An investigation of redundant genotype-phenotype mappings and their role in evolutionary search

Mark Shackleton
Admin 2-5
BT Labs, Adastral Park
Martlesham Heath
Ipswich, IP5 3RE, UK
mark.shackleton@bt.com

Rob Shipman
Admin 2-5
BT Labs, Adastral Park
Martlesham Heath
Ipswich, IP5 3RE, UK
rob.shipman@bt.com

Marc Ebner
Universität Würzburg
Lehrstuhl für Informatik II
Am Hubland
97074 Würzburg, Germany
ebner@informatik.uni-wuerzburg.de

Abstract- The neutral theory of evolution suggests that most mutations do not cause a phenotypic change. In this case the mapping from genotype to phenotype contains redundancy such that many mutations do not have an appreciable effect on the phenotype. This can result in neutral networks; sets of genotypes connected by single point mutations that map to the same phenotype. A population is able to drift along these networks, eventually encountering phenotypes of higher fitness, thus reducing the chance of becoming trapped in sub-optimal regions of genotype space. In this paper we explore the use and benefit of redundant mappings for evolutionary search. We investigate the properties of several genotype-phenotype mappings by performing random walks along the neutral networks in their genotype spaces. The properties are explored further by performing adaptive walks in which a concept of fitness is introduced. A mapping based on a random boolean network was found to have particularly interesting properties in both cases.

1 Introduction

Natural evolution differs in many respects from the evolutionary algorithms typically employed today. One such difference is highlighted by the neutral theory of evolution. According to this theory a considerable fraction of all mutations are phenotypically neutral with only a minute fraction of non-neutral mutations being beneficial [7]. However, most evolutionary algorithms use a fixed one-to-one mapping between genotype and phenotype with each genotype corresponding to exactly one phenotype. In this paper we explore the impact on evolutionary search of redundant genotype-phenotype mappings (in which each phenotype is represented by many genotypes). Such mappings offer the potential of increasing the efficacy of evolutionary search by allowing continued search on neutral networks.

If one considers an evolutionary search in which the only genetic operator available is single point mutation, then the number of genotypes directly reachable from a given genotype is $(A - 1)L$ where A is the number of alleles available and L is the length of that genotype [6]. For binary string genotypes the number of neighbors is simply L, the length of the genotype. Assuming a unique static fitness value is associated with every possible genotype, an individual genotype which represents a local optimum (with respect to these neighbors) would by definition be "trapped". That is, for the search to reach better fitness values the individual genotype would have to forego its current fitness value for a lesser fitness, in the hope of eventually locating a better fitness value. This is the case when there is no redundancy in the genotype-phenotype mapping. (Note that for the purposes of this discussion it is assumed that every distinct phenotype has a unique associated fitness.)

When the genotype-phenotype mapping exhibits redundancy there is the potential to continue the search by drifting along *neutral networks* in genotype space. If none of the point mutation neighbors is fitter than the current genotype, yet there exists a neighboring genotype that maps to the *same phenotype*, then it is possible to adopt that neighbor as the current search solution *without loss of fitness*. Provided such phenotypically neutral neighbors exist and form *extensive networks* throughout genotype space, they may offer a way of traversing genotype space without becoming "stuck" at local optima. There would be no need to pass through regions of lower fitness in order to reach regions of higher fitness.

The possible role of "neutral evolution" is discussed by Huynen [4] in the context of RNA sequence-structure mappings. This work showed that RNA folding algorithms predict the existence of extensive neutral networks in sequence (genotype) space for which the constituent sequences all map to an identical structure (phenotype). Further work by Huynen [5] considered aspects of a population searching such a genotype space. Fontana and Schuster [3] showed that the statistical topology organising the set of RNA shapes explains why neutral drift in sequence space is important in evolutionary search. Other work which has indicated the importance of such properties of the search space includes that of Ebner [1] and Shipman [8]. In this paper we consider redundant genotype-phenotype mappings which are far simpler than the RNA sequence-structure mapping but which might nevertheless exhibit extensive neutral networks which are helpful to an evolutionary search process. For comparison we also examine mappings which exhibit redundancy of a type which is *not* helpful for evolutionary search.

2 Genotype-phenotype mappings

A number of genotype-phenotype mappings are discussed below which exhibit varying degrees of redundancy, from none through to a high level of redundancy i.e. with many genotypes mapping to each possible phenotype. To allow the different aspects of each mapping to be explored, a single phenotypic representation is used consisting of 16 bits, resulting in 2^{16} = 65536 possible unique phenotypes. Other work (see Ebner [2] and Shipman [9]) has investigated smaller phenotype spaces exhaustively; here we investigate whether useful properties of redundancy extend to somewhat larger phenotype spaces.

2.1 Direct Binary Mapping

The Direct Binary one-to-one mapping maps a 16 bit genotype directly onto the 16 bit phenotype representation. As a result there is no redundancy in this mapping - exactly one genotype maps to each possible phenotype. This mapping provides a baseline for comparison with the other mappings.

2.2 Static Random Mapping

The Static Random mapping does exhibit redundancy. The genotype consists of a bit string of length 30, resulting in a genotype space of 2^{30} genomes, each mapping to a specific (16 bit) phenotype. The mapping is randomly initialised and remains static, with the same genotype always mapping to its corresponding phenotype. As the number of genomes is too large to conveniently store in a look-up table, this is accomplished by hashing into a fixed random number generator sequence. The degree of redundancy is quite high with on average 2^{14} genomes mapping to each phenotype.

2.3 Trivial Voting Mapping

The Trivial Voting mapping is redundant, and is illustrated in Figure 1 (top). The genotype has a total length of 48 bits. The first 3 genotype bits determine the value of the first phenotype bit, and consecutive sets of 3 genotype bits determine consecutive individual bits of the phenotype. Each phenotype bit is thus uniquely determined by a set of 3 genotype bits, with each of these sets being distinct (i.e. no two sets share any bits). A phenotype bit is set to unity if the *majority* of the 3 corresponding genotype bits are set, otherwise it is cleared to zero. Thus the genotype bits *vote* for the setting of the corresponding phenotype bit. This mapping exhibits a redundancy of 2^{32}:1.

2.4 Standard Voting Mapping

The Standard Voting mapping is again based on a voting approach where each bit of the phenotype is influenced by several bits from the genotype. (See Figure 1, bottom.) Each phenotype bit is determined by looking at all the bits of the genotype to which it is linked. A bit of the phenotype is set if

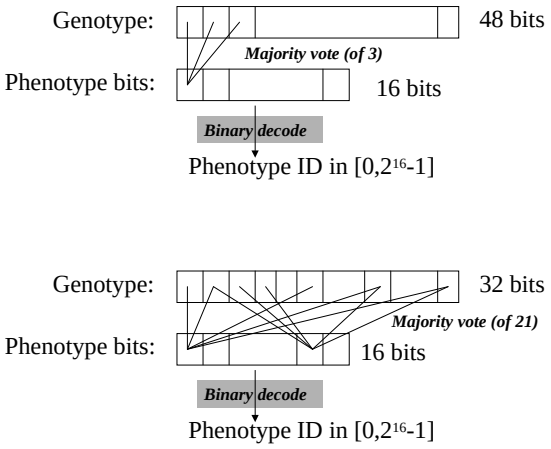

Figure 1: *Top:* Trivial voting mapping - each phenotype bit is uniquely determined by 3 genotype bits; if the majority of these bits are set then the corresponding phenotype bit will also be set. There is no overlap between the consecutive sets of 3 bits. *Bottom:* Standard voting mapping - each phenotype bit depends on a constant odd number of genotype bits which vote either for or against the corresponding phenotype bit being set. For this mapping there is considerable overlap between the sets of genotype bits.

the majority of connected bits in the genotype "vote" in favor of this. Thus, depending on the values of the other relevant bits, a point mutation may or may not have an effect on the phenotype. It is important to note that for the Standard Voting Mapping the set of genotype bits linked to a particular phenotype bit *will typically overlap* with the sets corresponding to other phenotype bits. This is the key distinction between this mapping and the *Trivial Voting* mapping. It is this aspect that permits multiple phenotype bits to potentially be changed simultaneously by a single point mutation. Alternatively, the redundant "majority voting" aspect of the mapping can sometimes result in the phenotype remaining unchanged, while still setting the scene for future transitions to different phenotypes.

The links between the genotype bits and the phenotype bits are determined in the following way. For each bit of the phenotype we randomly select a constant number of bits of the genotype which will vote for that phenotype bit. For each of the voting bits, we randomly choose whether a set bit will vote in favor of the corresponding phenotype bit being set, or against it being set. Thus there are positive and negative votes (links). The number of genotype bits which vote for each phenotype bit is fixed at a constant odd number. In the results reported later, a genotype of 32 bits was used with sets of 21 genotype bits being chosen for each of the 16 phenotype bits. Clearly there is significant overlap among the sets of voting bits.

2.5 Cellular Automaton Mapping

The Cellular Automaton mapping uses a developmental approach: the genotype encodes a set of cellular automaton rules which map a fixed initial value into the resulting phenotype value. (See Wolfram [10] for example, for more information on cellular automata.) A one-dimensional, non-uniform cellular automaton comprises an array of cells, each with an associated rule table which specifies how the state of that cell changes over time (see Figure 2). The cellular automaton consists of an array of 16 binary cells. For each cell, a rule table specifies the subsequent state of that cell, uniquely determined by its own state together with the state of its left and right adjacent cells. (Left and right extremes of the array are wrapped around.) The 3 input states result in $2^3 = 8$ possibilities, so each rule table has 8 entries, each of which specifies the new binary state resulting for that configuration. The genotype thus requires 8 bits to encode each rule table. This results in a total genotype length of 128 bits.

To perform the mapping from genotype to phenotype, the cellular automaton is first initialised with a fixed, arbitrary bit string. The cellular automaton is then iterated for a constant number of steps using the rule tables defined in the genotype. (For the results reported later, 20 iterations were used, but experiments suggested that increasing the number of iterations further had little effect on overall behavior.)

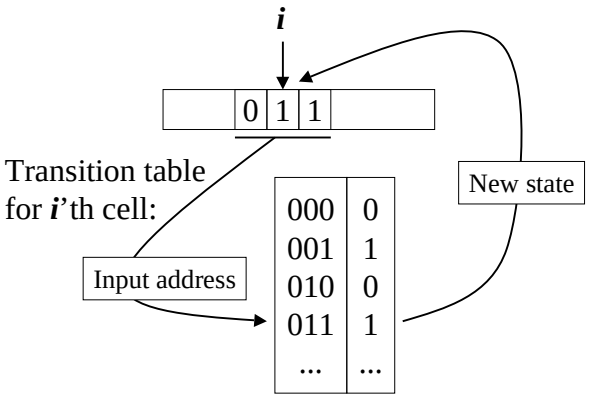

Figure 2: Cellular automata mapping: the phenotype consists of the final state of a cellular automaton whose state transition rules are encoded in the genotype. There is one cell for each of the 16 phenotype bits, and every cell has an associated transition table mapping the current bit state and neighbor states to a new state for that cell.

2.6 Random Boolean Network Mapping

The Random Boolean Network mapping is similar to the Cellular Automata mapping, but is rather more complex. The genotype again encodes rule tables describing how the state of each cell (in an array of 16 cells) changes over time. These rules are applied iteratively, as for the cellular automata, to derive the final phenotype value. For the Random Boolean Network used here the next state of each cell depends on 3 input bits, but these can be *any* 3 bits at *defined addresses* within the 16 bit cellular array. Thus in addition to the rule table, these addresses must also be specified: each address requires 4 bits to encode an address in the range *[0, 15]*. A further distinction from the cellular automaton defined above is that the initial state for each binary cell is directly encoded in the genotype. To summarise, the details for each cell include:

- initial state (1 bit);
- addresses for each input cell ($3 \cdot 4 = 12$ bits);
- the rule table defining the cell's next state ($2^3 = 8$ bits).

The above information is required for each of the 16 cells, which results in a total genotype length of 336 bits. The mapping from genotype to phenotype thus exhibits a high degree of redundancy.

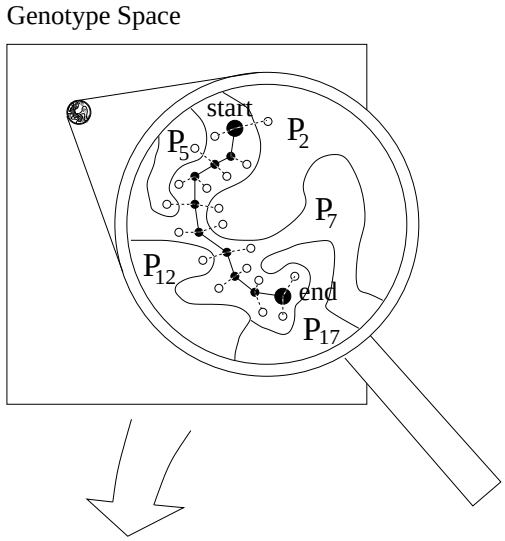

Different phenotypes encountered along random neutral walk:

P_2 P_5 P_{12}

Set of 10 equivalent genotypes found:

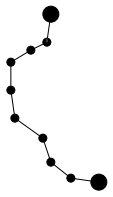

Figure 3: Random neutral walk through genotype space: the walk starts at a given genotype (and associated phenotype, P_7 in the figure above) and proceeds by randomly selecting successive point mutation neighbors which map to the same phenotype. During this walk different phenotypes are *encountered* which are *reachable by single mutations* of the genotypes along the walk.

3 Evaluation of the mappings on a random neutral walk

To evaluate various properties of the mappings we first performed an analysis of each mapping in the context of a " random neutral walk" through genotype space, as we are particularly interested in the ability to continue to explore genotype space without changing the current phenotype. (Later, in section 4, we discuss a fitness-adaptive walk.)

In this paper we report only those results which bear most strongly on later discussions of fitness-adaptive walks; for a more detailed statistical examination of these (and other) redundant mappings in the context of a random neutral walk we refer the reader to Shipman et al. [9]. Ebner et al. [2] considers the impact of redundant mappings on a population searching a space in the context of a changing environment. Here we consider a single point search to avoid the additional complications resulting from a population-based search.

3.1 Random neutral walk through genotype space

The following statistics were calculated in the context of a random neutral walk (Figure 3). Each random neutral walk starts from a given genotype and randomly selects one of the single mutation neighbors which maps to the *same phenotype*. The process is repeated for a fixed number of steps, each involving a transition in genotype space to a phenotype-neutral, adjacent neighbor. (If no such neighbor exists the walk will remain at the current position in genotype space.) The statistics reported here were all calculated for a series of 1024 random neutral walks, each starting at a genotype which was chosen randomly from the set of all possible genotypes.

3.2 Phenotypes reachable

The "phenotypes reachable" statistic measures the cumulative number of phenotypes encountered, averaged across a series of random neutral walks, and is plotted against walk length. A new phenotype is *encountered* when one of the single point mutant neighbors of the current genotype is found to map to that phenotype. (See Figure 3.)

The plotted value represents the mean over the entire series of random neutral walks. The number of phenotypes reachable via a neutral walk is an important statistic, as the slope gives an indication of the "innovation rate" i.e. the ability to find potentially better phenotypes while still maintaining the current phenotype. When we later associate a fitness value with each phenotype, the slope of this plot can indicate whether the search for fitter phenotypes should be continued.

3.3 Histogram of number of phenotypes encountered

The histogram of the number of phenotypes encountered indicates how many different phenotypes we can expect to encounter (as single point mutation neighbors) during a random neutral walk. The previous statistic, "phenotypes reachable", captures the mean of this distribution at the end of the neutral walk. In addition, this histogram indicates both the variance in number of phenotypes found over all walks and can also illustrate other biases (such as bimodality) in the distribution.

3.4 Results of neutral walks for each mapping

The above statistics were calculated for all of the mappings described earlier. The statistics were calculated over 1024 walks of length 500, each walk starting from a randomly selected genotype. The generated statistics are shown for the Standard Voting mapping (Figure 4), the Cellular Automata (CA) mapping (Figure 5), and the Random Boolean Network (RBN) mapping (Figure 6). Results from the other mappings are also discussed, but are not shown due to space limitations and because they can be simply described.

The Direct Binary mapping contains no redundancy, and every genotype maps to exactly one phenotype; consequently the number of phenotypes encountered during all "neutral" walks is equal to the number of single point mutation neighbors, namely 16.

The Static Random mapping contains a degree of redundancy, with the set of (approximately 2^{14}) genotypes which map to a given phenotype being scattered widely through genotype space. However, once again the number of phenotypes encountered is restricted to the number of single point mutation neighbors; this is because the set of genotypes corresponding to a given phenotype are typically *not connected by a neutral network* of single point mutations i.e. they are isolated in genotype space and cannot be reached one from another without foregoing the current phenotype. At higher levels of redundancy we might however expect neutral neighbors to exist. A Static Random mapping with a genotype length of 336 (equal to the largest genotype considered in this paper) was also investigated, but no neutral networks were observed.

The Trivial Voting mapping also exhibits redundancy, with on average 2^{32} genotypes corresponding to each unique phenotype. The mapping produces results very close to those of the Direct Binary mapping, encountering only 16 different phenotypes, with the only difference being that it takes several more steps to find these. In this case the redundancy simply slows down the search process, without allowing extensive search of genotype space. (A Trivial Voting mapping with a genotype of length 336 was also investigated and found to slow down the search even further.) Consideration of how the mapping is defined (Figure 1, top) indicates why this is the case: each genotype bit can only ever affect a single bit of the phenotype representation; thus the possible phenotype transitions resulting from a mutation are the same as those of the Direct Binary mapping. The Trivial Voting mapping illustrates that *redundancy alone is not enough* to help evolutionary search.

The Standard Voting mapping (Figure 1, bottom) was designed such that a single point mutation in the genotype could potentially affect multiple bits in the phenotypic representation simultaneously, or in some cases (due to the voting mech-

anism) may affect none of these bits. The results of the random neutral walks using this mapping are shown in Figure 4. The mean number of unique phenotypes encountered is 669, after a walk of 500 steps. This indicates that the mapping permits extensive phenotype-neutral walks through genotype space. The histogram illustrates that the walk length distribution is bimodal: the majority of walks encounter approximately 750 phenotypes but some walks encounter very few phenotypes. The latter observation indicates a weakness in the properties of the mapping which could lead to fitness-adaptive walks becoming "stuck" at locally fit phenotypes which lack an associated extensive neutral network.

The results for the Cellular Automata mapping are shown in Figure 5. The mean number of unique phenotypes encountered at the end of the neutral walk is 472. This is less than for the Standard Voting mapping but the gradient of the "phenotypes encountered" graph indicates that new phenotypes continue to be encountered and that extensive neutral networks exist in genotype space. In this case the histogram is unimodal, indicating that all walks sampled resulted in similar numbers of phenotypes being encountered. This property is more promising than that of the Standard Voting mapping as it implies that a fitness-adaptive walk is less likely to become trapped at a local optimum, in a restricted region of genotype space.

Results for the Random Boolean Network (RBN) mapping are shown in Figure 6. These statistics are particularly encouraging: the final mean count of phenotypes encountered over the random neutral walk is 4614, an order of magnitude greater than the maximum across the other mappings considered here. The gradient of the associated graph is also high, with little sign of an asymptote at the end of the walk. In addition, the histogram of number of phenotypes encountered is unimodal with no low values, suggesting the presence of many extensive neutral networks in genotype space.

4 Evaluation of the mappings on a fitness-adaptive walk

In this section we present results of using the redundant mappings in the context of a fitness-adaptive walk in genotype space. The statistics were calculated over 1024 walks of length 500 starting at randomly selected genotypes, as before. The walk conditions are as for the neutral walk, but with the key difference that a fitness value is randomly associated with each phenotype. The walk then proceeds by adaptive moves to the single-point mutation neighbor with the highest fitness, provided this fitness is higher than the current fitness. If such a neighbor does not exist, then the walk proceeds by a phenotype-neutral move to a single-point mutation neighbor with an equal fitness to the current fitness. If this is not possible either, then the walk ends.

Fitnesses are assigned to phenotypes randomly in the range *[0,1]* with higher values denoting higher fitness. However, the fitness assignment is not uniform; instead a uniform

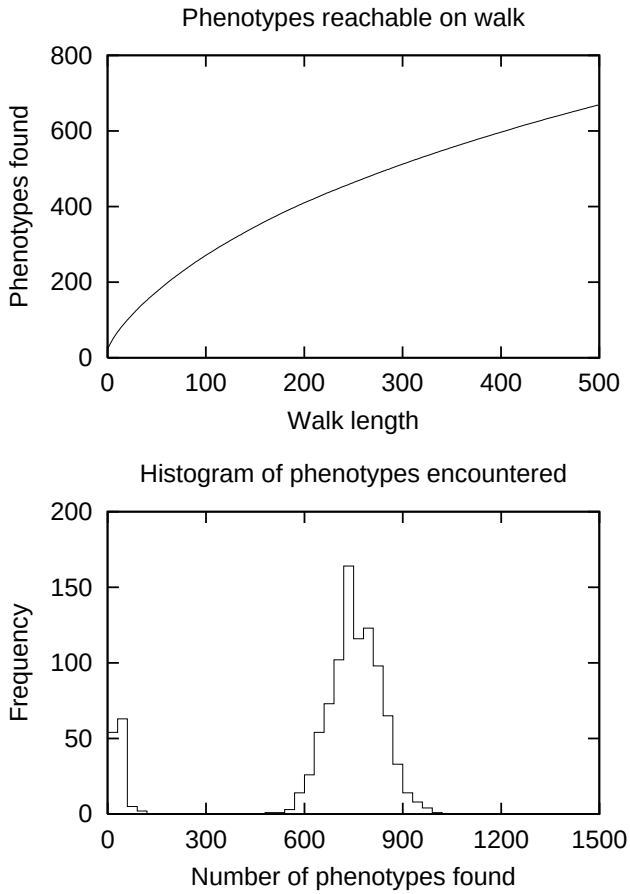

Figure 4: Standard voting mapping statistics: The graph shows the number of phenotypes found, averaged across all neutral walks. The redundancy in this mapping permits a neutral walk which continues to discover new phenotypes in its neighborhood. The eventual mean number of phenotypes found over the walks (669) is greater than for the cellular automata mapping, but the gradient (innovation rate) is gradually decreasing. The central portion of the histogram shows that most neutral walks encounter on average 750 different phenotypes, corresponding to extensive neutral networks in genotype space. The histogram also shows that some walks encounter very few phenotypes; this indicates a potential problem, as a fitness-adaptive walk could become "stuck" on a neutral network which is limited to a restricted region of genotype space.

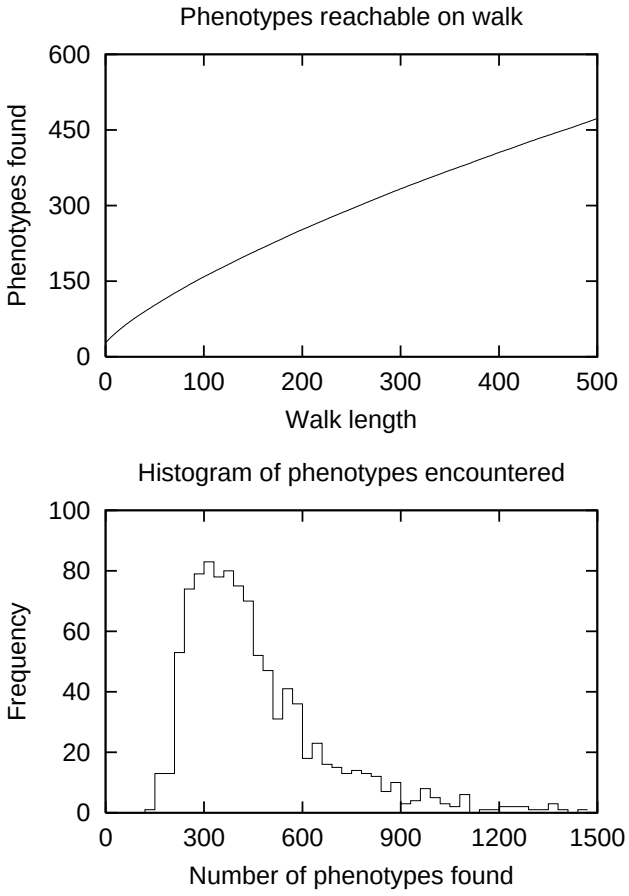

Figure 5: Cellular automata mapping statistics: the graph shows the number of different phenotypes encountered, averaged across all neutral walks. The redundancy in this mapping permits a neutral walk which continues to discover new phenotypes in its neighborhood. The eventual number of phenotypes found over the walk is less than for the standard voting mapping, but the gradient (innovation rate) has not decreased significantly indicating that new phenotypes continue to be encountered. The histogram indicates that all networks encountered are extensive, which is a desirable property, and may indicate that a fitness-adaptive walk is less likely to become "stuck" at a local optimum.

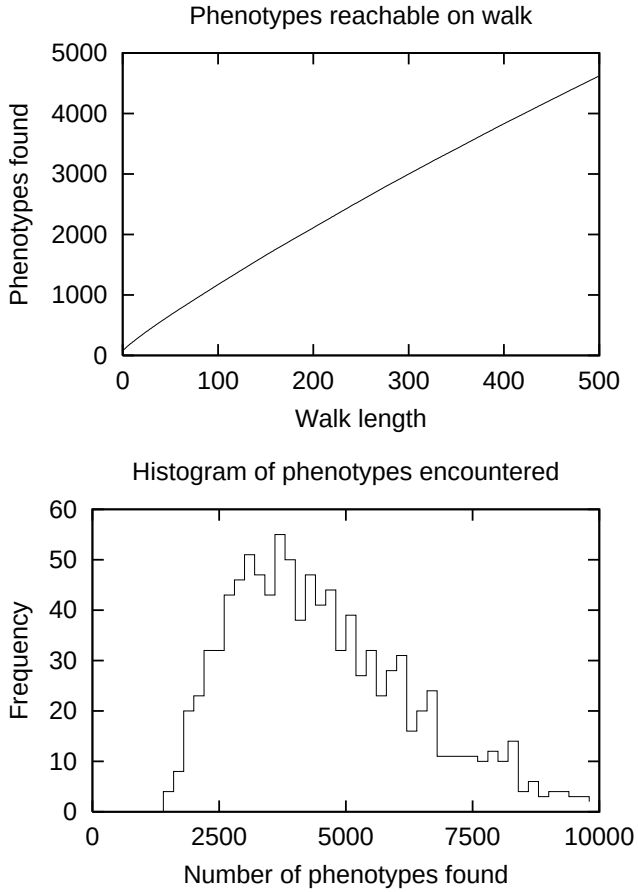

Figure 6: Random boolean network statistics: the graph shows the number of different phenotypes found, averaged across all neutral walks. The redundancy in this mapping permits a neutral walk which continues to discover new phenotypes in its neighborhood. The eventual number of phenotypes found over the walks is an order of magnitude greater than for the other mappings considered, with a steep gradient (innovation rate) showing little sign of decreasing. The histogram shows that all neutral walks encountered a substantial number of different phenotypes, indicating the presence of extensive neutral networks in genotype space.

random number *r* in *[0,1]* is remapped in order to produce many more low fitness values than high fitness values. The fitness is given by

$$f = e^{100(r-1)}. \quad (1)$$

The fitnesses are distributed in this way to generate a search space in which there are few good solutions among many poor solutions, which is a property of many search spaces of interest. It also allows greater distinction to be made between the efficacy of the different mappings and the fitness levels the evolutionary search is able to attain by searching the associated genotype spaces.

The results of the adaptive walks are shown in Figure 7. The Direct Binary one-to-one mapping has no redundancy and thus quickly becomes "trapped" in a local optimum, as indicated by the curve which rapidly flattens out at a low fitness level. The Trivial Voting mapping (not shown) performs no better than the Direct Binary mapping, rising more slowly to the same low fitness value.

The Static Random mapping behaves similarly to the Direct Binary mapping, as we would expect given the neutral walk performance reported earlier. No further fitness gains are made after the first few adaptive steps. The fitness reaches a higher level than the Direct Binary mapping, but this simply reflects the greater length of the genotype which results in more single point mutation neighbors and thus more phenotypes being sampled for higher fitness values in the few adaptive steps taken.

The Standard Voting mapping and Cellular Automata (CA) mapping do well over the entire walk and are still improving after 500 steps. The CA mapping in particular is continuing to discover new phenotypes as can be seen by the non-zero gradient. The Random Boolean Network (RBN) mapping does best of all, achieving highest fitness very early in the run. The results from the neutral walk presented earlier showed that the RBN mapping also had the greatest ongoing rate of phenotype discovery. The redundancy inherent in these mappings clearly helps prevent the adaptive search from becoming "trapped" in local optima when compared against the mapping with zero redundancy, or the "unhelpful" redundancy of the Trivial Voting mapping and the Static Random mapping. This is what we would expect, given the existence of extensive neutral networks in genotype space, and the neutral drift that was demonstrated in section 3 for these mappings.

Although the degree of redundancy alone is not sufficient to guarantee neutral networks, a minimum level of redundancy is likely to be helpful in the context of a population searching a space, in order to keep the population on the neutral network and avoid a situation where only a tiny fraction of mutants maintain fitness. The fraction of single point mutation neighbors which are phenotypically neutral for the key mappings were found to be approximately: 0.2 (Standard Voting mapping), 0.6 (CA), and 0.5 (RBN).

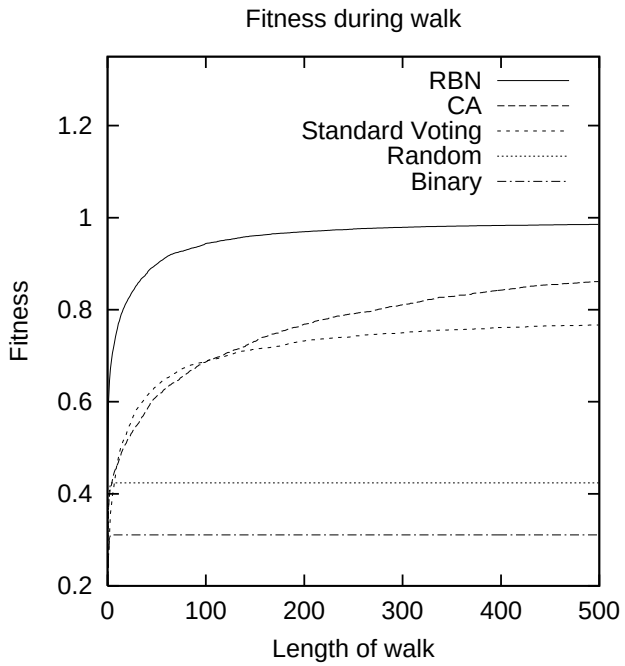

Figure 7: Fitness-adaptive walks. The graph shows the fitness levels achieved throughout adaptive walks through genotype space for each mapping, averaged over all 1024 walks. The Direct Binary one-to-one mapping clearly becomes trapped in local optima early in the adaptive walks, typically at a low fitness value. The Static Random mapping also becomes trapped after only a few adaptive steps, attaining a slightly higher fitness as a result of sampling more neighbors due to its longer genotype. The redundant mappings fare better with both the CA and Standard Voting mappings still increasing in fitness. The Random Boolean Network mapping significantly outperforms all other mappings, achieving a high fitness value early in the walk. The variance in this fitness value (not shown) is also very low, indicating consistency in locating good fitness values. Each walk consists of a combination of fitness-adaptive changes interspersed with fitness-neutral drift in genotype space. The redundancy inherent in the mappings clearly helps prevent the fitness-adaptive walks from becoming stuck in local optima. Eventual fitness values achieved are: 0.99 (RBN), 0.86 (CA), 0.77 (Standard Voting mapping), 0.42 (Static Random mapping) and 0.31 (direct Binary mapping).

5 Conclusions

In this paper we have explored the properties of several genotype-phenotype mappings with varying types, and degrees, of redundancy. Three of the redundant mappings were found to be beneficial for evolutionary search: the Standard Voting, Cellular Automata (CA) and Random Boolean Network (RBN) mappings.

These benefits were demonstrated in two ways. Firstly, random neutral walks were employed to show that neutral drift allowed for the discovery of many more, potentially better adapted, phenotypes than for a one-to-one mapping with no redundancy. Secondly, adaptive walks on fitness landscapes showed that the discovery of more phenotypes allowed for the attainment of higher fitness values. A mapping based on a random boolean network was shown to perform particularly well.

Care needs to be taken in the construction of redundant genotype-phenotype mappings. For redundancy to be of use it must allow for mutations that do not change the current phenotype, thus maintaining fitness, but which allow for moves to areas of genotype space where new phenotypes are accessible. Useless addition of redundancy was evidenced by a Trivial Voting mapping which had properties similar to a direct one-to-one mapping. In this case, no extensive exploration of genotype space via neutral mutation was possible.

This work has illustrated the potential use of redundant genotype-phenotype mappings to improve evolutionary search and prevent the search becoming trapped in local optima. Further work is required to determine whether these properties transfer to even larger phenotype spaces. Better measures are also needed to help characterise the nature of redundancy inherent in different genotype-phenotype mappings and the effect of that redundancy on evolutionary search. Such measures could be based on those developed to characterise RNA sequence-structure space properties, such as those discussed by Fontana et al. [3].

Acknowledgments

Part of this work was carried out at the Santa Fe Complex Systems Summer School 1999 by those involved in the "Redundancy" Special Interest Group. In particular we thank: Susan Ptak from Stanford University, Department of Biological Sciences and Richard Watson from Brandeis University, Volen Center for Complex Systems. We recognise their input in the form of discussions and helpful comments.

Bibliography

[1] Ebner, M. *On the search space of genetic programming and its relation to nature's search space.* In Proceedings of the 1999 Congress on Evolutionary Computation, Volume 2, Washington D.C, pp. 1357-1361. IEEE Press.

[2] Ebner, M; Shackleton, M; Shipman, R; Watson, R. *How redundant mappings influence the ability to adapt to a changing environment.* Submitted to the Journal of Evolutionary Computation, January 2000.

[3] Fontana, W; Schuster, P. *Shaping Space: the possible and the attainable in RNA genotype-phenotype mapping.* J. Theoretical Biology, 194, 491-515, 1998.

[4] Huynen, M. *Exploring phenotype space through neutral evolution.* Journal of Molecular Evolution, 43:165-169, 1996.

[5] Huynen, M; Stadler, P; Fontana, W. *Smoothness within ruggedness: the role of neutrality in adaption.* Proc. of the National Academy of Science, vol. 93, pp. 397-401, 1996.

[6] Kauffman, S. A. *The Origins of Order: self organization and selection in evolution.* Oxford University Press, 1993.

[7] Kimura, M. *The Neutral Theory of Molecular Evolution.* Cambridge University Press, Cambridge, UK, 1983.

[8] Shipman, R. *Genetic Redundancy: Desirable or problematic for evolutionary adaptation?* 4th International Conference on Artificial Neural Networks and Genetic Algorithms (ICANNGA), 1999.

[9] Shipman, R; Shackleton, M; Ebner, M; Watson, R *Neutral search spaces for artificial evolution: a lesson from life.* Submitted to Artificial Life VII, 2000.

[10] Wolfram, S. *Cellular Automata as Models of Complexity.* Nature, 311(4):419-424, October 1984.