

Lecture 4: Advanced Operators and Techniques

Inversion

- ▶ We have seen variations on the basic operators so far
 - ▶ Crossover
 - ▶ Mutation
- ▶ Holland's earliest work included a proposal for a further operator
 - ▶ Inversion
- ▶ Holland implicitly recognised the positional bias of 1X
- ▶ Inversion was his solution, intended to bring co-adapted alleles at distant loci closer together on the chromosome
- ▶ The biological interpretation of the inversion operator is that it maintains *linkage disequilibrium* due to selection in the face of disruption by crossover
 - ▶ *Definition:* two loci in a population are in *linkage equilibrium* if the frequency distribution of alleles at one locus is independent of the frequency distribution of alleles at the other locus

Inversion

Data: ℓ — length of chromosome

$i_1 \leftarrow$ random integer between 0 and ℓ inclusive;

$i_2 \leftarrow$ random integer $\neq i_1$ between 0 and ℓ inclusive;

if $i_1 > i_2$ **then**

 | swap i_1 and i_2 ;

end

for $l = i_1$ **to** $\lfloor (i_1 + i_2 - 1)/2 \rfloor$ **do**

 | swap allele and index at locus l with allele and index at locus
 | $i_1 + i_2 - 1 - l$;

end

Inversion

- ▶ An index is needed for each locus to preserve the meaning of the locus independent of its position on the chromosome
- ▶ Prior to crossover, both parents must be reordered if their ordering is different
 - ▶ Some ordering must then be assigned to the offspring, perhaps randomly selected from one of the parents
- ▶ Of course, inversion is redundant with operators such as UX, which do not have any positional bias

Population Selection Schemes

- ▶ We have previously looked at different selection operators
 - ▶ Roulette Wheel Selection (RWS)
 - ▶ Stochastic Universal Sampling (SUS)
- ▶ These operators all used Holland's original *generational* selection scheme
 - ▶ Each iteration of the GA completely replaces the population with the offspring of those selected to reproduce
- ▶ Other selection schemes exist
 - ▶ Overlapping
 - ▶ Steady state
 - ▶ (λ, μ) and $(\lambda + \mu)$
 - ▶ Elitism

Overlapping Generations and Steady State Selection

- ▶ Under selection let $1 \leq G \leq N$ individuals be selected for replacement
- ▶ These are replaced by G offspring from the reproduction phase
- ▶ If $1 < G < N$ we have overlapping generations
- ▶ G is sometimes referred to as the *generation gap*
- ▶ Clearly generational selection is a special case of overlapping selection where $G = N$
- ▶ At the other extreme $G = 1$
 - ▶ This is sometimes referred to as the *steady state GA*

(λ, μ) and $(\lambda + \mu)$

- ▶ Selection schemes from Evolutionary Strategies are also applicable
- ▶ (λ, μ)
 - ▶ λ parents are selected for reproduction
 - ▶ They produce $\mu \geq \lambda$ offspring
 - ▶ The best λ offspring are selected for the next generation
- ▶ $(\lambda + \mu)$
 - ▶ As for (λ, μ) , except the best individuals for the next generation are selected from the combined set of parents and offspring
- ▶ N.B. μ here is not the same as μ for mutation rate

Elitism

- ▶ Typical selection and genetic operators do not *guarantee* that the best individual in the population will be in the next generation
 - ▶ Fitness proportional, rank and tournament selection may not select the best individual
 - ▶ Even if selected, crossover and mutation are likely to destroy the best individual
- ▶ Thus the best solution found so far is frequently discarded by the GA
- ▶ *Elitism* avoids this by always preserving it in the next generation
 - ▶ Remaining $N - 1$ individuals are replaced by new strings (assuming a generational GA)

Further Measures of Selection

- ▶ In the last lecture we saw how to calculate selection probabilities and selection pressure for certain selection schemes
 - ▶ Linear rank selection
 - ▶ Soft tournament selection
 - ▶ ...
- ▶ Additional measures of selection for comparison of selection schemes exist
 - ▶ Selection intensity
 - ▶ Takeover time

Selection Intensity

- ▶ Selection *pressure* can be directly calculated from the selection probabilities under a particular selection scheme
- ▶ Selection *intensity* is a measure taken on the behaviour of a selection scheme

$$I(t) = \frac{\bar{f}_{sel}(t) - \bar{f}(t)}{\sigma_{f(t)}}$$

Selection Intensity

- ▶ Selection intensity is a *post hoc* measure of the difference in mean fitness of individuals selected for reproduction, and mean fitness of all individuals in the population
 - ▶ Scaled by population standard deviation to allow meaningful comparison across cases
 - ▶ Applies only to generational GAs in this form, but elitist and steady state versions have been derived
- ▶ Theoretical results have also been achieved, under the assumption of a normal fitness distribution

Takeover Time

- ▶ Selection can also be measured in terms of the time needed for the best individual to take over the entire population
 - ▶ Assuming no crossover and mutation, only selection
- ▶ Empirical results can be obtained to compare various selection schemes
- ▶ Theoretical results have been derived that agree well with empirical values
 - ▶ Takeover time for most selection strategies is $O(\log N)$
 - ▶ Takeover time for fitness proportional selection is $O(N \log N)$
- ▶ Takeover time is related to the rate at which diversity is lost through selection

Diversity Maintenance

- ▶ Diversity is crucial for any kind of evolution (as we will see in detail later in the course)
 - ▶ Natural
 - ▶ Artificial
 - ▶ Evolutionary computation
- ▶ Diversity can be maintained by various strategies
 - ▶ Modified selection / genetic operators
 - ▶ Modified representations
 - ▶ Fitness sharing schemes
- ▶ Loss of diversity can be a signal to terminate an evolutionary algorithm

Diversity Preserving Operators

- ▶ We can use ‘incest prevention’ to maintain diversity
 - ▶ Prevent crossover between individuals whose similarity lies above a threshold
 - ▶ Similarity can be measured as the reciprocal of Hamming distance
 - ▶ The threshold must be increased as selection results in population convergence
- ▶ Alternatively we can implement a ‘no duplicates’ policy
 - ▶ Individuals are not inserted into the population if they match an individual already in the population
 - ▶ This requires comparing each new individual against every individual already in the population
 - ▶ The naïve approach to this for a generational GA has $O(N)$ time complexity

Diversity Preserving Operators

- ▶ We can also force operators to produce novel offspring (i.e. not clones of either parent)
- ▶ E.g. 1X applied to 1101001 and 1100010 will always generate a clone if the crossover point is any of the first three positions
- ▶ We can identify suitable crossover points before applying the crossover operator
 - ▶ For binary chromosomes compute the XOR of the two parents
 - ▶ 0001011 for the example above
 - ▶ Select a crossover point lying between the outermost 1s in the XOR string

Fitness Sharing

- ▶ We can also maintain diversity through modifying the selection mechanism
- ▶ *Fitness sharing*, as its name suggests, shares fitness between individuals occupying the same *niche*
 - ▶ *Niche* is a term from biology denoting a particular set of environmental conditions, to which an organism or organisms may become adapted
 - ▶ E.g. herbivores and carnivores are adapted, or specialised, to eat plants and other animals respectively
- ▶ A simple linear *sharing function* such as

$$h(d_{ij}) = \begin{cases} 2 - \frac{d_{ij}}{D} & \text{if } d_{ij} < D \\ 1 & \text{otherwise} \end{cases}$$

- ▶ is evaluated for all pairs of individuals in the population
 - ▶ d_{ij} is the distance between individuals i and j
 - ▶ D is a tunable distance threshold

Fitness Sharing

- ▶ For every individual j we then compute the sum

$$\varsigma_j = \sum_{i \neq j} h(d_{ij})$$

- ▶ then divide its raw fitness by ς_j to give the adjusted fitness that will be used during selection
- ▶ Hence very similar chromosomes (under some distance metric) will have their fitnesses reduced more than unique chromosomes
- ▶ Care must be taken in choosing the distance metric
 - ▶ In particular, should it be based on genotypic or phenotypic distance?
 - ▶ E.g. if d_{ij} were Hamming distance then as we have seen genetically similar individuals can have very different phenotypes, and vice versa

Diploidy and Dominance

- ▶ All the representations we have seen in our examples so far allow one allele to occupy each locus
 - ▶ In biological terminology, this is *haploidy*
- ▶ Haploidy is not typical in nature, *diploidy* is far more common
 - ▶ Each locus carries *two* alleles, which encode for some trait
 - ▶ Often one allele is *dominant* over another
 - ▶ E.g. imagine that for eye colour in humans there are two alleles
 - ▶ The brown eye allele is *dominant* (B), the blue eye allele is *recessive* (b)
 - ▶ Then the *homozygote* bb encodes for blue eyes, the homozygote BB and the *heterozygotes* bB and Bb all encode for brown eyes
 - ▶ Alleles can also be *partially dominant*, or *additive*

Diploidy and Dominance

- ▶ Diploidy and dominance explain the ratio of offspring phenotypes that Mendel discovered
- ▶ Intriguingly, for any distribution of phenotypes where the underlying allele frequencies are m and $n = 1 - m$, one generation of random mating without selection results in a stable equilibrium

$$E(AA) = m^2, E(Aa) = 2mn, E(aa) = n^2$$

- ▶ This is the *Hardy-Weinberg* equilibrium, from population genetics
 - ▶ *Exercise:* for the *genotype* frequency ratio $p : 2q : r$, derive the genotype frequency ratio after one generation of random mating without selection
- ▶ The equilibrium demonstrates that, in the absence of selection, Mendelian genetics will not result in the spread or disappearance of dominant or recessive alleles
 - ▶ I.e. diploidy is diversity preserving in the absence of selection
 - ▶ It was originally thought that dominance would lead to elimination of a recessive allele even in the absence of selection

Diploidy and Dominance

- ▶ Diploidy and dominance are also a diversity preserving technique
 - ▶ If a dominant trait is fitter than its recessive counterpart(s), alleles for those recessive traits will still be preserved in the population despite negative selection
 - ▶ Dominance itself may evolve, so that the fitter trait is necessarily the dominant one
 - ▶ This could allow an improved evolutionary response to environmental change, where relative fitnesses of traits are reversed
- ▶ Hence diploidy and dominance could be useful in GAs
 - ▶ Particularly in GAs with non-static objective functions

Diploidy and Dominance

- ▶ Dominance and diploidy can be simply implemented in a GA
- ▶ Assume 3 alleles
 - ▶ 0: encodes for trait 0
 - ▶ 1: recessively encodes for trait 1
 - ▶ 2: dominantly encodes for trait 1
- ▶ This gives the following dominance map

	0	1	2
0	0	0	1
1	0	1	1
2	1	1	1

- ▶ The phenotypic ratio for dominants of 3 to 1 is achieved for allele pairings of 0 with 1, and 0 with 2
- ▶ Dominance can evolve through allele substitution of 1 for 2 and vice versa

Termination

- ▶ So far we have not considered when to stop the GA
 - ▶ We are probably applying the GA to a problem for which we do not know the global optimum
 - ▶ Hence we may not be able to specify a termination criterion in terms of quality of solution discovered
 - ▶ On the other hand we may be able to specify a termination criterion in terms of a solution that is 'good enough'
 - ▶ We can term this *satisficing*
- ▶ Given that evolution requires variability to act on, we could terminate the GA when population variability has fallen
 - ▶ E.g. when some statistical measure of population diversity falls below a threshold
 - ▶ E.g. when an attempt is made to cross an individual with a clone of itself
 - ▶ ...
- ▶ Other common termination criteria are fixed number of generations, or fixed processing time, etc.

Multiobjective Optimisation

- ▶ Many optimisation algorithms, including EC, seek to optimise a single objective function
- ▶ Many optimisation problems, however, have multiple conflicting objectives
 - ▶ I.e. multiple objective functions
- ▶ The simple approach to solving this is to assign fitness as a weighted sum of the value of each of the separate objective functions
 - ▶ This requires us *a priori* to decide on the relative importance of the different objectives
- ▶ As GAs (and most EC algorithms) are population based, we can take a different approach...

Pareto Optimality

- ▶ By using the concept of Pareto-optimality we can find a set of solutions that are all optimal compromises between the conflicting objectives
 - ▶ We can then examine this set and select one solution from it according to our needs
 - ▶ This approach is much more flexible
- ▶ Pareto-optimality is a concept used in economics, game theory, etc.
 - ▶ *Definition:* Solution A is *dominated* by solution B if solution B is better according to at least one objective, and no worse according to the other objectives
 - ▶ *Definition:* A *Pareto-optimal solution* is one that is not *dominated* by any other solution, i.e. it is one in which no objective can be improved without a deterioration in one or more of the other objectives
- ▶ N.B. do not confuse the definitions of *dominates* and *dominance*!

Multiobjective EC

- ▶ Several approaches to using Pareto-optimality in EC are possible, e.g.
 - ▶ Adjust the fitness function to favour non-dominated individuals
 - ▶ Rank-based selection assigning joint ranks to individuals
 - ▶ Assign rank 1 to all non-dominated individuals in the population, and remove
 - ▶ Assign rank 2 to all remaining non-dominated individuals, and remove
 - ▶ ...
 - ▶ Select only dominated individuals for deletion from the population (in overlapping selection)