

Two Ways to Grow Tissue for Artificial Immune Systems

Peter J. Bentley

University College London
Department of Computer Science
London WC1E 6BT, UK

P.Bentley@cs.ucl.ac.uk

Julie Greensmith

University of Nottingham
School of Computer Science (ASAP)
Nottingham NG8 1BB, UK

jqg@cs.nott.ac.uk

Supiya Ujjin

University College London
Department of Computer Science
London WC1E 6BT, UK

s.ujjin@cs.ucl.ac.uk

ABSTRACT

An immune system without tissue is like evolution without genes. Something very important is missing. Here we present the novel concept of tissue for artificial immune systems. Much like the genetic representation of genetic algorithms, tissue provides an interface between problem and immune algorithm. Two tissue-growing algorithms are presented with experimental results illustrating their abilities to dynamically cluster data and provide useful signals. The use of tissue to provide an innate immune response driving the adaptive response of conventional immune algorithms is then discussed.

Categories and Subject Descriptors

E.1 Data structures E.2 Data storage representations.

General Terms

Algorithms, Experimentation, Theory

Keywords

Artificial Immune Systems, Tissue growth, Innate immunity, Swarm intelligence, Dynamic clustering

1. INTRODUCTION

So the immune system exists to protect tissue from harm. In one sense, an immune system without tissue is meaningless. Yet in the field of artificial immune systems there is no real concept of tissue. Data is typically mapped directly to antigens. In many cases there is not even the concept of immune cells, let alone tissue cells. Both conceptually and technically, this can cause difficulties – for if every new artificial immune system is directly “wired” to a specific problem, then it becomes difficult to compare, analyse and even to apply the AIS to new problems.

Here we propose an alternative treatment for artificial immune systems. Instead of joining the AIS to its application directly,

it is proposed that an intermediary representation is employed, much like the genetic representation of the genetic algorithm. This intermediary will be a dynamic encoding of the current problem providing the equivalent of an innate immune response to support the adaptive response of an AIS. The encoding will be modified according to the problem like the genetic encoding of a GA [bentley99]. But regardless of the underlying data, it will present a consistent interface to an artificial immune system. That interface will be *tissue*, fig 1.



Figure 1 Tissue should act as the interface between problem and AIS.

1. BACKGROUND

The concept of artificial tissue is used extensively in cell modelling and simulation, with additional applications in electronics and biotechnology. One well-known example was the POETic project, which used the concept of cellular tissue and immune cell modelling within hardware devices [11]. In this architecture each cell is treated as an individual processing device, with the tissue performing the role of providing an interface between a biologically inspired processing mechanism and data provided by the environment. Similarly, in [3], fault tolerant electronic circuits were constructed and used a combination of embryonically grown cells coupled with immune-inspired negative selection. This model provides an immune inspired component and entity to protect, though the protected cells did not provide feedback signals to the AIS. The protected cells in this system were embryonically grown, sending out signals to support each other. The system partitioned the AIS and the cells into separate layers, providing communication between the two components. This architecture was implemented and applied to various hardware devices. Examples of developmental models that include aspects of tissue growth are becoming more popular; interested readers should consult [7].

In biology, tissue has long been known to be a crucial component of the immune system, and this role was highlighted further by Matzinger. The Danger Model, proposed by Polly Matzinger in 1994 [9], attempted to alter the perspective from which the immune system was viewed. This involved abandoning the belief that the immune system

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Conference '04, Month 1–2, 2004, City, State, Country.
Copyright 2004 ACM 1-58113-000-0/00/0004...\$5.00.

is conditioned at an early age to distinguish self from non-self proteins. Instead, this model proposes that the immune system contains cells sensitive to cellular damage. In her words: “The Danger model ... suggests that neither the innate nor adaptive immune systems are in ultimate control. This function belongs to the ancient innate responses of the normal bodily tissues themselves” [8]

The theory suggests that signals are innately released from cells under stress, due to damage, often derived from pathogens, physical disruption, radiation, extreme pHs or temperature. These signals may cause tolerance to proteins through regulatory cell activation or lead to the activation of effector cells [10]. This discrimination is based on the information gathered from proteins collected within the body, in combination with various signals derived from host tissue cells. The combination of antigens-plus-signals can give information regarding damage to a specific area of tissue. In order to understand what the signals are and under what conditions they arise, two important types of cell death have to be examined.

1. Apoptosis. Tissue cells can die in a number of different ways, forming part of the life cycle of a cell. It is essential for cells to die under controlled conditions to provide regulation of tissue growth and to remove defective and virally infected cells. This type of pre-programmed cell death is known as apoptosis. On receipt of an apoptotic signal the cell releases a number of degrading enzymes which have dramatic effects on the internal structure of the cell. The cell's DNA is fragmented into orderly portions, nuclear condensation is initiated and organelles are broken down. During this period of degradation, the integrity of the outer cell membrane remains intact, while expressing greater quantities of signalling molecules on the membrane surface. These molecules are detected by innate immune cells, such as macrophages, which are triggered to ingest the cell, ultimately resulting in removal of the apoptotic cell from the tissue[5].

2. Necrosis. In contrast, unexpected, chaotic cell death does not involve an intricate removal system. Unlike apoptotic cells, the necrotic cell swells up, the internal material is chaotically fragmented and the membrane integrity is lost. Ultimately, the cell explodes, releasing its contents into the fluid surrounding the cell. Cellular products released as a result of necrotic cell death are known as danger signals - endogenous activators of the innate immune system. This includes molecules derived due to cell degradation, inclusive of uric acid, adenosine-tri-phosphate, and heat shock proteins[12], in addition to an array of pro-inflammatory cytokines.

Without tissue there would be no endogenous danger signals, no innate immune activation and nothing to protect. Additionally it is thought that the absence of tissue derived danger signals is as equally important as their presence, through the generation of proteins that do not belong to the host, yet cause no damage, e.g. bacterial gut flora. The detection of an apoptotic signal is translated into the activation of the adaptive immune system's regulatory cells [10].

It is clear that tissue has been highlighted as an integral part of immune function. Danger signals released from cells dying under stressful conditions activate cells belonging to the innate immune systems. These cells ultimately control the effector cells, and giving direction to the immune response.

Yet, the concept of tissue has not been widely used within AIS. The question remains: is it possible to construct artificial tissue to provide an interface between an application and an artificial immune system?

1. DEFINING TISSUE

Focussing for now on the task of anomaly detection, it is proposed that tissue designed for artificial immune algorithms should comprise a series of linked cells, each cell “grown” in response to specific data, in a data stream being input to the system. Cells should grow and be supported by homogeneous data. Where data does not exist to support a cell, the cell dies. Where too much/too diverse data exists for a cell, the cell divides. Cells should exist in a dynamic network structure, with similar cells linked or placed near to each other. The use of a cellular representation is also intended to enable distributed processing and the support of multiple datastreams simultaneously.

In the ‘tissue paradigm’ all communication between a problem and AIS is mediated via the tissue. Tissue thus provides some functionality of the innate immune system, with the AIS performing the common role of adaptive immune system.

1.1 Uplinks

Given a data stream of temporally homogeneous data items, the tissue will quickly grow to form a specific shape, structure and size, which will be maintained indefinitely. The artificial immune system should consult all cells in the tissue, examining them and any corresponding danger signals. If the data changes, the tissue will change in response. Those aspects of the data that remain the same will continue to support the corresponding parts of tissue. Those aspects that differ will result in a restructuring or even cell death. An artificial immune system should thus be able to ignore static tissue and quickly cause an immune response on and near to the cells where the changes (and corresponding signals) are occurring. In this way the tissue provides more than an interface to the underlying data – it provides a spatial and temporal structure, enabling the AIS to specialise and focus to different extents, spatially and temporally.

It is recognised (and experiments confirm later) that the tissue will not perform perfectly as a clusterer and anomaly detector – if it did there would be no reason to have the AIS. Instead, the tissue provides useful data preprocessing, gathering similar data items together, and presenting gross, short-term anomalies to the AIS. (Specific, problem dependent knowledge can also be incorporated and exploited in the cells in order to present other innate signals to the AIS.) It is expected that critical anomalies will still occur within “normal” tissue. Thus the role of the AIS in the ‘tissue paradigm’ is now to consult cells within the tissue and identify fragments of data (antigens) presented by the cells that together may indicate a critical anomaly. Note that there is no real concept of a self/non-self division; here the concept is more one of stability/entropy. A stable tissue is considered ‘healthy’; unstable or entropic tissue is ‘unhealthy’ and will attract attention from the AIS.

1.1 Downlinks

The natural immune system is designed to both detect harmful anomalies and remove the causal agents. However, an artificial immune system using the ‘tissue paradigm’ cannot simply remove ‘infected’ cells from the tissue – this would only

prevent the tissue from presenting information about the anomaly to the AIS, it would not prevent the underlying anomaly in the application from reoccurring. Instead, the AIS should use the tissue as an interface to the application. If a critical anomaly is discovered, cells should be informed which antigens are responsible. The cells then pass this information down to the underlying application, where the information should be used to remove the cause of the anomaly. For example, in a computer network intrusion detection application, if the AIS identifies a specific antigen in one cell, the cell will then communicate this information to the network management software. This software might terminate a corresponding process and thus remove the 'infection' from the input data stream, or just inform the system administrator. If there is a one-to-one correspondence between cell and anomaly, then by identifying the anomalous antigen within the cell, and causing the subsequent prevention of the anomalous data in the input stream, the corresponding cell will no longer be supported by the data stream and will die. In other words, it is possible for the AIS to cause tissue cell death by interacting with the application via the tissue.

1. Tissue Algorithms

There are many ways in which tissue can be developed. Here we present two different approaches: a network tissue growing algorithm, and a swarm tissue growing algorithm. Both effectively act as dynamic clusterers, using danger signals as approximate alerts of anomalies in the input stream.

1.1 Network Tissue Algorithm

The network-based algorithm explicitly maintains cells in a dynamic network, with parent cells pointing to daughter cells, and link restructuring on cell death to maintain network coherence (e.g., the death of a parent cell results in the oldest daughter cell taking the parent's position in the network). In this algorithm, each cell may hold up to `maxantigenspercell` antigens before dividing into two. Figure 2 outlines the network tissue algorithm.

1.1.1 Biological Analogies

Figure 3 summarises the model with respect to natural biology. In this model a single cell may represent a particular cell type of a living organism. While there is data to support a cell (i.e., while the impact of the environment and genes results in a particular type of tissue structure), the cell will survive indefinitely (the tissue will have a certain cell type and structure indefinitely). If the input stream changes permanently (or for a sufficiently long duration), even if the change is dramatic, the new data will cause corresponding new tissue to develop and be supported (i.e. a long-term change in the environment causes long-term useful changes in tissue structure). But if an anomalous datum creates a cell, and there is insufficient subsequent similar data to support that cell, then the cell will die. (In an organism, cells can be created in response to the environment, affected by the existing tissue; but the environment might include some form of pathogen, which infects and destroys cells of that type). It is not necessary for apoptosis to be modelled explicitly – it is assumed that a single cell represents many cells of that type growing and dying to be replaced by new cells naturally. So should a cell die in the model, this can only be necrosis – and thus it causes the release of a danger signal, to be passed to the neighbouring cells in the tissue.

```

create zygote with first data point (antigen)
get next antigen from data steam
find nearest cell (cell with mean antigen
closest to current antigen)
if current antigen is sufficiently similar*
to nearest cell mean
then add antigen to nearest cell
    if nearest cell has number of antigens ==
        maxantigenspercell
    then split current cell into two linked cells
        s.t. antigens are shared equally**
    update cell means, danger signals and
        linked neighbours
else
    create new cell at current antigen;
    nearest cell is linked parent

for every cell
    for every antigen in the cell
        age antigen
        if antigen age > maxantigenage
            remove antigen
    if antigens in cell == 0
        cell dies (can no longer respond to input)
        create new dangersignal, origin = final antigen,
            range = cell stddev,
            strength = max (or inversely proportional
                to cell age)
        pass all danger signals of dying cell to
            linked neighbouring cells

for every dangersignal
    reduce strength
    if strength == 0
        delete dangersignal

```

*similarity measures depend on the matching function used and underlying application; in the experiments reported here, data values are normalised and the Euclidian distance between cell mean and antigen compared against a similarity threshold of 0.2 (default).
**the cell split function should use the same distance function to divide antigens into two groups; in the experiments reported here, all antigens greater than the mean are placed in one cell, all antigens less than the mean are placed in the other.

In addition to the similarity measures, there are 2 important constants:

`maxantigenage` - determines number of antigens held by tissue cells at any point in time.

`maxantigenspercell` - affects how many cells there will be in the tissue

Figure 2 The network tissue algorithm.

In an attempt to match biological characteristics of danger signals, in the model, danger signals emitted as a result of necrosis are general indicators of an anomaly, but are spatially and temporally specific. The danger signals from a dead cell are held by its neighbouring cells (which, through automatic network restructuring or swarming after necrosis, "fill the gap" left by the dead cell). It is possible for cells to hold many danger signals at once. Danger signals decay over time; they are removed once their strength falls to zero.




Figure 3. Left: Organic tissue grows according to its DNA, and interactions with its environment. To create a new cell, an existing cell divides into two. Necrosis results in “danger signals”. Immune cells consult antigens presented by tissue cells and respond by destroying infection. Right: AIS tissue grows according to its rules of growth and interactions with a problem. To create a new cell, an existing cell may split into two. Unnatural cell death results in danger signals. The AIS consults antigens presented by tissue cells and responds by signalling the underlying application via the cells, potentially resulting in the destruction of cells that are presenting anomalous antigens.

1.1.1 Experiments

A series of experiments were performed using the standard “breast cancer” UCI machine learning data set, comprising 240 ‘malignant’ items in class 1 and 460 ‘benign’ items in class 2. For each system setting, the same experiment was repeated 30 times. Implemented in ‘C’ and running on a Mac Powerbook G4, each run of 10,000 iterations (with one randomly picked data item presented to the tissue each iteration) lasted less than 5 seconds. Class 2 (benign) is treated as the “normal” class of data, with items from class 1 (malignant) being introduced into the datastream every 25 iterations (this value is investigated in the first 3 experiments).

Table 1 lists the different parameter settings used in the experiments.

Table 1 System setups for the nine experiments.

Experiment	Max antigen age	Max antigens/cell	Similarity threshold	Class 1 item freq.
1	40	10	0.2	25
2	40	10	0.2	10
3	40	10	0.2	5
4	40	10	0.1	25
5	40	10	0.3	25
6	40	5	0.2	25
7	40	15	0.2	25
8	20	10	0.2	25
9	60	10	0.2	25

1.1.1 Analysis

Table 2 shows the results for the nine experiments. As is to be expected from a deterministic algorithm (where the only stochastic element is the data item order in the data stream), the results for all experiments were very consistent across runs, as shown by the low standard deviation values. Experiments 1 to 3 indicate how the frequency of anomalous data items influences the accuracy of danger signals, i.e., more frequent items from class 1 reduces the tendency of the tissue to treat class 1 items as anomalous (true positive), while the percentage of items in class 2 treated as anomalous remains unchanged (false positive). Experiments 4 and 5 (also compare

with experiment 1) show how changing the similarity threshold affects danger signal accuracy. A smaller threshold produces near perfect detection of anomalies from class 1, but also increases the tendency for items in class 2 to be detected as anomalies. The opposite effect occurs when the threshold is increased. Experiments 6 and 7 (also compare with experiment 1) show how the number of antigens per cell affects danger signals. No real change occurs to the accuracy of detection of anomalous items from class 2, but a smaller number of antigens produces far less tolerance for different items in class 1 (the cells are more specialised, increasing the chances for even slightly different antigens to be treated as anomalous). Increasing the number of antigens has the opposite effect – causing a significant reduction in the number of items in class 2 that are treated as anomalous. Finally, experiments 8 and 9 (also compare with experiment 1) show the effect of varying the maximum antigen age. In the experiments, this has only a minor effect on danger signal accuracy, although the results suggest that the age should be set in relation to the expected frequency of anomalies in the datastream, i.e., a long age for frequent anomalies increases the tolerance of the tissue for the anomalies, while a short age causes infrequent but normal data items to be treated as anomalies.

Table 2 Results for experiments 1-9, showing mean number of danger signals per run, standard deviation, and percentage of danger signals that correspond to data items in class 1 and class 2.

exp	Class 1			Class 2		
	mean	stddev	%	mean	stddev	%
1	382.8	2.6	98.1%	1008.5	25.1	10.5%
2	933.4	5.8	95.6%	946.7	20.2	10.5%
3	1801.7	20.2	92.3%	842.6	22.0	10.5%
4	390.1	1.4	99.9%	1378	27.2	14.4%
5	373.5	3.5	95.7%	849.7	20.0	8.9%
6	383.7	2.4	98.3%	2165.8	27.3	22.6%
7	384.0	2.2	98.4%	796.7	24.8	8.3%
8	389.6	1.6	99.6%	1046.1	19.5	10.9%
9	378.7	3.0	97.2%	946.0	24.0	9.9%

1.1 Swarm Tissue Algorithm

The swarm-based algorithm is a second, alternative approach to tissue development. It is designed to follow much the same “tissue growing” principles as outlined previously, but now clusters in the tissue are formed by cell movement in a two-dimensional space (of size 1000 by 1000 units) which is unrelated to the data values, with similar cells moving together and dissimilar cells moving apart. In this algorithm, each cell holds just one data item; cells are created by new data and die constantly – thus apoptosis is modelled in this algorithm. If cells have not grouped themselves into a cluster by the time they die, they produce a danger signal, i.e. necrosis is modelled by the death of “abnormal” cells that do not participate in normal tissue development. Figure 4 outlines the algorithm, which uses the following swarming rules to drive the motion of cells:

$v_i = wv_i + C_1 r_1 (x_{pbest,i} - x_i)$ Rule 1
 $v_i = v_i - f_2 C_2 r_2 (x_{avoid,i} - x_i)$ Rule 2
 if $(|v_i| > v_{max})$ $v_i = (v_{max}/|v_i|)v_i$ Rule 3
 $x_i = x_i + v_i$ Rule 4

where:
 x_i is the current position of data item i
 $x_{pbest,i}$ is the current best position of data item i
 $x_{avoid,i}$ represents the current avoidance position of data item i .
 v_i is the velocity of data item i
 w is a random inertia weight between 0.5 and 1 [4]
 C_1 and C_2 are spring constants set to 1.494 [4]
 r_1 and r_2 are random numbers between 0 and 1 [2]
 f_2 is the repulsive factor (default value 2). Defines the effect of the repulsive force on velocity; the higher the value the more that dissimilar items repel each other.
 v_{max} is the maximum velocity (default value of 300)

Note: $x_{pbest,i}$ is either the central position of all items in the same cluster as i or the central position of all items in the swarming space (if i does not belong to a cluster)

$x_{avoid,i}$ represents the central position of all data items in i 's neighbourhood whose similarity value falls below the similarity threshold.

1.1.1 Biological Analogies

Figure 5 summarises the model with respect to natural biology. In this model cells are modelled more directly. New data generates new cells which all live for a fixed lifespan before dying. While they live they move with respect to each other, with similar cells clustering and dissimilar cells moving apart (i.e., the impact of the environment and genes results in a particular type of tissue structure, with similar cells adhering to each other and forming organs). As with the previous algorithm, if the input stream changes permanently (or for a sufficiently long duration), even if the change is dramatic, the new data will cause corresponding new tissue to develop and be supported (i.e. a long-term change in the environment causes long-term useful changes in tissue structure). But if an anomalous datum creates a cell, and there is insufficient subsequent similar data to produce similar cells, then the cell will be unable to form a cluster before it dies. (In an organism, cells can be created in response to the environment, affected by the existing tissue; but the environment might include some form of pathogen, which infects and alters cells of that type). In this algorithm, apoptosis is modelled explicitly – cells grow and die to be replaced by new cells naturally. So in this model, necrosis is modelled by a dying cell that has not formed part of a group of other cells – and thus it causes the release of a danger signal, to be passed to the neighbouring cells in the tissue.

```

get next antigen from data stream
create cell using antigen and place in swarm-tissue

for every cell, (current cell = C1)
  for every cell in neighbourhood* of C1
    (neighbour cell = C2)
    if C1 is sufficiently similar**
      to C2/cell-cluster***
      then C1 joins/makes cluster with C2/cell-cluster
      if C1 and C2 were in clusters with differences <
        current similarity
        then they form a new cluster together
    else
      C2/cell-cluster*** is added to
        cell avoidance list
  if C1 is in a cluster, C1 best position is mean pos.
    of cells in C1 cluster
  else C1 best position is mean tissue position

update velocity of C1 using best pos,
  mean avoidance values (Rules 1 to 3)
update C1 position based on velocity (Rule 4)

increase age of C1
if C1 age is greater than celllifespan
  remove C1 from swarm-tissue
  if C1 was not in a cluster
    create dangersignal, origin = C1,
      range = cell stdev,
      strength = max
  (pass all danger signals of dying cell
  to neighbouring cells)

for every dangersignal
  reduce strength
  if strength == 0
    delete dangersignal
  
```

* defined by radius around C_1 where radius = 300
 ** similarity measures depend on the matching function used and underlying application; in the experiments reported here, data values are normalised and the Euclidian distance between the two cell values are compared against a similarity threshold of 0.2 (default). In addition, the similarity measure between C_1 and C_2 where C_2 is in a cluster is scaled by the inverse of the number of cells in the cluster, making larger clusters more attractive.
 *** if C_2 is in a cluster, the mean value of cells in the cluster is used, otherwise the value of C_2 is used.

Figure 4 The swarm tissue algorithm.




Figure 5. Left: Organic tissue grows according to its DNA, and interactions with its environment. Right: AIS tissue grows by moving cells relative to each other according to their rules of growth and interactions with a problem. A cell that has not formed part of the tissue before it dies is necrotic and produces a danger signal. The AIS consults antigens presented by tissue cells and responds by signalling the underlying application via the cells.

1.1.2 Experiments

Again, a series of experiments were performed using the standard “breast cancer” UCI machine learning data set. As before, for each system setting, the same experiment was repeated 30 times. This time implemented in Java J2SE and running on a 2.4 Ghz Pentium 4 PC, each run of 10,000 iterations lasted between 30 and 145 seconds. The same parameter settings as listed in table 1 were used for the experiments, although experiments 6 and 7 could not be performed as each cell only holds one antigen in this model.

Table 3 Results for experiments 1-9, showing mean number of danger signals per run, standard deviation, and percentage of danger signals that correspond to data items in class 1 and class 2.

experiment	Class 1			Class 2		
	mean	stdev	%	mean	stdev	%
1	398.0	0.18	99.99%	1973.6	39.2	20.6%
2	996.0	0.18	100.0%	1903.1	38.4	21.2%
3	1991.9	0.37	100.0%	1783.4	46.0	22.4%
4	398	0	100.0%	4045.0	61.7	42.3%
5	397.3	0.92	99.82%	1434.4	44.8	15.0%
8	399	0.0	100.0%	6425.9	75.4	67.1%
9	396.9	0.25	99.98%	1422.4	44.5	14.9%

1.1.3 Analysis

Table 3 shows the results for the nine experiments. Accuracy of danger signals for class 1 is consistently high for all experiments, but the changes in parameter settings do appear to affect the percentage of items in class 2 treated as anomalous. Presenting items from class 1 more frequently (see results for experiments 1,2,3) produces a subtle increase in class 2 anomalies; this may be caused by a disturbance effect of more cells in class 1 disrupting the path of class 2 cells as they try to cluster. A lower similarity threshold allows fewer cells to cluster and so produces a considerably worsened percentage for class 2 anomalies, while a higher threshold has the reverse effect (see results for experiments 1,4,5). The same effect occurs when cell age is modified (see results for experiments 1,8,9) – a lower age produces fewer chances for clusters to form in time; a higher age increases the chance and thus reduces the class 2 signals. Further experiments showed that increasing the lifespan to 100 and using a threshold of 0.3 produced accuracy in class 1 of 99.8% and in class 2 of 8.9%, although execution times increased to 145 seconds for 10,000 items.

1. Discussion

Like immunobiology, the field of artificial immune systems has been obsessed with the workings of the adaptive immune system and its capabilities of specificity, diversity and memory, with little work spent on the innate immune system. This work attempts to lay the foundations of a more complete view of the immune system for AIS. We propose that the concept of tissue is important for several reasons:

- Tissue provides a generic data representation which interfaces between problem and conventional AIS, simplifying future AIS development.
- Tissue stores the current state of the application, providing a clearer concept of “organism” and enabling

the AIS to learn to detect changes in the organism, correct harmful changes and prevent future damage by similar agencies.

- For applications such as anomaly detection, tissue provides a dynamic window of the input data stream; the data is dynamically organized and spatially structured, encapsulating the important concepts of temporal and spatial variability. An AIS exploiting tissue would be able to specialize and focus on different areas of the problem, at different times, enabling a more precise response.
- Tissue encapsulates ideas of homeostasis – if the problem becomes heterogenous or chaotic, the tissue will reorganize its structure in response. An AIS collaborating with the tissue would be able to correct harmful changes and work to maintain homeostasis.
- Tissue is essential for the innate immune system, and tissue algorithms can be used to provide desirable “automatic” processing and signals from data.

It is proposed that an AIS will employ tissue by traversing its spatial representation and allocating resources according to the spatial and temporal requirements. A network-based AIS might form distinct and functionally diverse subnetworks to focus on tissue cells of different types. A population-based AIS would be able to allocate subpopulations of agents (e.g., antibodies, B-cells or T-cells) for specific regions of tissue. In all cases, all aspects of the problem should be presented to the AIS through tissue, and all AIS responses should be presented to the underlying application by the tissue.

In this work we have focused on the task of anomaly detection, and both tissue-growing algorithms were developed with this in mind. However, we propose that the concept of tissue should be employed for all AIS applications. This may inevitably involve different forms of tissue growth. For example, in a robot control application [6], sensor input might be used as an input data stream and the algorithms presented above could be used. Alternatively, the state of sensors and actuators might be represented by a fixed and predefined tissue structure (e.g. a cell for each sensor, and a cell for each motor). Such a structure would change if sensors or motors were lost through damage – obviously requiring a significant response from the controlling AIS. But normal control would occur through the robot presenting its changing state via antigens and signals from the cells, interpreted by the AIS, with responses made to the cells being mapped back to robot motor control.

Like the genetic representations of genetic algorithms, the exact tissue representation necessary is likely to be application-specific, but the AIS used to consult with the tissue and respond to it should be generic. It is conceivable that evolutionary computation could be employed to evolve useful innate tissue responses for a given application and AIS. Indeed if each tissue cell contained an evolving GP function [1], cells would be able to present one or more evolved interpretations (i.e., signals) derived from the raw data, in addition to the raw data.

1. Conclusions

In this work we have presented the novel concept of tissue for artificial immune systems. Much like the genetic representation of genetic algorithms, tissue provides an interface between problem and immune algorithm. From the

perspective of immunobiology, tissue provides an innate immune response, with the AIS providing an adaptive response. Two tissue-growing algorithms were presented with experimental results illustrating their abilities to dynamically cluster data and provide useful signals. Both algorithms are able to detect anomalous data items with accuracies up to 100% depending on the parameter settings. Future work will investigate the integration of these algorithms with artificial immune systems for intrusion detection.

Acknowledgements

Thanks to all the members of the “Danger Project” <http://www.dangertheory.com/> for their helpful comments, feedback and discussions which helped shape the ideas in this paper. This project is supported by the EPSRC (GR/S47809/01), Hewlet Packard Labs, Bristol, and the Firestorm intrusion detection system team. This work is ODDISIAC-friendly.

2. References

- [1] Bentley, P. J. 1999. An Introduction to Evolutionary Design by Computers. Chapter 1 in Bentley, P. J. (Ed.). *Evolutionary Design by Computers*. Morgan Kaufman Publishers Inc., San Francisco, CA, 1-73.
- [2] Blackwell, T.M. and Bentley, P.J. 2002. Dynamic Search with Charged Swarms. In Proceedings of the Genetic and Evolutionary Computation Conference 2002. New York.
- [3] Canham, R. and Tyrell, A.M. 2003. A Hardware Artificial Immune System and Embryonic Array for Fault Tolerant Systems. *Genetic Programming and Evolvable Machines*. vol 4, issue 4, pp.359-382.
- [4] Eberhart, R.C. and Shi, Y. 2001. Particle Swarm Optimization: Developments, Applications and Resources. In Proceedings of the 2001 Congress on Evolutionary Computation. vol.1, pp.81-86.
- [5] Edinger, A.L. and Craig B Thompson. 2004. Death by Design: apoptosis, necrosis and autophagy. *Current Opinion in Cell Biology*. vol.16, pp663-669.
- [6] Ko, A., Lau, H. Y. K., and Lau, T. L.. 2004. An Immuno Control Framework for Decentralised Mechatronic Control. In Proc. of the Third International Conference on Artificial Immune Systems (ICARIS 2004). Catania, Sicily. pp. 91-105.
- [7] Kumar, S. and Bentley, P. J. (contributing editors). 2003. *On Growth, Form and Computers*. Academic Press Ltd, London.
- [8] Matzinger, P. 1998. An Innate Sense of Danger. *Seminars in Immunology*.vol.10, pp.399-415.
- [9] Matzinger, P. 1994. Tolerance, Danger and the Extended Family. *Annual Reviews In Immunology* 12:991-1045.
- [10] Matzinger, P. 2002. The Danger Model: A Renewed Sense of Self. *Science*. 296, pp301-305.
- [11] Thoma, Y., Tempesti, G., Sanchez, E. and Arostegui, J.M.M. 2004. POEtic: an electronic tissue for bioinspired cellular applications. *Biosystems*. 76, pp.191-200.
- [12] Wallin, R.P.A., Lundquist, A., More, S.H., Von Bonin, A., Keissling, R. and Ljunggren, H.G. 2002. Heat Shock Proteins as Activators of the Innate Immune System. *Trends in Immunology*. 23, no. 3, pp.130-135.