Lecture 25

Reading

- Read Flake, ch. 17, "Competition & Cooperation"

Demonstration of GA: Finding Maximum of Fitness Landscape

Run Genetic Algorithms — An Intuitive Introduction
by Pascal Glauser
<homepage.sunrise.ch/
homepage/pglaus/gentore.htm>

Demonstration of GA: Evolving to Generate a Pre-specified Shape (Phenotype)

Run Genetic Algorithm Viewer
<www.rennard.org/alife/english/gavgb.html>

Demonstration of GA: Eaters Seeking Food

http://math.hws.edu/xJava/GA/

Why Does the GA Work?

The Schema Theorem

Schemata

A schema is a description of certain patterns of bits in a genetic string

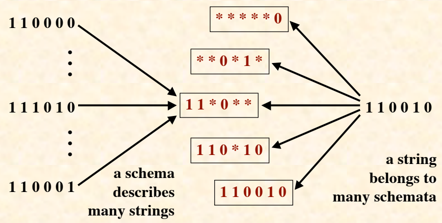

11/20/07 7

The Fitness of Schemata

- The schemata are the building blocks of solutions
- We would like to know the average fitness of all possible strings belonging to a schema
- We cannot, but the strings in a population that belong to a schema give an estimate of the fitness of that schema
- Each string in a population is giving information about all the schemata to which it belongs (implicit parallelism)

11/20/07 8

Effect of Selection

Let n = size of population

Let $m(S,t)$ = number of instances of schema S at time t

String i gets picked with probability $\frac{f_i}{\sum_j f_j}$

Let $f(S)$ = avg fitness of instances of S at time t

So expected $m(S,t+1) = m(S,t) \cdot n \cdot \frac{f(S)}{\sum_j f_j}$

Since $f_{av} = \frac{\sum_j f_j}{n}$, $m(S,t+1) = m(S,t) \frac{f(S)}{f_{av}}$

11/20/07 9

Exponential Growth

- We have discovered: $m(S, t+1) = m(S, t) \cdot f(S) / f_{av}$
- Suppose $f(S) = f_{av}(1 + c)$
- Then $m(S, t) = m(S, 0)(1 + c)^t$
- That is, exponential growth in above-average schemata

11/20/07 10

Effect of Crossover

1 … 0*
|←δ→|

- Let λ = length of genetic strings
- Let $\delta(S)$ = defining length of schema S
- Probability {crossover destroys S}: $p_d \le \delta(S) / (\lambda - 1)$
- Let p_c = probability of crossover
- Probability schema survives:

$$p_s \ge 1 - p_c \frac{\delta(S)}{\lambda - 1}$$

11/20/07 11

Selection & Crossover Together

$$m(S,t+1) \ge m(S,t) \frac{f(S)}{f_{av}} \left[1 - p_c \frac{\delta(S)}{\lambda - 1}\right]$$

11/20/07 12

Effect of Mutation

- Let p_m = probability of mutation
- So $1 - p_m$ = probability an allele survives
- Let $o(S)$ = number of fixed positions in S
- The probability they all survive is $(1 - p_m)^{o(S)}$
- If $p_m << 1$, $(1 - p_m)^{o(S)} \approx 1 - o(S)\, p_m$

11/20/07 13

Schema Theorem: "Fundamental Theorem of GAs"

$$m(S,t+1) \geq m(S,t)\frac{f(S)}{f_{av}}\left[1 - p_c\frac{\delta(S)}{\lambda - 1} - o(S)p_m\right]$$

11/20/07 14

The Bandit Problem

- Two-armed bandit:
  - random payoffs with (unknown) means m_1, m_2 and variances σ_1, σ_2
  - optimal strategy: allocate exponentially greater number of trials to apparently better lever
- k-armed bandit: similar analysis applies
- Analogous to allocation of population to schemata
- Suggests GA may allocate trials optimally

11/20/07 15

Goldberg's Analysis of Competent & Efficient GAs

11/20/07 16

Paradox of GAs

- Individually uninteresting operators:
  - selection, recombination, mutation
- Selection + mutation ⇒ continual improvement
- Selection + recombination ⇒ innovation
  - fundamental to invention: *generation* vs. *evaluation*
- Fundamental intuition of GAs: the three work well together

11/20/07 17

Race Between Selection & Innovation: Takeover Time

- Takeover time t^* = average time for most fit to take over population
- Transaction selection: population replaced by s copies of top $1/s$
- s quantifies selective pressure
- Estimate $t^* \approx \ln n / \ln s$

11/20/07 18

Innovation Time

- Innovation time t_i = average time to get a better individual through crossover & mutation
- Let p_i = probability a single crossover produces a better individual
- Number of individuals undergoing crossover = $p_c\, n$
- Probability of improvement = $p_i\, p_c\, n$
- Estimate: $t_i \approx 1 / (p_c\, p_i\, n)$

Steady State Innovation

- Bad: $t^* < t_i$
  - because once you have takeover, crossover does no good
- Good: $t_i < t^*$
  - because each time a better individual is produced, the t^* clock resets
  - *steady state innovation*
- Innovation number:

$$\text{Iv} = \frac{t^*}{t_i} = p_c p_i \frac{n \ln n}{\ln s} > 1$$

Feasible Region

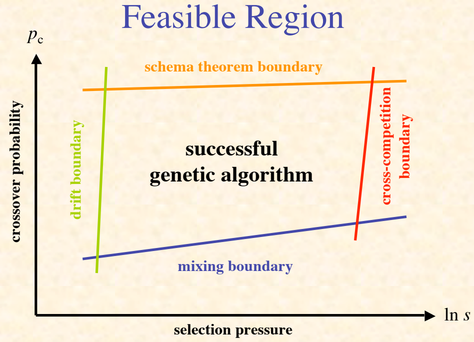

Other Algorithms Inspired by Genetics and Evolution

- Evolutionary Programming
  - natural representation, no crossover, time-varying continuous mutation
- Evolutionary Strategies
  - similar, but with a kind of recombination
- Genetic Programming
  - like GA, but program trees instead of strings
- Classifier Systems
  - GA + rules + bids/payments
- and many variants & combinations…

Additional Bibliography

1. Goldberg, D.E. *The Design of Innovation: Lessons from and for Competent Genetic Algorithms*. Kluwer, 2002.
2. Milner, R. *The Encyclopedia of Evolution*. Facts on File, 1990.

VI

VI. Cooperation & Competition

Game Theory and the Iterated Prisoner's Dilemma

The Rudiments of Game Theory

11/20/07 25

Leibniz on Game Theory

- "Games combining chance and skill give the best representation of human life, particularly of military affairs and of the practice of medicine which necessarily depend partly on skill and partly on chance." — Leibniz (1710)
- "… it would be desirable to have a complete study made of games, treated mathematically."
  — Leibniz (1715)

11/20/07 26

Origins of Modern Theory

- 1928: John von Neumann: optimal strategy for two-person zero-sum games
  - von Neumann: mathematician & pioneer computer scientist (CAs, "von Neumann machine")
- 1944: von Neumann & Oskar Morgenstern:*Theory of Games and Economic Behavior*
  - Morgenstern: famous mathematical economist
- 1950: John Nash: *Non-cooperative Games*
  - his PhD dissertation (27 pages)
  - "genius," Nobel laureate (1994), schizophrenic

11/20/07 27

Classification of Games

- Games of Chance
  - outcome is independent of players' actions
  - "uninteresting" (apply probability theory)
- Games of Strategy
  - outcome is at least partially dependent on players' actions
  - completely in chess
  - partially in poker

11/20/07 28

Classification of Strategy Games

- Number of players (1, 2, 3, …, *n*)
- Zero-sum or non zero-sum
- Essential or inessential
- Perfect or imperfect information

11/20/07 29

Zero-sum vs. Non Zero-sum

- Zero-sum: winnings of some is exactly compensated by losses of others
  - sum is zero for *every* set of strategies
- Non zero-sum:
  - positive sum (mutual gain)
  - negative sum (mutual loss)
  - constant sum
  - nonconstant sum (variable gain or loss)

11/20/07 30

Essential vs. Inessential

- Essential: there is an advantage in forming coalitions
  - may involve agreements for payoffs, cooperation, etc.
  - can happen in zero-sum games only if $n \geq 3$ (obviously!)
- Inessential: there is no such advantage
  - "everyone for themselves"

11/20/07 31

Perfect vs. Imperfect Information

- Perfect information: everyone has complete information about all previous moves
- Imperfect information: some or all have only partial information
  - players need not have complete information even about themselves (e.g. bridge)

11/20/07 32

Strategies

- Strategy: a complete sequence of actions for a player
- Pure strategy: the plan of action is completely determined
  - for each situation, a specific action is prescribed
  - disclosing the strategy might or might not be disadvantageous
- Mixed strategy: a probability is assigned to each plan of action

11/20/07 33

Von Neumann's Solution for Two-person Zero-sum Games

11/20/07 34

Maximin Criterion

- Choose the strategy that *maximizes* the *minimum* payoff
- Also called *minimax*: minimize the maximum loss
  - since it's zero-sum, your loss is the negative of your payoff
  - pessimistic?

11/20/07 35

Example

- Two mineral water companies competing for same market
- Each has fixed cost of $5 000 (regardless of sales)
- Each company can charge $1 or $2 per bottle
  - at price of $2 can sell 5 000 bottles, earning $10 000
  - at price of $1 can sell 10 000 bottles, earning $10 000
  - if they charge same price, they split market
  - otherwise all sales are of lower priced water
  - payoff = revenue – $5 000

11/20/07 Example from McCain's *Game Theory: An Introductory Sketch* 36

Payoff Matrix

		Perrier	
		price = $1	price = $2
Apollinaris	price = $1	0, 0	5000, –5000
	price = $2	–5000, 5000	0, 0

Maximin for A.

minimum at $1 · Maximin · minimum at $2

		Perrier	
		price = $1	price = $2
Apollinaris	price = $1	0, 0	5000, –5000
	price = $2	–5000, 5000	0, 0

Maximin for P.

		Perrier	
		price = $1	price = $2
Apollinaris	price = $1	0, 0	5000, –5000
	price = $2	–5000, 5000	0, 0

Maximin Equilibrium

		Perrier	
		price = $1	price = $2
Apollinaris	price = $1	0, 0	5000, –5000
	price = $2	–5000, 5000	0, 0

Implications of the Equilibrium

- If both companies act "rationally," they will pick the equilibrium prices
- If either behaves "irrationally," the other will benefit (if it acts "rationally")

Matching Pennies

- Al and Barb each independently picks either heads or tails
- If they are both heads or both tails, Al wins
- If they are different, Barb wins

Payoff Matrix

Minimum of each pure strategy is the same		Barb	
		head	tail
Al	head	+1, –1	–1, +1
	tail	–1, +1	+1, –1

Mixed Strategy

- Although we cannot use maximin to select a pure strategy, we can use it to select a mixed strategy
- Take the maximum of the minimum payoffs over all assignments of probabilities
- von Neumann proved you can always find an equilibrium if mixed strategies are permitted

Analysis

- Let P_A = probability Al picks head
- and P_B = probability Barb picks head
- Al's expected payoff:

$$E\{A\} = P_A P_B - P_A (1 - P_B) - (1 - P_A) P_B + (1 - P_A)(1 - P_B)$$

$$= (2 P_A - 1)(2 P_B - 1)$$

Al's Expected Payoff from Penny Game

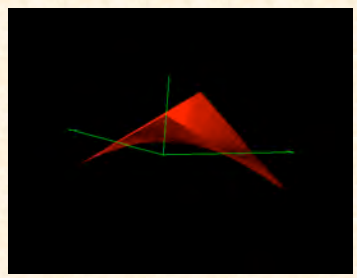

How Barb's Behavior Affects Al's Expected Payoff

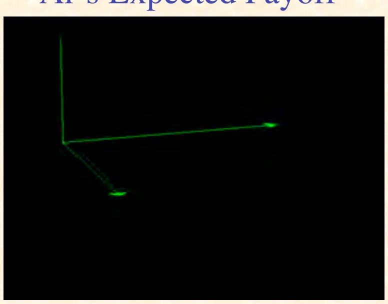

How Barb's Behavior Affects Al's Expected Payoff

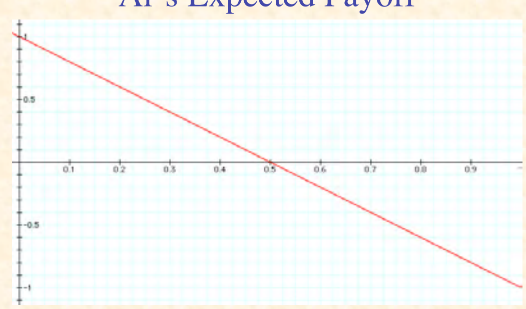

More General Analysis (Differing Payoffs)

- Let A's payoffs be:
  H = HH, h = HT, t = TH, T = TT
- $E\{A\} = P_A P_B H + P_A(1 - P_B)h + (1 - P_A)P_B t + (1 - P_A)(1 - P_B)T$
  $= (H + T - h - t)P_A P_B + (h - T)P_A + (t - T)P_B + T$
- To find saddle point set $\partial E\{A\}/\partial P_A = 0$ and $\partial E\{A\}/\partial P_B = 0$ to get:

$$P_A = \frac{T - t}{H + T - h - t}, \quad P_B = \frac{T - h}{H + T - h - t}$$

Random Rationality

"It seems difficult, at first, to accept the idea that 'rationality' — which appears to demand a clear, definite plan, a deterministic resolution — should be achieved by the use of probabilistic devices. Yet precisely such is the case."

—Morgenstern

Probability in Games of Chance and Strategy

- "In games of chance the task is to determine and then to evaluate probabilities inherent in the game;
- in games of strategy we *introduce* probability in order to obtain the optimal choice of strategy."

— Morgenstern

Review of von Neumann's Solution

- Every two-person zero-sum game has a maximin solution, provided we allow mixed strategies
- But— it applies only to two-person zero-sum games
- Arguably, few "games" in real life are zero-sum, except literal games (i.e., invented games for amusement)

Nonconstant Sum Games

- There is no agreed upon definition of rationality for nonconstant sum games
- Two common criteria:
  - dominant strategy equilibrium
  - Nash equilibrium

Dominant Strategy Equilibrium

- Dominant strategy:
  - consider each of opponents' strategies, and what your best strategy is in each situation
  - if the same strategy is best in all situations, it is the dominant strategy
- Dominant strategy equilibrium: occurs if each player has a dominant strategy and plays it

Another Example

Price Competition		Beta		
		$p = 1$	$p = 2$	$p = 3$
Alpha	$p = 1$	0, 0	50, –10	40, –20
	$p = 2$	–10, 50	20, 20	90, 10
	$p = 3$	–20, 40	10, 90	50, 50

There is no dominant strategy

Nash Equilibrium

- Developed by John Nash in 1950
- His 27-page PhD dissertation: *Non-Cooperative Games*
- Received Nobel Prize in Economics for it in 1994
- Subject of *A Beautiful Mind*

Definition of Nash Equilibrium

- A set of strategies with the property: No player can benefit by changing actions while others keep strategies unchanged
- Players are in equilibrium if any change of strategy would lead to lower reward for that player
- For mixed strategies, we consider expected reward

Another Example (Reconsidered)

Price Competition		Beta		
		$p = 1$	$p = 2$	$p = 3$
Alpha	$p = 1$	0, 0	50, –10	40, –20
	$p = 2$	–10, 50	20, 20	90, 10
	$p = 3$	–20, 40	10, 90	50, 50

better for Beta better for Alpha

Not a Nash equilibrium

The Nash Equilibrium

Price Competition		Beta		
		$p = 1$	$p = 2$	$p = 3$
Alpha	$p = 1$	0, 0	50, –10	40, –20
	$p = 2$	–10, 50	20, 20	90, 10
	$p = 3$	–20, 40	10, 90	50, 50

Nash equilibrium

Extensions of the Concept of a Rational Solution

- Every maximin solution is a dominant strategy equilibrium
- Every dominant strategy equilibrium is a Nash equilibrium

Cooperation Better for Both: A Dilemma

Price Competition		Beta		
		$p = 1$	$p = 2$	$p = 3$
Alpha	$p = 1$	0, 0	50, –10	40, –20
	$p = 2$	–10, 50	20, 20	90, 10
	$p = 3$	–20, 40	10, 90	50, 50

Cooperation

Dilemmas

- Dilemma: "A situation that requires choice between options that are or seem equally unfavorable or mutually exclusive"

  – *Am. Her. Dict.*
- In game theory: each player acts rationally, but the result is undesirable (less reward)