

Inexact Graph Matching with Genetic Search

Andrew D.J. Cross and Edwin R. Hancock
Department of Computer Science, University of York
York, YO1 5DD, UK

Abstract. This paper describes a novel framework for performing relational graph matching using genetic search. The fitness measure is Bayesian in origin. It gauges relational consistency at both the symbolic and attribute levels. The basic measure of symbolic consistency is Hamming distance, while attribute consistency is measured using Mahalanobis distance. We provide examples of the performance on synthetic graphs containing significant levels of clutter. We also demonstrate that the technique is capable of resolving multiple graphs with significant overlap. The performance advantages over deterministic hill climbing are also demonstrated.

1 Introduction

Although genetic search is new and imperfectly understood, it provides an attractive means of solving configurational optimisation problems [6, 11, 12]. Basic to genetic search is the idea of maintaining a population of alternative global solutions to the discrete optimisation problem in-hand. The initial population may be generated in a number of different ways, but should in some sense uniformly sample the feasible solution space. Associated with each of the different solutions is a cost function which in keeping with the evolutionary analogy is termed the “fitness” [6]. Genetic updates involve three distinct stages. Crossover maintains diversity by randomly selecting pairs of solutions from the current population and interchanging the symbols at corresponding configuration sites with a uniform probability [12]. Mutation aims to introduce new information into the population by randomly updating the component symbols for individual solutions with a uniform probability [12]. The net effect of modifying the population in this way is to randomly sample the landscape of the fitness function. Configurations generated by crossover and mutation are subjected to a stochastic selection process in order to avoid convergence to a local optimum [11]. The probability that a modified configuration enters the population is computed on the basis of the fitness measure. In many ways genetic search provides an interesting compromise between the continuous transformation of the discrete optimisation problem [5, 10] and its realisation by simulated annealing [1, 9, 7].

Our interest in this paper centres on exploiting genetic search in the matching of attributed relational graphs [13, 14, 16]. Here we aim to find a discrete matching configuration that optimises a Bayesian fitness measure which gauges relational consistency. This fitness measure is defined over connected subgraphs. The development of the consistency measure draws on the modelling of both

symbolic and attribute errors in the matching process. Symbolic differences are measured by Hamming distance. Attribute differences are measured by Mahalanobis distance. Genetic updates are aimed at locating the relational matches that maximise probability measure. In order to realise the matching process efficiently, we augment the standard genetic search operator in two ways. Firstly, we incorporate a hill-climbing step which ensures that the solutions in the population reside at the nearest local optimum prior to selection. The second algorithm refinement is to realise crossover at the subgraph level rather than selecting graph nodes at random. As recently demonstrated, the two algorithm refinements significantly accelerate convergence without compromising the global properties of genetic search [4].

The outline of this paper is as follows. In Section 2 we describe the basic graph formalism used in the remainder of the paper. Section 3 describes the development of our Bayesian fitness measure. Section 4 describes how the optimisation of this measure may be mapped onto a genetic search procedure. Finally, Section 6 offers some conclusions.

2 Relational Graphs

We abstract the matching process in terms of attributed relational graphs [2, 3, 13, 16]. 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 edges and $\mathcal{A} = \{\mathbf{x}_i, \forall i \in V\}$ 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. 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. $f \subseteq V_1 \times V_2$.

In performing the matches of the nodes in the data graph G_1 we will be interested in exploiting 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 structural 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 feasible relational mappings between representational subunits. We will strike a compromise by using subgraphs that consist of neighbourhoods of nodes interconnected by edges; for convenience we refer to these structural subunits or N -ary relations as super-cliques.

The super-clique of the node indexed j in the graph G_1 with edge-set E_1 is denoted by the set of nodes $C_j = j \cup \{i | (i, j) \in E_1\}$. The corresponding set of attributes is denoted by $\mathcal{R}_j = \{\mathbf{x}_u | u \in C_j\}$. The matched realisation of this super-clique is denoted by the symbolic relation $\Gamma_j = (f(u_1), f(u_2), \dots, f(u_{|C_j|}))$.

Our aim is to modify the match to optimise a measure of global relational consistency using the constraints provided by the model graph G_2 . There are two components to this consistency measure. The first of these is symbolic and involves comparing the matched relations, i.e. Γ_j , from the data graph with their exact counterparts in the model graph. The second component relates to the consistency of the attribute relations associated with the symbolic matches. Our aim is therefore to compare the matched configuration of symbols and attributes denoted by $\Upsilon_j = (\Gamma_j, \mathcal{R}_j)$ with their counterparts in the model graph.

Each of our constraint relations is formed on the super-cliques of the model graph G_2 . We compile the set of feasible relational matches, or structure preserving mappings, between the super-clique centred on the node j of the data graph and those of the model graph in a dictionary which we denote by Θ_j . Every entry in the dictionary is itself a relation defined over symbols and attributes. We denote the individual dictionary items by the mixed relation $\Lambda^\mu = (\Phi^\mu, S^\mu)$ where μ is an index over the dictionary. According to this notation $\Phi^\mu \subset V_2$ is the set of nodes which form the μ^{th} matchable super-clique. If $k \in V_2$ is the central node of the super-clique, then the set of symbols forming the μ^{th} structure preserving mapping is $\Phi^\mu = \{v_1^\mu, v_2^\mu, \dots, v_{|C_k|}^\mu\}$. The unary attribute set for the structure preserving mapping is given by $S^\mu = \{\underline{x}_l | l \in \Phi^\mu\} \subset \mathcal{A}_2$. With these ingredients $\Theta_j = \{\Lambda^\mu | \mu = 1, Z_j\}$.

3 Fitness

Our modelling of structural consistency is Bayesian and commences from the joint probabilities for the matched relations defined on the super-clique of the data graph, i.e. $P(\Upsilon_j)$. This model of the matching probability can be viewed as providing a means of imposing constraints on consistent relational matches. The available constraints residing in the dictionary are mixed N-ary relations defined over both symbols and attributes. We develop a Bayesian model of relational corruption. This results in a consistency metric which is a compound exponential function of two relational distance measures. Symbolic differences are gauged by Hamming distance and attribute differences are gauged by the Mahalanobis distance.

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

$$P(\Upsilon_j) = \sum_{\Lambda^\mu \in \Theta_j} P(\Upsilon_j | \Lambda^\mu) \cdot P(\Lambda^\mu) \quad (1)$$

In order to develop this expression further we recall that each structure preserving mapping Λ^μ is composed of both symbolic and attribute components. The different dictionary items are assumed to occur with a uniformly distributed prior, i.e. $P(\Lambda^\mu) = \frac{1}{|\Theta_j|}$. We further assume that the attribute and symbol information may be dichotomised in the matching process. Accordingly we apply

the Bayes theorem to the probability $P(\Upsilon_j|\Lambda^\mu)$ to separate the relations into their symbolic and attribute components in the following way

$$P(\Upsilon_j|\Lambda^\mu) = P(\Gamma_j|\Phi^\mu) \cdot p(\mathcal{R}_j|\mathcal{S}^\mu) \quad (2)$$

The development of a useful graph-mapping measure from this expression requires a two-component model of the processes at play in matching and of their roles in producing errors. According to our dichotomy, there are separate models for the symbolic matching constraints which are represented by the conditional probability $P(\Gamma_j|\Phi^\mu)$ and for the attribute matching constraints which are modelled by the conditional measurement density $p(\mathcal{R}_j|\mathcal{S}^\mu)$.

3.1 Symbolic Constraint Violation

Our model of the symbolic constraint process follows Wilson and Hancock [16]. Accordingly, we assume that the various types of matching error for nodes belonging to the same super-clique are memoryless. In direct consequence of this assumption, we may factorize the probability $P(\Gamma_j|\Phi^\mu)$ over the symbolic constituents of the relational mapping under consideration. As a result the conditional probability may be expressed in terms of a product over label confusion probabilities

$$P(\Gamma_j|\Phi^\mu) = \prod_{k=1}^{|C_j|} P(f(u_k)|v_k^\mu) \quad (3)$$

The matching errors at individual sites in the super-cliques are assumed to occur with a uniform and memoryless probability distribution. If the probability of matching errors is P_e , then the following distribution rule applies

$$P(f(u_k)|v_k^\mu) = \begin{cases} (1 - P_e) & \text{if } f(u_k) = v_k^\mu \\ P_e & \text{if } f(u_k) \neq v_k^\mu \end{cases} \quad (4)$$

As a natural consequence of this distribution rule the joint conditional probability is a function of the Hamming distance $H(\Gamma_j, \Phi^\mu) = \sum_{i \in C_j} (1 - \delta_{f(u_k), v_k^\mu})$ between the assigned matching and the feasible symbolic relational mapping Φ^μ [16]. This quantity counts the number of conflicts between the current matching assignment Γ_j residing on the super-clique C_j and those assignments demanded by the symbolic relational mapping Φ^μ . With these ingredients, the resulting expression for the joint conditional probability acquires an exponential character

$$P(\Gamma_j|\Phi^\mu) = K_{C_j} \exp[-k_e H(\Gamma_j, \Phi^\mu)] \quad (5)$$

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 probability distribution appearing in Equation (5) may be regarded as providing a natural way of softening the hard symbolic constraints operating in the model graph.

3.2 Attribute Constraints

Our modelling of the conditional measurement density $p(\mathcal{R}_j|S^\mu)$ is based on the assumption of Gaussian measurement errors in the attribute acquisition process. Accordingly we gauge differences between the attribute relations in the data graph and their counterparts in the dictionary using the Mahalanobis distance. The parameter of this distribution is the measurement covariance matrix Σ . The corresponding multivariate Gaussian density is

$$p(\mathcal{R}_j|S^\mu) = \frac{1}{(2\pi)^{\frac{|\mathcal{R}_j|}{2}}} \frac{1}{\sqrt{|\Sigma|}} \exp\left[-\frac{1}{2}(\mathcal{R}_j - S^\mu)^T \Sigma^{-1} (\mathcal{R}_j - S^\mu)\right] \quad (6)$$

With this ingredient the relational matching probability $P(\Upsilon_j)$ becomes

$$P(\Upsilon_j) = \beta_j \sum_{\Lambda^\mu \in \Theta_j} \exp\left[-\left(k_e H(\Gamma_j, \Phi^\mu) + \frac{1}{2}(\mathcal{R}_j - S^\mu)^T \Sigma^{-1} (\mathcal{R}_j - S^\mu)\right)\right] \quad (7)$$

where $\beta_j = \frac{K_{C_j}}{(2\pi)^{\frac{|\mathcal{R}_j|}{2}} |\Theta_j| \sqrt{|\Sigma|}}$.

In this way the separate roles of the symbolic and attribute components of the relations becomes explicit. All that now remains is to use the configurational probability $P(\Upsilon_j)$ to define a global fitness measure for use in genetic search for the optimal relational matches.

3.3 Global Fitness Measure

The configurational probability $P(\Upsilon_j)$ is the basic ingredient of our genetic search procedure. It represents the probability of a particular matching configuration evaluated over the state-space of feasible possibilities (i.e. the dictionary). We use as our *global measure of consistency* by the average clique matching probability enumerated over the cliques of the data graph i.e.

$$P_G = \frac{1}{|V_1|} \sum_{j \in V_1} P(\Upsilon_j) \quad (8)$$

With the ingredients described above, our global fitness measure draws on both attributes and symbols in gauging the quality of match. In this way it naturally provides an interesting compromise between the purely symbolic consistency measure of Wilson and Hancock [15, 16] on the one hand, and the attribute oriented relational consistency measures of Boyer and Kak [3] or of Yang and Kittler [17] on the other hand.

4 Genetic Search

Genetic search [6, 11, 12] provides a very natural way of locating the global optimum of the global consistency measure described in the previous section. In essence the approach relies on generating a population of random global matching configurations. These undergo cross-over, mutation and selection to locate the match that optimises a fitness measure. The main stages of the algorithm are outlined below and more detailed discussion can be found in [4].

Initial population generation: The initial population is generated either to be uniformly distributed or to be biased towards the initial matches suggested by unary measurement information. Whereas, the uniform distribution is appropriate if the attribute information is uncertain or unreliable, biasing may be employed to direct the search procedure towards favourable portions of the solution space.

Crossover: Crossover exchanges information between graph pairs in the population. Rather than using a uniform crossover [12], we realise the process at the level of disjoint subgraphs. This mixing of partially consistent subgraphs accelerates convergence [4].

Mutation: Mutation operations ensure that the fitness landscape is uniformly sampled by randomly swapping matches. This can be viewed as introducing uniform noise into the population. This not only introduces diversity, it also inhibits premature convergence.

Hill Climbing: One of the novel features of our genetic search process is the incorporation of a deterministic hill-climbing stage. This additional step is applied to the fitness measure once mutations have occurred and is used to accelerate convergence to the nearest optimum of the average consistency measure. The hill-climbing step ensures that each solution in the genetic population resides at a local optimum of the fitness measure. In this way local sub-optima may be rapidly rejected by the selection process.

Graph Editing: One of the critical ingredients in effective relational matching is the way in which unmatchable entities or clutter are accommodated. Here we follow a graph-edit philosophy which removes the clutter nodes and recomputes the edge-set of the graph as necessary [14]. This process is incorporated into the hill-climbing stage in the following way. Each node in turn is deleted from the graph and the edge-set recomputed. Our decision concerning node deletion or re-insertion is based on the value of P_G . If the value of P_G increases due to the deletion process, then the node is edited from the graph. If, on the other hand, the value of P_G increases as a result of node re-insertion at a later stage, then it is reinstated.

Selection: The final stochastic element of genetic search is the selection process. The aim here is to randomly admit the configurations refined by the hill climbing process to the population on the basis of their fitness measure. The probability distribution defined in equation (5) lends itself naturally to the definition of a population membership probability. Suppose that $P_G^{(i)}$ denotes the global configurational probability for the i^{th} member of the pool (population) of graphs. By normalising the sum of clique configuration probabilities over the population of matches, the probability for randomly admitting the i^{th} solution to the pool of graphs \mathcal{P} is $P_s = \frac{P_G^{(i)}}{\sum_{i \in \mathcal{P}} P_G^{(i)}}$.

5 Synthetic Matching Examples

Our aim in this Section is evaluate the behaviour of the genetic search procedure on synthetic data-sets with known ground-truth. The main goals here are to provide a study of the systematics of the method when structural corruption is a limiting factor. We also provide some examples to illustrate the effectiveness of the method at matching overlapped or highly corrupted graphs.

Figures 1a and 1b illustrate some typical matching results on synthetic graphs. Figure 1a shows the fittest solution from the initial population. The left-hand graph is the model while the right-hand graph is the data; lines between the two graphs indicate matches. The data graph has been obtained by adding random clutter to the model and perturbing the nodes with Gaussian position errors. Associated with each node in the model graph is a single unary attribute which has been generated at random from a uniform distribution. The corresponding attributes in the data graph have been obtained by adding Gaussian measurement errors to the unary attributes of the model. It should be noted that the unary attributes are entirely uncorrelated to either the absolute or relative positions of the nodes. In other words, we make no use of any kind of transformational information between the two scenes. The original model graph contains 20 nodes while the corrupted data graph contains 40 nodes. Figure 1b shows the fittest match from the genetic population after 3 iterations. There are two features worth noting. Firstly, the overall consistency of match has improved. The lines connecting the nodes in the data and model graphs are no longer randomly distributed. Secondly, the added clutter nodes have all been correctly identified and deleted from the data graph; they appear as disjoint points on the right-hand image of Figure 1b. The overall accuracy of match in this example is 100%.

The example described above is typical of the problem of matching a relational description that is subsumed in noise or clutter. Another common problem in computer vision is to match scenes containing multiple objects. Under particularly severe imaging conditions these objects may be significantly overlapped. The following two examples illustrate the capacity of our genetic search procedure to match under these two sets of conditions.

We commence with the simpler example which involves the matching of multiple non-overlapping models. Figure 2a shows the fittest initial match while Figure 2b shows the final match. The data graph, on the left hand side of Figures 2a and 2b, is a non-overlapping union of the three models on the right-hand side of the figures. Here the genetic search algorithm correctly partitions the data graph into three disjoint subgraphs. As indicated by the lines between the data and model, each of the subgraphs is correctly matched.

A more complex case in which the three graphs are overlapped is illustrated in Figures 3a and 3b. Here our genetic matching technique is again capable not only of correctly partitioning the nodes of the data graph into the three disjoint subgraphs but also of locating the consistent matches. In fact these results indicate that our matching technique has considerable potential as a tool for extracting relational clusters from highly overlapped data.




Fig. 1. a) An initial guess b) Recovered Solution




Fig. 2. a) An initial guess b) Recovered Solution

In order to illustrate the effectiveness of the genetic search technique, we have compared its performance with deterministic hill climbing. The deterministic algorithm aims to optimise the global cost function given in equation (8) by gradient ascent; in other words, the label update that results in the greatest increase in P_G is always accepted at a particular node. The comparison has been performed under conditions of controlled structural corruption. We have generated random graphs and added a controlled fraction of spurious noise. Figure 4 shows the fraction of the graph correctly recovered and matched as a function of the fraction of added noise nodes. The lower curve is the result obtained by iterating the deterministic method to convergence. The intermediate curve is the result after performing one iteration of genetic search with a population size of 100 graphs. After two iterations of genetic search the upper curve is obtained. The main conclusion from this study is that once the corruption level exceeds 20%, the gradient ascent technique is likely to become trapped in a local minimum. By exploring a much greater fraction of the search-space, genetic search is capable of finding good results even at very severe corruption levels. In other words, when combined with the probabilistic cost function, genetic search can recover significantly better results than its deterministic counterpart.





Fig. 3. a) An initial guess b) Recovered Solution



Genetic Hill Climbing vs. Gradient Ascent

Fig. 4. The effect of controlled structural corruption.

6 Conclusions

To conclude, we have shown how the optimisation of a Bayesian relational consistency measure naturally maps onto genetic search. This measure gauges relational consistency using both symbolic and attribute information. Symbolic differences are represented in terms of Hamming distance while attribute differences are represented by Mahalanobis distance. The genetic optimisation of our relational consistency measure is capable not only of correcting initialisation errors, but also of rectifying structural differences. Moreover, the technique can accurately partition merged or overlapping graphs into component model subgraphs.

References

1. Aarts E. and J. Korst, "Simulated Annealing and Boltzmann Machines", *John Wiley and Sons, New York*, 1989.
2. Barrow H.G. and R.J. Popplestone, "Relational Descriptions in Picture Processing", *Machine Intelligence*, **6**, 1971.
3. Boyer K. and A. Kak, "Structural Stereopsis for 3D Vision", *IEEE PAMI*, **10**, pp 144-166, 1988.
4. Cross A.D.J, R.C. Wilson and E.R. Hancock, "Genetic Search for Structural Matching", *Computer Vision -ECCV96*, Lecture Notes in Computer Science, 1064, pp. 514-525, 1996.
5. Faugeras O.D. and M. Berthod, "Improving Consistency and Resolving Ambiguity in Relaxation Labelling", *IEEE PAMI*, **3**, pp. 412-424, 1981.
6. Fogel D.B., "An Introduction to Simulated Evolutionary Optimisation", *IEEE Transactions on Neural Networks*, **5**, pp. 3-14, 1994.
7. Geman S. and D. Geman, "Stochastic relaxation, Gibbs distributions and Bayesian restoration of images," *IEEE PAMI*, **PAMI-6** , pp.721-741, 1984.
8. Hancock, E.R. and J. Kittler, "Discrete Relaxation," *Pattern Recognition*, **23**, pp. 711-733, 1990.
9. Kirkpatrick S., C.D. Gelatt and M.P. Vecchi, "Optimisation by Simulated Annealing", *Science*, **220**, pp. 671-680, 1983.
10. Peterson C. and B. Soderberg, "A New Method for Mapping Optimisation Problems", *International Journal of Neural Systems*, **1**, pp 2-33, 1989.
11. Qi, X.F. and F. Palmieri, "Theoretical Analysis of Evolutionary Algorithms with an Infinite Population in Continuous Space: Basic Properties of Selection and Mutation" *IEEE Transactions on Neural Networks*, **5**, pp. 102-119, 1994.
12. Qi, X.F. and F. Palmieri, "Theoretical Analysis of Evolutionary Algorithms with an Infinite Population in Continuous Space: Analysis of the Diversification Role of Crossover" *IEEE Transactions on Neural Networks*, **5**, pp. 120-129, 1994.
13. Shapiro L.G. and R.M. Haralick, "A Metric For Comparing Relational Descriptions", *IEEE PAMI*, **7**, pp 90-94, 1985.
14. Sanfeliu A. and K.S. Fu, "A Distance Measure Between Attributed Relational Graphs for Pattern Recognition", *IEEE SMC*, **13**, pp 353-362, 1983.
15. Wilson R.C and E.R. Hancock, "Graph Matching by Discrete Relaxation", *Pattern Recognition in Practice IV*, edited by E Gelsema and L Kanal, pp 165-176, 1994.
16. Wilson R.C., A.N. Evans and E.R. Hancock, "Relational Matching by Discrete Relaxation", *Image and Vision Computing*, **13**, pp. 411-421, 1995.
17. 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.