An Introduction to Artificial Immune Systems

Jon Timmis
University of Kent. UK.
ICARIS 2004.

What do I want to achieve?

- Give you a taster of what AIS are all about:
  - Why do we find the immune system useful?
  - Explain what AIS are;
  - Look at basic AIS algorithms;
  - Look at what we might do.
- I won't:
  - Talk about all areas of AIS and applications (Leandro will do that)
  - Talk too much about how AIS relate to other bioinspired ideas (although I will mention it)
  - Go into too much detail: this is an introduction

Outline

- **What are AIS?**
- Useful immunology
- Thinking about AIS
  - AIS Algorithms
  - Examples
- Examples of AIS

A Definition

AIS are adaptive systems inspired by theoretical immunology and observed immune functions, principles and models, which are applied to complex problem domains

Why the Immune System?

- Recognition
  - Anomaly detection
  - Noise tolerance
- Robustness
- Feature extraction
- Diversity
- Reinforcement learning
- Memory
- Distributed
- Multi-layered
- Adaptive

Some History

- Developed from the field of theoretical immunology in the mid 1980's.
  - Suggested we 'might look' at the IS
- 1990 – Bersini first use of immune algos to solve problems
- Forrest et al – Computer Security mid 1990's
- Hunt et al, mid 1990's – Machine learning

History

- Began to rapidly develop
- Not just security but all sorts of applications
- SMC, GECCO, CEC workshops 1999 onwards (Dasgupta)
- 1st ICARIS in 2002
- Now becoming accepted as a ‘sensible thing to do’
  - But we can do it much better than we are!

History

- Started quite immunologically grounded
- Bersini’s work
- Forrest's work with Perelson etc
- Kind of moved away from that, and abstracted more
- Now there seems to be a move to go back to the roots of immunology

Scope of AIS

- Fault and anomaly detection
- Data Mining (machine learning, Pattern recognition)
- Agent based systems
- Scheduling
- Autonomous control
- Optimisation
- Robotics
- Security of information systems

Outline

- What are AIS?
- **Useful immunology**
- Thinking about AIS
  - AIS Algorithms
  - Examples
- Examples of AIS

Immunity

- Protect our bodies from infection
- Innate and acquired immunity
  - Innate reacts to invaders but does not adapt
  - AIS has largely ignored this for some reason
- Acquired immunity has been of a great deal of interest:
  - Primary immune response
    - Launch a response to invading pathogens
  - Secondary immune response
    - Remember past encounters
    - Faster response the second time around

Multiple layers of the immune system

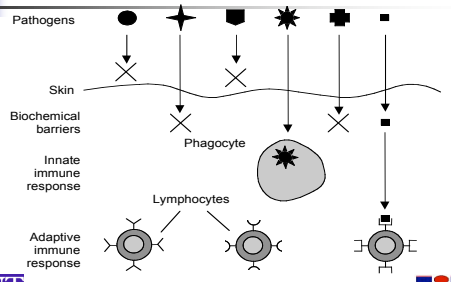

Immune Pattern Recognition

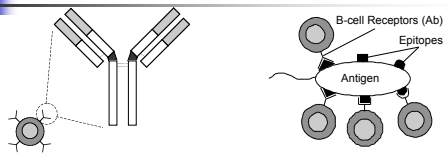

- Produces *antibodies* that recognise and attack *antigens*
- The immune recognition is based on the *complementarity* between the binding region of the receptor and a portion of the antigen called *epitope*.
- Recognition is not just by a single antibody, but a collection of them

Processes within the Immune System (very basically)

- Negative Selection
  - Censoring of T-cells in the thymus gland of T-cells that recognise *self*
- Positive Selection
  - If cells recognises an antigen it is activated and stimulates other components of the immune system
- Clonal Selection
  - Proliferation and differentiation of cells when they have recognised something

How does it work: A simplistic view

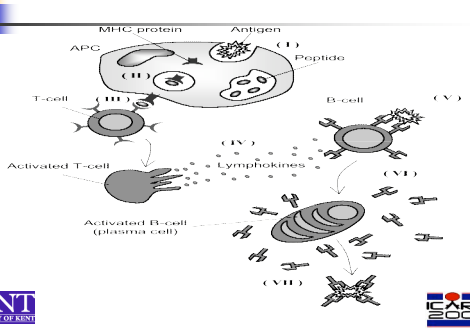

Immune Memory

- A number of theories exist:
  - Long-lived memory cells
  - Antigen ingesting cells
  - Idiotypic network
    - B cells co-stimulate each other
    - Treat each other a bit like antigens
- All go towards creating an immunological memory
- But immunologists are not sure, probably much more complicated than all of these!

Immune Network Theory

- Idiotypic network (Jerne, 1974)
- B cells co-stimulate each other
  - Treat each other a bit like antigens
- Creates an immunological memory

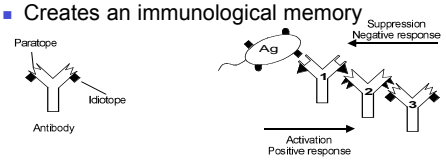

Shape Space Formalism

- Repertoire of the immune system is complete (Perelson, 1989)
- Extensive *regions of complementarity*
- Some threshold of recognition

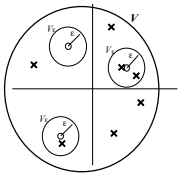

Danger theory (Matzinger 1994)

- it is not “non-self”, but “danger” that the IS recognises
  - dangerous invaders cause cell death or stress
  - these cells generate “danger signal” molecules
    - unlike *natural* cell death, ***apoptosis***
  - these stimulate an immune response *local* to the danger
    - to identify the “culprit”

Antibodies
Antigens
Cells
Damaged Cell

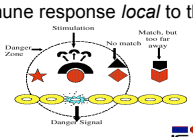

[Slide from S. Stepney, 2004]

Outline

- What are AIS?
- Useful immunology
- **Thinking about AIS**
  - AIS Algorithms
  - Examples
- What might we do?

What do want from a Framework?

- In a computational world we work with representations and processes. Therefore, we need:
  - To be able to describe immune system components
  - Be able to describe their interactions
  - Quite high level abstractions
  - Capture *general purpose* processes that can be applied to various areas

General Framework for AIS

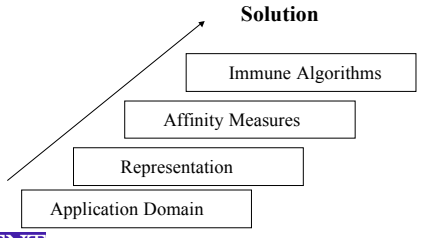

Representation – Shape Space

- Describe the general shape of a molecule

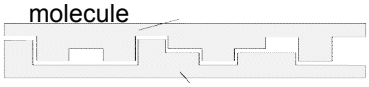

•Describe interactions between molecules

•Degree of binding between molecules

•Complement threshold

Representation

- Vectors

  Ab = $\langle Ab_1, Ab_2, ..., Ab_L \rangle$

  Ag = $\langle Ag_1, Ag_2, ..., Ag_L \rangle$
- Real-valued shape-space
- Integer shape-space
- Hamming shape-space
- Symbolic shape-space

Define their Interaction

- Define the term **Affinity**
- Affinity is related to distance
  - Euclidian $D = \sqrt{\sum_{i=1}^{L} (Ab_i - Ag_i)^2}$
- Other distance measures such as Hamming, Manhattan etc. etc.
- Affinity Threshold

Basic Immune Models and Algorithms

- Bone Marrow Models
- Negative Selection Algorithms
- Clonal Selection Algorithm
- Somatic Hypermutation
- Immune Network Models

Bone Marrow Models

- Gene libraries are used to create antibodies from the bone marrow
- Antibody production through a random concatenation from gene libraries
- Simple or complex libraries

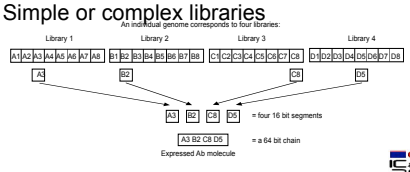

Illustration of NS Algorithm:

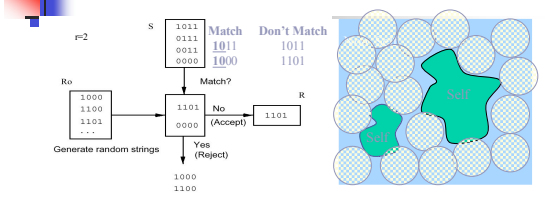

[Slide from Dasgupta and Timmis, 2003]

Negative Selection (NS) Algorithms

- Define *Self* as a normal pattern of activity or stable behavior of a system/process
  - A collection of logically split segments (equal-size) of pattern sequence.
  - Represent the collection as a multiset S of strings of length *l* over a finite alphabet.
- Generate a set R of *detectors*, each of which fails to match any string in S.
- Monitor new observations (of S) for changes by continually testing the detectors matching against representatives of S. If any detector ever matches, a change (or deviation) must have occurred in system behavior.

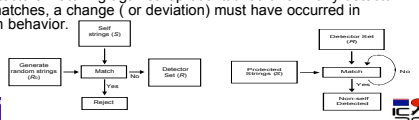

Clonal Selection

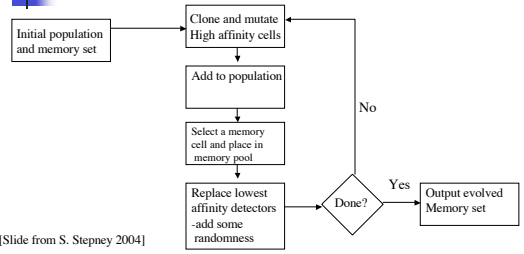

[Slide from S. Stepney 2004]

Clonal selection

- clone detectors that recognise patterns

```
                                  // S = set of patterns to be recognised
D := rnd ; m := ∅                 // initial random set of detectors; empty memory
    while not stopped do
      for each s in S do
        for each d in D do affinity(d,s) endfor
        C := clone(D,n1)          // n1 highest affinity d 's, rate ∝ affinity
        C := mutate C // rate ∝ 1\affinity
        D := D ∪ C
        m := m ∪ {highest affinity in D}
        D := D \ lowest_affinity(D,n2) ∪ rnd(n2)
      endfor
    endwhile
return D
```

[Taken from S. Stepney 2004]

Immune Networks

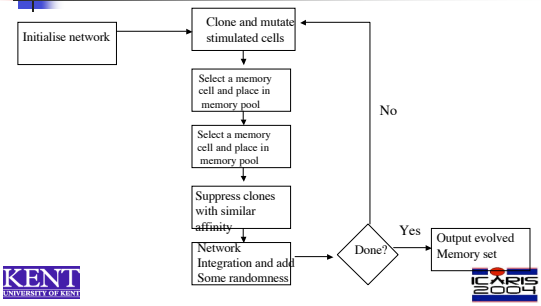

A Slight Aside

Inductive Bias

- Can affect choice of representation and affinity functions
- It is any (explicit or implicit) bias favouring one hypothesis over another;
  - All classification algorithms have it, otherwise it would only perform *rote* learning
- Bias is domain dependant
  - Therefore, can not always say **A** is better than **B**

Inductive Bias

- **Representation Bias**
  - Associated with the knowledge representation of the algorithm
- **Preference Bias**
  - Associated with the evaluation function employed
  - Use distance to measure affinity (this is important for AIS in particular)

Choice of Representation

- Assume the general case:
  Ab = $\langle Ab_1, Ab_2, ..., Ab_L \rangle$
  Ag = $\langle Ag_1, Ag_2, ..., Ag_L \rangle$
- Binary representation
  - Matching by bits
- Continuous (numeric)
  - Real or Integer, typically Euclidian
- Categorical (nominal)
  - E.g **female** or **male** of the attribute *Gender*. Typically no notion of *order*

Choice of Representation (2)

- Hybrid
  - Possible and desirable, as is quite common to have data of different types at the same time
  - More common in data mining literature, not really done (or appreciated) in AIS literature
- Typically, AIS seem to *adapt the data* to fit their particular algorithm rather than *adapting the algorithm* to making it more specific
  - From a problem orientated point of view, this could lead to throwing away useful data without realising it!

Choice of Affinity Functions

- Choice of function should take into account the data being mined as they will all have a **bias** …
- Binary Representation
  - Typically employ *Hamming* or *r-contiguous* rule
  - Argued that r-contiguous is more biologically plausible, therefore, use it … not so.
    - This assumes an ordering within the data that may not exist and will introduce a **positional** bias
    - In the data mining, quite common not to have unordered sets, representing the data when doing classification.
    - Therefore, a measure that takes into account position is not needed.

Choice of Affinity Functions (2)

- Continuous Representation
  - Vast majority of AIS use Euclidean .. Because … ?
  - Also is Manhattan. They will produce different results .. They have different inductive biases and are more effective for different data sets
  - $Dist(Ab, Ag) = (\sum (Ab_i - Ag_i)^2)^{1/2}$ (Euclidan)
  - $Dist(Ab, Ag) = \sum |Ab_i - Ag_i|$ (Mahantten)
- How do they differ?

Differences

- Which of the two antibodies is closer?

	Euc.	Man.
Ab_1 = **5.66**	8	
Ab_2 = 6.08	**7**	

- It depends …..

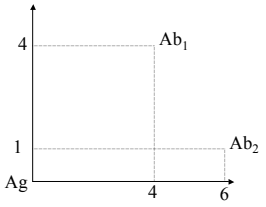

Why?

- Euclidean is more sensitive to noisy data
  - A single error in the coordinate of a vector could be seriously amplified by the metric
- Manhattan is more robust to noisy data and the differences tend not to be amplified
- So, results will be different and computational complexity is also different
- A rationale behind the choice is needed.

Categorical Representation

- AIS will be limited if all we can do is binary and numbers
- Data mining community make use of hybrids and have metrics for them
- Not just a transformation from one to another (ala AIS), but use the data
- Can measure the difference between categorical data. *Value Distance Metric.*
  - This assumes non-continuous learning though (as does the whole paper really)

Outline

- What are AIS?
- Useful immunology
- Thinking about AIS
  - AIS Algorithms
  - **Examples**
- What might we do?

Clonal Selection Based

- CLONALG [de Castro & von Zuben, 2001]
- DynamicCS [Kim and Bentley, 2002]
- AISEC [Secker et al, 2003]
- AIRS [Watkins et al, 2003]

AIRS Algorithm [Watkins, 2003]

- Data normalization and initialization
- Memory cell identification and ARB generation
- Competition for resources in the development of a candidate memory cell
- Potential introduction of the candidate memory cell into the set of established memory cells

Memory Cell Identification

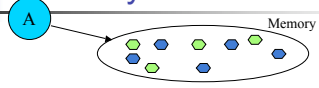

MCmatch Found

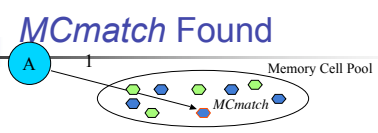

ARB Generation

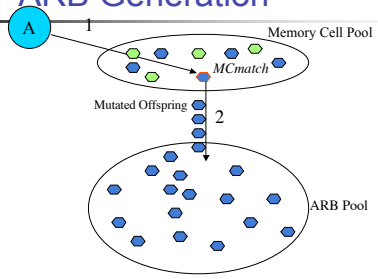

Exposure of ARBs to Antigen

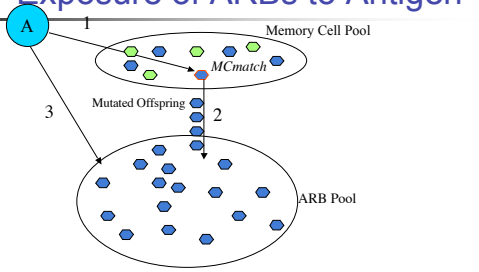

Development of a Candidate Memory Cell

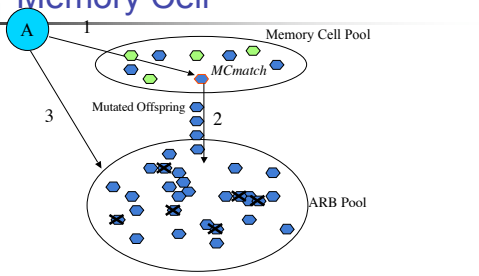

Comparison of *MCcandidate* and *MCmatch*

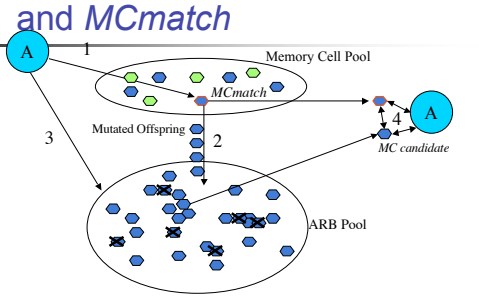

Memory Cell Introduction

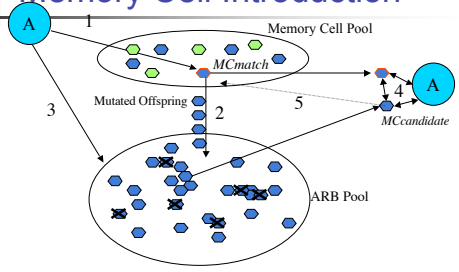

Artificial Immune Networks

- Many attempts at network-based immune system algorithms
- Based on a variety of ideas which include:
  - Jerne's network theory
  - Recognition regions
  - The properties of network interactions
  - Effectively clustering
  - One shot and continuous learning

Meta-Stable AIS [Neal, 2003]

- Clone when a data-item is not recognised (further than NAT from an existing ARB)
- Cloning just means placing an ARB directly on the unrecognised item
- The Algorithm:
- Initialise from data-set
- Do forever:
  - Stimulate
  - allocate,
  - clone,
  - cull

Outline

- What are AIS?
- Useful immunology
- Thinking about AIS
  - AIS Algorithms
  - Examples
- **What might we do?**

Interactions with Other Systems

- The immune system does not act in isolation
- All AIS work (and most bio-inspired approaches) to date have ignored this fact
- It is time that we did not, we need:
  - Complex, adaptive, autonomous and *self-organising* behaviour embodied within computational systems
  - Integrate (in a sound manner) from the biological domain, thus making use of the bioinspired techniques??
    - Look at homeostasis within an organism
  - A new radical look at how we do things:

Interacting Systems: Homeostasis

- Nervous System
  - Process and acts on stimuli from external sources
    - Vision, taste, touch etc
    - Develop over the lifetime of an organism
- Endocrine System
  - Hormones have a regulatory affect
  - Can affect behaviour
  - Production of hormone is linked to changes in state of the organism
  - Hormones typically decay over time
- Immune System
  - All cells within these system have receptors for each other, allowing them to interact in a wide variety of ways

The Way Forward ..

- Better interactions with biologists
  - Create better models of immunology
  - Develop and utilise models of interacting systems
- More accurate metaphors and development of a conceptual framework
- More mathematical analysis of algorithms
- More work on where AIS are good (and not)
- Work in demanding application areas
- Physiological view of the system

ARTIST

- ARTIST: A Network for Artificial Immune Systems (EPSRC funded network) http://www.artificial-immune-systems.org
- Work towards:
  - A theoretical foundation for AIS as a new CI
  - Extraction of accurate metaphors
  - Immune System Modelling
  - Application of AIS
- Train PhD students
- Fund workshops/meetings
- Coordinate and Disseminate UK based AIS research (links to Europe)

AIS Resources: Books

- Artificial Immune Systems and Their Applications by Dipankar Dasgupta (Editor) Springer Verlag, January 1999.
- Artificial Immune Systems: A New Computational Intelligence Approach by Leandro N. de Castro, Jonathan Timmis, Springer Verlag, November 2002.
- Immunocomputing: Principles and Applications by Alexander O. Tarakanov, Victor A. Skormin, Svetlana P. Sokolova, Springer Verlag, April 2003.
- http://www.artificial-immune-systems.org