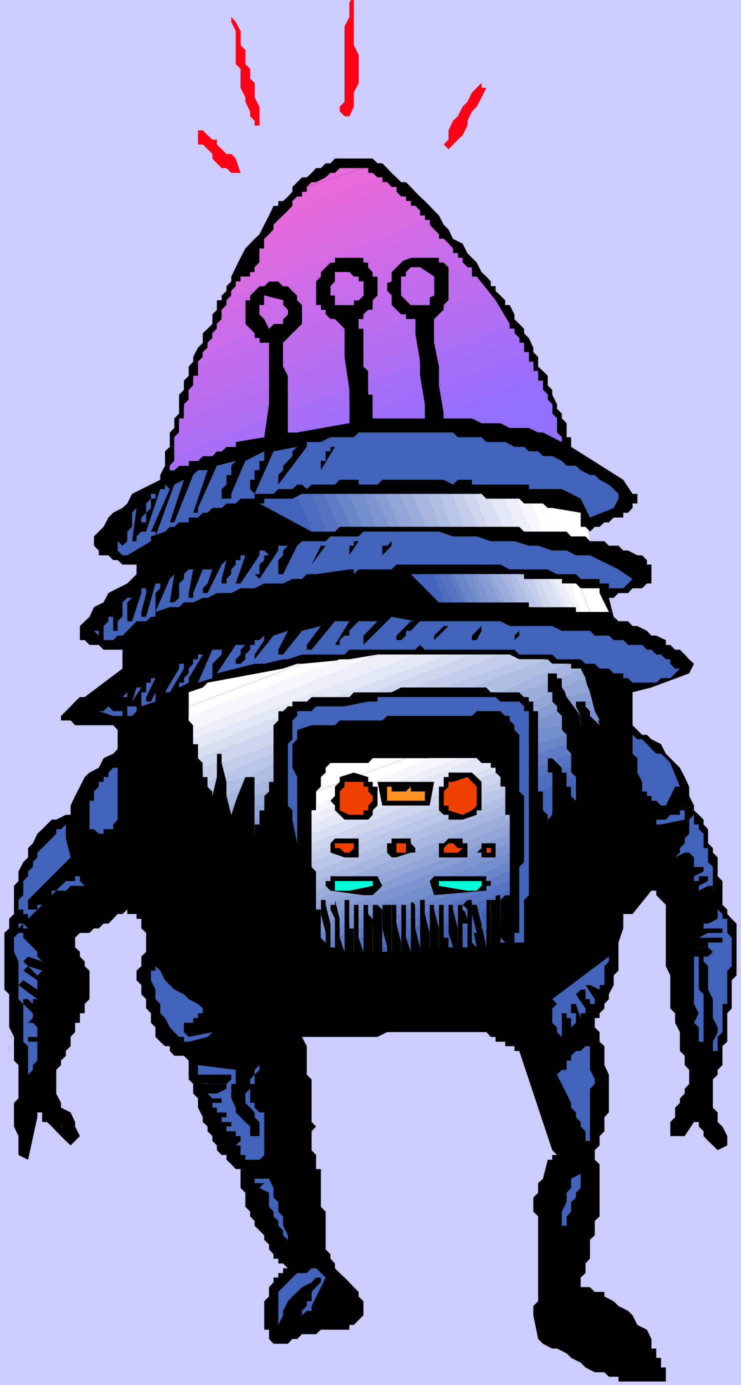

Intelligent Agent Systems

Intelligent Systems, Robotics and Artificial Intelligence

- Artificial intelligence and intelligent behavior of agents
- Comparing artificial and natural intelligence; what should be added to agents
- The role of society of individual agents
- Intelligent agents can learn, adapt and play games; role of personality and emotions

Intelligent Agent (IA)

Intelligent agents are software or programmable hardware entities.

They carry out some set of operations on behalf of a user or another program, or act for their own selfish benefit.

They have some degree of independence.

IA's employ some knowledge or representation of the user's goals or other agents goals

Soon there will be billions agents on Internet connected to various intelligent appliances.

Characteristics of IA

Characteristic	*Description*
Autonomous	Capable of acting on its own, being goal oriented and collaborative, able to alter its activity if needed (see Maes 1995).
Proactive response	The agents response must be corrective, i.e., they must exhibit goal-directed behavior by taking the initiative.
Unobstructive	Work without constant attention of its "master," being out of site (remote executions).
Module	Transportable across different systems and networks. Many agents are not mobile (e.g., Wizards in spreadsheets).
Dedicated and automated	An agent is usually designed to carry in a specific usually repetitive, normally difficult task. For multifaceted jobs we need a multiagent system.
Interactive	Agents are designed to interact with human, other agents or software programs (see opening case). This is critical for a multiagent system.
Conditional processing, practice	Using rule-based or pattern-matching logic (supplied by the user), the agent can make decisions in choosing contexts in which they perceive, or send alerts to the user in timely manner.
Friendly and dependable	An effective agent must be believable and exhibit easy interactivity with people
Able to learn	Only few agents can really do some learning, for example, observing the user and making predictions on its future behavior. Agent must be highly autonomous.

Major Tasks Performed by IA

- Information access and navigation
- Decision support and empowerment
- Repetitive office activity
- Mundane personal activity
- Search and retrieval
- Domain experts
- Intelligent Robots and Robot societies
- Robots with emotions and morality

Practical Applications of IA

- User interface
- Operating systems agents
- Spreadsheet agents
- Workflow and administrative management agents
- Software development
- Negotiation in electronic commerce
- Playing games

Questions of our interest

- How to formalize life situation of conflict and collaboration to be able to use them to program agents?
- What models of behavior can we learn in Nature?
- Why there exist cooperation?
- Is it useful to be altruistic rather than totally selfish?
- What are the best strategies to compete and cooperate with other agents?

Evolving Game Playing Strategies

Overview

1) Introduction

- Agent Based Modeling

2) Prisoner's Dilemma (PD)

3) Deterministic Strategies for PD

- Tournaments
- PD in a Natural Setting
- Downfall of Deterministic Strategies

4) Beyond Determinism

- Nowak and Sigmund

5) PD In Nature

1.1 Introduction

- Complexity Theory – Study of agents and their interactions
- Usually done by Computer Simulation
  - Agent Based Modeling
  - Bottom Up Modeling
  - Artificial Social Systems

1.2 Agent Based Modeling

- **Induction** – Patterns within Empirical Data
- **Deduction** – Specifying Axioms and Proving Consequences

1.3 How Does One do <u>Agent Based</u> Modeling?

- Begin with Assumptions
- Generate data which can be analyzed inductively
- Purpose is to aid Intuition
- Emergent Properties

1.4 Types of Agent Based Modeling

- **Rational Choice** Paradigm
  - Game Theory is based on Rational Choice

- **<u>Adaptive Behavior</u>**
  - Individual
  - Group

2. Prisoner's Dilemma (PD)

2.1) Background

2.2) Robert Axelrod

2.3) PD as a Model of Nature

2.4) Game Setup

2.5) Structure of the Game

2.6) Payoff Matrix

2.1 Background:

- The Prisoner's Dilemma was one of the earliest "games" developed in game theory.
- By simulating the Prisoner's Dilemma we are given an excellent method of studying the issues of conflict versus cooperation between individuals.
- Since the Prisoner's Dilemma is so basic, it can be used as a model for various schools of thought:
  - economics
  - military strategy
  - zoology
  - robotics
  - Artificial Intelligence.

2.2 Robert Axelrod

- Interested in political relationships and reproductive strategies in nature
  - Wanted to study the nature of cooperation amongst nations
  - He used the Prisoner’s Dilemma game as a model to help **explain the evolution of cooperating species** from an inherently selfish genetic pool

2.3 PD as a Model of Nature

- Accurate in the fact that an agent only cares about itself
  - (It is naturally selfish)
- Furthermore, *cooperation* can be mutually beneficial for all involved

2.4 Game Setup

- **The Game:**
  - Two people have been arrested separately, and are held in separate cells. They are not allowed to communicate with each other at all.
- **Each prisoner is told the following:**
  - We have arrested you and another person for committing this crime together.

- If you both confess, we will reward your assistance to us, by sentencing you both lightly: 2 years in prison.
- If you confess, and the other person does not, we will show our appreciation to you by letting you go. We will then use your testimony to put the other person in prison for 10 years.
- If you both don’t confess, we will not be able to convict you, but we will be able to hold you here and make you as uncomfortable as we can for 30 days.

- If you don't confess, and the other person does, that person's testimony will be used to put you in prison for 10 years; your accomplice will go free in exchange for the testimony.
- Each of you is being given the same deal. Think about it.

2.5 Structure of the Game

- If both players Defect on each other, each gets P (the Punishment payoff);
- If both players Cooperate with each other, each gets R (the Reward payoff);
- If one player Defects and the other Cooperates, the Defector gets T (the Temptation payoff), and the Cooperator gets S (the Sucker payoff);

Structure of the Game - Cont'd

- T > R > P > S and R > (T+S)/2.
  - These inequalities rank the payoffs for cooperating and defecting.
  - The condition of R > (T+S)/2 is important if the game is to be repeated. It ensures that individuals are better off cooperating with each other than they would be by taking turns defecting on each other.

Structure of the Game - Cont’d

- **Iterative** PD vs. **Single** PD
  - Single instance games of PD have a “rational” decision.
  - Always defect, since defecting is a **dominating strategy.**
  - However, with iterative PD always defecting is not optimal since an “irrational” choice of mutual cooperation will cause a net gain for both players.
  - This leads to the *“Problem of Suboptimization”*

2.6 Payoff Matrix

Subject B / Subject A	Cooperate	Defect
Cooperate	A: (R = 3) B: (R = 3)	A: (S = 0) B: (T = 5)
Defect	A: (T = 5) B: (S = 0)	A: (P = 1) B: (P = 1)

Iterative Prisoner's Dilemma Demo

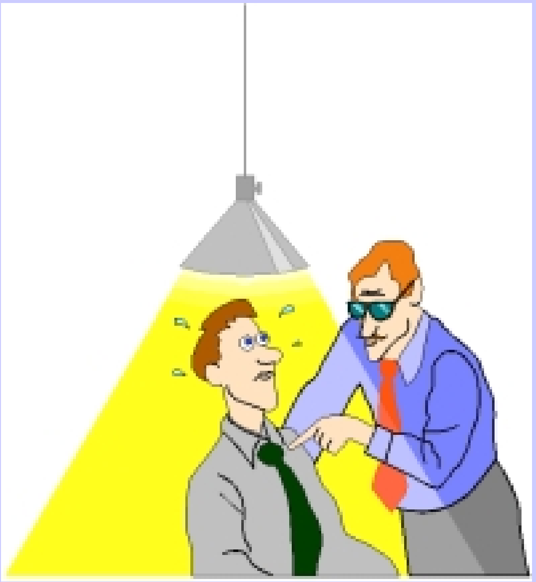

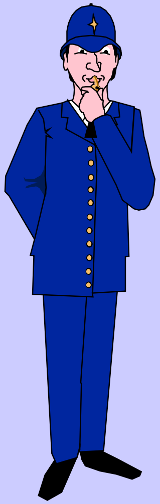

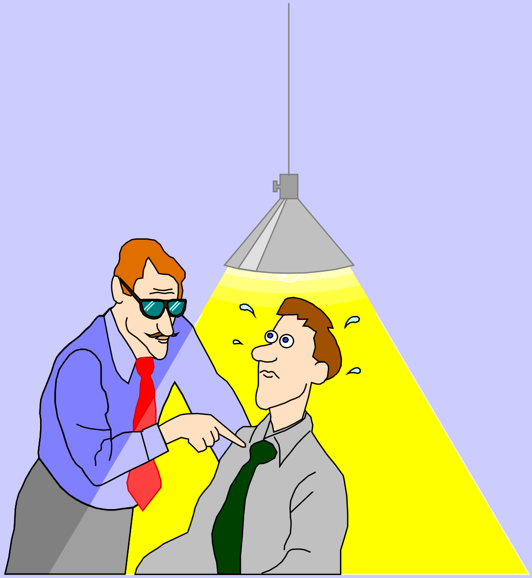

3 Deterministic Strategies for the Prisoner's Dilemma

3.1) Tit for Tat

3.2) Tit for Two Tat

3.3) Suspicious Tit for Tat

3.4) Free Rider

3.5) Always Cooperate

3.6) Axelrod's Tournament

3.7) PD in a Natural Setting

3.8) Downfall of Deterministic Strategies

3.1 Tit for Tat (TFT)

- The action chosen is based on the opponent's last move.
  - On the first turn, the previous move cannot be known, so always cooperate on the first move.
  - Thereafter, always choose the opponent's last move as your next move.

- **Key Points of Tit for Tat**

  - **Nice;** it cooperates on the first move.
  - **Regulatory;** it punishes defection with defection.
  - **Forgiving;** it continues cooperation after cooperation by the opponent.
  - **Clear;** it is easy for opponent to guess the next move, so mutual benefit is easier to attain.

3.2 Tit for Two Tat (TF2T)

- Same as Tit for Tat, but requires two consecutive defections for a defection to be returned.
  - Cooperate on the first two moves.
  - If the opponent defects twice in a row, choose defection as the next move.

- **Key Points of Tit for Two Tat**

  - When defection is the opponent’s first move, this strategy outperforms Tit for Tat
  - Cooperating after the first defection causes the opponent to cooperate also.
    - Thus, in the long run, both players benefit more points.

3.3 Suspicious Tit for Tat (STFT)

- Always defect on the first move.
- Thereafter, replicate opponent’s last move.

- Key Points of Suspicious Tit for Tat
  - If the opponent’s first move is defection, this strategy outperforms Tit for Tat
  - However, it is generally worse than Tit for Tat.
    - The first move is inconsequential compared to getting stuck in an infinite defection loop.

3.4 Free Rider (ALLD)

- Always choose to defect no matter what the opponent’s last turn was.
- This is a dominant strategy against an opponent that has a tendency to cooperate.

3.5 Always Cooperate (ALLC)

- Always choose to cooperate no matter what the opponent's last turn was.
- This strategy can be terribly abused by the Free Rider Strategy.
  - Or even a strategy that tends towards defection.

3.6 Axelrod's Tournaments

- Took place in the early 1980's
- Professional game theorists were invited by Axelrod to submit their own programs for playing the iterative Prisoner's Dilemma.
- Each strategy played every other, a clone of itself, and a strategy that cooperated and defected at random hundreds of times
- Tit for Tat won the first Tournament.
- Moreover, Tit for Tat won a second tournament where all 63 entries had been given the results of the first tournament.

3.7 PD in a Natural Setting

- All available strategies compete against each other (interaction amongst individuals as in nature)
- Recall that only strategies scoring above some threshold <u>will survive</u> to new rounds
- Surviving strategies then spawn new, <u>similar strategies</u>
- Success of a strategy depends on its <u>ability to perform well against other strategies</u>

3.8 Downfall of Deterministic Strategies

- Although Axelrod has argued reasonably well that TFT is the best deterministic strategy in the PD, they are inherently flawed in a natural setting
- **Theorem:** As proven by Boyd and Lorberbaum (1987) no deterministic strategy is evolutionarily stable in the PD.
  - In other words, they may die out in an evolution simulation

- Basic idea is that if two other strategies emerge that are just right, they can outperform and kill off another
- Consider TFT being invaded by TF2T and STFT
- TFT and TF2T both play STFT repeatedly
  - TFT falls into continual defection when it wouldn't have to.
    - They both score 1 each round
  - TF2T on the other hand, loses once and cooperates from then on
    - They both score 3 each round

4 Beyond Determinism

4.1) Nowak and Sigmund (1993)

4.2) Stochastic Strategies

- 4.2.1) Generous Tit For Tat
- 4.2.2) Extended Strategy Definition
- 4.2.3) Pavlov

4.3) Results: Nowak and Sigmund

- 4.3.1) Evolution Simulation

4.1 Nowak & Sigmund (1993): New Experiment

- Nowak and Sigmund extended the definition of a strategy slightly and performed large evolution simulations
- When populations can mutate, (as in an evolution simulation) we get <u>noise</u>
  - Suppose a strategy that always cooperates defects once due to mutation
  - Deterministic strategies (TFT in particular) can't handle this well as it could cooperate

4.2 Stochastic Strategies

- By definition, they involve an element of randomness
- Generous Tit For Tat (GTFT)
  - Instead of immediately defecting after an opponent does, there is a probability (q) that it will forgive the defection by cooperating on the next move
  - $q = \min[1-(T-R)/(R-S), (R-P)/(T-P)] = 1/3$
    - Should have about 1/3 chance of forgiveness

4.2.1 GTFT

- As we have seen, TFT is too severe in an evolution simulation
- In such simulations however, it is interesting to note that TFT needs to be present at some point to suppress defectors
- After the suppression, GTFT often emerges and stabilizes in the population, replacing TFT

4.2.2 Extended Strategy Definition

- Strategy takes not only opponent into consideration, but itself as well
  - There are 4 possible outcomes from each round
  - A probability for cooperating can be defined after each possible round outcome
  - Thus, a strategy can be given as a 4 dimensional vector (p1, p2, p3, p4) for cooperating after R, S, T, and P
  - So, TFT would be (1,0,1,0)

4.2.3 Pavlov

- The strategy (1, 0, 0, 1) was investigated and named Pavlov by Nowak and Sigmund
- It cooperates after both mutual cooperation and mutual defection
  - Can exploit a TF2T strategy by apologizing once TF2T starts defecting
  - Also exploits generous cooperators well by continuing to defect if it gets payoff T

- Deals well with noise by defecting once to punish a defection, but then by apologizing if both start defecting
- Has a weakness where it alternates between cooperating and defecting with ALLD
  - Thus, it is not evolutionarily stable against ALLD
- In a simulation however, Pavlov emerges as (.999, 0.001, 0.001, 0.995) which can survive against ALLD

4.3 Results: Nowak & Sigmund

- They ran 40 simulations
  - Started with random strategies of (0.5, 0.5, 0.5, 0.5)
  - The duration of the simulation was about 500,000 generations
  - Every 100th generation a new strategy was introduced (one of 10,000 predefined strategies)

- Typical development started with a chaotic period
- Followed by dominance of defectors as they take advantage of cooperators
- Eventually fairly strict TFT strategies choke out defectors
- Finally, TFT is too strict and is replaced from time to time by GTFT or more often (about 80% of the time) by Pavlov

4.3.1 Evolution Simulation

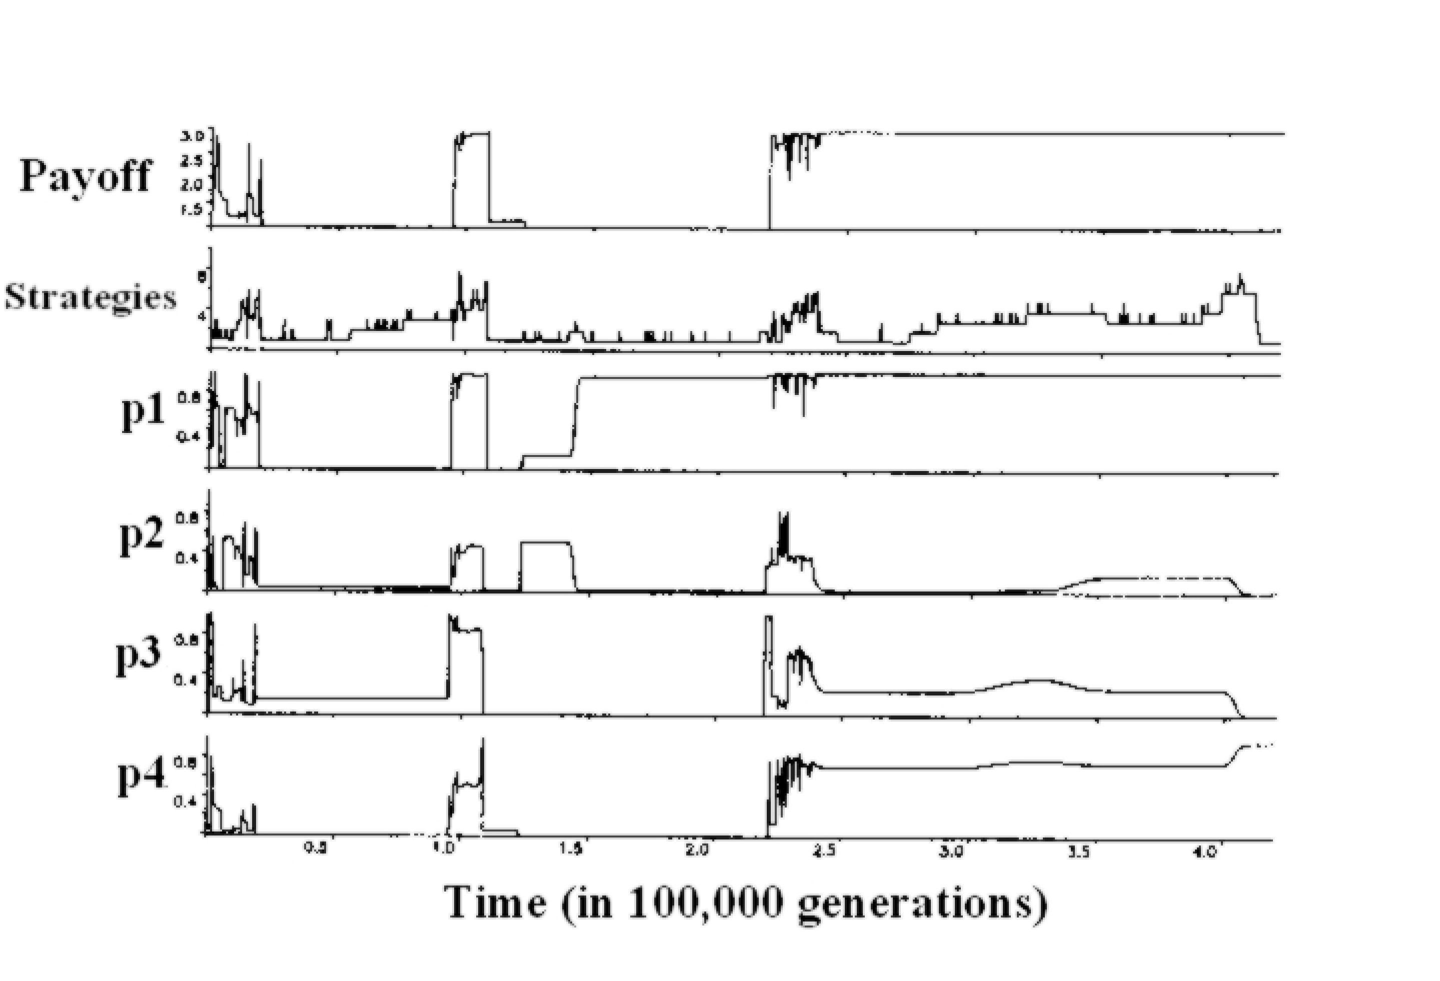

5 PD in Nature

5.1) Spatial Chaos

5.2) Case Study: Shoaling Fish

5.3) Conclusion: PD as an Agent Based Model

5.1 Spatial Chaos

- In some simulations the proximity of individuals is considered
  - strategies only compete with neighbors on a 2-D board
- At the end of a round, an individual will adopt the strategies of a successful neighbor
- In this scenario, a cluster of ALLC can even invade ALLD

5.2 Case Study: Shoaling Fish

- So far an exact representation of the conditions of the Prisoner’s Dilemma has not been identified in nature
- Predator inspection on shoaling fish is close, but the scenario is debatable
  - A pair of fish can break from the group to swim near and inspect the predator chasing the shoal
  - They get a payoff in the form of gaining knowledge about the predator

- Two fish can move closer to the predator, so they benefit from cooperation
- In addition, one can “defect” by not moving so close as the other and then gets the temptation payoff which is the knowledge without risk
- Thus T>R>P>S is satisfied but …
- Can they recognize previous defectors in order to punish them?
- Do they really prefer to approach in pairs?
- Does an inspector share information with the group regardless?

- Despite these shortcomings, some have claimed that guppies use a TFT strategy when approaching a predator

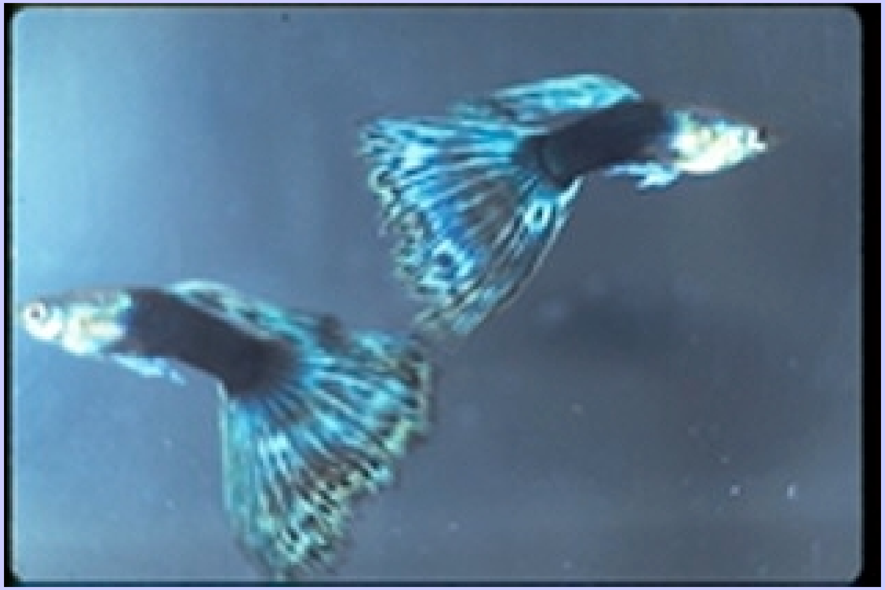

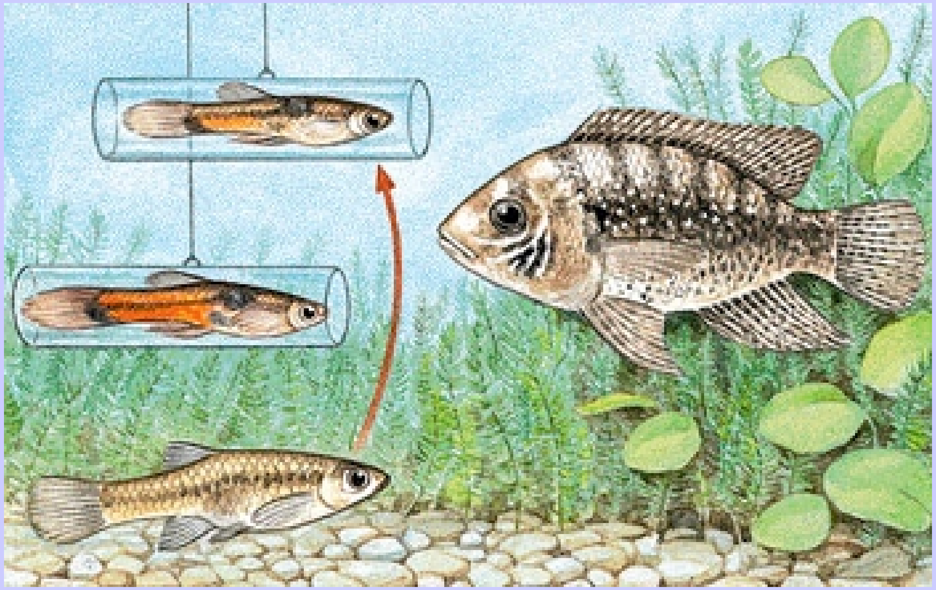

Evolution of Behavior in Nature

Levels of selection

- **Individual selection- each individual behaves so as to maximize its own fitness**. (Behavior that increases the donor's fitness at the expense of the recipient's fitness is termed "selfish" behavior.)
- **Group selection- each individual behaves so as to increase the fitness of its group**. (Behavior that reduces the fitness of the donor while it increases the fitness of the recipient, is termed "altruistic" behavior.)

Levels of selection

- Would it therefore be prudent for an individual to avoid reproduction so that its group does not become overpopulated?
- *Analysis*: what happens to “cheaters”?
- ***Observation***: groups that overpopulate their environment and overexploit their resources have high rates of extinction.

- Altruistic behavior is frequently directed towards relatives. Parents feed and care for their offspring; siblings support each other, grandparents care for grandchildren while the parent work, etc.
- Behavior that increases the fitness of kin at the expense of the donor’s fitness produces ***kin selection***.
- Kin selection occurs when the inclusive fitness of the altruistic individual is increased by the behavior.
- The gain in inclusive fitness is a function of the degree of relatedness of the donor and recipient, as well as the benefits and costs of the altruism.

Calculation of coefficients of relationship

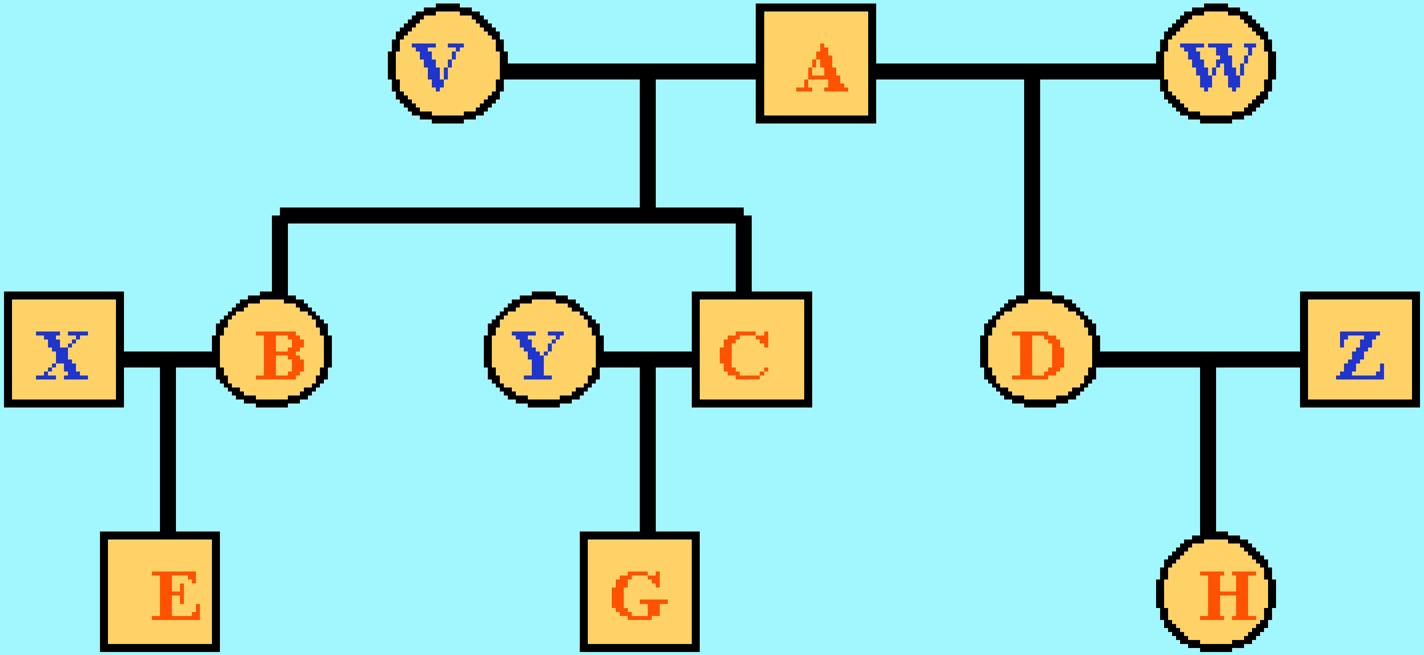

r_{JK} = the coefficient of relationship of individuals J and K is the probability that a randomly chosen gene from J is identical by descent to a second randomly chosen, homologous gene from K.

NOTE: $r = \sum (1/2)^n$ (sum is taken over all paths of n steps each).

A → D → H

B → A → C

B → V → C

$r_{AH} = 1/4$ $r_{BC} = 1/2$ $r_{CD} = 1/4$ $r_{EG} = ?$ $r_{GH} = ?$

Cooperative Breeding

From Purves, Orians, & Heller

Cooperative Breeding

- 1. Seychelles warbler
- 2. Florida scrub jay- up to 6 helpers per territory, each staying 1-3 yrs. More young survive in nests with helpers than without. Males stay longer than females.

Ornaments are weapons; usually great sexual dimorphism in male-female sizes

Sexual Selection

A) Trial by combat

Peafowl

Peacock displaying to peahen

Intersexual selection- charm

- Individuals having characters attractive to the opposite sex are more likely to mate and produce offspring.
- ***Why are some characters attractive?***
  - **1. Because they may be indicators of vitality and thus correlated with genes conferring high fitness.**
  - **2. Alternatively, they may be selected to ensure that the female's offspring will have the same characters and therefore enjoy greater mating success.**
- Mate choice is usually made by the gender with the most to loose *if a mistake is made*; this is usually, but not always, the female.

Asian jungle fowl

Gallus gallus

ancestor of the domestic chicken

Hens choose males with bright eyes and large red combs and wattles as shown (study of Marlene Zuk)

Male tungara frog

more complex calls are preferred by females… and bats

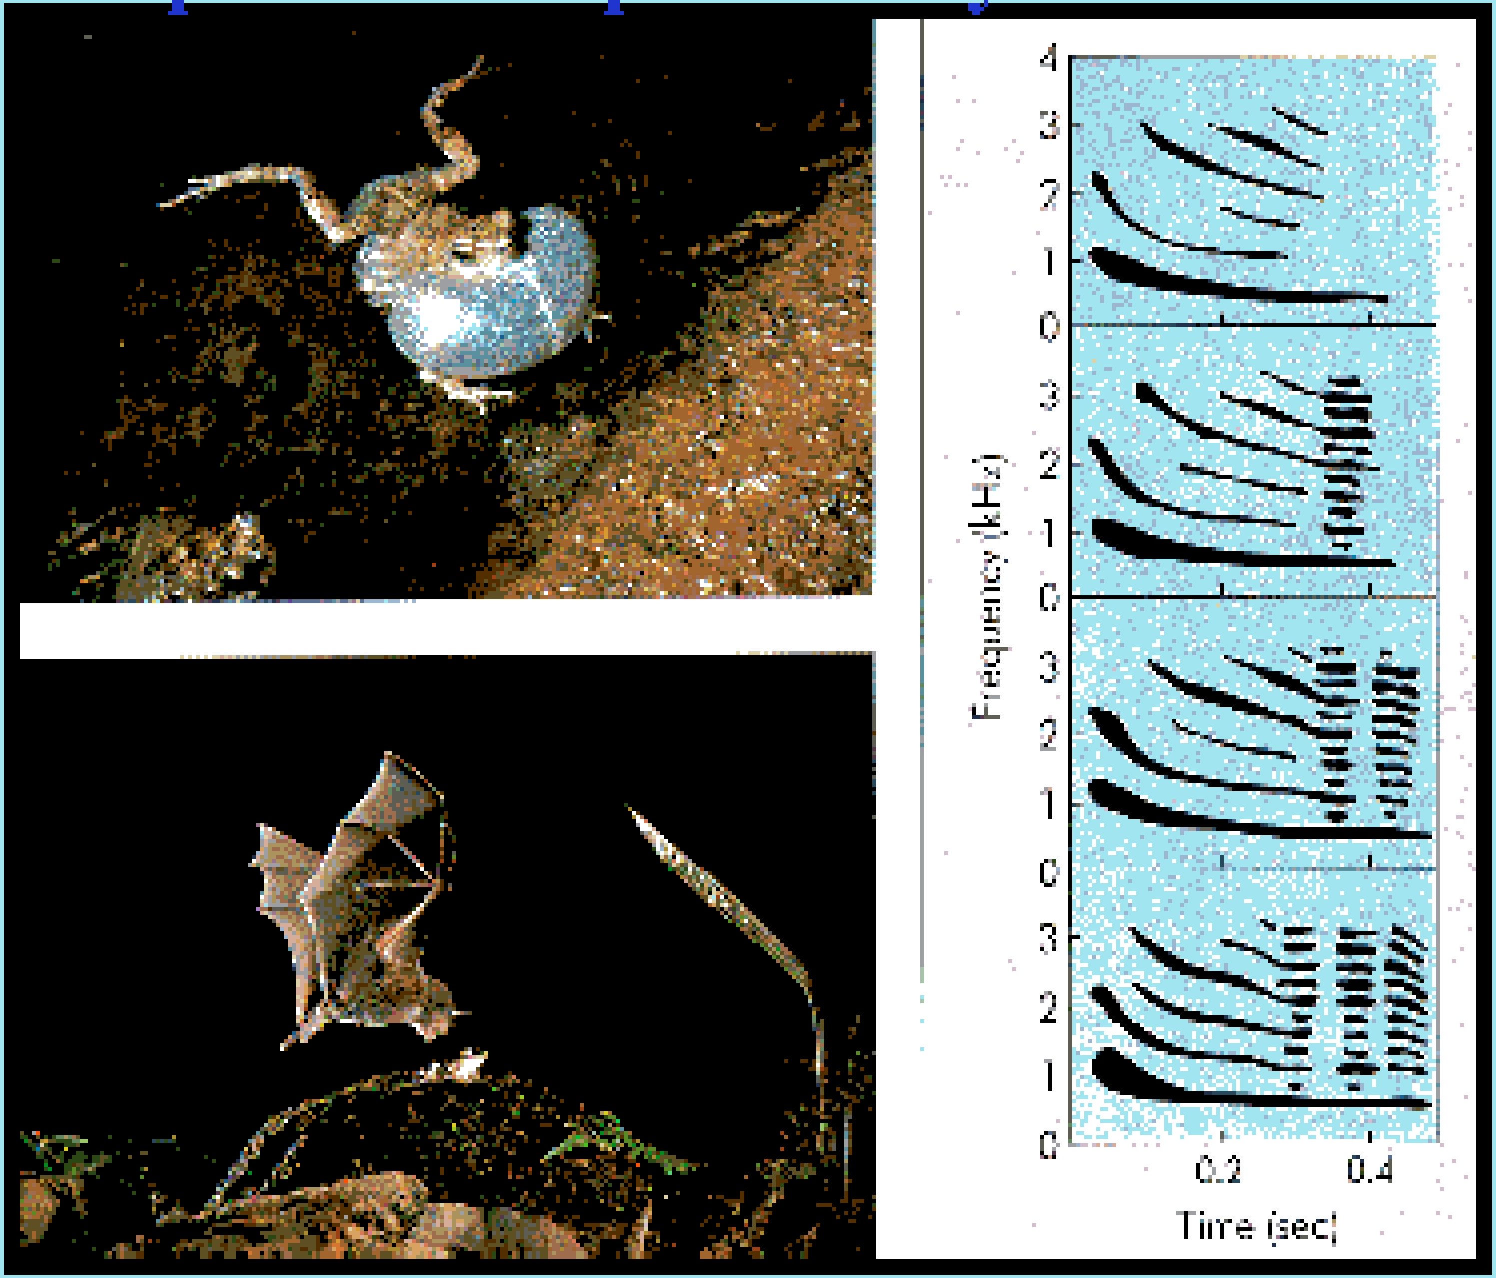

Humpback whales (*Magaptera*)

Does gene flow occur between whales living in different pods and geographical areas?

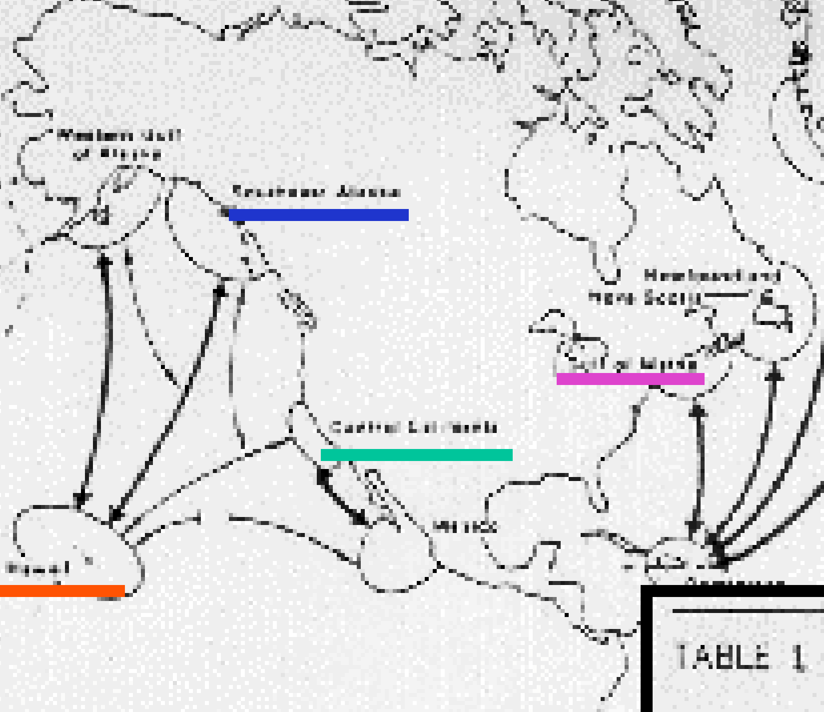

~20 individuals from each site were each scored for 468 nucleotides (which is 468/16,500 = 2.8 % of the mitochondrial genome).

Baker et al, 1990

SEA is the common Pacific summer feeding ground, CCA and Hi are winter breeding grounds; GOM is the outgroup

TABLE 1. Nucleotide divergence (%) of mtDNA within and between subpopulations of humpback whales

	SEA	HI	CCA	GOM
Southeastern Alaska (SEA)	**0.000**	0.000	0.176	0.282
Hawaii (HI)	0.014	**0.028**	0.156	0.262
Central California (CCA)	0.274	0.267	**0.195**	0.112
Gulf of Main (GOM)	0.380	0.374	0.308	**0.196**

Diagonal shows nucleotide diversity within a geographic region; values below the diagonal are genetic distances between subpopulations uncorrected for within-region divergence. Values above the diagonal show genetic distances between regions after correcting for within-region divergence.

The Green Sea Turtle (*Chelonia mydas*)

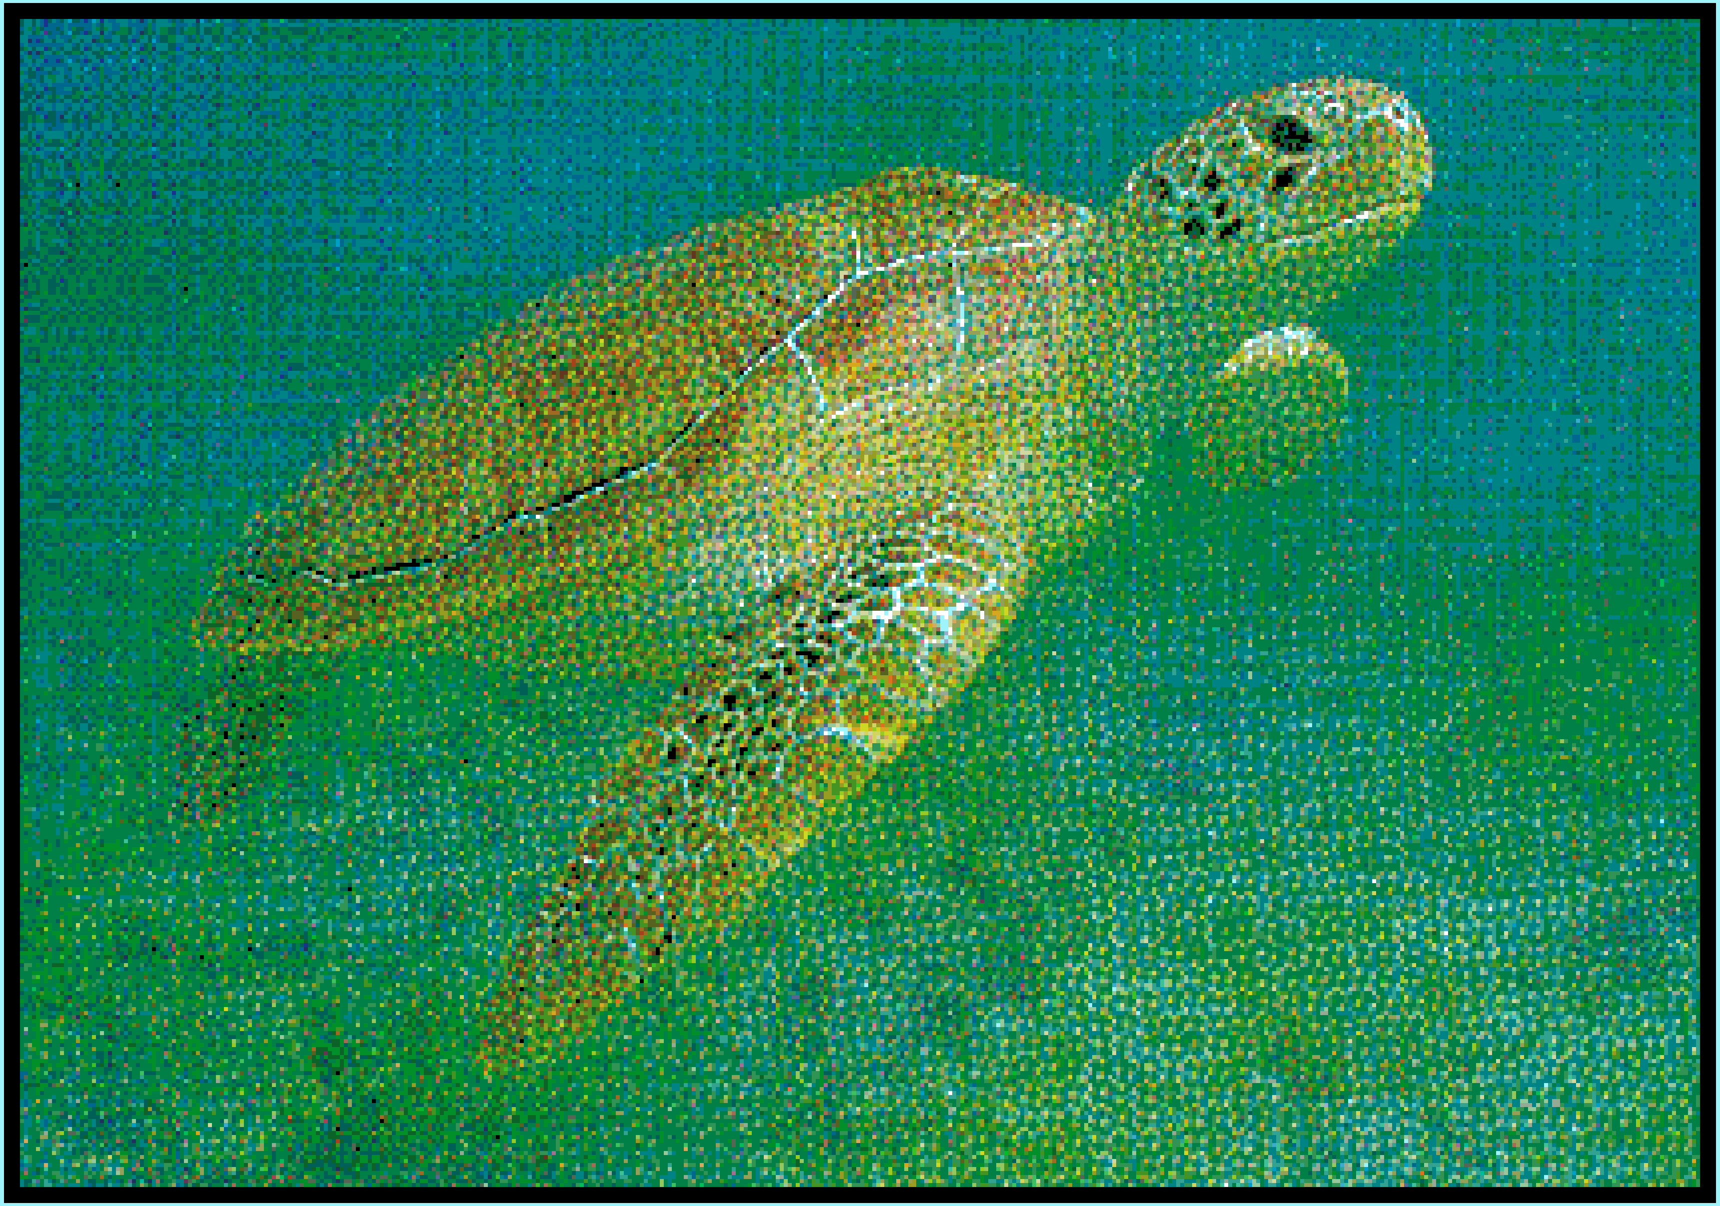

Do mature sea turtles return to their birth sites to lay eggs?

Oviposition and feeding sites

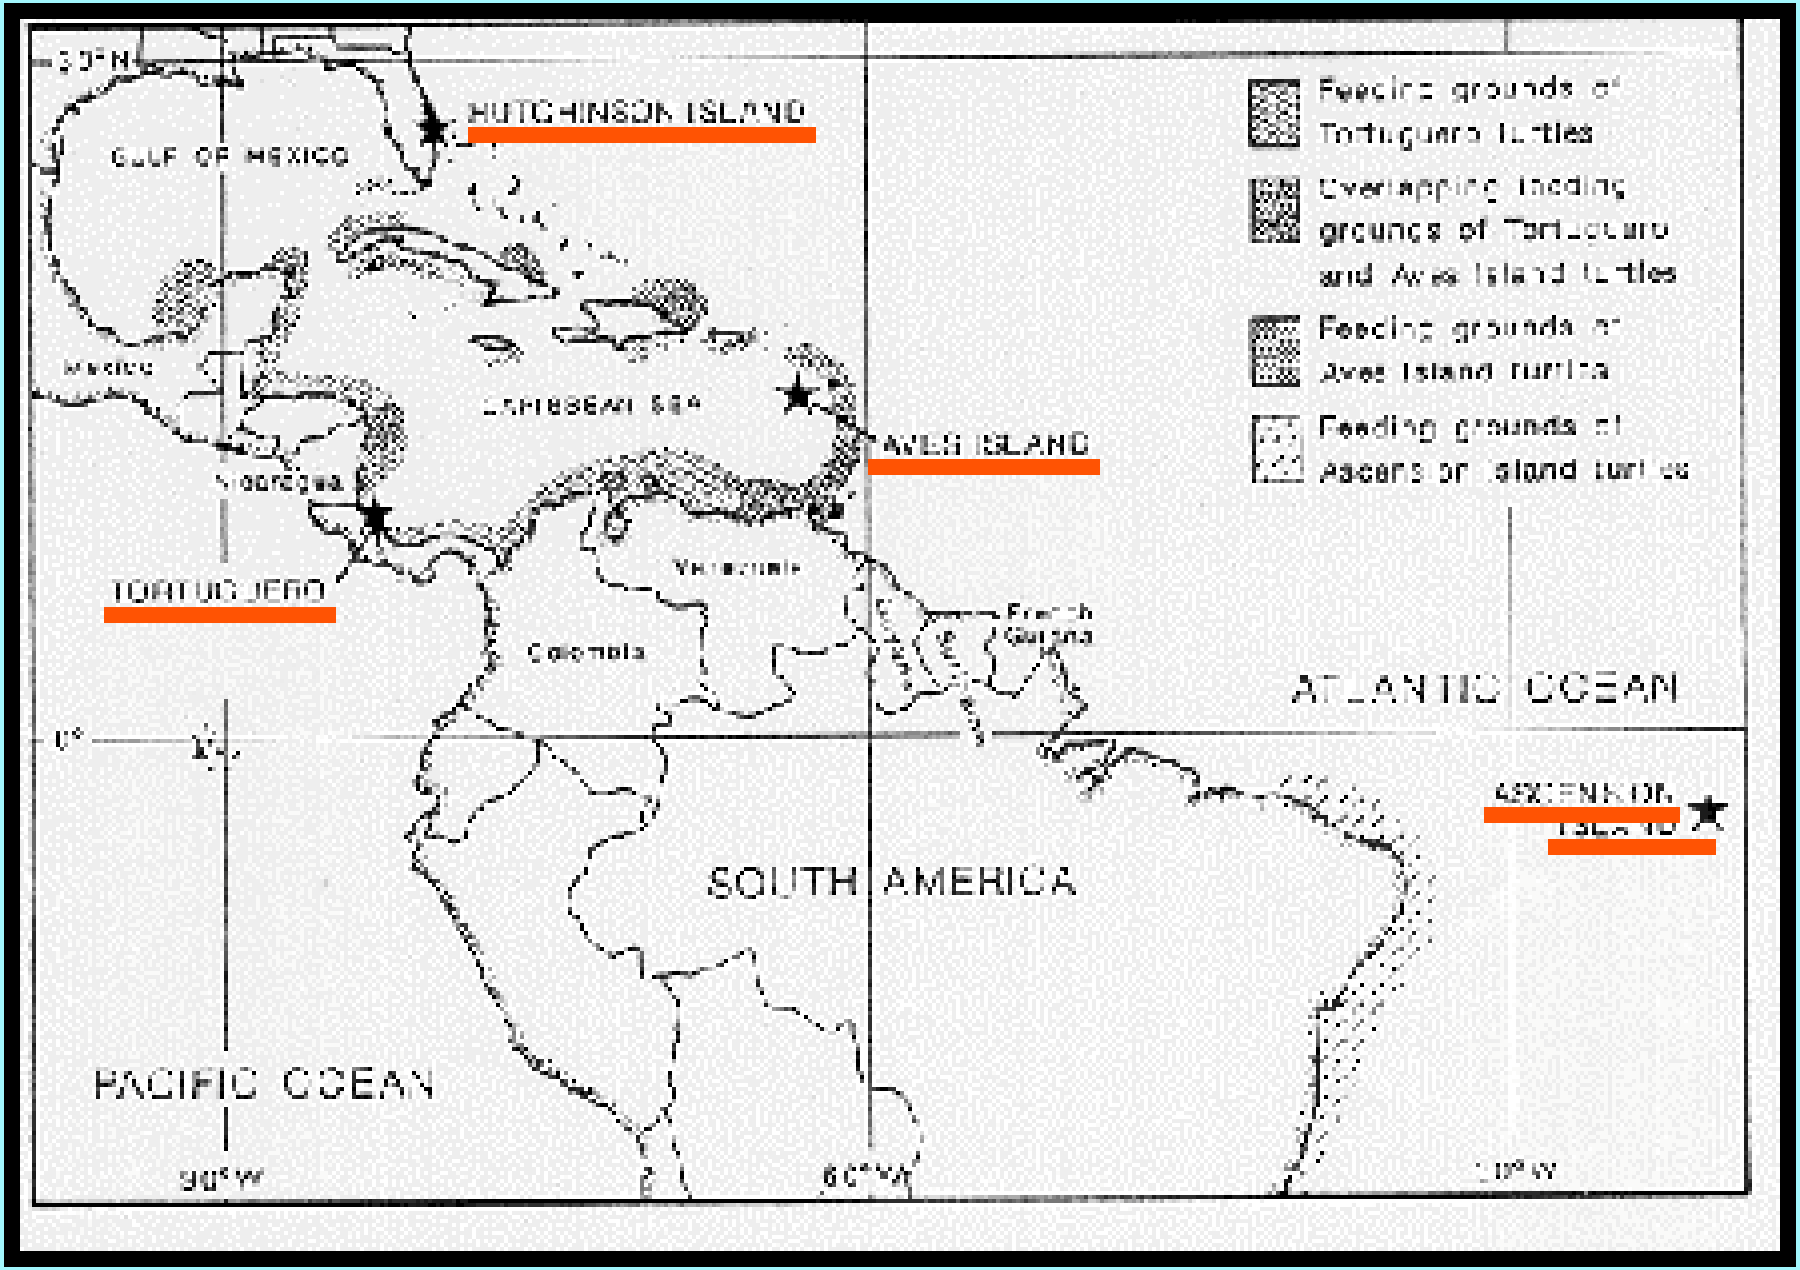

Alternative behavioral models of dispersal

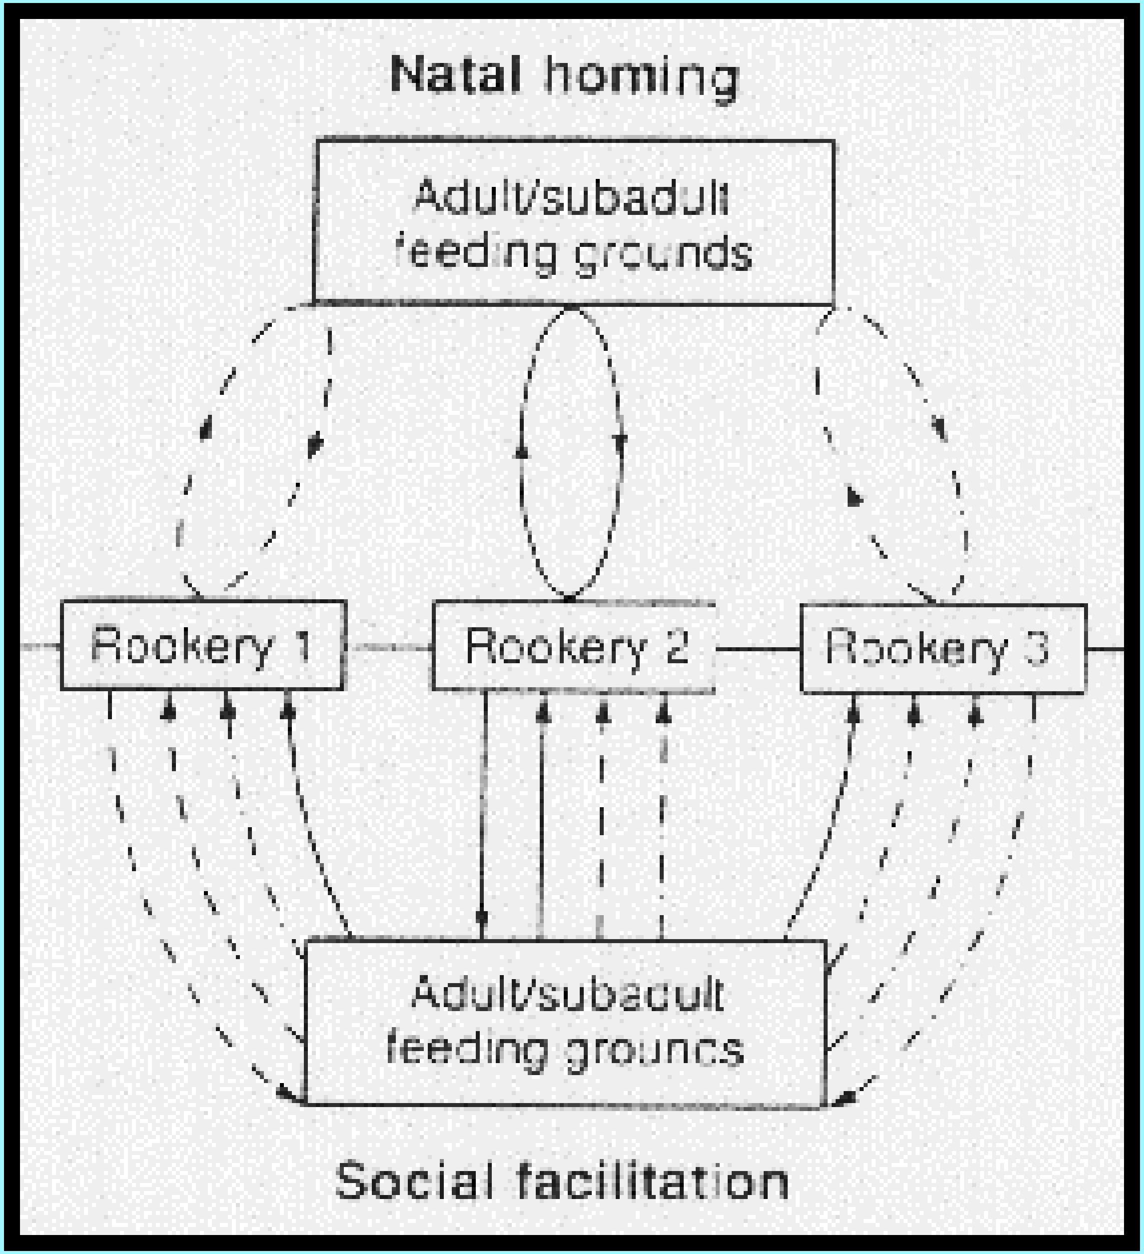

Adults can be tagged and it is known that they return to the same beach many times to lay their eggs.

But, what determines which beach is selected the first time?

Almost 30 years is required for a newborn turtle to reach reproductive maturity. Tagging is not feasible.

98 mtDNA restriction sites yielding three haplotypes (A, B, and C) were obtained using 14 enzymes

Numbers of individuals showing each genotype:

	A	B	C
Aves Is (Venezuela)	1	7	0
Tortuguero (Costa Rica)	15	0	0
Ascension Is (U.K.)	0	0	16
Hutchinson Is (Fla)	10	0	0

Conclusion: these haplotype distributions indicate that natal homing occurs, but there is a small amount of error providing the opportunity for gene flow and colonization of virgin areas.

Meylan et al. 1990

5.3 Conclusion: PD as an Agent Based Model

- In an abstract form, the PD simulations have proven valuable and powerful
  - Properties of successful strategies have been identified (nice, retaliatory, forgiving, etc.)
  - New strategies previously not considered were found (such as GTFT and Pavlov) and shown to be very good
  - Simulation has allowed a progression from Deterministic to Stochastic strategies

- However, the lack of natural systems corresponding to the PD clearly demand the development of new models
- The PD has weaknesses that need to be addressed when developing new models:
  - Individuals cannot alter their environment
  - Other forms of cooperation (by-product mutualism) are ignored
  - There is no information exchange between individuals

- Varying degrees of cooperation and defection are not taken into account
- Proximity of individuals sometimes matters
- N-player situations (group behavior) is ignored

- New paradigms need to be developed that take these variables into account
- "The aim, of course, is to combine such new paradigms to a model, that would provide a powerful tool to investigate, under which precise conditions, which forms of cooperation could evolve"

  - Bjorn Brembs

- **Other game theory paradigms than PD should be created to explain phenomena existing in Nature**

Review of Literature

- Anderson, Rules of the Mind, 1993, Erlbaum Associates
- Hayes-Roth, Barbara, BB1: An architecture for blackboard system that control, explain, and lean about their own behavior, No. 84-16 BB1, Department of Computer Science Stanford University
- Johnson, Proceedings of the First International Conference on Autonomous Agents (Autonomous Agents '97), 1997, ACM Press

Review of Literature

- Maes, Designing Autonomous Agents, 1990, MIT Press, reprinted as a special issue, Robotics and Autonomous Systems 6(1-2), 1990.
- Newell, Unified Theories of Cognition, 1990, Harvard University Press
- Nii, H. Penny, AI Magazine: Blackboard Systems, 7(2-3), 1986.
- Russell and Norvig, Artificial Intelligence: A Modern Approach, 1995, Prentice Hall

- VanLehn, Architectures for Intelligence, 1991, Lawrence Erlbaum Associates
- Wooldridge and Jennings, Intelligent Agents: ECAI-94 Workshop on Agent Theories, Architectures, and Languages, 1995, Springer-Verlag
- Wooldridge, Muller, and Tambe, Intelligent Agents II: Agent Theories, Architectures and Languages, 1996, Springer-Verlag

Sources

Darren Gerling

Jason Gerling

Jared Hopf

Colleen Wtorek

Eric Sparling