of the nodes in the graph when padding with dummy insertions is allowed. To overcome this problem, we turned to string edit distance or Levenshtein distance [77] as a means of computing the similarity of labelled neighbourhoods. There had been previous attempts to extend the idea of edit distance to graphs. The idea of doing this had already been explored by Eshehera and Fu [53], and by Sanfeliu and Fu [61] in their early work on structural matching. However, although effective the extension to graphs, lacks some of the formal neatness of the string treatment. Hence, we adopted an approach in which each neighbourhood is represented by a string, and the different permutations are implicitly traversed in the edit matrix. In this way we lift the exponential overhead on dictionary enumeration [21]. The location of the optimal string edit path is computed using Dijkstra's algorithm. Since, the Hamming distance of the original probability distribution is a special case of the Levenshtein distance, we again use a family of exponential distributions to model the probability of match between neighbourhoods.

Recently, we have taken this work one step further. We make use of the recent observation by Bunke [49] that the size of the maximum common subgraph and the edit distance are related to one another. Hence, by locating the maximum common subgraph, we may estimate edit distance. The recent work of Pelillo [60] provides a means of locating the max-clique using a quadratic programming method. We have made use of this result to efficiently compute tree edit distance [37].

## 4 Search and Optimisation

Once a measure of graph similarity is to hand, then a means of locating the matching configuration of maximum similarity is one of configurational optimisation. In our initial experiments with the methods described in the previous section we used simple hill-climbing methods. These are prone to convergence to local optima. Hence, an important theme in our work has been to explore the use of global optimisation methods for graph matching.

### 4.1 Genetic Search

The simplest way to overcome local convergence problems is to adopt a stochastic optimisation framework. There are a number of well known techniques here, including the extensively studied process of simulated annealing. However, genetic search offers a topical method which has furnished us with an effective route to graph matching.

Evolutionary or genetic search is a population-based stochastic optimisation process which draws its inspiration from population genetics. The process relies on three different operations to modify a population of possible solutions to discretely defined problems which have been encoded in a string of symbols which resemble the arrangement of genes on a chromosome. Mutation is a background operator which randomly switches individual symbols. Crossover or recombination, exchanges randomly selected subparts of pairs of strings residing in the population to produce new ones. Finally, selection uses a fitness measure for the solutions candidates which can be retained in the population. In a series of papers [8,12,22,24,26,28] we have investigated how genetic search can be tailored to our probabilistic framework for graph-matching. The reason for doing this is that there are several features of our formulation that have a natural assonance with genetic search. The first of these is that we have a probabilistic measure of graph similarity, which naturally plays the role of a fitness function [8]. Second, the development of this model relies on the specification of the assignment error probability. This probability can be used to control the re-assignment of the labels to nodes via mutation operator. However, despite these appealing conceptual isomorphisms, there are a number of more difficult problems to be solved in mapping the graph-matching problem onto a genetic search procedure. The most obvious of these is that GA's usually rely on a string encoding, and a graph is a relational structure. Although the encoding issue is not one of major importance, the main difficulty arises when considering how to implement crossover since we no longer have the convenience of strings. To overcome this problem we have developed a graph-based crossover mechanism. This involves performing a graph-cut and recombining the two subgraphs so formed. Normally, this would be rather intractable since we would have to produce cuts in different graphs which traversed the same number of edges so that the subgraphs could be recombined. We circumvent this problem by working with the Delauney graphs of point-sets. Here we bisect planar point-sets with a straight line and recombine the point-sets at the cut line. The recombined point-sets can be re-triangulated

to produce new Delaunay graphs. Structural errors may be corrected by deleting and re-inserting points in the recombination step. We have performed both theoretical and empirical analysis on the resulting graph-matching process. First, we have established the convergence properties of the resulting algorithm using a Gaussian model of the distribution of population fitness. Second, we have embarked on an exhaustive investigation of the best choice of mutation, crossover, and selection operators for use in conjunction with the graph-matching problem. Finally, we have exploited the population architecture of the optimisation process to develop a least commitment graph-matching method which allows multiple ambiguous solutions to be maintained.

## 4.2 Softened Representations

One way to escape local optima in discretely defined optimisation problems such as graph matching, consistent labelling and constraint satisfaction is to soften the representation. In other words, the aim is to replace the symbolic labels by real-valued assignment variables or by label probabilities. Probabilistic relaxation is such a scheme, but it only possesses local convergence properties. A method which possesses global convergence properties is mean field annealing. There have been several attempts to cast the graph-matching problem into a mean field setting. These include the early work of Simic [68] and Suganathan [78]. The idea here is to locate an energy function for the assignment or mean field variables through a process of local averaging over the configuration space of the matching problem. The mean field variables are updated by exponentiating the derivatives of the energy function. As the process proceeds, the constant of the exponentials is annealed, so as to increasingly penalise departures from consistency. These methods are also closely akin to the soft-assign method of Gold and Rangarajan [54].

However, these early mean-field methods are often referred to as “naive” since they commence from an ad hoc energy function rather than a principled probability distribution. More recently, the apparatus of variational inference has been deployed to develop mean field equations [58,79]. The idea here is to find the set of mean field variables which have minimum Kullback divergence to a prior distribution for the discrete entities for the problem in hand. We have adopted this approach to develop soft assign equations for graph matching [11,13]. Our approach is midway between the naive one and the full variational treatment. The reason for the compromise is, of course, that of tractability. Rather than using the variational approach to find minimum Kullback divergence update equations for the assignment variables, we use it to find an equivalent free energy for the dictionary-based probability distribution. This energy is couched in terms of assignment variables. We follow Gold and Rangarajan [54] and update the assignment variables by exponentiating the derivatives of the free energy. The result is a non-quadratic energy function which proves more robust than its quadratic counterpart.

### 4.3 Tabu Search

In addition to stochastic and mean-field methods for optimisation methods there are also deterministic methods which can yield good results for the graph-matching problem. Recently, tabu search has been used with great effect for a number of path-based planning and scheduling problems. The method uses an aspiration criterion to intensify and diversify the search procedure. We have used this search procedure with our probabilistic similarity measure to perform a deterministic brushfire search for graph matches [6,17]. We commence from the most confident matches and propagate the search for matches across the edges of the graph to adjacent nodes. This method provides a search procedure which is both efficient and effective.

## 5 Matrix Methods

One of the disappointments of the early matrix-based methods for graph-matching is that while they are extremely elegant, they can not cope with graphs of different size and are not robust to structural error. For this reason we have recently attempted to extend our probabilistic framework for graph-matching to the matrix domain. Our aim here has been to combine probabilistic methods with ideas from spectral graph theory [51].

### 5.1 Singular Value Decomposition

The method of Umeyama attempts to match graphs by comparing the singular vectors of the adjacency matrix. One of the problems which limits this method is that there is no clear way to compare the singular vectors of matrices of different size. To overcome this problem, we have recently presented a study which aims to cast the Umeyama [70] algorithm into the setting of the EM algorithm. Commencing from a simple model for the correspondence process in which assignment errors follow a Bernoulli distribution, we have developed a mixture model for graph-matching [27]. This leads to a utility measure for the graph-matching problem which is obtained from the trace of the weighted correlation of the adjacency matrices for the two graphs being matched. The weight matrix allows for both correspondence errors and difference in the size of the two adjacency matrices.

We find optimal matches in a two-step iterative EM algorithm. In the M or maximisation step, the optimal matches are found by performing singular value decomposition on the weighted adjacency correlation matrix. In the E-step, the weight matrix used to perform adjacency matrix correlation is updated.

### 5.2 Adjacency Matrix Markov Chains

Spectral methods hold out another promising route to graph-matching. It is well known that the leading eigenvector of the transition probability matrix for



a Markov chain is the steady state random walk on the graphical model of the Markov chain. Viewed slightly differently, if we commence with the adjacency matrix for a graph and convert it into a transition probability matrix using row and column normalisation, then the leading eigenvector provides us with a string representation of the graph. By converting graphs into strings in this way, then graph-matching can be effected using string-matching methods. With this goal in mind, we have recently reported a method which aims to match graphs using the Levenshtein distance between steady state random walks [43].

## 6 Segmentation and Grouping

To apply the matching methods listed above, which are entirely generic in nature, we require graph structures to be segmented from image data. The ideas of error modelling and graph editing also provide theoretical tools which can be used to develop segmentation and grouping algorithms.

### 6.1 Regions and Arrangements

The idea of using a Bernoulli distribution to model structural errors also has applications to the pairwise grouping of image entities. To do this we have used the Bernoulli distribution to model the distribution of edge weights or grouping strength between image tokens [35,36,42]. These may be viewed as the elements of a weighted adjacency matrix. We have developed a simple EM algorithm based on this Bernoulli model, which encourages blocks to form in the permuted adjacency matrix. This algorithm uses the eigenvectors of the adjacency matrix to update the connectivity weights, and shares much in common with the normalised cuts method of Shi and Malik [66].

### 6.2 Meshes and Surfaces

The explicit graph edit operations used to correct structural errors in the graph-matching process can also be used to modify relational structures so as to improve their consistency with raw image data. One application in which this arises is the control of surface meshes. Here the aim is to have a fine mesh in the proximity of high curvature surface detail, and a coarse grained mesh in the proximity of relatively unstructured surface regions. We have developed an adaptive surface mesh which uses split and merge operations to control the placement of mesh facets so as to optimise a bias-variance criterion [23]. Here facets may be added or subtracted so that they minimise the sum of variance and squared bias. In this way we achieve good data-closeness while avoiding over fitting with too many mesh facets.

## 7 Learning

Our probabilistic framework has provided us with a means of measuring the similarity of graphs. With good similarity measures to hand, then we can address a number of problems that involve learning from structural representations.

### 7.1 Graph Clustering

The simplest task that arises in machine learning with structural representations is how to cluster similar graphs together. This problem is more complicated than that of using feature or pattern vectors, since it is notoriously difficult to define the mean and covariance of a set of graphs. Hence, central clustering methods can not easily be used to learn classes of graphs. However, there are several methods which allow data to be clustered using pairwise or relational attributes rather than ordinal ones. Two such methods are multidimensional scaling and pairwise clustering. We have recently applied both methods to the clustering of shock-tress using edit distances between pairs of graphs [38,44]. The results prove promising and indicate that shape classes can be learnt from structural representations.

### 7.2 Structural Variations

One of the shortcomings of simply using a pattern of distances, is that it does not capture the variations in detailed structure associated with a particular class of graph-structure. In a recent study we have turned to spectral graph theory to furnish us with a representation suitable for this purpose. One of the problems with computing the mean and variance of a set of graphs is that they may potentially contain different numbers of node and edges, and even if the structures are the same then correspondences between nodes are required. To overcome this problem we have used the leading eigenvectors of the adjacency matrices of sets of graphs to define fuzzy clusters of nodes. From the fuzzy clusters, we compute features including volume, perimeters and Cheeger constants. By ordering the clusters according to the magnitude of the associated eigenvalue, we construct feature-vectors of fixed length. From these vectors it is possible to construct eigenspaces which trace out the main modes of structural variation in the graphs [45].

## 8 Applications

The research described in the previous sections is of a predominantly methodological and theoretical nature. However, we have applied the resulting techniques to a large variety of real world problems. In this section we summarise these.

## 8.1 Model Alignment

Graph-matching provides correspondences between image entities. The entities may be points, lines, regions or surfaces depending on the abstraction adopted. However, once correspondence information is to hand then it may be used for the purposes of detailed model alignment. We have explored a number of different approaches to using the correspondence information delivered by our graph-matching methods for model alignment.

**Rigid Models** The first class of alignment process are those which involve rigid transformations. The rigid geometries most frequently used in computer vision are the similarity, affine and perspective transformations. To exploit our probabilistic framework for measuring the consistency of graph correspondences, we have developed a novel dual-step EM algorithm [14,18]. Here we work with a mixture model over the space of correspondences. The aim is estimate the parameters of any the three rigid transformations that bring noisy point sets of potentially different size into correspondence with one another. If this is done using the EM algorithm then the aim is to find the set of parameters that optimise an expected log likelihood function. If the alignment errors between the different point sets follow a Gaussian distribution then the expected log-likelihood is a weighted squared error function, where the weights are the a posteriori alignment probabilities. Starting from first principles, we have shown how the correspondence probabilities delivered by our dictionary-based structural error model can be used to impose additional constraints on the alignment process. This takes the form of an additional weighting of the contributions to the expected log-likelihood function. We develop a dual step EM algorithm which interleaves the tasks of finding point correspondences and estimating alignment parameters. The method has been applied to recovering plane perspective pose. We have also shown how the method relates to the Procrustes alignment process for point-sets subject to similarity transformations [31].

**Non-rigid Models** For non-rigid alignment we have adopted a slightly different approach. Here we have posed the alignment and correspondence problems in a setting which involves maximising the cross entropy between the probability distributions for alignment and correspondence errors [59]. We have applied the resulting matching process to the problem of aligning unlabelled point distribution models to medical image sequences. In this way we are able to impose relational constraints on the arrangement of landmark points, and improve the robustness of the matching process to both measurement and structural errors [30,41].

## 8.2 Object Recognition

Another important application of the graph matching methods is to object recognition, and in particular that of recognising an object from a large number of image instances. This problem is one of practical significance since it arises in the retrieval of images from large data-bases.

**Line-Patterns** The main application vehicle for the object recognition has been the retrieval of trade marks abstracted in terms of line pattern arrangements. Here we have used a coarse to fine search procedure. We commence with a representation of the line patterns based on a so-called relational histogram. This is a histogram of affine invariant relative angle and length ratio attributes. However, only those pairs of lines that are connected by a line-centre adjacency graph can contribute to the histogram [20]. The next step is to refine the candidate matches using a set-based representation of the edge attributes. Here we use a robust variant of the Hausdorff distance to find the closest matching sets of attributes [29]. Finally, the pruned set of matches is searched for the best match using a full graph-based matching scheme [19].

**Surface Topography** The framework described above can be extended from line patterns to surface patches. One of our complementary interests is in the process of shape-from-shading [15]. This is concerned with recovering surface patches from local shading variations. We have shown how the robust Hausdorff distance can be used to recognise 3D objects from 2D views using the arrangement surface patches together with their shape-index attributes [25].

**View-Based Object Recognition** The final application studied has been view based object recognition. Here the aim is to a large number of views of a 3D object as the viewpoint is slowly varied.

## 9 The Future

One of the main challenges facing the structural pattern recognition community is that of developing methods for shape modelling that are comparable in performance and versatility to those developed by the computer vision community. Here eigenspace methods have proved powerful for both modelling variations in boundary shape and appearance attributes. Although the vision community have recently been pre-occupied with the use of geometric methods, the recent use of shock-graphs have increased their confidence and curiosity in the use of structural methods. However, while recent work has demonstrated the utility of graph-based methods for matching and recognition, it is fair to say that the available methodology for learning structural descriptions is limited. In many ways the graphical models community has much from which we can learn. They have developed principled information theoretic methods for learning the structure of both Bayes nets and other graphical structures such as decision trees. It must be stressed that these are inference structures rather than relational models. Hence, they are rather different from the relational structures encountered in image analysis and computer vision which are predominantly concerned with the representation of spatially organised data. What we do have to offer is a methodology for measuring the similarity of relational structures at a relatively fine-grained level. Moreover, if spectral techniques can be further tamed then there is scope for considerable convergence of methodology.

## References

1. J. Kittler and E. R. Hancock, "Combining Evidence in Probabilistic Relaxation," *International Journal of Pattern Recognition and Artificial Intelligence*, **3**, pp.29–52, 1989.
2. E. R. Hancock and J. Kittler, "Discrete Relaxation," *Pattern Recognition*, **23**, pp.711–733, 1990. [34](#)
3. E. R. Hancock and J. Kittler, "Edge-labelling using Dictionary-based Probabilistic Relaxation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **12**, pp 165–181 1990.
4. R. C. Wilson, A. N. Evans and E. R. Hancock, "Relational Matching by Discrete Relaxation", *Image and Vision Computing*, **13** , pp. 411–422, 1995. [34](#), [35](#)
5. R. C. Wilson and E. R. Hancock, "A Bayesian Compatibility Model for Graph Matching", *Pattern Recognition Letters*, **17**, pp. 263–276, 1996. [34](#)
6. M. L. Williams, R. C. Wilson, and E. R. Hancock, "Multiple Graph Matching with Bayesian Inference", *Pattern Recognition Letters*, **18**, pp. 1275–1281, 1997. [38](#)
7. R. C. Wilson and E. R. Hancock, "Structural Matching by Discrete Relaxation", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **19**, No.6, pp.634-648, 1997. [34](#), [35](#)
8. A. D. J. Cross, R. C. Wilson and E. R. Hancock, "Genetic Search for Structural Matching", *Pattern Recognition*, **30**, pp.953-970, 1997. [36](#)
9. A. M. Finch, R. C. Wilson and E. R. Hancock, "Matching Delaunay Graphs", *Pattern Recognition*, **30**, pp. 123–140, 1997. [34](#)
10. A. M. Finch, R. C. Wilson and E. R. Hancock, "Symbolic Graph Matching with the EM Algorithm", *Pattern Recognition*, **31**, pp. 1777–1790, 1998.
11. R. C. Wilson, A. D. J. Cross and E. R. Hancock, "Structural Matching with Active Triangulations", *Computer Vision and Image Understanding*, **72**, pp. 21–38, 1998. [35](#), [37](#)
12. A. D. J. Cross and E. R. Hancock, "Matching Buildings in Aerial Stereogramms using Genetic Search and Matched Filters", *ISPRS Journal of Photogrammetry and Remote Sensing*, **53**, pp. 95–107, 1998. [36](#)
13. A. M. Finch, R. C. Wilson and E. R. Hancock, "An Energy Function and Continuous Edit Process for Graph Matching", *Neural Computation*, **10**, pp. 1873–1894, 1998. [37](#)
14. A. D. J. Cross and E. R. Hancock, "Graph Matching with a Dual Step EM Algorithm", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **20**, pp. 1236–1253,, 1998. [41](#)
15. P. L Worthington and Edwin R Hancock, "New Constraints on Data-Closeness and Needle Map Consistency for Shape-from-Shaping", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, IEEE Computer Society Press, **vol. 21**, no. 12, pp 1250-1267, 1999. [42](#)
16. R. C. Wilson and E. R. Hancock, "Graph matching with hierarchical discrete relaxation", *Pattern Recognition Letters*, **20**, pp 1041-1052, 1999.
17. M. L. Williams and R. C. Wilson and E. R. Hancock, "Deterministic Search For Relational Graph Matching", *Pattern Recognition*, **32**, pp. 1255-1271, 1999. [38](#)
18. S. Moss and Richard C. Wilson and E. R. Hancock, "A mixture model of pose clustering", *Pattern Recognition Letters*, **20**, pp 1093-1101, 1999. [41](#)
19. B. Huet and E. R. Hancock, "Shape recognition from large image libraries by inexact graph matching", *Pattern Recognition Letters*, **20**, pp. 1259-1269, 1999. [42](#)

20. B. Huet and E. R. Hancock, "Line Pattern Retrieval Using Relational Histograms", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **21**, pp. 1363-1370, 1999. 42
21. R. Myers, R. C. Wilson and E. R. Hancock, "Bayesian Graph Edit Distance", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **22**, pp 628-635, 2000. 35
22. A. D. J. Cross and E. R. Hancock, "Convergence of a Hill Climbing Genetic Algorithm for Graph Matching", *Pattern Recognition*, **33**, pp 1863-1880, 2000. 36
23. R. C. Wilson and E. R. Hancock, "Bias Variance Analysis for Controlling Adaptive Surface Meshes", *Computer Vision and Image Understanding* **77**, pp 25-47, 2000. 39
24. R. Myers, E. R. Hancock, "Genetic Algorithms for Ambiguous Labelling Problems", *Pattern Recognition*, **33**, pp. 685-704, 2000. 36
25. P. L. Worthington and E. R. Hancock, "Object Recognition using Shape-from-shading", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **33**, pp. 535-562, 2001. 42
26. R. Myers and E. R. Hancock, "Least Commitment Graph Matching with Genetic Algorithms", *Pattern Recognition*, **34**, pp 375-394, 2001. 36
27. B. Luo and E. R. Hancock "Structural Matching using the EM algorithm and singular value decomposition", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **23**, pp. 1120-1136, 2001. 38
28. R. Myers and E. R. Hancock, "Empirical Modelling of Genetic Algorithms", *Evolutionary Computation*, **9**, pp. 461-493, 2001. 36
29. B. Huet and E. R. Hancock, "Relational Object Recognition from Large Structural Libraries", *Pattern Recognition*, 2002. 42
30. M. Carcassoni and E. R. Hancock, "Spectral Correspondence for Point Pattern Matching", *Pattern Recognition*, to appear 2002. 41
31. B. Luo and E. R. Hancock, "Iterative Procrustes Analysis using the EM Algorithm, *Image and Vision Computing*, to appear 2002. 41
32. B. Huet and E. R. Hancock, "Object Recognition from Large Structural Libraries", *Advances in Structural and Syntactic Pattern Recognition, Springer, Lecture Notes in Computer Science*, **1451**, pp. 190-199, 1998.
33. S. Moss and Edwin R. Hancock, "Structural Constraints for Pose Clustering", *Computer Analysis of Images and Patterns, Springer Lecture Notes in Computer Science*, **1689**, Franc Solina and Ales Leonardis eds.,pp. 632-640, 1999.
34. A. Torsello and E. R. Hancock, "A Skeletal Measure of 2D Shape Similarity", *Springer Lecture Notes in Computer Science*, **2059**, Edited by C. Arcelli, L. P. Cordella and G. Sannitii di Baja, pp. 260-271, 2001.
35. A. Robles-Kelly and E. R. Hancock, "An Expectation-Maximisation Framework for Perceptual Grouping", *Springer Lecture Notes in Computer Science*, **2059**, Edited by C. Arcelli, L. P. Cordella and G. Sannitii di Baja, pp. 594-605, 2001. 39
36. A. Kelly and E. R. Hancock, "A maximum likelihood framework for grouping and segmentation", *Springer Lecture notes in Computer Science*, **2134**, pp. 251-266, 2001. 39
37. A. Torsello and E. R. Hancock, "Efficiently computing weighted tree edit distance using relaxation labeling", *Springer Lecture notes in Computer Science*, **2134**, pp. 438-453, 2001. 35
38. B. Luo, A. Robles-Kelly, A. Torsello, R. C. Wilson, E. R. Hancock, "Discovering Shape Categories by Clustering Shock Trees", *Springer Lecture notes in Computer Science*, **2124**, pp. 151-160, 2001. 40

39. B. Huet and E. R. Hancock, "Fuzzy Relational Distance for Large-scale Object Recognition", *IEEE Computer Society Computer Vision and Pattern Recognition Conference, IEEE Computer Society Press*, pp. 138–143, 1998.
40. S. Moss and E. R. Hancock, "Pose Clustering with Density Estimation and Structural Constraints", *IEEE Computer Society Conference on Computer Vision and Pattern Recognition, IEEE Computer Society Press*, Fort Collins, vol II, pp. 85–91, 1999.
41. M. Carcassoni and E. R. Hancock, "Point Pattern Matching with Robust Spectral Correspondence", *IEEE Computer Society Conference on Computer Vision and Pattern Recognition, IEEE Computer Society Press*, I, pp. 649–655, 2000. 41
42. A. Robles-Kelly and E. R. Hancock, "A Maximum Likelihood Framework for Iterative Eigendecomposition", *Eighth International Conference on Computer Vision*, IEEE Computer Society Press, Vancouver, Canada, pp. 654–661, 1999. 39
43. A. Robles-Kelly and E. R. Hancock, "Graph Matching using Adjacency Matrix Markov Chains", *British Machine Vision Conference*, pp. 383–390, 2001. 39
44. A. Torsello, B. Luo, A. Robles-Kelly, R. C. Wilson and E. R. Hancock, "A Probabilistic Framework for Graph Clustering", *IEEE Computer Vision and Pattern Recognition Conference*, pp. 912–919, 2001. 40
45. A. Torsello and E. R. Hancock, "Matching and Embedding through Edit Union of Trees", *ECCV 2002*, to appear. 40
46. H. G. Barrow and R. J. Popplestone. Relational descriptions in picture processing. *Machine Intelligence*, VI:377–396, 1971. 31
47. K. Boyer and A. Kak. Structural Stereopsis for 3D Vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10:144–166, 1988. 32, 33
48. H. Bunke and K. Shearer. A graph distance metric based on the maximal common subgraph. *Pattern Recognition Letters*, 19:255–259, 1998. 32
49. H. Bunke. Error correcting graph matching: On the influence of the underlying cost function. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21:917–922, 1999. 32, 35
50. W. J. Christmas, J. Kittler, and M. Petrou. Structural matching in computer vision using probabilistic relaxation. *IEEE PAMI*, 17(8):749–764, 1995. 32
51. F. R. K. Chung. *Spectral Graph Theory*. American Mathematical Society Ed., CBMS series 92, 1997. 33, 38
52. A. D. J. Cross, R. C. Wilson, and E. R. Hancock. Inexact graph matching with genetic search. *Pattern Recognition*, 30(6):953–970, 1997. 32
53. M. A. Eshera and K. S. Fu. An image understanding system using attributed symbolic representation and inexact graph-matching. *Journal of the Association for Computing Machinery*, 8(5):604–618, 1986. 32, 35
54. S. Gold and A. Rangarajan. A graduated assignment algorithm for graph matching. *IEEE PAMI*, 18(4):377–388, 1996. 32, 37
55. M. Fischler and R. Elschlager. The representation and matching of pictorial structures. *IEEE Transactions on Computers*, 22(1):67–92, 1973. 31
56. L. Herault, R. Horaud, F. Veillon, and J. J. Niez. Symbolic image matching by simulated annealing. In *Proceedings of British Machine Vision Conference*, pages 319–324, 1990. 32
57. R. Horaud and H. Sossa. Polyhedral object recognition by indexing. *Pattern Recognition*, 28(12):1855–1870, 1995. 33
58. M. I. Jordan and R. A. Jacobs. Hierarchical mixture of experts and the EM algorithm. *Neural Computation*, 6:181–214, 1994. 37
59. B. Luo and E. R. Hancock. Relational constraints for point distribution models. *Springer Lecture Notes in Computer Science*, 2124:646–656, 2001. 41

60. M. Pellilo. Replicator equations, maximal cliques, and graph isomorphism. *Neural Computation*, 11(8):1933–1955, 1999. [32](#), [35](#)
61. A. Sanfeliu and K. S. Fu. A distance measure between attributed relational graphs for pattern recognition. *IEEE Trans. Systems, Man and Cybernetics*, 13(3):353–362, May 1983. [32](#), [35](#)
62. G. L. Scott and H. C. Longuet-Higgins. An Algorithm for Associating the Features of 2 Images. *Proceedings of the Royal Society of London Series B-Biological*, 244(1309):21–26, 1991. [33](#)
63. K. Sengupta and K. L. Boyer. Modelbase partitioning using property matrix spectra. *Computer Vision and Image Understanding*, 70(2):177–196, 1998. [33](#)
64. L. G. Shapiro and R. M. Haralick. A metric for comparing relational descriptions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 7(1):90–94, January 1985. [31](#)
65. L. S. Shapiro and J. M. Brady. Feature-based Correspondence - An Eigenvector Approach. *Image and Vision Computing*, 10:283–288, 1992. [33](#)
66. J. Shi and J. Malik. Normalized cuts and image segmentation. *PAMI*, 22(8):888–905, August 2000. [39](#)
67. A. Shokoufandeh, S. J. Dickinson, K. Siddiqi, and S. W. Zucker. Indexing using a spectral encoding of topological structure. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, pages 491–497, 1999. [33](#)
68. P. D. Simic. Constrained nets for graph matching and other quadratic assignment problems. *Neural Computation*, 3:268–281, 1991. [37](#)
69. S. áTirthapura, D. Sharvit, P. Klein, and B. B. Kimia. Indexing based on edit-distance matching of shape graphs. *Multimedia Storage And Archiving Systems III*, 3527:25–36, 1998. [32](#)
70. S. Umeyama. An eigen decomposition approach to weighted graph matching problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10:695–703, 1988. [32](#), [33](#), [38](#)
71. M. L. áWilliams, R. C. Wilson, and E. R. Hancock. Deterministic search for relational graph matching. *Pattern Recognition*, 32(7):1255–1271, 1999. [32](#)
72. A. K. C. Wong and M. You. Entropy and distance of random graphs with application to structural pattern recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 7:509–609, 1985. [32](#), [33](#)
73. A. L. Yuille, P. Stolorz, and J. Utans. Statistical Physics, Mixtures of Distributions, and the EM Algorithm. *Neural Computation*, 6:334–340, 1994. [32](#)
74. M. Pelillo, K. Siddiqi, and S. W. Zucker. Matching hierarchical structures using association graphs. *IEEE PAMI*, 21(11):1105–1120, 1999. [32](#)
75. W. J. Christmas, J. Kittler, and M. Petrou. Structural matching in computer vision using probabilistic relaxation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(8):749–764, 1995. [33](#)
76. J. Kittler and E. R. Hancock. Combining evidence in probabilistic relaxation. *IEEE PRAI*, 3:29–51, 1989. [34](#)
77. V. Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics-Doklady*, 10:707–710, 1966. [35](#)
78. P. N. Suganthan, E. K. Teoh, and D. P. Mital. Pattern-recognition by graph matching using the potts mft neural networks. [37](#)
79. T. Hofmann and J. M. Buhmann. Pairwise data clustering by deterministic annealing. *PAMI*, 19(2):192–192, February 1997. [37](#)