

Open the Box

Inman Harvey
School of Cognitive and Computing Sciences
University of Sussex
Brighton BN1 9QH, UK
inmanh@cogs.susx.ac.uk
<http://www.cogs.susx.ac.uk/users/inmanh>

1 Introduction

SAGA or Species Adaptation Genetic Algorithms have been developed over the last 8 years as the modification of standard GAs necessary when one is using them not as function optimisers, but rather as incremental adaptation algorithms. This is inevitably associated with variable-length genotypes. I here give a brief background survey.

2 Why use Variable-Length Genotypes?

Many standard GA problems involve minimising/maximising a function of n variables, and for these a fixed-length direct-encoding genotype representation is usually appropriate. There are at least two possible ways in which var-len GAs can be used:

- Separate parts of the genotype directly encode separate components of a phenotype, and the latter can have a variable number of components.
- The genotype only indirectly ‘encodes’ the phenotype via some morphogenetic process, with no direct relationship between genotype length and size/complexity of phenotype.

These two classifications need not be distinct; Genetic Programming where the genotype is a concatenation of programming primitives could fit into either classification. What is common is the possibility of a phenotype of variable, potentially unlimited complexity.

This is particularly appropriate for design problems — e.g. design of robot control systems, of electronic circuits, or pharmaceutical molecules — and for incremental evolution, where rather than a single target problem there is a sequence of problems of increasing size or complexity.

3 GAs as Adaptive Improvers

GAs were not originally intended by John Holland as function optimisers, but rather — by analogy with natural evolution — as adaptive improvers. An adaptive species of organisms can be considered as ‘solving’ the problems posed by its particular environmental niche, and evolution allows the species to change over generations as and when its environmental circumstances change. In the natural world, a significant part of one species’ environment will be other species coevolving with it.

This aspect of evolution may be particularly relevant for some engineering design purposes. In practice the design requirements for the 1998 specification for artefact X will be a modification or complexification of the 1997 specification for X. This looks like coevolution between species X and the design requirements for X.

If the design requirements change at intervals of a certain length, and it takes less than this length of time to produce a solution from scratch, then each new specification can be treated as a brand new problem to optimise. However, when using artificial evolution to design complex artefacts, the bottleneck in expense and time is the evaluation of each member of the population, over successive generations. At some stage the evolution-from-scratch

time will exceed the interval between specification changes, and at this stage one cannot afford to start from scratch each time; one must turn to adaptive improvement.

4 Genetic Convergence as a Consequence

The end-point of a normal GA run is normally seen as convergence onto a (near-)optimum, and often there is much effort and worry on the topic of ‘avoiding premature convergence’. This seems to imply that there are problems continuing a GA run with a converged population. Wrong.

There is a confusion here between ‘asymptotic convergence of fitness levels to a supposed optimum’ and ‘genetic convergence such that genetic variation in the population is minimised’. These are totally different, and it can easily be demonstrated [4, 5] that for many complex real problems genetic convergence in the latter sense takes place in the first few generations, whilst fitness continues to climb despite genetic convergence. Genetic variation within the population will not drop to zero in the presence of mutation, but rather drop rapidly to a level dictated by the balance between selection and mutation pressures. This balance can be deliberately manipulated to optimal values to allow maximal continued evolution, as has been developed in SAGA [2, 3].

Proponents of ‘one-shot GAs’ — a bounded GA run searching for an optimal solution to a single problem — are not so concerned about mutation/selection balance. In a typical GP run with a population of 640,000 for 50 generations, the variety lies in the original population, and mutation may not be necessary; however even with such a large population the variation very quickly disappears. If statistics are maintained on typical GA runs then you can see that genetic convergence falls rapidly in the first few generations until the selection/mutation balance is reached.

5 Ideal Selection/Mutation Balance

Too much selection → no further exploration.

Too much mutation → loss of accumulated information in the genotypes of the population.

For normal selection pressures, the ideal mutation rate is around one mutation per functional (non-junk) parts of the genotype — see cited references for background to this. By expressing such a mutation rate as **mutations per genotype** rather than **mutations per locus**, the same mutation rate is appropriate regardless of the length of genotype. It is no coincidence that mutation rates in natural evolution have themselves evolved to be typically around this level — independent of length of genotype, from bacteria of order 10^3 to humans of order 10^9 .

To fit this into a picture of a fitness landscape: a genetically converged population will be sitting on and around a local hill, where selection tends to focus it on the hilltop, and mutation spreads the population down the hillsides into the valleys. The optimal rate retains a hold on the current hilltop, while maximising search in particular along high ridges which may lead to even higher hills.

6 Recombination with Variable Lengths

Recombination is disruptive when it exchanges non-homologous sections of the two parent genotypes, and this is particularly a problem when genotypes are of variable lengths. The problem for recombination becomes then, given a randomly selected crossover point in one parent genotype, how to identify an appropriate place to break the other parent genotype so as to exchange homologous sections as far as is possible.

As a matter of practical concern, therefore, an algorithm needs to be developed which can determine on ‘syntactic’ grounds rather than ‘semantic’ ones how to exchange homologous segments of genotypes of varying lengths. This can be quantified as maximising the similarity (under some appropriate measure) of the segments exchanged. This is of course a problem which nature, at the level of molecular biology, has found its own method of tackling, so an investigation of the relevant literature is suggested. It turns out that molecular biologists have developed algorithms for their own rather different, but related purposes. They are interested in quantifying on ‘syntactic’ grounds the similarities between two given nucleotide or amino-acid sequences, and doing so with computational efficiency, and it turns out that their algorithms can be adapted and extended for our present purposes. This is done in [6], where the algorithm is explicitly presented.

7 Neutral Networks

Under many circumstances there can be loci on a genotype where mutations do not change the phenotype; or alternatively, do change the phenotype without altering fitness. Mutations at such loci are *neutral*. A connected pathway through genotype space of genotypes with the same fitness constitutes a *neutral network*. Examples from RNA evolution [1] show that in some circumstances such neutral networks can percolate throughout genotype space. A population moving on such a neutral network has potentially an enormous number of neighbours, and is therefore very unlikely to get stuck at a local optimum.

One desirable possibility that variable-length genotypes offer is that of extending the dimensionality of a genotype search space, and thereby creating hyper-dimensional bypasses to avoid entrapment on local optima. Here is an explicit mapping (due to Nick Jakobi) whereby a maximally rugged fitness landscape of binary genotypes of length N (with many local optima) can be transformed into a different landscape of dimensionality $2N + 1$ which has a single global optimum. In the new space with genotypes of length $2N + 1$, treat the first locus as a switch; when it is 0, interpret the next N loci as a genotype within the first space, and when it is 1, interpret the last N loci in the same way. There are now 2^N neutral networks, and any two networks neighbour each other in a couple of places. Thus one can reach the global optimum from any point without going downhill, there are no local optima apart from the global one.

The optimum selection/mutation balance needs to be adjusted to take into account the amount of neutrality in the genotypes. The form which neutral networks take may be relevant to the appropriateness of different types of recombination.

8 Implicit Variable Length

Under some circumstances a fixed-length genotype may have the proportion of loci which are functional (non-neutral) varying during an evolutionary run. An example is in work by Adrian Thompson in hardware evolution [7], where genotypes of 1800 bits specified the hardware configuration of an FPGA, a reconfigurable silicon chip. A 10×10 array of cells was available on the silicon for wiring (as genetically specified) into a circuit to perform the signal discrimination task, though the final successful configuration was shown to use only some third of these. The two-thirds of the genotype that encoded for the unused parts of silicon was effectively ‘junk DNA’, and mutations there were (in the context of the rest of the successful genotype) not subject to selection and hence ‘junk DNA’. Nevertheless analysis of the evolutionary pathway [5] indicates that an earlier stage in the evolutionary run virtually all the loci were under selection. In other words the ‘effective length’ of the genotype varied between some 600 bits and the full 1800 bits during the run.

Given that the optimal selection/mutation balance depends on setting the expected number of mutations per ‘effective’ part of the genotype — that part which is non-neutral — then such implicit changes of genotype length are just as significant as more immediately visible changes in length.

References

- [1] C.V. Forst, C. Reidys, and J. Weber. Evolutionary dynamics and optimization: Neutral networks as model-landscapes for rna secondary-structure folding-landscapes. In F. Moran, A. Moreno, J.J. Merelo, and P. Cachon, editors, *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life (ECAL95). Lecture Notes in Artificial Intelligence 929*, pages 128–147. Springer Verlag, 1995.
- [2] I. Harvey. Species adaptation genetic algorithms: The basis for a continuing SAGA. In F. J. Varela and P. Bourguine, editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 346–354. MIT Press/Bradford Books, Cambridge, MA, 1992.
- [3] I. Harvey. Evolutionary robotics and SAGA: the case for hill crawling and tournament selection. In C. Langton, editor, *Artificial Life III, Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. XVI*, pages 299–326. Addison Wesley, 1993.
- [4] I. Harvey, P. Husbands, and D. T. Cliff. Genetic convergence in a species of evolved robot control architectures. In S. Forrest, editor, *Genetic Algorithms: Proceedings of Fifth Intl. Conference*, page 636, San Mateo CA, 1993. Morgan Kaufmann.

- [5] I. Harvey and A. Thompson. Through the labyrinth evolution finds a way: A silicon ridge. In T. Higuchi, editor, *Proc. of The First International Conference on Evolvable Systems: From Biology to Hardware (ICES96)*. Springer-Verlag, 1996.
- [6] Inman Harvey. The SAGA cross: the mechanics of crossover for variable-length genetic algorithms. In R. Männer and B. Manderick, editors, *Parallel Problem Solving from Nature, 2*, pages 269–278. North-Holland, 1992.
- [7] A. Thompson. An evolved circuit, intrinsic in silicon, entwined with physics. In T. Higuchi, editor, *Proc. of The First International Conference on Evolvable Systems: From Biology to Hardware (ICES96)*. Springer-Verlag, 1996.