

Self-configuration in Autonomic Electronic Institutions

Eva Bou¹ and Maite López-Sánchez² and J. A. Rodríguez-Aguilar¹

Abstract. Electronic institutions (EIs) have been proposed as a means of regulating open agent societies. EIs define the rules of the game in agent societies by fixing what agents are permitted and forbidden to do and under what circumstances. And yet, there is the need for EIs to adapt their regulations to comply with their goals despite coping with varying populations of self-interested agents. In this paper we focus on the extension of EIs with autonomic capabilities to allow them to yield a dynamical answer to changing circumstances through the adaptation of their norms.

1 Introduction

The growing complexity of advanced information systems in the recent years, characterized by being distributed, open and dynamical, has given rise to interest in the development of systems capable of self-management. Such systems are known as self-* systems [14], where the * sign indicates a variety of properties: self-organization, self-configuration, self-diagnosis, self-repair, etc. A particular approximation to the construction of self-* systems is represented by the vision of autonomic computing [11], which constitutes an approximation to computing systems with a minimal human interference. Some of the many characteristics of autonomic systems are: it must configure and reconfigure itself automatically under changing (and unpredictable) conditions; it must aim at optimizing its inner workings, monitoring its components and adjusting its processings in order to achieve its goals; it must be able to diagnose the causes of its eventual malfunctions and reparate itself; it must act in accordance to and operate into a heterogeneous and open environment.

In what follows we argue that are EIs [5] a particular type of self-* system. When looking at computer-mediated interactions we regard Electronic Institutions (EI) as regulated virtual environments wherein the relevant interactions among participating agents take place. EIs have proved to be valuable to develop open agent systems [10]. However, the challenges of building open systems are still considerable, not only because of the inherent complexity involved in having adequate inter-operation of heterogeneous agents, but also because the need for adapting regulations to comply with institutional goals despite varying agents' behaviors. Particularly, when dealing with self-interested agents.

The main goal of this work consists in studying how to endow an EI with autonomic capabilities that allow it to yield a dynamical answer to changing circumstances through the adaptation of its regulations. Among all the characteristics that define an autonomic system we will focus on the study of self-configuration as pointed out in [11] as a second characteristic: "An autonomic computing system must configure and reconfigure itself under varying (and in the future, even unpredictable) conditions. System configuration or "setup" must occur automatically, as well as dynamic adjustments to that configuration to best handle changing environments".

The paper is organized as follows. In section 2 we introduce the notion of autonomic electronic institution as an extension of the classic notion of electronic institution along with a general model for norm adaptation. Section 3 details a case study to be employed as a scenario wherein to test the model presented in section 2. Section 4 provides some preliminary, empirical results. Finally, section 5 summarizes some conclusions and outlines paths to future research.

2 Autonomic Electronic Institutions

The idea behind EIs [17] is to mirror the role traditional institutions play in the establishment of "the rules of the game"—a set of conventions that articulate agents' interactions— but in our case applied to agents (be them human or software) that interact through messages whose (socially relevant) effects are known to interacting parties. The essential roles EIs play are both descriptive and prescriptive: the institution makes the conventions explicit to participants, and it warrants their compliance. EIs involve a conceptual framework to describe agent interactions as well as an engineering framework [1] to specify and deploy actual interaction environments.

Although EIs can be regarded as the computational counterpart of human institutions for open agent systems, there are several aspects in which they are nowadays lacking. According to North [18] human institutions are not static; they may evolve over time by altering, eliminating or incorporating norms. In this way, institutions can adapt to societal changes. Nonetheless, neither the current notion of EI in [5] nor the engineering framework in [1] support norm adaptation so that an EI can self-configure. Thus, in what follows we study how to extend the current notion of EI in [5] to support self-configuration.

First of all, notice that in order for norms to adapt, we believe that a "rational" view of EIs must be adopted (likewise the rational view of organizations in [6]) and thus consider

¹ IIIA - CSIC Artificial Intelligence Research Institute, Campus UAB 08193 Bellaterra, Spain, email: {ebm, jar}@iia.csic.es

² WAI, Volume Visualization and Artificial Intelligence, MAiA Dept., Universitat de Barcelona, email: maite@maia.ub.es

that *EIs seek specific goals*. Hence, EIs continuously adapt their norms to fulfill their goals. Furthermore, we assume that an EI is *situated* in some environment that may be either totally or partially observable by the EI and its participating agents.

With this in mind, we observe that according to [5] an EI is solely composed of: a dialogic framework establishing the common language and ontology to be employed by participating agents; a performative structure defining its activities along with their relationships; and a set of norms defining the consequences of agents' actions. From this follows that further elements are required in order to incorporate the fundamental notions of goal and norm transition as captured by the following definition of *autonomic electronic institution*.

Definition 2.1 *Given a finite set of agents A , we define an Autonomic Electronic Institution (AEI) as a tuple $\langle PS, N, DF, G, P_i, P_e, V, \delta \rangle$ where:*

- PS stands for a performative structure;
- N stands for a finite set of norms;
- DF stands for a dialogic framework;
- G stands for a finite set of institutional goals;
- $P_i = \langle i_1, \dots, i_s \rangle$ stands for the values of a finite set of institutional properties, where $i_j \in \mathbb{R}$, $1 \leq j \leq s$ contains the value of the j -th property;
- $P_e = \langle e_1, \dots, e_r \rangle$ stands for the values of the environment properties, where each e_j is a vector, $e_j \in \mathbb{R}^{n_j}$ $1 \leq j \leq r$ contains the value of the j -th property;
- $P_a = \langle a_1, \dots, a_n \rangle$ stands for the values that characterize the institutional state of the agents in A , where $a_j = \langle a_{j1}, \dots, a_{jm} \rangle$ $1 \leq j \leq n$ stands for the institutional state of agent A_j ;
- V stands for a finite set of reference values; and
- $\delta : PS \times N \times G \times V \rightarrow PS \times N$ stands for a normative transition function that maps a performative structure and a set of norms into a new performative structure and a new set of norms given a set of goals and a set of values for the reference variables.

Notice that a major challenge in the design of an AEI is to learn a *normative transition function*, δ , that ensures the achievement of its institutional goals under changing conditions. Next, we dissect the new elements composing an AEI.

An AEI employs norms to constrain agents' behaviors and to assess the consequences of their actions within the scope of the institution. Although there is a plethora of formalizations of the notion of norm in the literature, in this paper we adhere to a simple definition of norms as effect propositions as defined in [9]:

Definition 2.2 *An effect proposition is an expression of the form*

$$A \text{ causes } F \text{ if } P_1, \dots, P_n$$

where A is an action name, and each of F, P_1, \dots, P_n ($n \geq 0$) is a fluent expression. About this proposition we say that it describes the effect of A on F , and that P_1, \dots, P_n are its preconditions. If $n = 0$, we will drop it and write simply A causes F . From this definition, changing a norm amounts to changing either its pre-conditions, or its effect(s), or both.

Agents participating in an AEI have their social interactions mediated by the institution according to its norms. As a

consequence of his interactions, only the *institutional (social) state* of an agent can change since an AEI has no access whatsoever to the inner state of any participating agent. Therefore, given a finite set of participating agents $A = \{A_1, \dots, A_n\}$ where $n \in \mathbb{N}$, each agent $A_i \in A$ can be fully characterized by his institutional state, represented as a tuple of observable values $\langle a_{i1}, \dots, a_{im} \rangle$ where $a_{ij} \in \mathbb{R}$ $1 \leq j \leq m$. Thus, the actions of an agent within an AEI may change his institutional state according to the institutional norms.

The main objective of an AEI is to accomplish its goals. For this purpose, an AEI will adapt its norms. We assume that the institution can observe the environment, the institutional state of the agents participating in the institution, and its own state to assess whether its goals are accomplished or not. Thus, from the observation of environment properties (P_e), institutional properties (P_i), and agents' institutional properties (P_a), an AEI obtains the reference values required to determine the fulfillment of goals. Formally, the reference values are defined as a vector $V = \langle v_1, \dots, v_q \rangle$ where each v_j results from applying a function h_j upon the agents' properties, the environmental properties and/or the institutional properties; $v_j = h_j(P_a, P_e, P_i)$, $1 \leq j \leq q$.

Finally, we can turn our attention to institutional goals. An example of institutional goal for the Traffic Regulation Authority could be to keep the number of accidents below a given threshold. In other words, to ensure that a reference values satisfies some constraint.

Formally we define the goals of an AEI as a finite set of constraints $G = \{c_1, \dots, c_p\}$ where each c_i is defined as an expression $g_i(V) \triangleleft [m, M]$ where $m, M \in \mathbb{R}$, \triangleleft stands for either \in or \notin , and g_i is a function over the reference values.

2.1 Norm Adaptation

A major challenge in the design of an AEI is to learn a *normative transition function* that allows to accomplish institutional goals under changing situations. In this work, we concentrate on norm adaptation and therefore we consider that there is no definition of performative structure. Thus, institutional goals must be accomplished through norms, which will be the only means of regulating agents' actions. We are considering the *normative transition function* defined in 2.1 in a more simple way, $\delta : N \times G \times V \rightarrow N$, namely as a normative transition function that maps a set of norms into a new set of norms.

From the definition 2.2 of norm, changing a norm amounts to changing either its pre-conditions, or its effects, or both. Norms can be parameterized, and therefore we propose that each norm, $N_i \in N$, has a set of parameters $\langle p_{i1}, \dots, p_{im} \rangle \in \mathbb{R}^{im}$. Notice that when the parameters of the norms are associated to the pre-conditions and/or to the effects, changing the values of these parameters means changing the norm. When we refer to change the norms or to adapt the norms we are referring to change or to adapt the values of the parameters of the norms. Norms have associated parameters that can be changed to increase its persuasiveness depending on the agent population behavior. We propose to learn the *normative transition function* by exploring the space of parameter values in search for the ones that best accomplish goals for a given population of agents. In this manner, if we can automatically adapt norms to the global behavior of an agent population, then, we can repeat it for a number of different agent popu-

lations and thus characterize the overall *normative transition function*.

Figure 1 describes how this learning process is performed for a given population of agents (A) using an evolutionary approach. We have an initial set of individuals ((I_1, \dots, I_n)), where each individual represents a set of norm parameters ($\{\langle p_{11}, \dots, p_{1m} \rangle, \dots, \langle p_{i1}, \dots, p_{im} \rangle\}$). The institution performs a simulation for each individual with the population of agents A , so that the norms represented by each individual can be evaluated according to the institutional goals (*Norm evaluation*). Finally, the AEI compiles the evaluations of all individuals in order to perform the *Norm adaptation* process that results with a new set of individuals (*New norms*) to be used as an initial set of individuals for next step in the learning process.

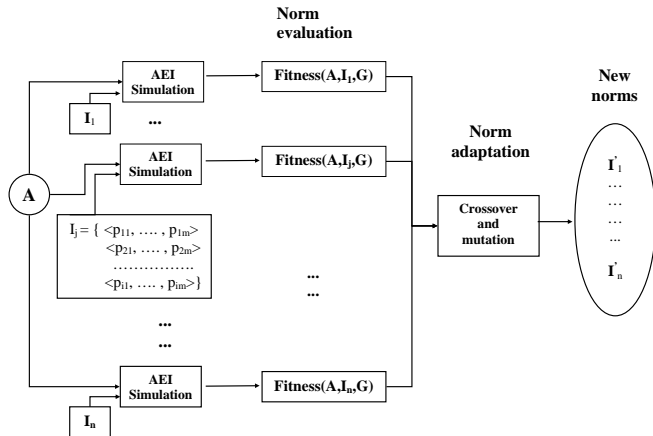


Figure 1. Example of an step in norm adaptation using an evolutionary approach.

3 Case Study: Traffic Control

Traffic control is a well-known problem that has been approached from different perspectives, which range from macro simulation for road net design [20] to traffic flow improvement by means of multi-agent systems [15]. We tackle this problem from the Electronic Institutions point of view, and therefore, this section is devoted to specify how traffic control can be mapped into Autonomic Electronic Institutions.

In this manner, we consider the Traffic Regulation Authority as an Autonomic Electronic Institution, and cars moving along the road network as agents inside the institution. Considering this set-up, traffic norms regulated by Traffic Authorities can therefore be translated in a straight forward manner into norms belonging to the Electronic Institution. Norms within this normative environment are thus related to actions performed by cars (in fact, in our case, they are always restricted to that). Additionally, norms do have associated penalties that are imposed to those cars refusing or failing to follow them. In our case study, we assume that the Traffic Authority is always aware of norm violations: cars may or may not respect rules, but they are not able to avoid the consequences of their application. Furthermore, our Electronic

Institution is able to change norms based on its goals – just as traffic authorities do modify their traffic rules– and, therefore, it is considered to be autonomic.

Our AEI sets up a normative environment where cars do have a limited amount of credit (just as some real world driving license credit systems) so that norm offenses cause credit reductions. The number of points subtracted for each traffic norm violation is specified by the sanction associated to each norm, and this sanction can be changed by the regulation authority (that is, our AEI) if its change leads –or contributes to– the accomplishment of goals. Eventually, those cars without any remaining points are forbidden to circulate. On the other hand, we assume a non-closed world, so expelled cars are replaced by new ones having the total amount of points.




Figure 2. Grid environment representation of a 2-lane road junction.

Getting into more detail, we focus on a two-road junction. It is a very restrictive problem setting, but it is complex enough to allow us to tackle the problem without losing control of all the factors that may influence the results. In particular, no traffic signals (neither yield or stop signals nor traffic lights) are considered, therefore, cars must only coordinate by following the traffic norms imposed by the AEI. Our institution is required to define these traffic norms based on general goals such as minimization of the number of accidents or deadlock avoidance.

We model the environment as a grid composed by road and field cells. Road cells define 2 orthogonal roads that intersect in the center (see figure 2). Discretization granularity is such that cars have the size of a cell. As section 3.2 details, our model has been developed with the Simma tool [13]. Although the number of road lanes can be changed parametrically, henceforth we assume the 2-lane case. Next subsections are devoted to define this “toy problem” and present our solution proposal in terms of it. But before that, we introduce some nomenclature definitions:

- A_i : an agent i , agents correspond to cars.
- t : time step. Our model considers discrete time steps (ticks).
- (J_x, J_y) : size in x, y of our road junction area.
- J : inner road junction area with (x_0^J, y_0^J) as top left cell inside it
 $J = \{(x, y) \mid x \in [x_0^J, x_0^J + J_x - 1], y \in [y_0^J, y_0^J + J_y - 1]\}$
 Considering the 4 J cells in the junction area of Figure 2:
 $J = \{(x_0^J, y_0^J), (x_0^J + 1, y_0^J), (x_0^J, y_0^J + 1), (x_0^J + 1, y_0^J + 1)\}$.
- J_{BE} : Junction Boundary Entrance, set of cells surrounding the junction that can be used by cars to access it. They correspond to cells near by the junction that belong to incoming lanes. Figure 2 depicts $J_{BE} = \{(x_0^J, y_0^J - 1), (x_0^J -$

$1, y_0^J + J_y - 1), (x_0^J + J_x - 1, y_0^J + J_y, (x_0^J + J_x, y_0^J))\}$.

Nevertheless, the concept of boundary is not restricted to adjacent cells: a car can be also considered to be coming into the junction if it is located one –or even a few– cells away from the junction.

- (x_i^t, y_i^t) : position of car A_i at time t , where $(x, y) \in \mathbb{N} \times \mathbb{N}$ stands for a cell in the grid.
- (h_{ix}^t, h_{iy}^t) : heading of car A_i , which is located in (x, y) at time t . Heading directions run along x, y axes and are considered to be positive when the car moves right or down respectively. In our orthogonal environment, heading values are: 1 if moving right or down; -1 if left or up; and 0 otherwise (i.e., the car is not driving in the axe direction). In this manner, car4’s heading on the right road of figure 3 is $(-1, 0)$.

3.1 AEI specification

3.1.1 Environment

As mentioned above, we consider the environment to be a grid. This grid is composed of cells, which can represent roads or fields. The main difference among these two types is that road cells can contain cars. Indeed, cars move among road cells along time.

Figure 2 depicts a 8×8 grid example. The top left corner of the grid represents the origin in the x, y axes. Thus, in the example, cell positions range from $(0, 0)$ in the origin up to $(7, 7)$ at the bottom-right corner. Additionally, a cell is a road if one of its x, y coordinates belong to J inner junction area (see previous definition).

We define this grid environment as:

$$P_e = \langle (x, y, \alpha, r, d_x, d_y) \mid 0 \leq x \leq \max_x, 0 \leq y \leq \max_y, \alpha \subseteq P(A) \cup \emptyset, r \in [0, 1], d_x \in [-1, 0, 1], d_y \in [-1, 0, 1] \rangle$$

being x and y the cell position, α defines the set of agents inside the grid cell (x, y) (notice that $\alpha \subseteq A$), r indicates whether this cell represents a road or not, and, in case it is a road, d_x and d_y stand for the lane direction, whose values are the same as the ones for car headings. Notice that the institution can observe the environment properties along time, we use P_e^t to refer the values of the grid environment at a specific time t . This discretized environment can be observed both by the institution and cars. The institution observes and keeps track of its evolution along time, whilst cars do have locality restrictions on their observations.

3.1.2 Agents

We consider $A = \langle A_1, \dots, A_n \rangle$ to be a finite set of n agents in the institution. As mentioned before, agents correspond to cars that move inside the grid environment, with the restriction that they can only move within road cells. Additionally, agents are given an account of points which decreases with traffic offenses. The institution forbids agents to drive without points in their accounts.

The institution can observe the $P_a = \langle a_1, \dots, a_n \rangle$ agents’ institutional properties, where

$$a_i = \langle x_i, y_i, h_{ix}, h_{iy}, speed_i, indicator_i, offenses_i, accidents_i, distance_i, points_i \rangle$$

These properties stand for: car A_i ’s position within the grid, its heading, its speed, whether the car is indicating a trajectory change for the next time step (that is, if it has the intention to turn, to stop or to move backwards), the norms being currently violated by A_i , whether the car is involved in an accident, the distance between the car and the car ahead of it; and, finally, agent A_i ’s point account. Notice that the institution can observe the agent properties along time, we use a_i^t to refer the agent A_i ’s properties at a specific time t .

3.1.3 Reference values

In addition to car properties, the institution is able to extract reference values from the observable properties of the environment, the participating agents and the institution. Thus, these reference values are computed as a compound of other observed values. Considering our road junction case study, we identify different reference values:

$$V = \langle num_collisions, num_crashed, median_offenses, dispersion_offenses, num_blocked \rangle$$

where **num_collisions** indicates total number of collisions for last t_w ticks ($0 \leq t_w \leq t_{now}$):

$$num_collisions = \sum_{t=t_{now}-t_w}^{t_{now}} \sum_{e \in P_e^t} f(e_{\alpha^t})$$

being P_e^t the values of the grid environment at time t , e_{α^t} the α^t component of element $e \in P_e^t$ and

$$f(e_{\alpha^t}) = \begin{cases} 1 & \text{if } |e_{\alpha^t}| > 1 \\ 0 & \text{otherwise} \end{cases}$$

Furthermore, **num_crashed** counts the number of cars involved in accidents for last t_w ticks, **median_offenses** indicates the median of the offenses accumulated by all agents during last t_w ticks, similarly, **dispersion_offenses** computes the dispersion of the offenses, and finally, **num_blocked** shows how many cars have been blocked by other cars for last t_w ticks.

3.1.4 Goals

Goals are institutional goals. The aim of the traffic authority institution is to accomplish as many goals as possible. The institution tries to accomplish these goals by defining a set of norms (see subsection 3.1.5).

Institutional goals are defined as constraints upon a combination of reference values. Considering our scenario, we define restrictions as intervals of acceptable values for the previous defined reference values (V) so that we consider the institution accomplishes its goals if V values are within their corresponding intervals. In fact, the aim is to minimize the number of accidents, the number of traffic offenses, as well as the number of blocked cars by establishing the list of institutional goals G as:

$$G = \langle num_collisions \in [0, MaxCollisions], num_crashed \in [0, MaxCrashed], Median_Off \in [0, MaxOffences], Disp_Off \in [0, MaxDispOffences], num_blocked \in [0, MaxBlocked] \rangle$$

3.1.5 Norms

Autonomic Electronic Institutions try to accomplish goals by defining norms. Norms have associated penalties that are imposed to those cars refusing or failing to follow them. These penalties can be parameterized to increase its persuasiveness depending on the agent population behavior.

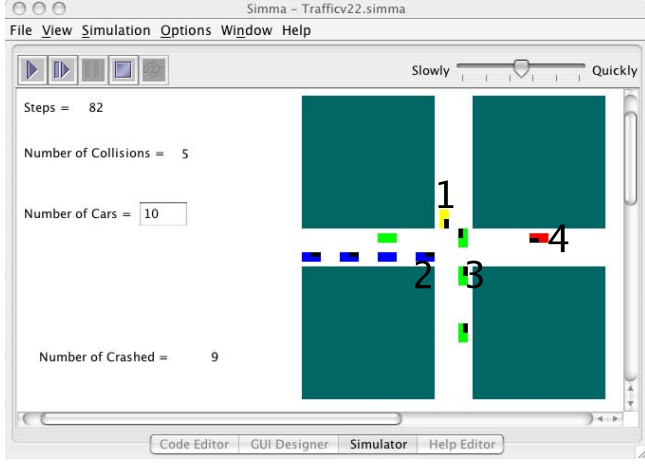


Figure 3. Priority to give way to the right (Simma tool screenshot).

Considering a road junction without traffic signals, priorities become basic to avoid collisions. We consider, as in most continental Europe, that the default priority is to give way to the right. This norm prevents a car A_i located on the Junction Boundary Entrance (J_{BE}) to move forward or to turn left whenever there is another car A_j on its right. For example, car 1 in figure 3 must wait for car 2 on its right, which must also wait for car 3 at the bottom J_{BE} . The formalization in table 1 can be read as follows: “if car A_i moves from a position in J_{BE} at time $t - 1$ to its next heading position at time t without indicating a right turn, and if it performs this action when having a car A_j at the J_{BE} on its right, then the institution will fine A_i by decreasing its points by a certain amount” (see figure 4).

Table 1. Right priority norm.

Action	$in(a_i, J_{BE}, t - 1) \wedge in(a_i, (x_i^{t-1} + h_{ix}^{t-1}, y_i^{t-1} + h_{iy}^{t-1}), t) \wedge \neg indicator(a_i, right, t - 1)$
Pre-conditions	$right(a_i, a_j, t - 1)$
Consequence	$points_i^t = points_i^{t-1} - fine$

Where the predicate $\mathbf{in}(a_i, Region, t)$ in table 1 is equivalent to $\exists(x, y, \alpha^t, r, d_x, d_y) \in E^t$ so that $(x, y) \in$




Figure 4. Priority to give way to the right.

Region and $a_i \in \alpha^t$ and $\mathbf{right}(a_i, a_j, t)$ is a boolean function that returns true if car a_j is located at J_{BE} area on the right side of car a_i . For the 2-lane J_{BE} case in Figure 2, it corresponds to the formula: $(x_i^t - h_{iy}^t + h_{ix}^t J_x, y_i^t + h_{ix}^t + h_{iy}^t J_y) = (x_j^t, y_j^t)$.

Other norms, such as deadlock avoidance or junction blocking prevention have been considered and implemented. Nevertheless, due to the lack of space, we cannot detail them.

3.2 Experimental Settings and Design

As a proof of concept of our proposal in section 2.1, we have designed an experimental setting that implements the traffic case study. In this preliminary experiment we consider a single normative goal ($num_collisions$) and the right priority norm in table 1, which is parameterized by its fine (i.e., points to subtract to the car falling to follow the norm).

The 2-road junction traffic model has been developed with Simma [13], a graphical MAS simulation tool shown in Figure 3, in such way that both environment and agents can be easily changed. In our experimental settings, we have modeled the environment as a 16×16 grid where both crossing roads have 2 lanes with opposite directions. Additionally, the environment is populated with 10 cars, having 40 points each.

Our institution can observe the agents properties for each tick and can keep a record of them in order to refer to past ticks. In fact, the institution usually determines traffic offenses by analyzing agent actions along time. Agent actions are observed through consecutive car positions and indicators (notice that the usage of indicators is compulsory for cars in this problem set up). During our discrete event simulation, the institution replaces those cars running out of points by new cars, so that the cars' population is kept constant. Cars follow random trajectories at a constant 1-cell/tick speed and they collision if two or more cars run into the same cell. In that case, the involved cars do remain for two ticks in that cell before they can start following a new trajectory.

Cars correspond to agents without learning skills. They just move based on their trajectories and institutional norms. Agents have local information about their environment (i.e., grid surrounding cells) and know whether their next moves will violate a norm and what fine will be thus applied. Agents decide whether to comply with a norm based on three parameters: $\langle fulfill_prob, high_punishment, inc_prob \rangle$. Being $fulfill_prob \in [0, 1]$ the probability of complying with norms that is initially assigned to each agent, $high_punishment \in \mathbb{N}$ the fine threshold that causes an agent to consider a fine to be high enough to reconsider the norm compliance, and

$inc_prob \in [0,1]$ the probability increment that is added to $fulfill_prob$ when the fine threshold is surpassed by the norm being violated. In summary, agents decide whether they keep moving regardless of violated norms or they stop in order to comply with norms based on a probability that is computed as: $final_prob = fulfill_prob + inc_prob$ when $fine > high_punishment$.

Our goal is to adapt norms to agent behaviors by applying Genetic Algorithms (GA)³ to accomplish the institutional goal, to minimize the total number of collisions. We propose learn the norms by different agent populations behavior by simulation. Once specified what are the different agent populations behavior, a genetic algorithm is running by each population of agents. We use 10 individuals in each step of the genetic algorithm, where each individual is a set of parameters. Therefore, norm adaptation is implemented as a learning process of the “best” norm parameters. To evaluate an individual we run 10 times the simulator with the set of parameters of the individual. The simulator run the AEI model explained above during 5000 ticks. Thus, norm quality is given by a fitness function that considers the number of collisions, which is computed as an average of 10 different 5000-tick-long simulations for each model setting.

4 Results

From the experimental settings specified above, we have run experiments for three different agent populations. These populations are characterized by their norm compliance parameters, being $fulfill_prob = 0.5$ and $inc_prob = 0.4$ for the three of them whereas $high_punishment$ varies from 5 for the first, to 10 for the second, up to 14 for the third (see table 2).

Since the right priority norm contributes to reduce accidents, our AEI must learn how to vary its fine parameter to increase its persuasiveness for agents, and eventually, to accomplish the normative goal of minimizing the total number of collisions. As to shows table 2, our experiments have resulted in that our AEI learned a fine of 14, 12, and 15 for each respective population. In all three cases, the learned fine is larger than the population’s $high_punishment$ value, and therefore, the goal is successfully reached⁴. In this manner, we can state the AEI success in learning the norms that better accomplish its goal.

Next figure 5 gives some more detail about the performance of agent populations for different norm fine values. First chart compares the number of collisions per 100 ticks when the fine is 4 with the resulting number of collisions when it is 14, which is the learnt value for agents with a $high_punishment$ threshold equal to 5. Analogously, second and third charts compare results between value 4 and learnt values 12 and 15 respectively (which, again, are learnt when the corresponding agent populations have 10 and 14 threshold values). For all three cases, we can observe that the number of collisions for fine 4 keep above the ones for learnt fines. It is so both in average and along the curve that results from a simulation of 5000 ticks. As expected, the reason is that value 4 is smaller than the $high_punishment$ values for all three agent populations.

³ We use GALib [19], a C++ library of genetic algorithm components

⁴ Notice that, due to the agent’s behavior, any fine value higher than the population $high_punishment$ value will be equally successful.

Additionally, we can also observe that the deviation in the number of collisions is smaller as well.

Table 2. Learning results for three different agent populations.

Parameters	population1	population2	population3
$fulfill_prob$	0.5	0.5	0.5
$high_punishment$	5	10	14
inc_prob	0.4	0.4	0.4
Learned fine	14	12	15

5 Discussion and Future work

Within the area of Multi-Agent Systems, adaptation has been usually envisioned as an agent capability. In this manner, works such as the one by Excelente-Toledo and Jennings [7] propose a decision making framework that enables agents to dynamically select the coordination mechanism that is most appropriate to their circumstances. On the other hand, it has been long stated [2] that agents working in a common society need norms to avoid and solve conflicts, make agreements, reduce complexity, or to achieve a social order. Both approaches –i.e. adaptation and norms– have been considered together by Lopez-y-Lopez et al. [12], where agents can adapt to norm-based systems and they can even autonomously decide its commitment to obey norms in order to achieve associated institutional goals. This adaptation from the point of view of agents in these related works is the most remarkable difference with the approach presented in this paper, which focuses on adapting the institution –that is, the authority issuing norms– rather than adapting the agents. Institution adaptation is accomplished by changing norms autonomously. Therefore, we do not select norms at design stages as it is done by Fitoussi and Tennenholtz [8], who do it so by proposing the notions of minimality and simplicity as selecting criteria. They study two basic settings, which include Automated-Guided-Vehicles (AGV) with traffic laws, by assuming an environment that consists of (two) agents and a set of strategies available to (each of) them. From this set, agents devise the appropriate ones in order to reach their assigned goals without violating social laws, which must be respected.

Regarding the traffic domain, MAS has been previously applied to it [15] [4] [3]. But traffic has been also widely studied outside the scope of MAS, for example, the preliminary work by [16] used Strongly Typed Genetic Programming (STGP) to control the timings of traffic signals within a network of orthogonal intersections. Their evaluation function computed the overall delay.

Preliminary results in this paper provide soundness to our AEI approach. We plan to perform the same experiments with other norms and with more goals. Nevertheless, we plan to extend both our traffic model and the institutional adaptation capabilities so that the AEI will not only learn the most appropriate norms for a given agent population, but it will be able to adapt to any change in the population.

ACKNOWLEDGEMENTS

This work was partially funded by the Spanish Science and Technology Ministry as part of the Web-i-2 project (TIC-

2003-08763-C02-01). The first author enjoys an FPI grant (BES-2004-4335) from the Spanish Education and Science Ministry.

REFERENCES

- [1] Josep L. Arcos, Marc Esteva, Pablo Noriega, Juan A. Rodríguez-Aguilar, and Carles Sierra, 'Engineering open environments with electronic institutions', *Engineering Applications of Artificial Intelligence*, (18), 191–204, (2005).
- [2] R. Conte, R. Falcone, and G. Sartor, 'Agents and norms: How to fill the gap?', *Artificial Intelligence and Law*, (7), 1–15, (1999).
- [3] Arnaud Doniec, Stéphane Espié, René Mandiau, and Sylvain Piechowiak, 'Dealing with multi-agent coordination by anticipation: Application to the traffic simulation at junctions.', in *EUMAS*, pp. 478–479, (2005).
- [4] Kurt Dresner and Peter Stone, 'Multiagent traffic management: An improved intersection control mechanism', in *The Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, pp. 471–477, New York, NY, (July 2005). ACM Press.
- [5] Marc Esteva, *Electronic Institutions: from specification to development*, IIIA PhD Monography, Vol. 19, 2003.
- [6] A. Etzioni, *Modern Organizations*, Prentice-Hall, Englewood Cliffs, NJ, 1964.
- [7] Cora Beatriz Excelente-Toledo and Nicholas R. Jennings, 'The dynamic selection of coordination mechanisms', *Autonomous Agents and Multi-Agent Systems*, 9(1-2), 55–85, (2004).
- [8] David Fitoussi and Moshe Tennenholtz, 'Choosing social laws for multi-agent systems: Minimality and simplicity', *Artificial Intelligence*, 119(1-2), 61–101, (2000).
- [9] Michael Gelfond and Vladimir Lifschitz, 'Representing action and change by logic programs', *Journal of Logic Programming*, 17, 301–321, (1993).
- [10] Nicholas R. Jennings, Katia Sycara, and Michael Wooldridge, 'A roadmap of agent research and development', *Autonomous Agents and Multi-agent Systems*, 1, 275–306, (1998).
- [11] J. O. Kephart and D. M. Chess, 'The vision of autonomic computing', *IEEE Computer*, 36(1), 41–50, (2003).
- [12] Fabiola López-López, Michael Luck, and Mark d'Inverno, 'Constraining autonomy through norms', in *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pp. 674–681, New York, NY, USA, (2002). ACM Press.
- [13] M. López-Sánchez, X. Noria, J. A. Rodríguez-Aguilar, and N. Gilbert, 'Multi-agent based simulation of news digital markets', *International Journal of Computer Science and Applications*, 2(1), 7–14, (2005).
- [14] Michael Luck, Peter McBurney, Onn Shehory, and Steve Willmott, *Agentlink Roadmap*, Agenlink.org, 2005.
- [15] Sean Luke, Claudio Cioffi-Revilla, Liviu Panait, and Keith Sullivan, 'Mason: A new multi-agent simulation toolkit', in *Proceedings of the 2004 SwarmFest Workshop*, p. 8, (2004).
- [16] David J. Montana and Steven Czerwinski, 'Evolving control laws for a network of traffic signals', in *Genetic Programming 1996: Proceedings of the First Annual Conference*, pp. 333–338, Stanford University, CA, USA, (28–31 July 1996). MIT Press.
- [17] Pablo Noriega, *Agent-Mediated Auctions: The Fishmarket Metaphor*, IIIA PhD Monography, Vol. 8, 1997.
- [18] D. North, *Institutions, Institutional Change and Economics Performance*, Cambridge U. P., 1990.
- [19] Matthew Wall, *Galib, A C++ Library of Genetic Algorithm Components*, Massachusetts Institute of Technology (MIT). <http://lancet.mit.edu/ga/>.
- [20] Qi Yang, *A Simulation Laboratory for Evaluation of Dynamic Traffic Management Systems*, Ph.D. dissertation, MIT, 1997.

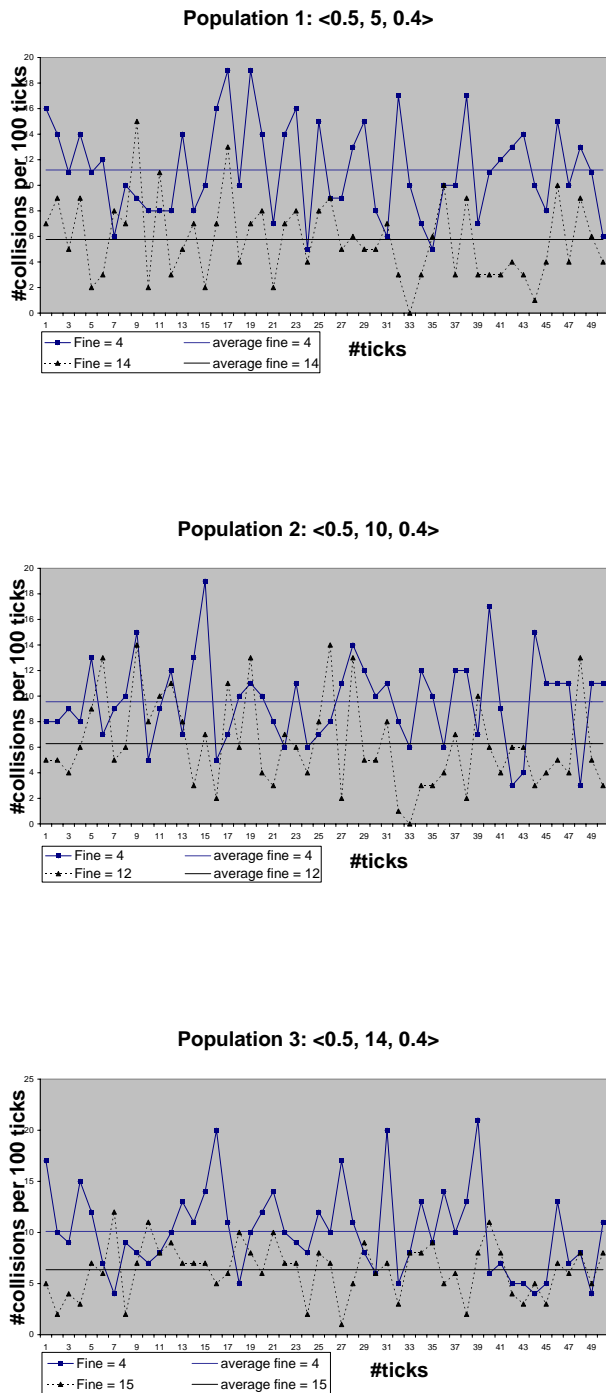


Figure 5. Number of collisions per 100 ticks along a 5000-tick simulation.