

Niching and Speciation

Fernando Lobo

University of Algarve

Introduction

- We will look at techniques that are useful to maintain diversity in the population.
- Especially useful to find multiple (different) solutions in multimodal problems.
- Also useful in multi-objective optimization, dynamic environments and simulation of complex adaptive systems.

Multimodal problems

- In problems with multiple optima, a standard EA eventually focuses its attention on a single optima.
 - Even if there's many equally good ones.
- Why is that?
 - Due to chance variation alone, the population will end up being in the vicinity of a single solution.

Methods

- We will be looking at two methods inspired by what happens in Nature:
 - Niching
 - Speciation

Niching

- Niching methods allow EAs to maintain a population of diverse solutions through time.
- EAs with niching are capable of obtaining multiple optimal solutions in a single population in a single run.
- Motivation from Nature:
 - natural evolutionary processes maintain a variety of species, each occupying its own ecological niche.

Niching

- Two major classes of methods:
 - Fitness sharing
 - Crowding

Fitness sharing

- Proposed by Goldberg and Richardson (1987).
- Idea:
 - Resources in Nature are limited.
 - Similar individuals (those occupying the same niche) must share those resources (e.g., food) with each other.
 - In EAs, food = fitness

Fitness sharing (cont.)

- Sharing works by degrading an individual's fitness by an amount related to the number of similar individuals contained in the population.
- The shared fitness f' is equal to the original fitness f , divided by its niche count.
- An individual's niche count is the sum over the entire population, of the sharing function (sh) values between itself and every individual in the population.

$$f'(i) = \frac{f(i)}{\sum_{j=1}^n sh(d(i, j))}$$

Fitness sharing (cont.)

- The sharing function is a function of the distance between two individuals:
 - returns 1 if the two individuals are identical
 - returns 0 if the the distance is above a certain threshold (σ_{share} , a user defined parameter).
 - returns something in between 0 and 1 depending on the level of similarity.

Fitness sharing (cont.)

- A common sharing function is:

$$sh(d) = \begin{cases} 1 - \left(\frac{d}{\sigma_{share}}\right)^\alpha, & \text{if } d < \sigma_{share} ; \\ 0, & \text{otherwise .} \end{cases}$$

- α is a constant (typically set to 1).
- distance d can be computed at the genotype level (Hamming distance for bitstrings) or at the phenotype level.
- σ_{share} should be set to allow discrimination between desired peaks. (In practice this usually unknown a priori.)

Crowding

- First proposed by De Jong (1975).
- Idea: new individuals replace similar individuals in the population.
- Also uses a distance function to measure similarity.
- Unlike sharing, crowding methods do not allocate individuals proportional to each niche (peak) fitness.
- By replacing similar individuals, crowding promotes pre-existing diversity of a population.

Crowding (cont.)

- There's various ways to implement this mechanism.
- Let's look at the original method proposed by De Jong.
- The algorithm is a steady-state EA: only a fraction of the population reproduces and dies each generation.

De Jong' s original crowding model

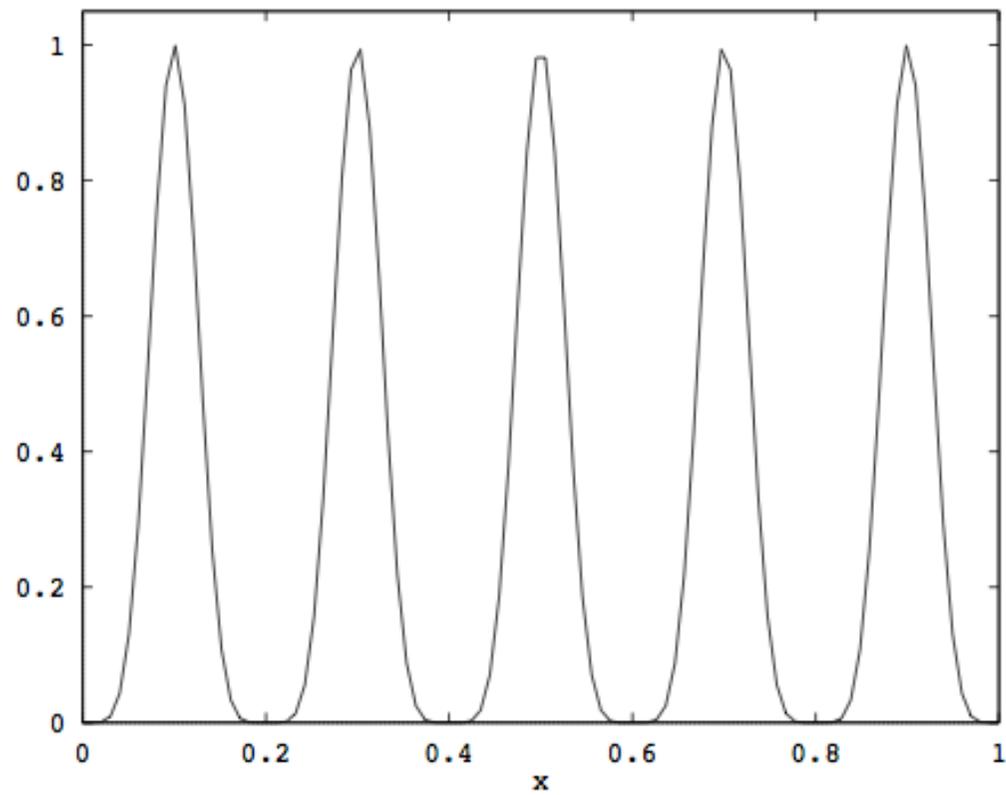
- A proportion GG (*generation gap*) of the population is chosen via fitness-proportionate selection to undergo crossover and mutation.
- For each new individual, take a random sample of CF (*crowding factor*) individuals from the population. The new individual replaces the one that is most similar to it in that random sample.
- (GG was typically set around 0.1, and CF around 2-5)

Simulation results: Sharing vs Crowding

- Results taken from Deb and Goldberg (*An investigation of niche and species formation in genetic function optimization*, ICGA 1989).
- Two test functions: F1 and F2
 - F1: equal peaks
 - F2: non-equal peaks

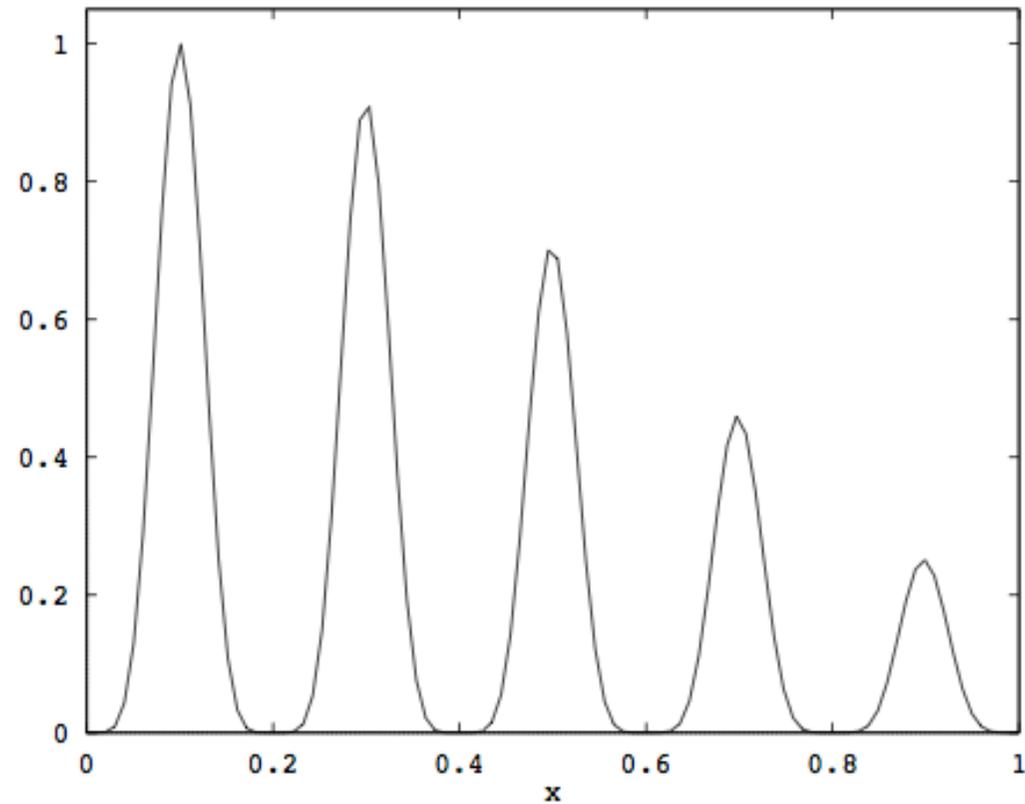
F1: equal peaks

$$F1(x) = \sin^6(5\pi x)$$



F2: unequal peaks

$$F2(x) = e^{-2(\ln 2)\left(\frac{x-0.1}{0.8}\right)^2} \sin^6(5\pi x)$$

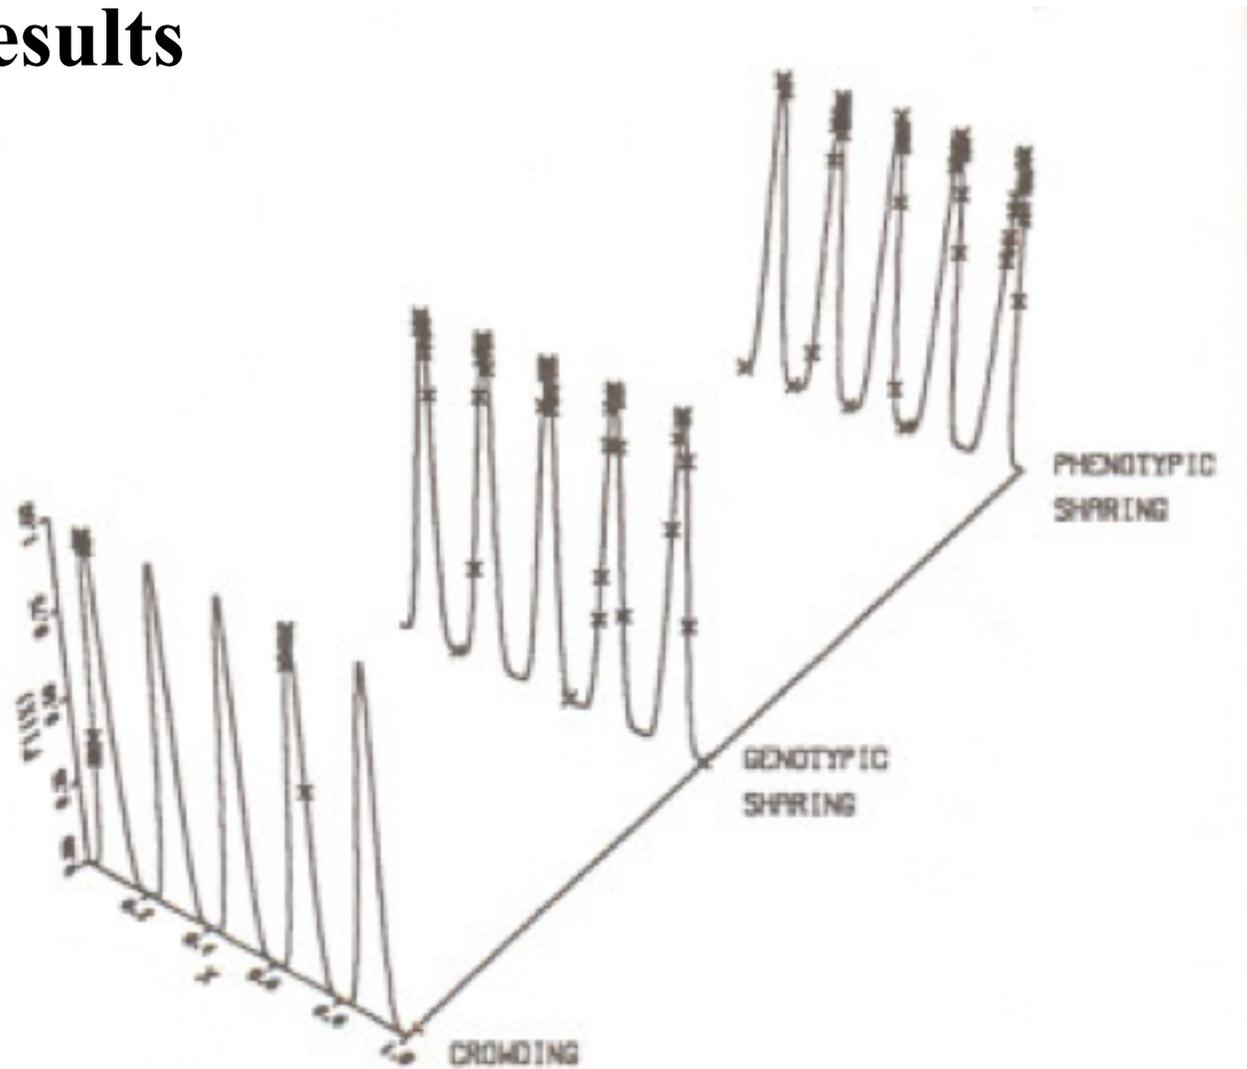


GA Parameters

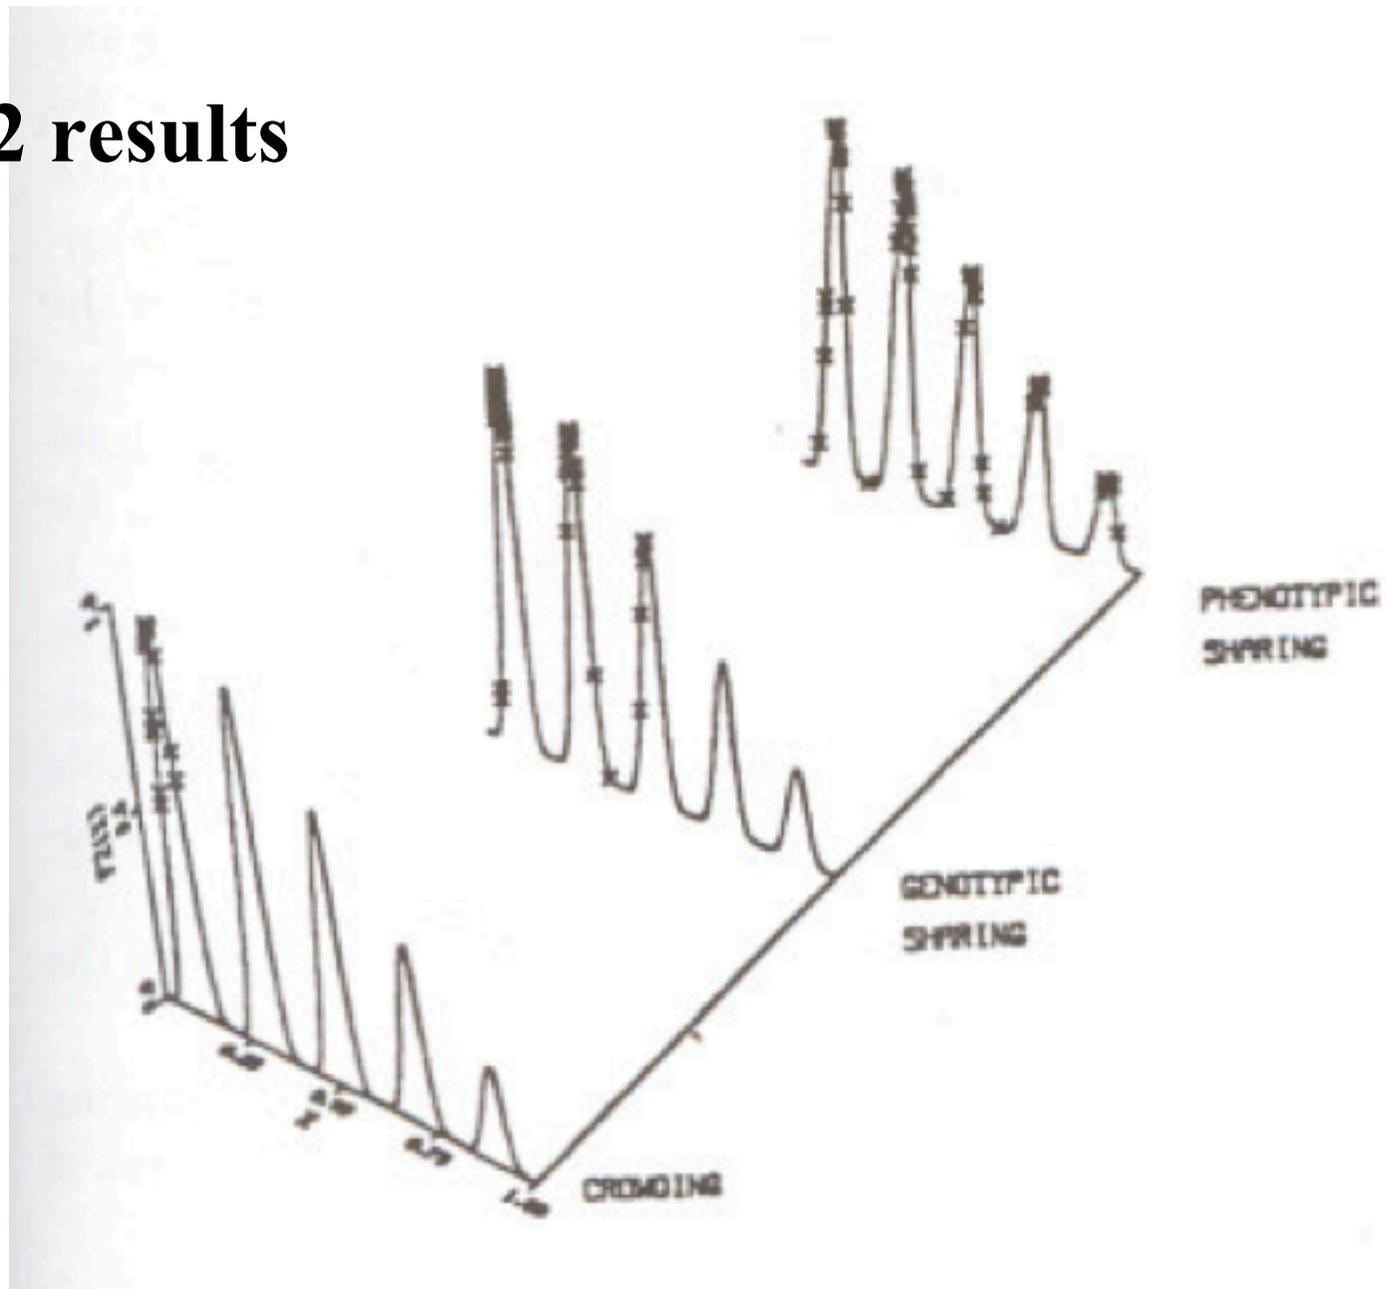
- σ_{share} set according to Deb and Goldberg's recommendations (see paper for details).

maximum generation	:	200
population size	:	100
string length (binary coded)	:	30
probability of crossover	:	0.9
probability of mutation	:	0.0
crowding factor	:	3
generation gap	:	0.1

F1 results



F2 results



Results (cont.)

- Crowding is unable to maintain the multiple peaks (maintains only 2 for F1 and only 1 for F2).
- On peaks with unequal value, genotypic sharing is not able to maintain the lower peaks.

Restricted Tournament Selection (RTS)

- RTS is a crowding-like method proposed by Harik (1995).
- RTS is a modification of tournament selection that restricts an individual from competing with others that are very different from it.

RTS works as follows

- Select two solutions at random, A and B , from the population and perform crossover and mutation, resulting in two new solutions, A' and B' .
- For each new solution (A' and B'), scan w individuals (randomly chosen) from the population and pick the one that is most similar to it. Call them A'' and B'' .
- A' competes with A'' . If A' is better, then it replaces A'' in the population. (Do the same for B' and B'')
 - The algorithm is steady-state (solutions are introduced in the population incrementally).

RTS

- RTS performs much better than the standard crowding mechanism.
- Also much simpler to use than fitness sharing.
- Harik suggested setting w as 4 times the number of desired peaks to be found, and obtained good results.

Speciation

- Finding multiple solutions can be improved with speciation methods.
- A species is a collection of individuals which resemble each other more closely than they resemble individuals of another species.
- Artificial species can be created in EAs by only allowing mating (crossover) between similar individuals.
 - In Nature, lions don't mate with elephants!

Speciation (cont.)

- Niching distributes individuals among multiple peaks.
- But niching does not prevent an individual from one peak (niche) to mate with an individual from another peak.
 - Such matings often result in so-called *lethal* solutions, representing none of the peaks.
- Can solve this problem with speciation methods.

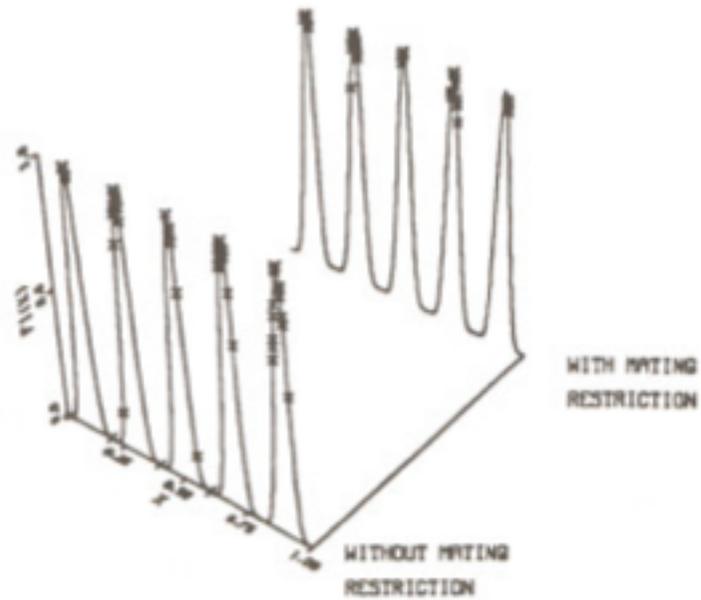
Speciation methods

- Deb and Goldberg proposed two mating restricting schemes based on phenotypic and genotypic distances between mating individuals.
- Idea: if the distance is closer than a parameter σ_{mating} , they participate in the crossover operation. Otherwise another individual is chosen at random as a candidate mate.
 - This process continues until a proper mate is found. If no such member exists, then a random individual from the population is chosen as mate.

Speciation methods

- Niching with σ_{share} is implemented in the selection operator.
- Mating restriction with σ_{mating} is implemented in the crossover operator.
- Deb and Goldberg used $\sigma_{mating} = \sigma_{share}$
- Results on F1 and F2 were improved.

F1 with and without mating restriction



Other speciation method

- Using tag bits (Spears, 1994)
- Tag bits are appended to every individual.
- Each species correspond to a particular configuration of those bits.
- Mating is restricted to individuals containing the same set of tag bits.
- Spears allowed the tag bits to be mutated.