

Environment classification in multiagent systems inspired by the adaptive immune system

Danesh Tarapore^{1,2}, Anders Lyhne Christensen³, Pedro U. Lima¹ and Jorge Carneiro²

¹Institute for Systems and Robotics (ISR), Instituto Superior Técnico (IST), Lisbon, Portugal

²Instituto Gulbenkian de Ciência, Oeiras, Portugal

³Instituto de Telecomunicações & Instituto Universitário de Lisboa (ISCTE-IUL), Lisbon, Portugal
daneshtarapore@gmail.com

Abstract

The adaptive immune system in vertebrates is a complex, distributed, adaptive system capable of effecting collective multicellular responses. Our study introduces many of the desirable properties of this biological system to decentralized multiagent systems. We adopt the crossregulation model of the adaptive immune system involving interactions between effector and regulatory cells. Effector cells can mount beneficial immune responses to microbial antigens as well as pathologic autoimmune responses to self-antigens. Deleterious autoimmunity is prevented by regulatory cells that suppress the effectors to tolerate the self-antigens. We redeploy the crossregulation model within a multiagent system by letting each agent run an ODE-based instance of the model. Results of extensive simulation-based experiments demonstrate that a distributed multiagent system can mount different responses to distinct objects in their environment. These responses are solely a result of the dynamics between virtual cells in each agent and interactions between neighboring agents. The collective dynamics gives rise to a meaningful “self”-“nonself” classification of the environment by individual agent, even if these categories were not prescribed a priori in the agents.

Introduction

Multiagent systems (MAS) comprise a large number of research domains, ranging from software agents to multirobot systems, and play an important role in several applications, such as supply chain management, transportation logistics and network routing. The coordination of agents in a MAS is a major challenge because agent behavior depends not only on interactions with their immediate environment but also on the behavior of other agents. A centralized control approach may not always be feasible due to computational and/or communication constraints on agents (e.g., Crespi et al. (2008); Mermoud et al. (2010)). Distributed control, on the other hand, is often complicated to realize because the behavioral rules for the individual units cannot be easily derived from a desired macroscopic behavior (e.g., Parker (2000); Yamins and Nagpal (2008); Hamann (2010)). In the design of large scale distributed systems, several researchers have therefore taken inspiration from nature e.g., aggregation of amoeba into slime mold (Payton et al., 2003), quorum sensing and communication in bacteria (Sahin, 2005),

division of labor in social insects such as ants and honey bees (Parker et al., 2003; Waibel et al., 2009; Hauert et al., 2009; Tarapore et al., 2010; O’Grady et al., 2010).

The cell collective that constitutes the adaptive immune system has been extremely successful during the course of evolution as evidenced by its presence in all jawed vertebrate species (Janeway et al., 1997). Central to the success of these cells is the important role they play in establishing and maximizing the capabilities of the immune system, by allowing an exquisite “self-nonself” discrimination that is not present in invertebrates. The cell collective is able to recognize and mount specific immune responses to microbial agents that the organism and its ancestors had never faced before. It does this immersed in the constant presence of diverse and abundant body antigens, which are molecularly similar to the microbial antigens. In normal healthy individuals, sporadic microbial invaders are specifically eliminated by immune responses and, at the same time, pathologic autoimmune responses to the abundant body antigens is prevented, i.e. natural tolerance to “self” is maintained. Experimental evidence indicates that natural tolerance results from the dynamics and interactions between specific regulatory and effector T-cells (e.g., Sakaguchi (2004)). Interestingly, the decentralized nature of the interactions may impart a high degree of robustness for natural tolerance, without the need of maintaining a specific, genetically hardwired, “memory” of self-antigens.

The decentralized and adaptive nature of the immune system is a source of inspiration for designers of large scale MAS. In particular, the ability of the system to dynamically maintain natural tolerance has many industrial applications. Some typical studies that take inspiration from this “self”-“nonself” discrimination capability of the immune system include, distributed intrusion detection systems (Nino and Beltran, 2002; Kim and Bentley, 1999), and fault tolerance systems (Bradley and Tyrrell, 2000, 2001; Canham and Tyrrell, 2002). However, most of these models assume which particular antigens or features are prescribed as “self”, and consequently the system is trained to tolerate them. While this approach does provide some in-

interesting results of robust feature classification, it does not fully incorporate the dynamics and adaptive nature of the immune system. This led us to propose the use of the *cross-regulation model* (CRM) for the maintenance of tolerance. The CRM (Leon et al., 2000, 2003, 2004; Carneiro et al., 2007) suggests a dynamics of interactions between cells of the immune system, that allows the system to discriminate between antigens based solely on their density and persistence in the environment. The system is able to tolerate body antigens (i.e. “self”) that are characteristically persistent and abundant, and to mount an immune response to foreign pathogens, that are characterized as being neither persistent nor abundant. The model has been used successfully in spam detection (e.g., Abi-Haidar and Rocha (2008)) and document classification (e.g., Abi-Haidar and Rocha (2010, 2011)) scenarios, making it a good candidate for MAS for environment classification.

In this study, we propose a CRM-based approach to replicate the capability of the immune system in maintaining tolerance. We use an agent-based simulator to model a situation where individuals have to tolerate certain features, while mounting an immune response against others. The different environmental features are represented by different sensory stimuli in the environment, and their nature (“self” or “nonself”) are not known by the agents beforehand. We demonstrate the capacity of the system to tolerate specific environmental features that may be characterized as persistent and abundant (“self”), while mounting an immune response against others (“nonself”). In addition, the system response is resilient to sensory noise, and can respond correctly under varying environmental conditions.

The rest of the paper is organized as follows: In the following section, we describe the CRM. We then present the application of the CRM in a MAS. We go on to report the results of our experiments in different environmental conditions and under varying levels of perceptual noise. Finally, we discuss our approach to environment classification and highlight the conclusions of this study.

The Crossregulation Model

Two general principles are essential for the viability of multicellular organisms. Firstly, the persistence of any cell lineage requires that its cells recurrently interact with other cell types in the organism. Cells that fail to interact with other cells eventually die. Secondly, the growth of a cell population involves density-dependent feedback mechanisms controlling individual cell proliferation. These feedback mechanisms may involve (i) indirect interactions among cells (such as a competition for limited growth factors) and (ii) direct interactions, such as contact inhibition. These two principles of multicellular organization are the foundation of the crossregulation model, and have been justified extensively in Carneiro et al. (2007). Below, we outline the model and highlight its interesting properties that are later replicated

with a cell recruitment mechanism.

The CRM describes the population dynamics of cells of the adaptive immune system, based on three mutually interacting cell types: (i) Antigen presenting cells (APCs) that display the antigen on their surface. Individual APCs have a fixed number of sites (s) on which effector and regulatory cells can form conjugates; (ii) effector cells T_E that can potentially mount an immune response which, depending on receptor specificity, can be directed to foreign pathogens or to self-antigens; and (iii) regulatory cells T_R that suppress proliferation of T_E cells with similar specificities. Furthermore, the APCs are classified into different sub-populations of equivalent APCs, with each APC in a sub-population presenting the same antigen on its surface. Similarly effector and regulatory cells are also classified into different sub-populations or clones according to their specificity.

The dynamics of T-cell population is regulated by the following density-dependent feedback mechanisms. (i) Effector and regulatory cells that are unable to interact with APCs are slowly lost by cell death. (ii) The proliferation of effector and regulatory cells requires interactions with APCs and depends on interactions these T-cells make with each other. Proliferation of the T_E cell population is promoted by the absence of regulatory cells on the APC. In contrast, T_R can only proliferate following co-conjugation with effector cells on the same APC. Additionally, T_E and T_R cells interact indirectly by competition for access to conjugation sites on APCs.

Behavior of cell population

Considerable work has focused on analyzing the properties of the CRM, and the underlying dynamics between T_E , T_R and APCs (Leon et al., 2000, 2003). An interesting characteristic of the CRM is the ability to discriminate between antigens based on their density. At low concentrations of APCs, the system evolves into a stable state composed only of effector cells (immune response). In contrast, at higher values of APCs, the system demonstrates bistable behavior. At these concentrations of antigens, the system can evolve either into an equilibrium state consisting predominantly of effector cells (immune response), or into a state composed largely of regulatory cells (tolerant response). The system develops into the regulatory cell dominated state, provided that the seeding population has sufficient T_R cells. By contrast, if T_R cells are initially underrepresented, T_E cells will competitively exclude the former from the system. Consequent to the antigen density dependent response, the effector cells are made tolerant to antigens that are persistent and abundant. In addition, the effector cells are free to mount immune responses to antigens that are not persistent or not abundant.

Table 1: Parameters of the crossregulation model.

Param.	Description	Value (a.u.)
A_j	Density of APCs of population j	—
s	Maximum number of T-cells that can bind to an APC	3
E_0	Seed density of effector cells	10
R_0	Seed density of regulatory cells	100
E_i	Density of effector cells of clone i	—
R_i	Density of regulatory cells of clone i	—
T_i	Density of T-cells of clone i	$E_i + R_i$
C_{ij}	Density of conjugates between T_i and A_j	—
γ_c	Conjugation rate of T-cells to APCs	10^{-1}
γ_d	Deconjugation rate of T-cells from APCs	10^{-1}
σ_E	Influx rate of new effector cells	10^{-3}
σ_R	Influx rate of new regulatory cells	0.6×10^{-3}
π_E	Proliferation rate of effector cells	10^{-3}
π_R	Proliferation rate of regulatory cells	0.5×10^{-3}
δ	Death rate of effector and regulatory cells	10^{-5}

Mathematical formulation of the model

The dynamics of the interactions between effector and regulatory cells, with APCs is described by a set of ordinary differential equations in the following variables: (i) The number of effector E_i and regulatory R_i T-cells of clonal type i , where $i \in \{1, 2, \dots, N\}$ and N is the number of T-cell clones. (ii) The number of APCs A_j , where $j \in \{1, 2, \dots, M\}$ and M is the number of different antigen types. (iii) The number of conjugates C_{ij} formed between effector and regulatory cells from clone i and APC from population j .

For the effector E_i and regulatory R_i cells of clone i , we have:

$$\frac{dE_i}{dt} = \sigma_E + \pi_E E_i^* - \delta E_i \quad (1)$$

$$\frac{dR_i}{dt} = \sigma_R + \pi_R R_i^* - \delta R_i \quad (2)$$

where the involved quantities are defined in Table 1.

The equations for E_i (eq 1) and R_i (eq 2) have three terms. The first term represents the influx of new cells, which is assumed to be constant. The second term accounts for the proliferation of activated effector and regulatory cells. Finally, the death of T-cells is represented by the third term of the equations. In the simulations, we generate all T-cell clones with similar initial conditions i.e., $\forall i$, $E_i(0) = E_0$ and $R_i(0) = R_0$.

The density of activated T_E and T_R cells of each clone are computed in a stepwise manner. Let us consider the interactions between the i -th T-cell clone and the j -th APC population. The dynamics of the conjugates C_{ij} is described by the following equation:

$$\frac{dC_{ij}}{dt} = \gamma_c \theta_{ij} \left(T_i - \sum_{j=1}^M C_{ij} \right) \left(A_j s - \sum_{i=1}^N C_{ij} \right) - \gamma_d C_{ij}$$

where $T_i = E_i + R_i$, and γ_c and γ_d involve the conjugation and deconjugation rates between APCs and T-cells, respectively (parameters in Table 1). In the above equation, new conjugates are formed by the free T-cells of clone i with the available sites on APCs of population j at rate γ_c . The conjugation rate is also controlled by the affinity (θ_{ij}) between the T-cells and APCs. The existing conjugates dissociate at rate γ_d . The conjugation and deconjugation of T-cells from the APCs is a fast process with respect to the overall T-cell clone dynamics. Consequently, we solve at each time step, the steady state values of the conjugates by the Euler-Heun adaptive step method (Butcher, 2003).

The density of activated effector E_i^* and regulatory R_i^* cells can now be calculated (for details see Appendix A). Conjugated effector cells are activated in the absence of regulatory cells on the same APC. In contrast, conjugated regulatory cells can only be activated if at least one effector cell is simultaneously conjugated to the same APC.

The population dynamics behavior exhibited by the CRM is governed by two key composite parameters representing the effective growth rates of T_E and T_R cell populations (Leon et al., 2000). These two parameters are directly proportional to the basic parameters controlling population growth i.e., conjugation constant (γ_c), affinity between T-cell and APCs (θ_{ij}), influx rate of new effector and regulatory cells (σ_E and σ_R), proliferation rates of these two types of T-cells (π_E and π_R), and the density of APCs (A_j). The effective growth rates of the T-cells is also inversely proportional to the death rate (δ) of the corresponding population. The composite T_E and T_R growth parameters define four parameter regimes according to the resulting cell population behavior. Three parameter regimes result in a single stable state that may correspond to either: (i) extinction of all T-cells ($T_E = 0, T_R = 0$), (ii) immune state ($T_E > T_R$), or (iii) tolerant state ($T_E < T_R$). The fourth parameter regime corresponds to a bistable system where both immune and tolerant states are stable. A detailed analysis of these parameter regimes is provided in Leon et al. (2000). For our present study, the parameter values have been set so that at low APC densities, the system evolves into a single state composed only of effector cells. By contrast, at relatively high density of APCs, the system is bistable and can evolve either into an immune or tolerant equilibrium state.

CRM in a Multiagent System

In this section, we demonstrate how the CRM can be implemented on a distributed embodied multiagent system in order to give the system the capacity to classify different features in the environment based on their concentrations. Features that are persistent and abundant are to be tolerated, while features that are present at a low density are not. We show that the multiagent system is able to adapt online and that it is resilient to perceptual noise.

We use a stochastic, spatial, discrete-time simulator. The

simulated environment is toroidal and has a size of 10×10 units. The MAS is composed of 50 point-sized agents that perform a random walk: each agent move at a constant speed of 0.01 units/time-step, and has a probability of 0.01 of changing to a new random direction each simulation step. The agents detect features of static objects within their sensory range (1 unit) and run an internal and individual instance of a CRM in order to determine if the objects should be tolerated or not (see details below).

Individual features of the static objects in the environment are encoded in Boolean form (present= 1, absent= 0), and then concatenated to form a binary string, the *feature vector*. At the start of each time-step, an agent computes the density of each feature vector (FV_j) within its sensory range. In the agent's internal CRM instance, APCs are then generated corresponding to each of the feature vectors perceived. Each APC presents an individual feature vector to the T-cells. The number of each type of the APCs generated $APC_j = FV_j$, for $j \in \{1, \dots, M\}$, where M is the number of different feature vectors perceived by the agent.

The T-cell clones (T_1, T_2, \dots, T_N), each have a different receptor encoded as a binary string, which determines their affinity to the APC population. The affinity between T cell clonal i and APC population j is denoted by θ_{ij} :

$$\theta_{ij} = \exp\left(-\frac{H(i, j)}{c}\right) \quad (3)$$

where H is the Hamming distance between the receptor of T_i and the feature vector presented by A_j , and c is the cross-reactivity between T-cells and APCs. A high value of c would result in all T-cell clones having a high affinity to all APC populations. By contrast, at low c , each T-cell clone would have a high affinity to only one distinct APC population.

At the start of the simulation, the number of effector and regulator cells on each agent is initialized to E_0 and R_0 respectively. Following this, Algorithm 1 (parameters in Table 2) is performed by the agents in each simulation time-step, allowing the agents to execute the behavior designed in the CRM. The agents begin by sensing their local environment and computing the density of feature vectors. Perceptual noise is modeled by randomly flipping the binary representation of one of the feature with probability x . The CRM is then numerically integrated for time S , allowing the system to respond to the different APCs. After computing the number of effector and regulatory cells at time S , the cells diffuse among agents. In this communication phase, each agent selects a neighboring agent within its communication range. The selection is random following a linear distribution on the total number of T-cells associated with each agent in communication range. Following the selection, each agent sends and receives d of its effector and regulatory cells. Finally, the agent decides the nature of each feature vector FV_j sensed, as follows:

Table 2: Parameters of the stochastic simulator

Param.	Description	Value (a.u.)
N	Number of T-cell clones	4
M	Number of different feature vectors	4
c	Cross-reactivity between T-cells and APCs	0.4
x	Probability to add noise on a feature	0.1 – 0.5
S	Time CRM instance is executed, in a single simulation-step	10^5
d	Proportion of T-cells diffused to neighboring agents	0.5

$$E = \sum_{i=1}^N \theta_{ij} E_i \quad R = \sum_{i=1}^N \theta_{ij} R_i$$

where the feature vector is accepted as tolerant if $R > E$, else the object associated with the feature vector is removed from the environment by the agent.

Algorithm 1 An agent's control loop (simulation of an CRM instance)

- 1: {Perceive static objects}
 - 2: Compute density of feature vectors (FV_j) in sensory range of agent
 - 3: For each of the sensed feature vectors, add noise to one of the features with probability x
 - 4: Assign feature vectors to APCs i.e., $\forall j, A_j = FV_j$
 - 5: {Run instance of CRM}
 - 6: $time \leftarrow 0$
 - 7: **while** $time \leq S$ **do**
 - 8: $\forall i \in \{1, 2 \dots N\}$ and $\forall j \in \{1, 2 \dots M\}$, compute the number of conjugated cells C_{ij} in steady state, integrating using the Euler-Heun adaptive step method
 - 9: Using the number of conjugated cells, compute the updated number of effector and regulatory cells with the Euler-Heun adaptive step method. The adaptive step size is stored in h
 - 10: $time \leftarrow time + h$
 - 11: **end while**
 - 12: {Diffuse cells across neighboring agents}
 - 13: Randomly select one of the agents in the communication range following a linear distribution and weighted by the total number of cells on the respective neighboring agents
 - 14: Exchange cells with agent
 - 15: {Decide if feature vectors are to be tolerated or not}
 - 16: For each feature vector, compute the sum of effector and regulatory cells, weighted by their affinity.
 - 17: Tolerate the FV if total regulatory cells exceeds effectors, else mount an immune response i.e., remove the static object associated with the feature vector from the environment.
-

Experiments

We set up a series of experiments in order to evaluate the classification capabilities of a multiagent system operating according to the model described above. In a first set of experiments, we distributed two different types of static objects in the environment: one with a high density ($10/\text{unit}^2$) and one with a low density ($1/\text{unit}^2$). Both types of static objects were placed at random positions drawn from a uniform distribution. In each replication of the experiment, the feature vectors of the two types of static objects were picked at

random in such a way that one would be the complement of the other. Within the CRM conceptual framework the abundant objects are interpreted as body/self-antigens, while the low density objects are foreign or “nonself”. We endowed agents with the capacity to remove objects and therefore tolerance to “self” was interpreted as the persistence of the objects. We show that the MAS is, under some specific non-trivial conditions, able to tolerate abundant objects that will persist and to remove less abundant objects.

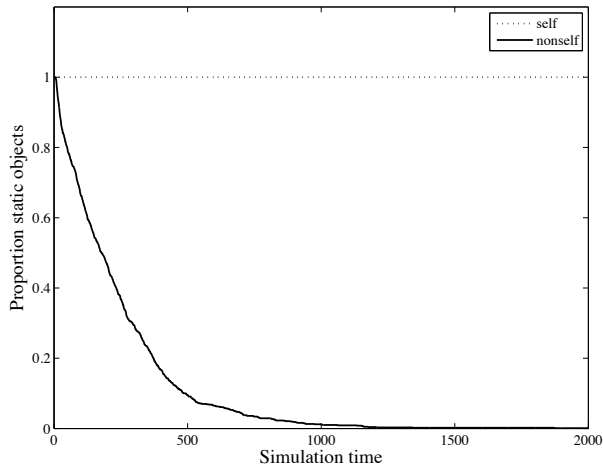


Figure 1: Mean proportion of static objects across 10 replicates. Individual agents had 30% probability perceptual noise.

In Fig. 1, we have plotted the mean proportions of static objects (with respect to the initial quantities) across 10 replicates, each for 2000 simulation steps and with 50 agents in the environment. The object density variance across simulation time was similar irrespective of the level of perceptual noise added ($x = 0.1 - 0.5$), and was therefore illustrated for a single case ($x = 0.3$, Fig. 1). After 2000 simulation time-steps, there was very little variation in the objects associated with “self”. The density of self-objects remaining at 10 for $0.1 - 0.4$ probability of perceptual noise, while at higher level of perceptual noise ($x = 0.5$), tolerance was maintained in all but two replicates (less than 0.2% of self-objects destroyed in each replicate). By contrast, the system exhibited an absence of tolerance to objects associated with “nonself” (Fig. 1 and 2). An immune response to these objects was mounted irrespective of the level of noise. However, the response was more effective at lower levels of noise (Fig. 2).

We set up a second series of experiments in order to evaluate the capabilities of a multiagent system to maintain tolerance under varying environmental conditions. These experiments were designed to assess the requirement for communication between agents. In this set of experiments, we divided the environment into three regions, with two con-

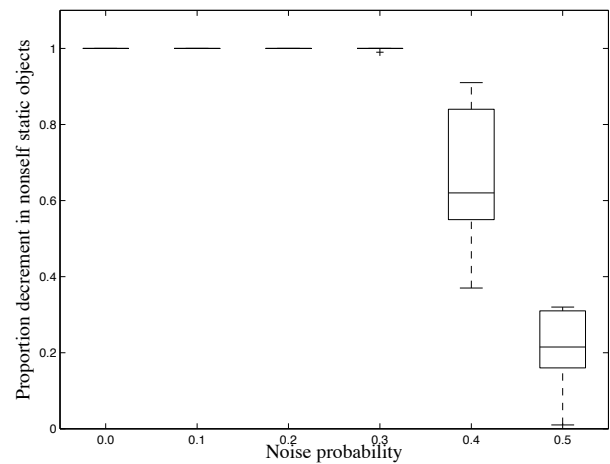


Figure 2: Proportion decrement of “nonself” static objects with different amounts of perceptual noise.

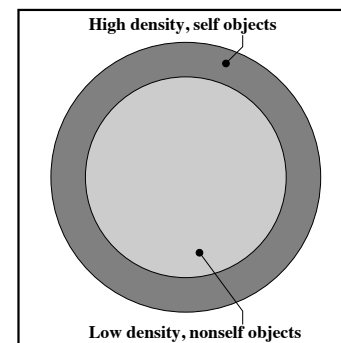


Figure 3: The heterogeneous environment used to investigate environment classification under varying environmental conditions.

centric circle of radii 4 and 5 units (Fig 3). Two different types of static objects were distributed in the environment in two different locations: one with a low density ($1.98/\text{unit}^2$) was distributed within the inner circle, and one with a high density ($70.7/\text{unit}^2$) was distributed between the inner and outer circles.

In Fig. 4 and 5, we have plotted the mean proportions of static objects (with respect to the initial quantities) with intra-agent communication suppressed and enabled respectively. Experiments were replicated 10 times, each for 2000 simulation steps and with 50 agent in the environment.

The communication of T-cells between agents had a strong effect on the maintenance of tolerance. In the absence of communication, the system was unable to maintain tolerance (Fig. 4). At 2000 simulation time-steps, the abundant “self” objects were removed from the environment in all 10 replicates. By contrast, in the presence of communication between agents, almost 100% of abundant “self”

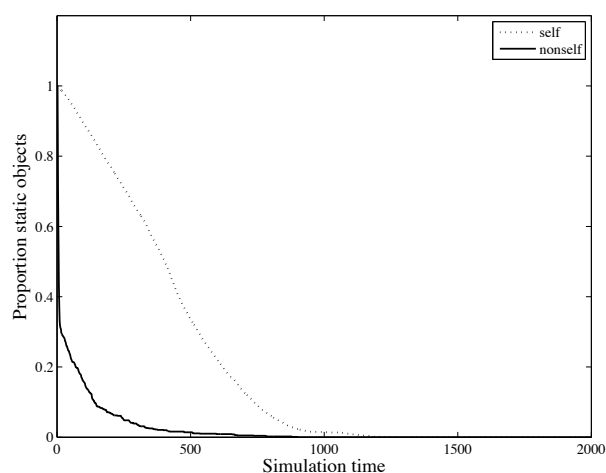


Figure 4: Mean proportion of static objects across 10 replicates, with heterogeneous distribution of static objects (“self” and “nonself”), and inter-agent communication suppressed.

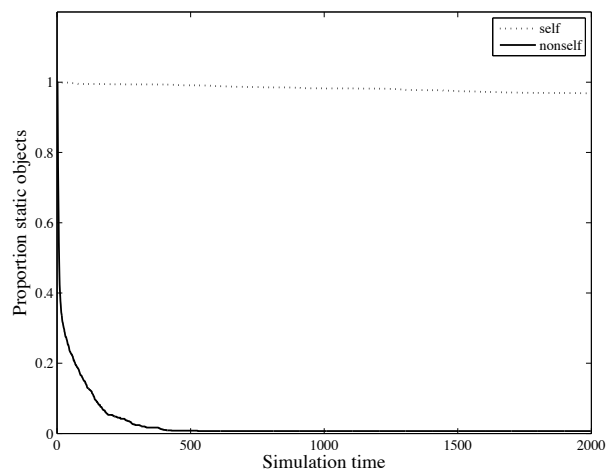


Figure 5: Mean proportion of static objects across 10 replicates, with heterogeneous distribution of static objects (“self” and “nonself”), and inter-agent communication enabled.

objects persisted in the environment at the end of the simulation (Fig. 5). In addition, the agents were also able to mount an effective immune response, such that all the “nonself” objects had been removed from the environment at the end of the simulation (Fig. 5).

Discussion

Our study revealed a robust maintenance of tolerance to “self”, understood as abundant antigens or features, irrespective of the level of perceptual noise on individual agents. Interestingly, even at a 50% chance to distort a sensed fea-

ture, the abundant “self” was largely tolerated. This resiliency to noise exhibited by the system was a consequence of the cross-reactivity between T-cells and APCs. At our level of cross-reactivity, regulatory cells with a high affinity to feature vectors of the “self”, were able to react with and consequently suppress effectors associated with a mis-read “self” feature vector (low Hamming distance apart) and consequently prevent their destruction. Separate experiments investigating the influence of this parameter, indicated a complete absence of tolerance at low values of cross-reactivity. By contrast, at very high levels of cross reactivity, regulatory cells suppressed effectors associated with all the sensed features thus preventing any discrimination by the system. Interestingly, the ability of our system to tolerate noise distinguishes it from a simple response threshold-based model for environment classification, wherein different feature vectors are assigned distinct tolerance thresholds, and the system response is governed strictly by the density of each feature vector type being above or below its corresponding threshold.

In simulated environments with a heterogeneous distribution of objects, the agents continued to classify environmental features correctly, despite the variations in their local environmental conditions. Our results revealed the requirement of communication of T-cells between neighboring agents in order to maintain the tolerance to abundant “self” objects. In the absence of communication, agents were unable to tolerate “self” objects when entering regions consisting of them. By contrast, in the presence of communication, regulatory cells communicated from agents already in the “self” associated region allowed the entering agents to respond faster to environmental changes and consequently, greatly improved their tolerance. The diffusion of T-cells between agents allows the agents to share information of their local environments and to perform better as a collective.

In our simulations, APCs are generated corresponding to each of the feature vectors. Each APC presents an independent feature vector present at that instance. Consequently, APCs related to a newly generated feature vector may not react to the existing T-cells in the agents’ history. This is because the reaction would be dependent on the feature vector chosen for this new event and its affinity to the existing T-cells. We illustrate this point with the following example: Consider an agent in an environment with FV_j presented by A_j at a density resulting in a tolerant response. The agent has in its history, T clonal-type T_i with $\Theta_{ij} = 1$. Consequent to the density of A_j , $R_i > E_i$. Now let us consider the agent moving into an environment resulting in another APC type A_k . However, the existing cells in the agents’ history may or may not react to this new APC, and the decision is stochastic and dependent on the choice of the new feature vector FV_k . In this system, for the existing T-cells to react with the new feature vector, $\theta_{ik} > 0$ and this is a direct consequence of the (preexisting) affinity mapping between

feature vectors and T-cell clonal types. Another possible approach wherein the history of the system could be explicitly taken into account would be the generation of APCs to represent various combinations of feature vectors. Based on the above example, APCs would present feature vectors of type $\{FV_j, FV_k, FV_j FV_k\}$. In this condition, existing cells in the agents' history would be able to respond to new feature vectors. Additionally, the system response would not solely be a consequence of the feature vector specific topology. However, the outcome of this scenario needs to be explored further.

In our experiments, we used a relatively abstract stochastic simulation in which mobile agents performing random walk perceived features on static objects present in the environment. The agents task, to distinguish between what is persistent and abundant and what is not, is a metaphor for a large class of detection and identification tasks in the field of MAS, and more specifically in multirobot systems (MRS): novelty detection, fault detection, intrusion detection, and so on. In our model, features were associated with external and immediately observable perceptual cues, but features may be computed based on other qualities, such as the behavior of nearby agents, proprioceptive sensory input, and environmental attributes. In this way, our CRM-based approach to classification in multiagent systems, could for instance give robots the capacity to distinguish between normal behavior and abnormal behavior. Since tolerance and its absence is determined online and does not require an initial training step, we expect that our CRM-based approach is particularly suitable to MRS operating in dynamic environments in which the task attributes may change over time and to MRS that adapt and change their behavior during task-execution. In this regard, we are currently investigating approaches to reduce the computational complexity of running the CRM on individual robots of a MRS.

Conclusions

In this study, we proposed an approach inspired by the capability of the adaptive immune system to maintain tolerance in multiagent systems. We further investigated the utility of this approach in task involving environment classification. Different environmental features were represented by different sensory stimuli in the environment, and their nature ("self" or "nonself") was not known by the agents beforehand i.e. it was not built into the individual agent's behavior. Our simulations revealed the capability of the collective of agents to tolerate features characterized as abundant and persistent, while mounting an immune response against specific features that were neither persistent nor abundant. Furthermore, the agent decision making was robust to perceptual noise and variations in their environmental conditions.

These encouraging results of our study provides a good stepping stone of our CRM-based approach for more detailed multiagent system experiments involving a broad

range of tasks.

Supplemental Data: Movies of MAS simulations are available online at <http://home.iscte-iul.pt/~alcen/alife2012/>.

Acknowledgment: This study was supported by the FCT grant PTDC/EEACRO/104658/2008.

References

- Abi-Haidar, A. and Rocha, L. (2008). Adaptive spam detection inspired by the immune system. In *Proceedings of the 11th International Conference on the Simulation and Synthesis of Living Systems, Artificial Life XI*, pages 1–8. MIT Press, Cambridge.
- Abi-Haidar, A. and Rocha, L. (2010). Biomedical article classification using an agent-based model of T-cell cross-regulation. In *Proceedings of the 9th International Conference on Artificial Immune Systems (ICARIS)*, pages 237–249. Springer-Verlag, Berlin, Germany.
- Abi-Haidar, A. and Rocha, L. (2011). Collective classification of textual documents by guided self-organization in T-cell cross-regulation dynamics. *Evolutionary Intelligence*, 4(2):69–80.
- Bradley, D. and Tyrrell, A. (2000). Hardware fault tolerance: an immunological solution. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, volume 1, pages 107 – 112. IEEE Press, Piscataway, NJ.
- Bradley, D. W. and Tyrrell, A. M. (2001). Immunotronics: Hardware fault tolerance inspired by the immune system. In *Proceedings of the 3rd International Conference on Evolvable Systems: From Biology to Hardware (ICES)*, pages 11–20. Springer-Verlag, Berlin, Germany.
- Butcher, J. (2003). *Numerical methods for ordinary differential equations*, chapter 23. John Wiley & Sons, West Sussex, England, second edition.
- Canham, R. and Tyrrell, A. (2002). A multilayered immune system for hardware fault tolerance within an embryonic array. In *Proceedings of the 1st International Conference on Artificial Immune Systems (ICARIS)*, pages 3 – 11. University of Kent at Canterbury Printing Unit.
- Carneiro, J., Leon, K., Caramalho, I., Van Den Dool, C., Gardner, R., Oliveira, V., Bergman, M., Sepúlveda, N., Paixão, T., Faro, J., and Demengeot, J. (2007). When three is not a crowd: a crossregulation model of the dynamics and repertoire selection of regulatory CD4⁺ T cells. *Immunological Reviews*, 216(1):48–68.
- Crespi, V., Galstyan, A., and Lerman, K. (2008). Top-down vs bottom-up methodologies in multi-agent system design. *Autonomous Robots*, 24(3):303–313.
- Evans, M., Hastings, N., and Peacock, B. (2000). *Statistical Distributions*, chapter 27, pages 134–136. John Wiley & Sons, New York, third edition.
- Hamann, H. (2010). *Space-Time Continuous Models of Swarm Robotic Systems: Supporting Global-to-Local Programming*. Cognitive Systems Monographs. Springer-Verlag, Berlin, Germany.

- Hauert, S., Zufferey, J., and Floreano, D. (2009). Evolved swarming without positioning information: an application in aerial communication relay. *Autonomous Robots*, 26(1):21–32.
- Janeway, C., Travers, P., Walport, M., and Shlomchik, M. (1997). *Immunobiology: The Immune System in Health and Disease*. New York: Garland Science.
- Kim, J. and Bentley, P. (1999). The human immune system and network intrusion detection. In *Proceedings of the 7th European Conference on Intelligent Techniques and Soft Computing (EUFIT)*. ELITE-Foundation, Aachen, Germany.
- Leon, K., Farob, J., and Carneiro, J. (2004). A general mathematical framework to model generation structure in a population of asynchronously dividing cells. *Journal of Theoretical Biology*, 229:455–476.
- Leon, K., Lage, A., and Carneiro, J. (2003). Tolerance and immunity in a mathematical model of t-cell mediated suppression. *Journal of Theoretical Biology*, 225:107–126.
- Leon, K., Perez, P., Lage, A., Farob, J., and Carneiro, J. (2000). Modelling t-cell-mediated suppression dependent on interactions in multicellular conjugates. *Journal of Theoretical Biology*, 207(2):231 – 254.
- Mermoud, G., Upadhyay, U., Evans, W., and Martinoli, A. (2010). Top-Down vs Bottom-Up Model-Based Methodologies for Distributed Control: A Comparative Experimental Study. In *Proceedings of the 12th International Symposium on Experimental Robotics (ISER)*, Springer Tracts in Advanced Robotics. Springer-Verlag, Berlin, Germany.
- Nino, F. and Beltran, O. (2002). A change detection software agent based on immune mixed selection. In *Proceedings of the 2002 Congress on Evolutionary Computation (CEC)*, pages 693–698. IEEE Computer Society, Washington, DC.
- O’Grady, R., Groß, R., Christensen, A. L., and Dorigo, M. (2010). Self-assembly strategies in a group of autonomous mobile robots. *Autonomous Robots*, 28:439–455.
- Parker, C., Zhang, H., and Kube, C. (2003). Blind bulldozing: multiple robot nest construction. In *Proceedings of the 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 2, pages 2010–2015. IEEE Press, Piscataway, NJ.
- Parker, L. (2000). Current state of the art in distributed autonomous mobile robotics. In *Proceedings of the 5th International Symposium on Distributed Autonomous Robotic Systems (DARS)*, pages 3–12. Springer, Berlin, Germany.
- Payton, D., Estkowski, R., and Howard, M. (2003). Compound behaviors in pheromone robotics. *Robotics and Autonomous Systems*, 44(34):229 – 240.
- Sahin, E. (2005). Swarm Robotics: From Sources of Inspiration to Domains of Application. In *Swarm Robotics*, volume 3342, pages 10–20. Springer, Berlin, Germany.
- Sakaguchi, S. (2004). Naturally arising CD4⁺ regulatory T cells for immunologic self-tolerance and negative control of immune responses. *Annual Review of Immunology*, 22:531–562.
- Tarapore, D., Floreano, D., and Keller, L. (2010). Task-dependent influence of genetic architecture and mating frequency on division of labour in social insect societies. *Behavioral Ecology and Sociobiology*, 64(4):675–684.
- Waibel, M., Keller, L., and Floreano, D. (2009). Genetic Team Composition and Level of Selection in the Evolution of Cooperation. *IEEE Transactions on Evolutionary Computation*, 13(3):648–660.
- Yamins, D. and Nagpal, R. (2008). Automated global-to-local programming in 1-D spatial multi-agent systems. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 615–622. IFAAMAS, Toronto, Canada.

Appendix A: Equations for activated T-cells

The section details the equations to calculate density of activated effector E_i^* and regulatory R_i^* cells, for all T-cell clones i . Given the conjugate density $C_{ij}(t)$ at steady state, the density of conjugated effector and regulatory cells is calculated proportional to the relative frequency of T_E and T_R cells in the clone. For the conjugated effector Ec_{ij} and regulatory Rc_{ij} cells of clone i at APC population j , we have:

$$Ec_{ij}(t) = \frac{C_{ij}(t)E_i(t)}{T_i(t)} \quad \text{and} \quad Rc_{ij}(t) = \frac{C_{ij}(t)R_i(t)}{T_i(t)}$$

Finally, for the number of activated effector E_i^* and regulatory R_i^* cells, we have:

$$E_i^* = \sum_{j=1}^M P_e(A_j, Ec_i, Rc_i) Ec_{ij} \quad (4)$$

$$R_i^* = \sum_{j=1}^M P_r(A_j, Ec_i, Rc_i) Rc_{ij} \quad (5)$$

where function P_e is the probability that an effector cell is conjugated with no neighboring regulatory cell at the same APC. P_r is the probability that a regulatory cell is conjugated with an APC that has at least one effector cell conjugated simultaneously. Additionally, Ec_i and Rc_i are the total number of conjugated effector and regulatory cells of clone i :

$$Ec_i = \sum_{j=1}^M Ec_{ij} \quad \text{and} \quad Rc_i = \sum_{j=1}^M Rc_{ij}$$

The probability functions P_e and P_r can be reduced to the following expressions, based on a multinomial approximation (Evans et al., 2000) that is valid given that the total number of sites (summed over all the APCs) is much larger than the number of sites per APC. For 3 binding sites ($s = 3$) on each APC, we have:

$$P_e(A_j, Ec_i, Rc_i) = \frac{(Rc_i - 3A_j)^2}{9A_j^2} \quad (6)$$

$$P_r(A_j, Ec_i, Rc_i) = \frac{(6A_j - Ec_i)Ec_i}{9A_j^2} \quad (7)$$

Utilizing the probability functions P_e and P_r , the density of effector and regulator cells can be calculated (eq 4 and 5).