Navigation Using Inverting Genetic Algorithms: Initial Conditions and Node-Node Transitions

Steven J. Simske[1] and David C. Matthews[2]

[1]Hewlett-Packard Company, 3404 East Harmony Road, Mailstop 85
Fort Collins, CO 80528, USA
steven.simske@hp.com
[2]Colorado State University, Computer Science Department,
Fort Collins, CO 80523-1873, USA
davematt@cs.colostate.edu

Abstract. Navigation path optimization derives its cost function from the total cost of travel. These costs can accrue from distance, traffic patterns, preferred order of node sequencing, maximum preferred distance between nodes, and other pragmatic considerations (node-node costs). For a traveling salesman problem (TSP), these costs are usually distances. This paper considers the effects of the initial states in the "genome"—focusing on the use of rule-based and clustering techniques for initial conditions. It also considers the effects of weighting the node-node transition costs on algorithm convergence and final path cost. The best tested combination of initial states and rules for recombination involves weighting each by the distances between nodes. In addition to the mean of the residual error μ, the metric designated algorithmic efficacy (AE) is introduced as a useful comparative metric for navigational optimization algorithms. Four novel and six published TSP problems are investigated: μ is shown to improve in every case, while AE has a wide range of results.

1 Introduction

The utility of genetic algorithms (GA's) for solving difficult optimization problems efficiently benefits navigational systems [1]. This is due in part to the role of mutation, crossover and inversion in preventing futile iterating around non-global minima [2]. However, when large numbers of nodes (positions that are traversed between) are involved, or when node-node transitions are difficult to represent with the traditional loci set of GA's [3], GA's can fail to explore the problem space sufficiently to find the global minimum cost. For navigational systems, node-node transitions can rarely be viewed in isolation, but instead practical issues around traversing particular node-node pathways must be accounted for. GA researchers have looked at gene linkage [4,5], punctuated crossover [6,7] and "messy" GA's [8] as means of allowing groups of co-adapted genes to be inherited together during the recombination/crossover phase of the GA. Multiple recombination strategies [9] also can provide greater variety in offspring. In this paper, the concept of "relative gene linkage" via weighted node-node transition costs is introduced. Relative gene linkage is imposed directly

(via initial pathway specification) and indirectly (via probabilistic inversion, or crossover via reversal of a substring of the gene). A means for comparing the algorithmic efficacy (AE) of different initial conditions and node-node probability schemes is introduced.

2 Initial Conditions, Node-Node Transitions and Test Cases

The experiments performed here are roundtrip traveling salesman problems (TSPs) in which the starting node is also the ending node, and every other node in the set is traversed. Nodes thus correspond to physical locations (points, or cities in three of the examples). Node-node transitions are the costs (herein, simply the distances) associated with moving between one node and another. A particular node-node transition will be called a state, and each experiment limits the number of states to N, the number of nodes. That is, for each node reached, the next node must be one that has not yet been visited, with the exception of the last node visited, for which the subsequent state returns to the starting node.

2.1 Initial Conditions

Because convergence efficacy is increased when (a) the initial genes are close to the optimal genes, and (b) at least some of the initial genes are evaluated (and selected) before the first crossover, the following initial conditions were explored:

1. Naïve. Here the N nodes in a navigational pathway are chosen randomly, one after the other, until all N are exhausted. This is a reasonable starting point when N is small or when node-node distances are relatively similar for all node-node transitions.
2. Weighted Nearest Neighbor. Here, for each node, the costs to each of the other N-1 nodes is calculated, and all "#1" choices, where possible, awarded. From the remaining pool of node-node transitions, all "#2" choices, where possible, are awarded. This continues until all node-node pathways have been assigned. For example, in a 4-node problem, if the following node-node transition distances are recorded:

Table 1. Sample 4-node problem with node-node transition distances as shown

	Node 1	Node 2	Node 3	Node 4
Node 1	--	40	30	35
Node 2	40	--	50	55
Node 3	30	50	--	70
Node 4	35	55	70	--

Node 1-Node 3 is assigned first. Node 4-Node 1 is assigned next. Because Node 1 is now complete, Node 2-Node 1 cannot occur in an exhaustive complete traversal, and so the following traversal is obtained: {Node 4-Node 1-Node 3-Node 2-Node 4}.

3. Lowest Remaining Distance. Here, the lowest remaining distance between any "nonexhausted" nodes is always assigned. This and the preceding method are generally reasonable when the nodes are relatively evenly spaced; however, when there are a small number of outlying points, this can results in extraneous large distances to/from/between such outliers.

4. Centroid + Clockwise/Counterclockwise Traversal. In this method, the starting point is selected at random, and then the path follows from one node to the next along a clockwise/counterclockwise path around a reference point. The centroid of all nodes is the reference point for the rotation.

5. Clustering + Method 4. If the distances between nodes vary considerably, then the nodes themselves are candidates for clustering. Traditional clustering techniques (e.g. K-means) can be used to assign the nodes to distinct clusters, along with a method found to be useful herein for cluster definition: A cluster C composed of 2 or more nodes is valid if the minimum distance of a path between all of its nodes is less than the distance to any other nodes in the set of all nodes S.

2.2 Node-Node Transitions

Node-node transitions can be differentially weighted to affect the initial generation of genes, to impact the location of crossover, or both. Some important means of assigning node-node transition probabilities (or costs, normalized to sum to 1.0 for all possible transitions from a given node) include:

1. Naïve. Here, all probabilities are identical, and so each node-node transition cost, or probability, is defined as: $P_i = 1/(N-1)$, where N=#nodes.

2. By Relative Distance. The node-node costs are weighted by relative distance. One such method is to use a step size of ΔP, so that $P_i = P(Min)+((i-1) * \Delta P)$, where P(Min) is the minimum cost, and i=1…N. This is also "ranked" distance.

3. By Cluster. If nodes are assigned to clusters, then the weighting within the clusters can be higher than the weighting outside of the cluster (many possible methods). That is, $p(C_i,C_j) > p(C_i,N_j)$, where C is the cluster, and N = S-C is the set of all other nodes (outside the cluster C).

4. By Direction. If a clockwise/counterclockwise traversal is occurring, then the next nodes in the direction being traversed can be weighted more highly. This Bayesian technique is advantageous for determining dynamically the probabilities for crossover (however, due to space limitations and its modest benefits, this is not considered further herein).

5. By Distance. An inverse weighting scheme is used here, where the weight/probability $P_i = C_i/d_{ij}$ for all nodes i to all other nodes j. Since the weighting is inversely proportional to distance, then $d_{ij}/d_{ji} = P_j/P_i$, and thus $C_j/C_i = 1$, $P_i = k/ d_{ij}$, and since $\Sigma_{i=1\ldots N}P_i = 1.0$, then $k = 1/(\Sigma_{j=1\ldots N}(1/d_{ij}))$, and so:

$$P_i = 1 / (d_i * \Sigma_{j=1\ldots N}(1/d_{ij})). \quad (1)$$

This can be extended to any power X of distance readily as given here:

$$P_i = 1 / (d_i^X * \Sigma_{j=1\ldots N}(1/d_{ij}^X)) \quad (2)$$

It can also be readily extended to any function f(i):

$$P_i = f(i) / \Sigma_{j=1...N}(f(i))) \quad (3)$$

In real-world applications, these node-node transition costs are based on, among other things, distance and traffic patterns between nodes, and intelligent clustering of nodes based on similarity in the navigator's intents at each destination (e.g. cities to visit may be clustered by language in Europe), maximum distance preferred between nodes, and other pragmatic considerations.

2.3 Test Cases Investigated

First we describe four novel test cases we designed to provide a wide gamut of tests. We developed a Java-based toolkit allowing for control of: (1) number of genes, (2) crossover rate and type, (3) mutation rate and type, (4) test for performance asymptote, (5) test for optimal fitness intransigence, (6) fitness (cost), (7) selection rate, (8) initial conditions, and (9) node-node transition costs.

Of the four custom test cases, the first is the "Iceland" challenge, wherein 9 cities in Iceland (Anglicized to: Akranes, Akureyri, Borgarnes, Egilstaddir, Hofn, Isafjordur, Reykjavik, Selfoss and Vik, distances garnered from [10]) were traversed in a roundtrip TSP path as described above. For this type of problem, the number of possible "optimum pathways" is at least 2N, where N=#nodes, since the optimum single traversal optimum pathway can start and end at any node, and traverse either clockwise or counterclockwise. Thus, the odds of randomly selecting an optimum pathway are 2N/N! = 2/(N-1)! For the Iceland case, this is 1/20160.

The second case used 20 cities in the United States and the reported distances between them. The cities were selected to ensure that no obvious optimal pathway could be obtained (unlike the Iceland case, in which it is in hindsight obvious that the "Centroid + Clockwise/Counterclockwise Traversal", among others, provides the optimal path). For the "USA" case, the odds of randomly selecting an optimum pathway are $2/19! = 1/(6.082255x10^{16})$.

These two ("Iceland" and "USA") test cases were used to explore the effects of initial conditions and node-node transitions. The optimal set of initial condition and node-node transitioning methods were then applied to two further cases. The "Europe" test case uses 30 cities in Europe and the distances between them. The cities are all on the continent and north of the Alps in an attempt to prevent obvious outliers and increase the complexity of the overall task. For the "Europe" test case, the odds of randomly selecting an optimal pathway are $2/29! = 1/(4.42088x10^{30})$. Finally, a synthetic case ("Grid") was considered. The "Grid" case consists of a four-by-four grid of "squares", simulating the inner and bordering intersections of a 3x3 set of city blocks. Because of the geometry of the grid, the odds of randomly selecting a minimum pathway is $192/16! = 1/(1.08973x10^{11})$.

For all of the test cases, except where noted otherwise, the following specifications were employed: (1) 50 genes were used. (2) Crossover rate was set at 90%, with crossover consisting of a roulette-wheel selection of two splicing locations (that is, two intra-loci transitions), followed by reversal of the pathways between the nodes.

This is known as **inversion**. For example, if the original pathway was: 1234567891 for cities 1-9, and the splicing transitions were found to be 2-3 and 6-7, then our "crossed-over" result is 12**6543**7891. The Java Math.random() function was used for determining the splicing locations; therefore, the **true inversion** rate was 80% = (0.889*90%) due to the possibility of identical splice points for the "Iceland" case, and 85.5%, 87% and 84.4% for the "USA", "Europe" and "Grid" cases, respectively, for the same reason. (3) Mutation rate was set at 2%, and employed as a swap of two loci. For the previous example, this results in 12**6**45**3**7891. True mutation rate was 1.78%, 1.9%, 1.93% and 1.88%, respectively, for the "Iceland", "USA", "Europe" and "Grid" cases. (4) Rapidly-converging (asymptotic) runs of the algorithm were determined by comparing after each iteration of the GA the standard deviation (STD) of the fitness (or "cost", which was total distance) for the first half of the iterations to the STD of the fitness over the last half of the iterations. If the ratio was greater than four, then the particular run of the algorithm was declared "converging". (5) If converging, the run was either terminated and the optimal fitness and path/s recorded or else the mutation rate was increased to 5% for one iteration and the run allowed to proceed. This "spot" increase in mutation rate was termed "jiggling." (6) Fitness cost was simply the sum of all node-node distances. (7) Selection rate was based on relative cost for each gene. Suppose two genes G1 and G2 had cost 2500 and 3000, respectively. Then, the survival weight for G1 was proportional to 1/2500 and that for G3 proportional to 1/3000. Summing these for all genes and normalizing to 1.0 allowed selection of survival using iterations of Math.random(). For example, if G1 ended up with 3% of the total fitness, then, for example, in the 50 iterations of Math.random(), each value in the interval [0.00, 0.03) would select for one G1 offspring (expected value 1.5 offspring out of 50). (8) Initial conditions could be assigned by any of the means discussed above. (9) Finally, node-node transition costs offered a unique opportunity to improve the expected fitness and performance of the GA. An example of the "naïve" assignment of these transitional probabilities is shown here (Probability, P, = 0.125 for all 8 Transitions from "Hofn"):

<Transition Source="Hofn" Destination="Akranes" Value="493" P="0.125"/>

The weighting scheme of Eq. 1 is shown here (Probability, P, = 0.23 from Boston to nearby New York, but is only 0.01 to far-away Seattle):

<Transition Source="Boston" Destination="New York" Value="211" P="0.23"/>

<Transition Source="Boston" Destination="Seattle" Value="3088" P="0.01"/>

In addition to these four novel test cases, we then performed tests on 2 symmetrical and 4 asymmetrical cases published on [11]. The asymmetrical tests can have differing distances from A to B than from B to A, thus complicating the effects of inversion.

3 Results and Discussion

For the "Iceland" and "USA" cases, three sets of 1000 runs (computing time: 1-15 min/set) of the genertic GA software were performed, and the following data computed: (1) minimum (C_{min}), maximum (C_{max}) and range (C_r) of "lowest cost", or "optimal" pathways obtained in each run; (2) mean (μ) of the optimal pathways obtained

(or "error"), and the standard deviation (σ_μ) of the means for the three sets of 1000 runs; and (3) the number of iterations (N_i) to converge on the optimal value so obtained (along with the standard deviation of N_i, σ_N). Since the true (global) optimum cost (C_{opt}) was known (2009 km for the "Iceland" case, 9271 miles for the "USA" case), C_{min}, C_{max}, C_r, μ and σ_μ were normalized by C_{opt} (and are thus presented as percentages of the optimal cost) and 100% subtracted from them (excepting C_r and σ_μ) so that they show incremental percentage over the optimum. For the "Europe" and "Grid" tests, sets of 10,000 (1-6 hours) or 100,000 (1-4 days) runs were performed, affording a very accurate set of data. For the published cases [11], sets of 1,000 (1-4 days) runs were performed. The data reported here represent more than a year's worth of 733 MHz CPU time altogether.

3.1 Iceland Case

Node-node probability regimens were used to determine the initial set of genes (Table 2). The first was a naïve assignment in which the next node was randomly selected from the remaining (legitimate) pool. The second used relative distance weighting for all nodes (normalized to 1.0 after each assignment for the remaining nodes) where step size ΔP=0.05 and P(Min)=0.02, so P(Max)=0.23. The third used the results of clustering to pool together the nodes (Akranes, Borgarnes, Reykjavik and Selfoss) and thus exclude the other five nodes (within cluster weighting was increased to 3, 5, 15 and 30 times the extra-cluster weighting, with peak effectiveness at 15X, the reported value in Table 2). The fourth used inverse distance weighting (Eq. 1).

Table 2. Results for the "Iceland" case using different node-node transition methods for determining the initial genome. All groups are significantly different ($p<.01$) comparing (μ +/- σ_μ) and (N_i +/- σ_N). C_{min} = 0.00 for all sets of runs, and so $C_T=C_{max}$

	Naïve	Rel. Dist	Cluster	Distance
C_{max} (%)	11.50	8.41	6.57	8.16
μ (%)	**0.247**	**0.085**	**0.130**	**0.054**
σ_μ (%)	0.012	0.021	0.029	0.017
N_i	**8.049**	**5.930**	**6.434**	**4.592**
σ_N	0.027	0.076	0.165	0.116

The naïve assignment (random gene sequencing) results in 0.247% expected increase in best solution cost over true optimum ("error"), with a mean of 8.05 iterations (representing 50*8.05 ~ 402 genes) to reach convergence. For the simple clustering technique, the mean optimal cost is 0.130% above true minimum (a 47.4% relative decrease) with a concomitant decrease to 6.434 (a 20.1% relative decrease) iterations (~ 322 genes) to reach convergence. Further improvement is obtained by the "Relative Distance" method, in which longer distances from a node are incrementally weighted less (for probability): the mean optimal cost is only 0.085% above true optimum (a 65.6% relative decrease from "Naïve") while reducing further (to 5.950, or ~298 genes, a 26.1% relative decrease from "Naïve") the iterations to reach con-

vergence. Lastly, the "Distance" mean optimal cost reduces to 0.054% above true optimum (a 78.1% relative decrease from "Naïve") and iterations to reach convergence reduce to 4.592 (~ 230 genes, a 42.9% relative decrease from "Naïve"). To compare these results relatively, "Algorithmic Efficacy" (AE) can be defined as:

$$AE = k/(\mu * N_i) \quad (4)$$

where k is a normalizing constant for the particular case (k=1 for the "Iceland" case). AE (not shown in Tables) for the "Naïve", "Rel. Distance", "Cluster", and "Distance" algorithms is 0.50, 1.98, 1.20 and 4.03, respectively. This value indicates that the "Distance" method is eight times as effective as the "Naïve" method.

Due to the simplicity of the "Iceland" case, the "Naïve" and "Clustering" initial conditions were the only ones that did not produce the global optimum immediately. "Naïve" results are presented above, and for the "Clustering" a maximum AE of only 0.56 was obtained, obviating any significant advantage of clustering for a set of 9.

The last set of experiments on the "Iceland" case focused on the utility of the four node-node probability schemes described earlier (Table 2) when deployed for inversion locations in addition to determining the initial genome (Table 3).

Table 3. Results for the "Iceland" case using different node-node transition methods for determining both the initial gene set and the inversion loci. "Rel. Distance" and "Distance" groups are statistically significantly different from all other groups in comparing (μ +/- σ_μ) and (N_i +/- σ_N). C_{min} = 0.00 for all sets of runs, and so C_T=C_{max}

	Naïve	Rel. Dist	Cluster	Distance
C_{max} (%)	8.46	8.16	8.51	4.74
μ (%)	**0.261**	**0.121**	**0.329**	**0.0066**
σ_μ (%)	0.058	0.010	0.040	0.0056
N_i	**8.167**	**5.371**	**6.196**	**3.056**
σ_N	0.052	0.068	0.130	0.138

The "Cluster" method (AE=0.49) applied to initial gene set and crossover provides no improvement in mean optimal cost compared to "Naïve" (AE=0.47 or 0.50 in Tables 2 and 3). For "Rel. Distance", AE = 1.54, slightly worse than when it is used in determining the initial genome only. However, the AE for the "Distance" technique is 49.6, or 100 times the "Naïve" value (and 25 times the value obtained when "Distance" is used for the initial gene set only). Thus, the "Distance" technique is the most effective of the techniques investigated. Because σ_μ and σ_N are relatively small for all of the other cases investigated, they will be shown no further. In Tables 2-5, C_{min} = 0.00 for all sets of runs, and so C_T=C_{max}.

3.2 USA Case

For relative distance weighting of the "USA" case, the step size ΔP=0.05 and P(Min)=0.00763, so P(Max)=0.09763. The clustering technique pooled the nodes (Los Angeles, San Diego), (Dallas, Houston, San Antonio), (Atlanta, Charlotte) and

(Boston, New York, Philadelphia and Washington D.C.), thus excluding the other 9 cities (a relative cluster-to-noncluster weighting of 25X is reported in Table 4). The values for the more complicated "USA" are higher for C_{max}, C_r, μ and N_i than for "Iceland" (Table 4). As a consequence, for AE (Eq. 4) in the "USA" case, k was set to 1000. For the "Naïve" initial gene set assignment, AE = 1.58. "Rel. Distance" AE = 1.66, "Cluster" AE = 1.60, and "Distance" AE = 1.62. These values are similar, and indicate that for more complex cases, initial genome assignment by these methods does not significantly affect AE.

Table 4. Results for the "USA" case using different node-node transition methods for determining the initial genome. All groups are statistically significantly different in comparing (μ +/- σ_μ) and (N_i +/- σ_N).

	Naïve	Rel. Dist	Cluster	Distance
C_{max} (%)	33.00	29.08	28.63	28.51
μ (%)	**9.866**	**8.094**	**7.108**	**6.234**
N_i	**63.99**	**74.57**	**87.69**	**99.04**

The effect of initial conditions ("Naïve", "Weighted Nearest Neighbor", "Lowest Remaining Distance", "Centroid + Clockwise / Counterclockwise Traversal" and "Clustering") were considered. When the "Lowest Remaining Distance" initial condition was used, the values for (μ +/- σ_μ) and (N_i +/- σ_N) obtained were (6.223+/-0.095%) and (99.61+/-3.76), respectively, yielding an AE of 1.61. However, the first noticeable improvement in algorithmic efficacy for the "USA" case occurred when the "Centroid + Clockwise / Counterclockwise Traversal" initial condition was used for 10% of the initial gene set. When this pathway (which has a supra-optimal traversal distance of 11100 miles, or +19.7%) was used together with "Naïve" node-node probabilities, the values for (μ +/- σ_μ) and (N_i +/- σ_N) obtained were (3.107+/-0.089%) and (52.45+/-2.73), respectively, yielding an AE of 6.14, or a four-fold improvement over the previous methods. In the next set of runs, this pathway was used together with "Distance" node-node probabilities, and the values for (μ +/- σ_μ) and (N_i +/- σ_N) obtained were (2.270+/-0.016%) and (81.69+/-0.97), respectively, yielding an AE of 5.39, with a significantly lower μ than for any previous methods.

For inversion loci determination, the same set as described in Table 3 was used (Table 5). Note that the first columns in Tables 4 and 5 (as for Tables 2 and 3) should be statistically similar/equivalent (and are), since they represent the same "Naïve" protocols. Unlike the results in Table 4, the results in Table 5 show compelling changes in μ without offsetting increases in N_i. Thus, AE (with k=1000) is 1.64, 3.99, 4.63 and 11.87, respectively, for the four columns, and the "error" μ improves significantly by 2.46, 2.81 and 6.48 times for the last three columns, respectively, when compared to the "Naïve" results.

Table 5. Results for the "USA" case using different node-node transition methods for determining both the initial genome and the inversion loci. All groups are statistically significantly different from all other groups in comparing (μ +/- σ_μ); and the "Distance" group is statistically significantly lower than the other groups for (N_i +/- σ_N).

	Naïve	Rel. Dist	Cluster	Distance
C_{max} (%)	52.00	23.55	17.58	11.50
μ (%)	**9.557**	**3.886**	**3.407**	**1.475**
N_i	**63.97**	**64.52**	**63.46**	**57.11**

The "USA" case was also investigated for its convergence sensitivity to mutation rate and to survival of the "best" initial state on the value for AE (using the "Naïve" method for node-node transitions). Mutation rate was varied from 2% to 40%, with AE very stable at 2%, 3%, 4% and 5% (AE=1.60-1.63 for runs of 10,000 in this range). Above 6%, the mutation rate deleteriously affected AE, reducing it to 1.42 at 8%, 1.00 at 20% and 0.58 at 40%. From this it was validated that a low mutation rate (2-6%) is appropriate for a task the size of the "USA" problem. When the "Distance" method for node-node transition is used, however, even modestly increased rates of mutation are more deleterious. When the mutation rate is 2%, AE varies from 11.27-11.87; when it is increased to 4%, AE varies from 7.31-7.63, a decrease of more than 35% (mean error also increased from approximately 1.5 to 1.85, a 23% increase). Thus, mutation rate should be set lower for "Distance" than other methods.

Selecting the 4% mutation rate (mid-optimal band), the "Centroid + Clockwise Traversal" initial state was then preserved for from 2%-100% of the initial states while using the "Distance" method for node-node transitions and the rest of the initial genome determination (3 separate sets of 10,000 runs were performed at each percentage). For this set of conditions, a very high AE was obtained for the "USA" set—nearly twice as high as the 11.87 value for the 0% preservation of this initial state shown in Table 5. For 2%, 4%, 6%, 8% and 10% preservation, AE values were 20.3, 20.8, 21.5, 22.3 and 23.3, respectively. The value for AE continued to climb until approximately 30% preservation of the original state, after which values from 28-31 were obtained. This series of tests emphasizes that, if possible to estimate effectively, initial state can significantly enhance even the "Distance" method for initial genome and node-node transitions thereafter.

3.3 Europe and Grid Cases

The "Europe" and "Grid" cases were tested on multiple runs of 100,000 (Table 6), comparing the naïve GA to the use of distance w. For the "Europe" case, the "Distance" technique provided a 3.5-fold improvement in mean error (μ), a moderate in iterations (N_i), and thus the ratio of AE("Distance")/AE("Naïve")=3.037. For the "Grid" case, the value of μ improved by a factor of 6.09, the value of N_i decreased by 12.1%, and thus the ratio of AE("Distance")/AE("Naïve")=6.83.

Table 6. Results for the "Europe" [E] and "Grid" [G] cases using "Naïve" or "Distance" node-node transition methods for determining the initial genome. Paired groups are significantly different ($p<.01$) comparing (μ +/- σ_μ) and (N_i +/- σ_N)

	[E]-Naïve	[E]-Distance	[G]-Naïve	[G]-Distance
C_{min} (%)	1.52	0.0	0.0	0.0
C_{max} (%)	54.79	29.56	37.5	25.0
μ (%)	**26.6**	**7.6**	**11.18**	**1.94**
N_i	**202.7**	**233.7**	**28.2**	**25.1**

The "Europe" Naïve case did not reach the minimum value in 800,000 runs of the GA (C_{min} = % in Table 6), while the "Distance" case reached C_{min} 22 times in 800,000 runs. This corresponds to 200*50*800,000, or $8x10^9$, genes evaluated, a large number but still far less than the $4.42088x10^{30}$ ratio of super-optimal to optimal paths. The extraordinary power of the "Distance" GA is evidenced by the ratio of $(22x4.42088x10^{30})/(8x10^9)$, which means the GA was $1.2x10^{22}$ times more effective than random pathway assignment at finding an optimal pathway for the "Europe" case.

3.4 TSPLIB Cases

Table 7. Results for the symmetrical (upper two data rows) and asymmetrical (lower 4 data rows) TSP problems [11].

	μ-Naïve	N_i-Naïve	μ-Distance	N_i-Distance
ei151	56.2%	360.7	24.6%	423.7
ei1101	102.9%	1323.6	53.4%	999.9
br17	43.7%	39.2	31.2%	79.4
ft70	27.3%	166.5	24.7%	1636.3
p43	4.39%	532.3	1.65%	1109.6
ry48p	49.6%	310.5	24.4%	466.4

We have also tested large sets (1000 runs each) of the "Naïve" and "Distance" cases on 2 symmetrical (ei151 and ei1101) and 4 asymmetrical (br17, ft70, p43 and ry48p) published TSPs [11]. The symmetrical problems show results (Table 7) consistent with those of our large novel cases: significant (2X) reduction in mean error using the "Distance" method compared to the "Naïve" method, with a modest change one way or the other in iterations. The asymmetrical problems, however, show less predictability: mean error drops by a mere 9.5% for ft70; 28.6% for br17; 50.8% for ry48p; and 62.4% for p43. Number of iterations (N_i) to converge, however, increases when using the "Distance" method compared to the "Naïve" method, such that for br17 and ft70, the value for AE actually increases. The poorer relative performance of the "Distance" method on the asymmetrical problems is likely due in part to the fact that in-

version of sub-paths may result in unpredictable changes in distance, removing some of the effectiveness of the “Distance” method.

4 Summary and Future Work

This paper introduces a metric useful for making comparisons between different GA strategies: “Algorithmic Efficacy”, AE, which incorporates both the (residual) error and the number of iterations required for convergence. For all of the cases presented herein, the use of “Distance” weighting of node-node transitions for both initial gene sequencing and for dictating inversion splicing locations results in compelling improvements in “error” and lesser so iterations. Initial conditions can also significantly improve the value of AE. The relative value of the other (simpler) distance-based methods depends on the nature of the case: the “Iceland” case, for example, is optimized by a clockwise traversal around the centroid while the other cases are not. From these data, it appears that the inversion distance weighting is a significant means of optimizing both the initial genome and the inversion decisions. Mean error μ is improved by 3.5 to 37.4 times for our novel cases, decreasing as the complexity of the case increases, and AE is improved by from 3 to 99 times for the “Distance” compared to “Naïve” method. In the most complex case, “Europe”, the “Naïve” method did not reach the minimum value in 800,000 runs, while the “Distance” method reached the minimum value 22 times. Clearly, if minimizing the cost of the path is important, the “Distance” method for both initial conditions and node-node transitions is preferable to any other methods investigated here. The modest increase in Ni observed for the “Europe” case need be tempered by the fact that after an equal number of runs (203) the value for μ was improved by 3.4 for the “Distance” compared to “Naïve” method—in other words, the additional 31 runs contributed only another 3% to the improvement in μ of the “Distance” method.

On the standard symmetrical and asymmetrical problems we explored [11], the inverse distance weighting generally reduced mean error appreciably, except in an extremely asymmetrical case (ft70), in which the node-node distances were highly bimodal, and so “naïve” methods were roughly equally effective. AE actually increased for inverse distance weighting techniques on several asymmetric problems (br17, ft70), but for similar reasons to those cited above for the “Europe” case: limiting the “Distance” method to the same number of runs as the “Naïve” method, in each case, resulted in less than a 5% change in mean error, μ.

The results are applicable to many broad areas of navigation: visiting a set of destinations, finding best paths when a particular destination must be replaced or a particular node-node transition is missing (e.g. detours, closed roads, traffic jams, etc.), and even extra-navigational process optimizations. Detours can be represented in at least two ways: (1) implicitly as prohibitively high node-node costs (or, equivalently, zero-valued node-node probabilities), and (2) explicitly as “missing” transitions.

Gene linkage and related co-inheritance techniques [4-7] can be viewed as similar to the clustering techniques described herein where $p(C_i,N_j) = 0.0$ for all nodes i in the cluster C and all nodes j in the nonclustered space S = N-C (N is the set of all nodes). However, the work presented here automates the linkage of loci/genes in a

non-binary and adaptive fashion, and its performance is significantly enhanced compared to clustering schemes. Loci pairs with lower costs of transition are innately linked to a greater degree than loci pairs with higher transition costs. In other words, the "gene linking" comes for free.

Future work includes the need for separate testing of the effect of inversion rate, mutation rate, and more advanced "jiggling" strategies and cost functions. Conditional (Bayesian) methods of assigning initial state and crossover splicing locations were introduced but not evaluated here. It is likely that in many navigational problems error and convergence rate improvements can be garnered applying these methods to the initial state, crossover, and mutation parameters. Rule-based direction of initial state and selection (e.g. spectacularly unlikely genes are discarded) should also be evaluated. Finally, since the technique described here is innately related to cost optimization, a more mature cost representation than simple node-node transition weighting is appropriate. When non-roundtrip TSP pathways are legitimate, the nodes themselves must be weighted, and cost optimization will involve several (often competing) factors: node value, node-node transition cost, replacement node value, multi-node value schema (i.e. where visiting a particular set of nodes sequentially or during the entire pathway increases their summed value), and possible additional value for re-visiting a node, among others.

References

1. Rao, R.P.N., Fuentes, O.: Hierarchical Learning of Navigational Behaviors in an Autonomous Robot Using a Predictive Sparse Distributed Memory. Autonomous Robots 5:3-4 (1998) 297-318 and Machine Learning 31:1-3 (1998) 87-113
2. Goldberg, D.E.: Genetic Algorithms in Search, Optimization, and Machine Learning. Addison-Wesley Publishing Company, Inc., Reading, Mass. (1989)
3. Laane, L.A.: Development of Navigational Controllers for Vehicles in Highway Traffic Situations via Genetic Programming. In Koza, J.R. (ed.): Genetic Algorithms and Genetic Programming at Stanford 1995, Stanford Bookstore (1995) 171-180
4. Smith, J.: On Appropriate Adaptation Levels for the Learning of Gene Linkage. Genetic Prog. Evolvable Machines 3 (2002) 129-155
5. Salman, A., Mehrota, K., Mohan, C.: Linkage Crossover for Genetic Algorithms. In Banzhaf, W., et al. (eds.): Proc. Genetic Evolutionary Comp. Conf. (1999) 564-571
6. Schaffer, J., Morishima, A.: An Adaptive Crossover Distributed Mechanism for Genetic Algorithms. In Grefenstette, J.J., et al. (eds.): Proc. Second Internat. Conf. Genetic Algorithms (1987) 36-40
7. Schaffer, J.D., Eshelmann, L.J.: On Crossover as an Evolutionary Viable Strategy. In Belew, R., Booker, L. (eds.): Proc. Fourth Internet. Conf. Genetic Algorithms (1991) 61-68
8. Goldberg, D.E., Korb, B., Deb, K.: Messy Genetic Algorithms: Motivation, Analysis and First Results. Complex Systems 3:5 (1989) 493-530.
9. Beyer, H.: Toward a Theory of Evolution Strategies: On the Benefit of Sex-the (μ/μ-λ)-Strategy. Evolutionary Computation 3 (1995) 81-111
10. Ísland Fer∂akort, Mál og menning, su∂urlandsbraut 12, IS-108 Reykjavik, http://www.edda.is (2002)
11. TSPLIB Home Page, library of sample instances for the TSP from various sources and of various types, http://www.iwr.uni-heidelberg.de/groups/comopt/software/TSPLIB95/ (2004)