# ECM3412 Nature-Inspired Computation

Revision Lecture

# Overview

- Module Overview
- Sections in More Detail
- Examinable Reading
- Past Exam Questions

# Module Overview

- Evolutionary algorithms
  - Introduction to evolution
  - Search and Optimisation, Easy & Hard Problems
  - Landscapes and Hillclimbing
  - Genetic Operators in more detail
  - Encodings and Representations
  - Genetic Programming
  - MultiObjective EAs
- Swarm Intelligence I
  - Ant Colony Optimisation
  - Applications of ACO
- Swarm Intelligence II
  - Flocking Behaviours
  - Particle Swarm Optimisation
- Neural Networks
  - The Brain and Artificial Neurons
  - MultiLayer Perceptrons
  - Applications of ANNs
  - Self-Organising Maps
- Artificial Life
  - ALife
  - Cellular Automata & Applications

# Evolutionary Algorithms (EAs): Introduction to evolution

- Evolution
  - Small random changes in individuals makes them more able to survive and replicate
  - Over time, individuals which are better at surviving in the environment prevail
  - If we look at the environment as a problem, we can do problem solving (Gecco).
- Magic Ingredients
  - Random mutation
  - Population of individuals

# EAs: Search and Optimisation

- Search usually involves generating solutions to a problem to which there exists a fitness function to evaluate ‘how good’ that solution is

- Problems (including search) can be grouped into tractable and intractable

- Almost all interesting search/optimisation problems are hard (intractable)

- Problem type:
  - Tractable – there are known fast algorithms to give us a solution
  - Intractable – need to find approximate solutions

# EAs: Landscapes & Hillclimbing

- Fitness ‘landscapes’ are theoretical views of the progress of algorithms

- There are various features of a landscape
  - Unimodal – single hill – hillclimbing will find the best
  - Multimodal – multiple hills, some better than others
  - Plateaus/Deceptive Peaks – can be difficult to find optimum

- Hillclimbing works by evaluating some local solutions and going ‘uphill’.
  - Gets stuck in local maxima

# EAs: Operator Details

- Steady State vs Generational Algorithms
- Selectors:
  - Tournament
  - Roulette Wheel
- Crossover
  - Single Point
  - N-Point
  - Uniform
- Mutation (encoding dependent)
  - Single point
  - M-Point
  - Swap

# EAs: Representations and Encodings

- Representations
  - Binary
  - Real values
  - Integers
  - Symbols
- Encodings
  - Direct – GA modifies directly the variables of the problem
  - Indirect – GA modifies variables heuristically – "first clash- free slot" in the exam timetabling example.

# EAs: Genetic Programming

- Evolves trees instead of flat chromosomes
- Trees represent functions or programs rather than explicit solutions to a problem
- Requires specialised crossover and mutation operators
- Has found success in many areas deemed to require human intelligence
- Have been described as ‘human competitive’

# EAs: MultiObjective EAs

- A lot of real-world problems (esp. in Engineering) have more than one objective (fitness).
- EAs can be used to optimise these with a MOGA.
- Uses the principle of 'domination'.
- Some modification of the algorithm required (e.g. selection operator)
- The end product of a MOGA is a pareto set of solutions representing the optimal trade off.

# Swarm Intelligence 1: Ant Colony Algorithms

- Based on experiments with real ants
- Ants soon find the shortest path to from the nest to food

- Ants need to communicate – do so through stigmergy (the modification of the environment)
  - Pheromones are laid to tell following ants which way to go

- ACO Algorithms – need a pheromone update procedure related to the fitness of the ant’s path through the problem

# Swarm Intelligence 1: Ant Colony Applications

- Some problems can be solved by ACO algorithms which are not topology-based.
- Most problems can be converted into a construction graph e.g.
  - Each variable choice can be represented as a node and the ant can find a path through the nodes
- Specific Applications:
  - Travelling Salesman
  - University Timetabling
  - Clustering (corpse aggregation)

# Swarm Intelligence 2: Flocking Behaviours

- Flocking
  - Natural phenomenon seen in animals in all media – air, land & sea
  - Even humans do it somewhat
  - Normally gives benefits to all members of the flock which outweighs the disadvantages
    - Increase in range for migrating geese
    - Predator avoidance for shoaling fish
- Craig Reynold's Boids – artificial representation of flocking

# Swarm Intelligence 2: Particle Swarm Optimisation

- Particles represent solutions in the solution space

- Each solution has:
  - A position in the search space
  - A velocity
  - Knowledge of the previous best encountered by the particle
  - Knowledge of the best solution found globally

- Each iteration, a particle computes its new position based on its velocity

- Velocity based on a weighted sum of pbest and gbest

# Artificial Immune Systems: Applications (CA2)

- Artificial Immune Systems use principles taken from the human immune system

- Generally applied to those fields which require the identification of anomalous events:
  - Computer security – virus and anomalous behaviour detection using immune system metaphors.

# Neural Networks: Inspiration and Early Work

- Differences in human/machine capabilities and architecture
  - Serial vs Parallel Processing

- Neuronal function
  - Axon
  - Synapses

- Artificial Neuron (McCulloch and Pitts)
  - Information processing system based on biological neuron

- Perceptron (Rosenblatt)
  - Learning via weighted connections (Hebb)

# Neural Networks: Learning & Applications

- Simple learning in perceptrons
  - Weight adaptation dependent on result
  - Limitations (XOR)

- Learning in Multi-Layer Perceptrons
  - Feedforward backpropagation
  - Batch and Online updating

- Neural network properties
  - Ability to generalise
  - Ability to gracefully degrade

- Data Representation & Overfitting

# Neural Networks: Variations

- Recurrent Neural Networks
  - Not needed in detail
  - Use context units & delayed connections
- Radial Basis Function (RBF) Neural Networks
- Unsupervised Learning
  - Self Organising Map
  - Unsupervised Learning Algorithm
  - Applications (e.g. phoneme recognition, gene clustering, webpage classification)

# Artificial Life and Cellular Automata

- ALife is the study/simulation/ synthesis of life in a computer
- Most of the Nature Inspired Techniques can be considered Alife
- Concentrates on emergent behaviour of collections of simple elements, in contrast to ‘reductionist’ biology
- Cellular automata embody the emergent princple

# Cellular Automata

- Represented by a grid or lattice of ‘cells’, which can be in a number of ‘states’
- States update using
  - State transition rules
  - States of cells in the neighbourhood
- CAs have been used for
  - Enzyme simulation
  - Epidemiology
  - Optimisation

# Examinable Reading

Specific papers that:

- Relate fairly closely in lectures
- Not (in fact, far from) fully covered in lectures
- For you to read, hopefully enjoy, and try to understand
- May turn up in the exam
- May not turn up in the exam

# What I might ask

Summarize the main algorithm, experiments done and the main findings of the paper *The Design of Analog Circuits by Genetic Programming by John Koza* [12 marks]

- Note this is *part* of a question.
- Note that this would be part of a very unimpressive answer:
  - *It's about an algorithm like the one in the lectures but a bit different …*

    *They ran a few experiments to test the algorithm …*

    *It seemed to work …*

# What I might ask II

- This may be part of a good answer, showing understanding and a grasp of the important points:
  - *The paper describes the design of six analog circuits using genetic programming, in which the trees represent instructions to 'draw' components onto a blank circuit board. The functions and terminals relate to the placement of electrical components on the board (e.g. a capacitor creating function) etc…*
  - *The fitness was computed by simulating the circuit using SPICE and various required features of the circuits in question (e.g. travel times for the robot controller) etc…*
  - *Many of the results obtained satisfied the requirements of the problem or displayed other desirable properties such as low distortion for an amplifier etc…*
  - *The system performs comparably against human designed circuits and demonstrates real human-competitiveness*

# How Much I Might Ask

- Possibly one or two part questions.
- These questions will be explicitly about a paper (or two)
- No `hidden' questions concerning the papers. E.g. part of another question may be usefully answered by using some knowledge, idea or example from one of these papers. If so, then it is excellent to make use of that knowledge, but this was not deliberate in my setting of the question.

# General Notes on examinable reading

- If you revise a paper so thoroughly that you can reproduce it in the exam word for word, you are wasting your revision time!

- If you read it, understanding the general gist, but have trouble with some of the detail, then that’s what I expect.

- If some of your trouble with the detail hampers your overall understanding of the main points, then I would expect you to work harder at understanding those bits – via a bit of personal research (e.g. google, or emailing me a clear question).

# Questions?