

Don't Touch Me, I'm Fine: Robot Autonomy Using an Artificial Innate Immune System

Mark Neal¹, Jan Feyereisl², Rosario Rascunà³, and Xiaolei Wang⁴

¹ Computer Science, University of Wales, Aberystwyth, UK

² School of Computer Science, University of Nottingham, UK

³ CCNR, University of Sussex, UK

⁴ Electrical and Communications Engineering, Helsinki University
of Technology, Espoo, Finland

ARTIST Network: Student Spring School, Aberystwyth, UK
April 2006

Abstract. A model for integration of low-level responses to damage, potential damage and component failure in robots is presented. This model draws on the notion of inflammation and introduces an extensible, sub-symbolic mechanism for modulating high-level behaviour using the notion of artificial inflammation. Preliminary results obtained via simulation are presented and demonstrate the potential benefits of such a scheme. Additionally the system maps the robot's physiological state-space, which is defined in terms of the levels and sources of inflammatory response. This is achieved using Kohonen's Self-Organizing Map algorithm to arrange the states experienced during the lifetime of the robot. The future use of this map for diagnosis and localization of faults and for the generation of specific high-level remediation behaviour is also discussed.

Keywords: Artificial Immune Systems, Human Immune Systems, Innate Immunity, TLR, PAMPs, Inflammation, SOM, Robot.

1 Introduction

With a few rare exceptions such as [10,4], the innate immune system has been neglected in artificial immune systems [3], especially in the field of robotics which appears to have much to be gained from such an approach. The functions making up this part of the immune system, offer a number of useful analogies that can be exploited in a robotic system. In the quest for autonomy an artificial innate immune system can be applied in order to create systems which are aware of their own state. This could allow them to maintain a "healthy", homeostatic balance and achieve self sufficiency. In order to achieve this a robot must contain a number of proprioceptive¹ sensors, monitoring various state measures across

¹ proprioceptive: sensing internal body state.

the physical domain of the robotic system [2]. An analogy emerges here with Toll-Like Receptors (TLRs) as sensors of potentially problematic signals within the body. Such signals are known as Pathogen Associated Molecular Patterns (PAMPs) [9]. In robotic systems simple sensors capable of detecting problematic circumstances (eg. “motor3 overheating”) can often be used locally to help remediate the problem without recourse to high-level software and control systems. This is directly analogous to the types of action taken by innate immune system components (such as macrophages) endowed with TLRs. Difficulties arise in engineering complex robotic systems (or other electro-mechanical systems) which attempt to integrate the input from large numbers of such local sensing and remediation devices into high-level control systems. It rapidly becomes impossible to predict all possible combinations of problem and remediation action, and computationally expensive to process all this information in the high-level controller. A number of approaches to robot control have addressed this problem with varying degrees of success, the best known being [1]. The notion of artificial inflammation allows the integration of information about low-level response patterns into a small number of global signals which represent the “state of health” of the system throughout time. These simple inflammatory signals can then be used via schemes such as neuro-endocrine control [7,8] to modulate high-level control systems appropriately.

The representation of the states of the robotic system using Kohonen’s Self-Organizing Maps (SOM) [6] allows the sources of the inflammation to be localized within individual nodes in order to both diagnose problems at intermediate levels (eg. “motor compartment 2 overheating”) and to allow higher-level remediation to be appropriately targeted on the components that directly affect the inflamed parts of the robot.

A description of the physiology of a robot follows, including the analogy drawn from the innate immune system. Next, a step-by-step description of the model is used to show exactly how it works both in this specific case and how the scheme works in general. A proof of concept experiment is described, supported by the results obtained and a commentary on what the results show. This is followed by some conclusions, including advantages and disadvantages of the proposed model.

2 Robot Physiology

In general a robot is a complex system made up of numerous interacting components that can fail or malfunction alone as well as in combination. Typical components also include automatic damage protection functions and circuits such as locally switched cooling fans and automatic overheat cut-outs. Analogies between such components and the innate immune system are presented here. Firstly the function of TLRs in the innate immune system is the detection of PAMPs. In a robot the proprioceptive sensors which monitor the state of the robot can be considered to be analogous to TLRs. For example a temperature sensor, measuring the temperature of a motor within a robot might have a TLR

associated with it containing a function (see Figure 1) which determines if the TLR gets triggered and by what amount. PAMPs, in terms of robotics are signals received by the robot's proprioceptive sensors (TLRs). These can trigger the TLRs starting the immune response in order to prevent possible damage in the long run. For example a temperature which exceeds various predefined thresholds might trigger responses designed to limit or prevent damage. In the natural innate immune system the action of TLRs leads to the generation of an inflammatory response via a number of pathways and mechanisms. This response is initially characterized by the generation, accumulation and diffusion of cytokines through the local tissues and into the bloodstream. In the longer term, continued inflammation results in a sustained "stress response" which has wide-ranging and diverse effects at a number of levels. This can affect physiological responses, behaviour and psychological state. These responses might vary from protection of an inflamed area, to the reduction of use of a limb due to localized pain through to increased sleep periods in severe cases. These varied responses can be incorporated into an innate artificial immune system with the help of the SOM. This can be achieved by activating the SOM using the current state vector of the robot (represented by the states of activation of all TLRs in the robot) and responding appropriately to affect the high-level controller, by releasing hormone into a neuro-endocrine controller for example. Whilst not implemented here assignment of remediation activities to particular nodes of the SOM (such as specification of which hormone to release) could be achieved automatically by examining which components of the robot are the source of the inflammation and selectively suppressing control system components which access those components. In the first instance this is a reasonable assumption, but in those cases where this response is insufficient to prevent further inflammation the SOM can be used to "spread" the inflammatory response to neighbouring nodes in order to suppress activity of components in closely related states. The gradual spreading of inflammation through the SOM ensures that the dependence on "engineered-in" relationships between component failures and remediation activities is only used in the first instance. When such relationships are incompletely or incorrectly assigned, the spreading to other closely related remediation activities improves the likelihood of an appropriate response being elicited in a computationally inexpensive and extensible manner.




Fig. 1. Schematic of the TLRs' functions

For example, this could enable the system to locally engage in an activity in an inflamed area in order to prevent damage. In the case of an overheating motor this localisation feature ensures that a nearby fan gets switched on, rather than a fan in a distant part of the robot.

The inflammatory response is simply accumulated over time from the individual TLR response levels. The sum of TLR activity is calculated at each time step and added to the current inflammation level. Also at each time step the inflammation level is geometrically decayed. Thus the formula for updating the inflammation at each time step is as follows:

$$inf_{t+1} = decay \times (inf_t + \sum_{x=1}^n f(TLR_x)) \quad (1)$$

where inf_t is the inflammation level at time t , $decay$ is a scalar in the range $0 < decay < 1$ and $f(TLR_x)$ is the activation level of the x 'th individual TLR from a set of n in the complete system.

3 Innate Autonomy

A detailed description of the functionality of the model follows, presenting a framework for robot autonomy based on the innate immune system.

Assuming a simplistic robot comprising of four motors, two fans and four sensors each measuring the temperature of one motor, a description of each step of the model is given. The robot is initially in a stable, homeostatic state from which it will deviate over the duration of the description of the model. The homeostatic state is defined to be when the four motors are operating continuously with the fans switched off. All four motors can be switched on and off at any point in time, according to the activity of TLRs (based on the motor temperatures). Both fans can also be triggered (also by the TLRs) to cool the motors.

3.1 PAMPs

At regular intervals sensors within the robot collect data about the robot's state and convert this data into signals. These signals are analogous to PAMPs in the human body. These signals are collected by the corresponding TLRs in order to monitor and respond locally to the state of the system. In our example robot these are simply the temperatures of each individual motor.

3.2 TLRs

If one motor is overheating while the others are functioning correctly, the associated sensor generates a PAMP which is passed to the related TLR. A PAMP is defined to be a sensor reading that deviates from normal according to a predefined function, which operates as part of the TLR. Each TLR has a predefined

set of responses as shown in Figure 1. Once the TLR receives the signal (ie. the temperature reading), it evaluates it according to predefined condition and response pairs shown in Table 1. The TLR also returns the inflammation level associated with the particular response:

$$y = f(TLR_x) \tag{2}$$

where x denotes the TLR in question and y denotes the inflammation level associated with the action that should be performed when the robot is in that state. Such functions can be implemented in terms of simple mathematical functions, lookup tables, fuzzy logic operators or any other appropriate technique. An example input could be the value 50, which represents the temperature of one of the motors and a response is generated according to the following lookup table (for example):

Table 1. Function table

Condition	Response
$T_x < 40^\circ C$	\emptyset
$40^\circ C < T_x < 80^\circ C$	Fan On
$T_x > 80^\circ C$	Fan On, Motor Off

This means, that the outcome of the function f will be the action *Fan On*. This is a local immediate response to the trigger of a single TLR. If the temperature is within the acceptable range, no action will be taken.

	Response ₁	Response ₂	...	Response _n
TLR ₁	1.0	1.0	...	$f(TLR_1)$
TLR ₂	0.0	0.0	...	$f(TLR_2)$
...
TLR _x	$f(TLR_x)$
$\sum_{m=1}^x f(TLR_m)$

Fig. 2. Input Feature Vector

In this model implementation, the stable state inflammation level is represented with the real value 0.0, while the TLR triggered state is 1.0. This is the contribution to the inflammation level described above. Once the model collects the outputs of the TLR functions of each individual TLR, a vector is created from the responses as shown in Figure 2. This vector is used as input to the SOM and the sum of its components is used to update the inflammation level according to equation 1.




Fig. 3. Self-Organizing Map trained using data acquired from the simulated robot control system. The upper left corner represents normal operation states, and the dark patch in the lower right quadrant represents states where many or all of the TLRs are responding. Other regions of the map represent states where fewer TLRs are responding.

3.3 Self-organizing Maps

The higher level state representation of the robot is encoded using a SOM [5].

The strength of a SOM algorithm in the context of this work is the way it deals with multidimensional input vectors. The algorithm is able to cope with large amounts of n -dimensional data and find correlations between them. This means that a system incorporating a SOM is highly scalable, as large numbers of input sensors can be dealt with. Upon finding a correlation between input vectors, the algorithm locates an appropriate neuron within the SOM, which consequently gets activated. This process is performed in an unsupervised manner, thus avoiding tedious and possibly inaccurate supervised methods, which would only allow a limited set of states to be represented within the map. A SOM is a low dimensional representation of the input data which preserves the topological properties of the input and explicitly represents multiple relationships between similar states. This feature enables the proposed system to evolve the map in a way which can be exploited for the purpose of inflammation. Neurons within the SOM which are topologically in close proximity represent states with certain similarities and thus result in only slightly different responses when activated. This is in contrast to most traditional statistical analysis methods such as cluster analysis or minimal spanning trees which do not unambiguously and explicitly represent such rich relationships between data items. The SOM also allows the possibility of learning on-the-fly without requiring discontinuous reorganisations of the state map which can result using statistical analyses such as cluster analysis.

The SOM contains all possible states of the robot, distributed across the map in a topologically ordered fashion and clustered according to similarity of the states. Initially the SOM is trained on a set of known problematic as well as stable states. This gives the map an informed starting point, from which it can evolve and adapt over the lifetime of the robot. A major feature of a SOM is the clustering effect which means that general robotic states can be identified in the maps produced when trained in this way. An example of this is the stable/homeostatic state; this state will be represented within the SOM by a cluster of similar nodes in which most of the TLR responses are zero. This can be seen in figure 3 in the top left corner of the map. In contrast the dark region in the lower right quadrant of the map has clustered all the states in which two motors are overheating and can be considered to be a stressed state of the robot, and if the robot remains in this state for long periods then inflammation will result and spread the activation throughout the map.

The input into the SOM is the TLR vector, which contains all TLR responses. This vector is presented to the SOM and the algorithm finds the node within the map which is closest to the input feature vector. In our case this is measured using the Euclidean distance.

3.4 Neuro-endocrine Control

The system then passes on the responses, which correspond to the winning node within the SOM, in order to influence the higher level control mechanism's behaviour. This response could be achieved in a number of ways, but perhaps a good candidate would be using a neuro-endocrine control system [7,8] where the artificial hormone is simply the inflammation level. These neuro-endocrine controllers rely on standard multi-layer perceptron neural networks with the simple addition of sensitivity to hormone concentrations built into their synapses. Thus the neural networks in the control system could be selectively (selection being performed by the SOM) suppressed by the application of the inflammation level as an artificial hormone at their synapses in the (now standard) neuro-endocrine way:

$$u = \sum_{i=0}^{nx} w_i \cdot x_i \cdot inf_t \quad (3)$$

where n is the number of synapses at the artificial neuron, w_i is the weight associated with the i 'th synapse, x_i is the input to that synapse and inf_t is the inflammation level at the time t . This new activation level is then used with the standard output function:

$$o = \frac{1}{1 + e^{-u}} \quad (4)$$

where o is the output from the neuron in question. This provides a simple but effective way of affecting the higher level control systems of the robot.

3.5 Spreading Inflammation

The clustering effect of the SOM offers a way of dealing with local as well as more widespread problems in a way which is analogous to inflammation. The robot is in a stable/homeostatic state if all its actuators are working correctly. Once problems start to occur, the nodes which become activated within the SOM fall outside the cluster of the stable behaviour. Once in such an unstable state the artificial innate immune system first deals with the problem locally at the level of TLRs. In case this local prevention does not return the robot to a stable state within a short period of time, inflammation starts to spread to neighbouring nodes of the current state node. This way the system deals with the problem by performing similar, yet slightly different responses, until the problem is rectified and the robot is returned to a stable state (a node within the SOM is activated which belongs to the cluster of stable/homeostatic behaviour).

4 Proof of Principle

A proof of principle implementation has been developed to demonstrate the key features of the operation of the model as described above. The model contains a small number of TLRs and uses inflammation responses generated by them to modify behaviour of a very simple high level control system. The inflammation response is integrated across the system and is decayed in the manner indicated above. Simple physical models of heating and cooling of motors are included in the simulation. The SOM component is not currently integrated into the system, but the vectors representing the system state were collected during the execution of the model and were used to train an SOM to prove the principle. This implementation has been performed as a simulation containing the important parts of the robot's functionality. The following results were obtained, supporting the proposed principle and its viability in a future physical system implementation.

4.1 Description of the Model

The simulated robot has two motor compartments: one for the front two wheels and one for the rear two wheels. Each wheel has a separate motor as is common in all-terrain robots. Each compartment also has a single cooling fan which is responsible for cooling the pair of motors in that compartment. Each motor has a TLR associated with it which monitors the motor's temperature. Each TLR has three possible states. The "normal" state is that the motor is enabled and the fan is switched off. When the motor reaches a predefined threshold temperature the TLR will activate and switch on the fan in that compartment. If the motor reaches a second, higher threshold temperature which endangers the motor then the TLR will activate a thermal cut-out which cuts all current to the motor in question in order to allow it to cool. This disables the motor and thus deprives the high-level control system and the robot as a whole of the use of that motor. The simulation ensures that the temperature of the motors increases proportionally to the current passing through it. The motor model also includes a simplistic but

sufficiently realistic cooling curve, the effects of which can be discerned in figure 5. The current applied to the motor is controlled by a high-level control system, which in the simulation is a simple fixed sequence of instructions. The purpose of the simulation is to demonstrate the action of the innate immune system components, and thus the implementation of a neuro-endocrine controller was not deemed necessary.

It is important to note that decisions are taken by the TLRs without the intervention of the high level control system, and they have to be considered as the first response of the immune system. The high level control system might then be influenced to change its behaviour depending on the inflammation present in the system through a scheme such as the neuro-endocrine controller outlined above (see section 3.4). In this model a more simplistic high-level control mechanism is used, but importantly it *is* affected by the inflammation level and modifies the requested current taking this inflammation level into account. This is a very simplistic remediation mechanism.

5 Results

Figure 4 shows how increasing current causes an increase in inflammation. The oscillations in the inflammation are due to the action of the TLRs switching the cooling fans and the motors themselves on and off. The effect of the inflammation is also to reduce the currents requested by the high level control system using a simple inversely proportional relationship (see Figure 6). The high-level control system is at the same time always attempting to return the motor currents to the requested levels.

Figure 5 shows the temperature of one motor over a period of time varying with the current. For a current of 0.1, after reaching the limit temperature of 40 (this value was fixed arbitrarily) a response is performed by the TLR which causes the fan to switch on. This operation causes the temperature of the motor to decrease. However the high level control system is trying to return current to the requested level. Considering a current of 0.1 the fan is always able to control the temperature. This pattern can also be seen when the current is 0.2. A different case occurs when the current is 0.5, this means that the high level control system is driving the motor at a high rate in order to fulfill its aim. This causes the TLR to activate the fan and frequently switch off the motor to prevent damage.

Figure 6 shows the effects of varying current over time in different motors and the resultant inflammation level. In this experiment motor1 simulates the occurrence of a fault, resulting in excessive current at time step 500. This causes the inflammation level to rise in steps as the requested current increases at time steps 1000 and 1500. The dramatic increase in inflammation at time step 2000 is due to the failure of the fan to cool motor1 and subsequent coincidental failure of motor2 and motor3. This inflammation comes from the activity of TLRs 2 and 3 as they activate the other fan and switch off the motors when required. At time step 2500 the faults are removed from the motors and the system returns to normal operation. This type of over-current condition can result from sticky




Fig. 4. Increasing current causes increasing inflammation level




Fig. 5. Impact of varying motor current over the time

motor bearings or fouling of axles by long grass and is relatively common in drive motors of all-terrain robots. The figure illustrates the way in which the inflammation level varies and responds to the state of the robot and how it can rapidly return to “normal” when faults are dealt with.




Fig. 6. Effect of the current change on the inflammation level

Figure 3 shows the SOM as generated using the state input feature vectors taken from the above experiments. A clear cluster, representing the homeostatic state, can be seen in the upper left region of the table. This cluster comprises of states which contain value 0 for all TLR responses. This value represents no triggering activity of the TLRs. By contrast, the dark region in the lower right quadrant represents triggering of TLRs both to switch on a fan and to switch two motors off to prevent damage. The region in the centre of the bottom row represents triggering of a single fan, and is bordered by regions to the left and right which represent switching on the fan in the other motor compartment (left) and switching off a motor in the overheating compartment (right). These adjacent regions can be used to highlight what might happen if inflammation caused by the single fan in the first motor compartment persists and is required to spread through the SOM. Activation of the adjacent regions mentioned will trigger preventative high-level actions appropriate for these closely related states. For example reducing current in the affected motors is likely to be one of the actions taken in order to pre-empt the triggering of the TLRs in the other components.

6 Conclusion

A scheme for incorporating low-level damage prevention and maintenance activities into a coherent biologically inspired control paradigm has been proposed, based on an innate immune system. Three important aspects of the innate immune

system have been applied with clear analogies between a robotic and a human immune system. These are the notion of TLRs, inflammation and localisation. The system has been developed with the help of a SOM as an adaptive state representation of the robot, which enables local as well as global failure prevention and ratification. A model has been implemented to support the above given principles. Results from performed experiments show that the activity of TLRs causes an incremental inflammatory response over time, in case the robot is not returned to a stable state in a reasonable period of time. This inflammatory response can be used along with the SOM to locate the affected area of the robot in order to deal with it on a more global level. The presented preliminary results support the described principles and encourage future development of a real robot implementation incorporating immune, neural and endocrine control components.

Some of the potential advantages of this scheme are highlighted throughout the earlier parts of the paper, but perhaps one of the most significant is that it offers a relatively simple mechanism for integrating existing engineering knowledge of how to deal with particular problems locally with the higher level and less well defined parts of the control system. Some potential disadvantages include: that the engineer must still manually assign fault conditions and remediation activities for local conditions which leaves room for oversight and error; the overhead of maintaining a system-wide map of the robot's state may cause problems (whilst maintaining the SOM is unlikely to be computationally expensive, the gathering of its input data from all over the robot could be problematic); and last but by no means least, it is not yet clear how such an innate system might fit into a full multi-layer artificial immune system for a robot. Apart from the obvious next step (implementing the system as described on a real robot), a pressing piece of future work will be identification of how this might be achieved.

It is also interesting to consider the effect of the system on the combination of task achieving behaviour and survival behaviour. Whilst the mechanism here does not explicitly address this (potential) conflict, it does provide an interesting possibility when combined with the neuro-endocrine control systems described above and elsewhere. The "soft" switching, suppression and promotion of behaviours is precisely what this conflict requires in order to achieve the sorts of complex trade-offs that are observed in nature. The addition of an inflammation based driver for such behaviour mediation provides an additional homogeneous driver *specifically for maintenance of homeostasis*. This is an important step forward as it provides a truly integrated mechanism for promotion of survival behaviours within task achieving robot systems.

Acknowledgements

This paper arose as a direct result of the ARTIST Network² funded Student Spring School held at Aberystwyth on April 8-13th 2006. Whilst the direct contributors are listed as authors we would like to thank all of those who attended

² ARTIST is an EPSRC (UK) funded network to support artificial immune systems research.

for help, discussions and ideas. In particular we would like to thank Julie Greensmith for coaching us on the functioning of the innate immune system and the inflammatory response.

References

1. R Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1):14–23, 1986.
2. W. Clancey. *Situated Cognition: On Human Knowledge and Computer Representations*. Cambridge University Press, 1997.
3. L N de Castro and J Timmis. *Artificial Immune Systems: A New Computational Intelligence Approach*. Springer-Verlag, 2002.
4. J. Greensmith, U. Aickelin, and S. Cayzer. Introducing Dendritic Cells as a Novel Immune-Inspired Algorithm for Anomaly Detection. In *Proceedings of the 4th International Conference on Artificial Immune Systems*, volume 3627, 2005.
5. T Kohonen. Self-organised formation of topologically correct feature maps. *Biological Cybernetics*, 43:59–69.
6. Teuvo Kohonen. *Self-organising Maps*. Springer, 1995.
7. M. Neal and J. Timmis. Timidity: A Useful Mechanism for Robot Control? *Informatica*, 27(4):197–204, 2003.
8. M Neal and J Timmis. *Recent Advances in Biologically Inspired Computing*, chapter : Once more unto the breach... towards artificial homeostasis? IGP, 2004.
9. L Sompayrac. *How the immune system works*. Blackwell, 2002.
10. Bruce C. Trapnell Jr. A peer-to-peer blacklisting strategy inspired by leukocyte-endothelium interaction. In *Proceedings of the 4th International Conference on Artificial Immune Systems*, volume 3627, 2005.