



ELSEVIER

Available at  
[www.ComputerScienceWeb.com](http://www.ComputerScienceWeb.com)  
POWERED BY SCIENCE @ DIRECT®

Information Sciences 150 (2003) 33–57

INFORMATION  
SCIENCES  
AN INTERNATIONAL JOURNAL

[www.elsevier.com/locate/ins](http://www.elsevier.com/locate/ins)

# A hierarchical fuzzy–genetic multi-agent architecture for intelligent buildings online learning, adaptation and control

Hani Hagra<sup>s</sup> \*, Victor Callaghan, Martin Colley,  
Graham Clarke

*Department of Computer Science, University of Essex, Wivenhoe Park, Colchester CO43SQ, UK*

Received 5 August 2000; accepted 7 April 2001

---

## Abstract

In this paper, we describe a new application domain for intelligent autonomous systems—intelligent buildings (IB). In doing so we present a novel approach to the implementation of IB agents based on a hierarchical fuzzy genetic multi-embedded-agent architecture comprising a low-level behaviour based reactive layer whose outputs are co-ordinated in a fuzzy way according to deliberative plans. The fuzzy rules related to the room resident comfort are learnt and adapted online using our patented fuzzy–genetic techniques (British patent 99-10539.7). The learnt rule base is updated and adapted via an iterative machine-user dialogue. This learning starts from the best stored rule set in the agent memory (Experience Bank) thereby decreasing the learning time and creating an intelligent agent with memory. We discuss the role of learning in building control systems, and we explain the importance of acquiring information from sensors, rather than relying on pre-programmed models, to determine user needs. We describe how our architecture, consisting of distributed embedded agents, utilises sensory information to learn to perform tasks related to user comfort, energy conservation, and safety. We show how these agents, employing a behaviour-based approach derived from robotics research, are able to continuously learn and adapt to individuals within a building, whilst always providing a fast, safe response to any situation. In addition we show that our system learns similar rules to other offline supervised methods but that our system has the additional capability to rapidly learn and optimise the learnt rule

---

\* Corresponding author.

E-mail address: [hani@essex.ac.uk](mailto:hani@essex.ac.uk) (H. Hagra<sup>s</sup>).

base. Applications of this system include personal support (e.g. increasing independence and quality of life for older people), energy efficiency in commercial buildings or living-area control systems for space vehicles and planetary habitation modules.

© 2002 Elsevier Science Inc. All rights reserved.

*Keywords:* Fuzzy systems; Genetic systems; Intelligent buildings; Agents; Online learning

---

## 1. Introduction

The building industry uses the term intelligent, to describe the way the design, construction and management of a building can ensure that the building is flexible and adaptable, and therefore profitable, over its full life span. A definition which finds favour with many building managers and architects is that “*An intelligent-building is one that provides a productive cost-effective environment through the optimisation of four basic elements; systems, structures, services, management and the inter-relationship between them*” [20].

Computer scientists, however, hold a different view of intelligence. They are more concerned with giving machines management, analytic and control capabilities that are comparable to intelligent human activity. We prefer the definition—“*An intelligent-building is one that utilises computer technology to autonomously govern and adapt the building environment so as to optimise user comfort, energy-consumption, safety and work efficiency*”. In the context of a building, a system works by taking inputs from building sensors (light, temperature, passive infra-red, etc.), and using this and other information to control effectors (heaters, lights, electronically operated windows, etc.). If this system is to be intelligent, an essential feature must be its ability to learn and adapt appropriately. For this a system which can adapt and generate its own rules (rather than being restricted to simple automation) is required. According to Kasabov [16] an intelligent system should be able to learn quickly from large amounts of data therefore using fast training. An intelligent system should also adapt in real-time and in an on-line mode where new data are accommodated as it arrives. Also the system should be able to accommodate in an incremental way any rules that will become known about the problem. It should be memory-based, plus possess data and exemplar storage and retrieval capacities. It also should be able to learn and improve through active interaction with the user and the environment. It should have parameters to represent short and long term memory, age, forgetting, etc. It should also analyse itself in terms of behaviour, error and success. To the authors knowledge no system in the field of intelligent building has satisfied these conditions.

We view intelligent buildings as computer-based systems, akin to robots, gathering information from a variety of sensors, and using embedded intelligent agent techniques to determine appropriate control actions. In controlling such systems one is faced with the imprecision of sensors, lack of adequate models of many of the processes and the non-deterministic aspects of human behaviour. Such problems are well known and there have been various attempts to address them. The most significant of these approaches has been the pioneering work on behaviour-based systems from researchers such as Brooks [4] and Steels [23] who have had considerable success in the field of mobile robots. It might not seem obvious that a building can be looked upon as a machine; indeed “*a robot that we live within*”, but, in other work [6] we have justified the view that intelligent buildings, are analogous to robots, gathering information from a variety of sensors, and able to use behaviour-based techniques to determine appropriate control actions.

Fuzzy logic offers a framework for representing imprecise, uncertain knowledge. Similar to the way, in which people make their decisions, fuzzy systems use a mode of approximate reasoning, which allows it to deal with vague and incomplete information. In addition fuzzy controllers exhibit robustness with regard to noise and variations of system parameters. However, fuzzy systems have a well-known problem concerning the determination of their parameters. In most fuzzy systems, the fuzzy rules are determined and tuned through “trial and error” by developers, taking much iteration to determine and tune them. As the number of input variables increases (which is the case of intelligent buildings) the numbers of rules increase further magnifying the difficulty.

Evolutionary algorithms constitute a class of search and optimisation methods guided by the principles of natural evolution. Genetic algorithms (GA) are optimisation methods inspired by principles of natural evolution and genetics. GA has been successfully applied to solve a variety of difficult theoretical and practical problems by imitating the underlying processes of evolution such as selection, recombination and mutation. Their capability of learning enables a GA to adapt a system to deal with any task [12]. There are numerous reports in the scientific literature on designing fuzzy controllers using GA [2,10,15,17]. However, most work uses simulation to overcome the time overhead problem caused by the large number of iterations needed for a conventional GA to develop a good solution. Thus it is not feasible for a simple GA to learn online and adapt in real-time. The situation is worsened by the fact that most evolutionary computation methods developed so far assume that the solution space is fixed (i.e. the evolution takes place within a pre-defined problem space and not in a dynamically changing and open one), thus preventing them from being used in real-time applications [16].

Our work is concerned with utilising an intelligent *embedded-agent* approach based on a double hierarchical fuzzy-genetic system, similar to the approach

already taken in our previous work in mobile robotics [13,14], to create an integrated and semi-autonomous building control system. There are other research projects concerned with applying AI to buildings most notably [5,7,8,19]. Examples of such work include research in Sweden [8] that utilises multi-agent principles to control an intelligent building. Their primary goal is energy-efficiency, and although their system does adjust the heating and light level to suit individual preferences, these settings must be pre-defined. Their agents are built from traditional AI (i.e. not behaviour based) and their work does not address issues, such as occupant based learning. The system, so far implemented in simulation only, managed to achieve energy savings of 40% over the same building being controlled manually by occupants. A group in Colorado [19] are using a soft computing approach—neural networks—focusing solely on the intelligent control of lighting within a building, by anticipating when particular *zones* (regions in a room) will be occupied or unoccupied. Their system, implemented in a building with a real occupant, also achieved a significant energy reduction, although this was sometimes at the expense of the occupant's comfort. They use a centralised control architecture which differs from our multi-agent approach and which includes no adaptation, as is the case for us. A third group based at the MIT Artificial Intelligence Lab in Massachusetts is working on an intelligent room project. They employ a mix of cameras, microphones and multiplexer to enable people to interface with room-based systems in a natural way using speech, gesture, movements, and context information [5]. This primary focus on facility of the user interface differs to our work where ideally the agent remains more or less invisible to the user of the building. There have been also other research projects by [1] and [11] concerned with producing optimal models for buildings which will used later in control, however these techniques lacks flexibility as if the system characteristics changes the system has to repeat a time consuming offline learning cycle. Also these techniques ignores the fact that the building control is subject to human desires which can vary from a person to another, so there is a need for online adaptive and interacting learning. We believe our approach is unique in seeking to apply a hierarchical fuzzy control architecture with genetic learning. This approach allows the learning or adaptation to be performed through interaction with the environment (with the occupant of the room being seen as part of the environment). In this approach there is no need for simulation or direct human intervention into the rule setting system thus satisfying the definitions of intelligent autonomous agents in robotics [23] and intelligent systems as explained above.

Broadly speaking, the agent work described here is situated in the recent line of research that concentrates on the realisation of embedded real-time agents grounded in the physical world. Traditional approaches such as real-time control and expert systems rely on predictive models and thus have difficulty in environments that are of intractable complexity (in terms of the dynamics and

number of variables) being essentially (if not actually) non-deterministic. However, a number of robotic based researchers have had notable success in dealing with this type of problem based on techniques that can broadly be categorised by the term behaviour based systems. We take this approach as our starting point and have developed a method of implementing this as a hierarchical fuzzy system with numerous consequent advantages. Thus, our macro-control architecture belongs the “behaviour based control architecture” school pioneered by Rodney Brooks of MIT in the late 80’s [4]. In this approach a number of concurrent behaviours (mechanisms to attain goal or maintain state) are active (sensing environment, effecting machine) to a degree determined the relationship of the machine and environment. We have extended this approach by developing a number of mechanisms that address both behaviour-integration, in the form of a hierarchical fuzzy controller, and genetic learning, combined into a single architecture which we label an “*associative experience engine*” [13,14]. In addition to dealing with control problems that are difficult to model, the self-programming nature of the system has benefits in reducing software development costs throughout the lifetime of the product. In this paper we give an overview of both the hierarchical fuzzy control and the genetically based associative experience engine applied to develop an embedded-agent within a distributed intelligent buildings architecture.

## 2. Distributed architecture

The granularity of our computational distribution is room-based. Thus, each room contains an *embedded-agent*, which is then responsible, via sensors and effectors for the local control of that room as shown in Fig. 1. The logic behind this is that it mirrors the architects’ vision of the functionality of the building and thereby provides a natural segmentation of agent types and functions. All embedded-agents are connected via a high level network (IP-ethernet in our case), thereby enabling collaboration or sharing of information to take place where appropriate. Within a room, devices such as sensors and effectors are connected together using a building services network (Lontalk in our case) and IP at the higher level. This distributed artificial intelligence (DAI) architecture is illustrated in Fig. 1.

## 3. The embedded-agents

Fig. 2 shows the internal of the behaviour-based agent. Controlling a large integrated building system requires a complicated control function resulting from the large input and output space and the need to deal with many imprecise and unpredictable factors, including people. In our system we simplify

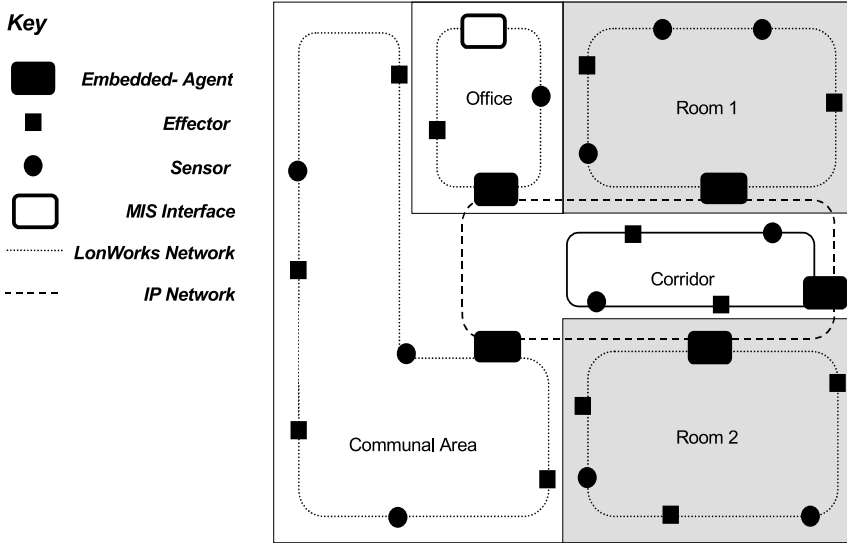


Fig. 1. The DAI building-wide architecture.

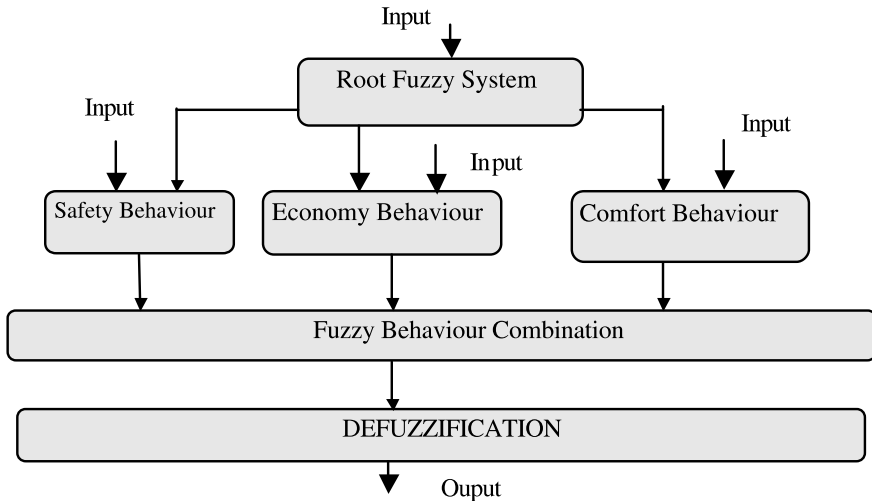


Fig. 2. The hierarchical fuzzy control system.

this problem by breaking down the control space into multiple behaviours, each of which responds to specific types of situations, and then integrating their recommendations.

### 3.1. The hierarchical fuzzy control architecture

The behaviour based approach, pioneered by Brooks, consists of many simple co-operating units and has produced very promising results when applied to the control of robotics, which we argue includes IB.

The problem of how to co-ordinate the simultaneous activity of several independent behaviour-producing units to obtain an overall coherent behaviour have been discussed by many authors [4,21]. The work described in this paper suggests a solution based on using fuzzy logic to both implement individual behaviour elements and the necessary arbitration (allowing both fixed and dynamic arbitration policies to be implemented). We achieve this by implementing each behaviour as a fuzzy process and then using fuzzy agents to co-ordinate them. In the resultant architecture, a hierarchical fuzzy logic controller (HFLC) takes a hierarchical tree structure form and is shown in Fig. 2. This hierarchical approach has the following advantages:

- It facilitates the design of the agent controller and reduces the number of rules to be determined. It uses the benefits of fuzzy logic to deal with imprecision and uncertainty.
- Using fuzzy logic for the co-ordination between the different behaviours which allows more than one behaviour to be active to differing degrees thereby avoiding the drawbacks of on–off switching schema (i.e. dealing with situations where several criteria need to be taken into account). In addition, using fuzzy co-ordination provides a smooth transition between behaviours with a consequent smooth output response.
- It offers a flexible structure where new behaviours can be added or modified easily. The system is capable of performing different tasks using identical behaviours by changing only the co-ordination parameters to satisfy a different high level objective without the need for re-planning.

Our room-based decomposition of behaviours consists of the following meta-functions. A *safety behaviour*, which ensures that the environmental conditions in the room are always at a safe level. An *emergency behaviour*, which in the case of a fire alarm or another emergency, might for instance open the emergency doors and switch off the main heating and illumination systems. In the case of an emergency this will be the only active behaviour. An *economy behaviour* that ensures that energy is not wasted so that if a room is unoccupied the heating and illumination will be switched to a sensible minimum value. All of the previous behaviours are fixed but settable, there is however a set of behaviours that the system learns from the occupant of the room and these are called *comfort behaviours*. These behaviours try to ensure that the conditions the occupant prefers (subject to being safe) are set. The learning process is done *interactively* using reinforcement where the controller takes actions and

monitors these actions to see if they satisfy the occupant or not, until a degree of satisfaction is achieved. This process would be acceptable in a hotel or apartment block but would probably require the intervention of a care assistant in housing for the elderly or those with learning difficulties. The behaviours, resident inside the agent, take their input from a variety of sensors in the room (such as occupancy, inside illumination level, outside illumination level, inside temperature, outside temperature etc.), and adjust device outputs (such as heating, lighting, blinds, etc.) according to pre-determined, but settable, levels. The complexities of training and negotiating satisfactory values for multiple use rooms would depend upon having a reliable means of identifying different users.

In our prototype system each agent has six inputs made up of four environmental variables—a room temperature (RTEMP), the outside natural temperature (ONTEMP), the room illumination (RILLUM) and the outside natural illumination (ONILLUM) each of which have the fuzzy membership functions shown in Fig. 3. Each input will be represented by three fuzzy sets as these were found to be the smallest number that give satisfactory results. The two remaining inputs to the system are the room occupancy flag and another input indicating whether there is an emergency alarm. The system has two outputs which are the room heating (RH) setting and a room illumination (RI) setting, these outputs have the membership functions shown in Fig. 4. Seven fuzzy sets were used as it was found to be smallest number to give a satisfactory result. If the alarm input is active then the emergency behaviour is dominant and the room illumination and heat are switched off. The economy behaviour is active to a fuzzy degree depending on occupancy and the outside temperature and light. Behaviour co-ordination is done in a fuzzy way as will be explained later and the fuzzy membership for behaviour co-ordination are shown in Fig. 5.

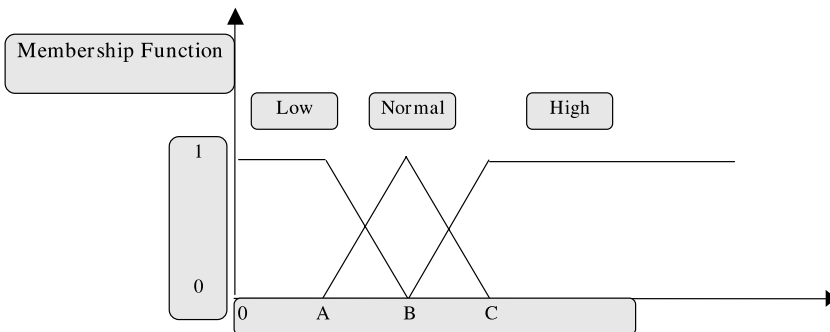


Fig. 3. The input membership functions for RTEMP and ONTEMP  $A = 10^\circ$ ,  $B = 20^\circ$ ,  $C = 30^\circ$ , for RILLUM and ONILLUM  $A = 300$  Lux,  $B = 400$  Lux,  $C = 500$  Lux.






Fig. 4. The output membership functions for RI  $A = 0\%$ ,  $B = 20\%$ ,  $C = 35\%$ ,  $D = 40\%$ ,  $E = 50\%$ ,  $F = 70\%$ ,  $G = 100\%$ , for RH  $A = 0\%$ ,  $B = 30\%$ ,  $C = 40\%$ ,  $D = 50\%$ ,  $E = 70\%$ ,  $F = 85\%$ ,  $G = 100\%$ .

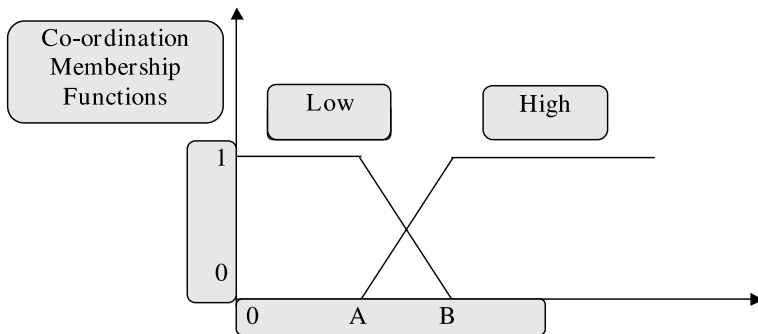


Fig. 5. The co-ordination parameters for ONILLUM  $A = 350$  Lux and  $B = 400$  Lux, for ON-TEMP  $A = 15^\circ$  and  $B = 25^\circ$ .

Each separate behaviour will use a singleton fuzzifier, triangular membership functions, product inference, max-product composition and height defuzzification. The selected techniques were chosen due to their computational simplicity. The equation that maps the system input to output is given by:

$$Y_t = \frac{\sum_{p=1}^M y_p \prod_{i=1}^G \alpha A_i p}{\sum_{p=1}^M \prod_{i=1}^G \alpha A_i p} \tag{1}$$

In this equation  $M$  is the total number of rules,  $y$  is the crisp output for each rule,  $\alpha A_i$  is the product of the membership functions of each rule input and  $G$  is the number of inputs. In this hierarchical architecture we will use a fuzzy operator to combine the preferences of different behaviour into a collective preference. According to this view, command fusion is decomposed into two

steps: preference combination, decision and in the case of using fuzzy numbers for preferences, product–sum combination and height defuzzification. The final output equation is [21]:

$$C = \frac{\sum_i (BW_i C_i)}{\sum_i BW_i} \quad (2)$$

Here  $i = \text{economy, comfort}$ ,  $C_i$  is the behaviour command output (room heat and temperature).  $BW_i$  is the behaviour weight. The behaviour weights are calculated dynamically taking into account the context of the agent. In Fig. 2 each behaviour is treated as an independent fuzzy controller and then using fuzzy behaviour combination we obtain a collective fuzzy output which is then defuzzified to obtain a final crisp output. The fuzzy values are used as inputs to the context rules which are suggested by the high level planner according to the mission to be performed by *the agent*.

*Which is in general:*

*IF ONTEMP IS HIGH AND ONILLUM IS HIGH AND THE ROOM IS OCCUPIED THEN ECONOMY  
IF ONTEMP IS LOW AND ONILLUM IS LOW THEN COMFORT  
IF THE ROOM IS VACANT THEN RH IS LOW AND RI IS LOW*

The context rules determine which behaviour is fired, and to what degree, depending on the fuzzy membership functions in Fig. 5. The final output is a mixture of the different behaviour outputs, each weighted by the degree of its importance, and the final output is calculated using Eq. (2).

#### 4. Overview of the genetic learning architecture

For learning and adapting the dynamic comfort rule base according to the occupant behaviours we use an evolutionary computing approach based on a development of novel hierarchical GA technique. This mechanism operates directly on the fuzzy controller rule-sets. We refer to any learning conducted without user interaction, in isolation from the environment as *offline* learning. In our case learning will be done *online* in real-time through interaction with the actual environment and user.

##### 4.1. The associative experience engine

Fig. 6 provides an architectural overview of what we term an *associative experience engine* (AEE) which forms the learning engine within the control architecture and is the subject of British patent 99-10539.7. Behaviours are

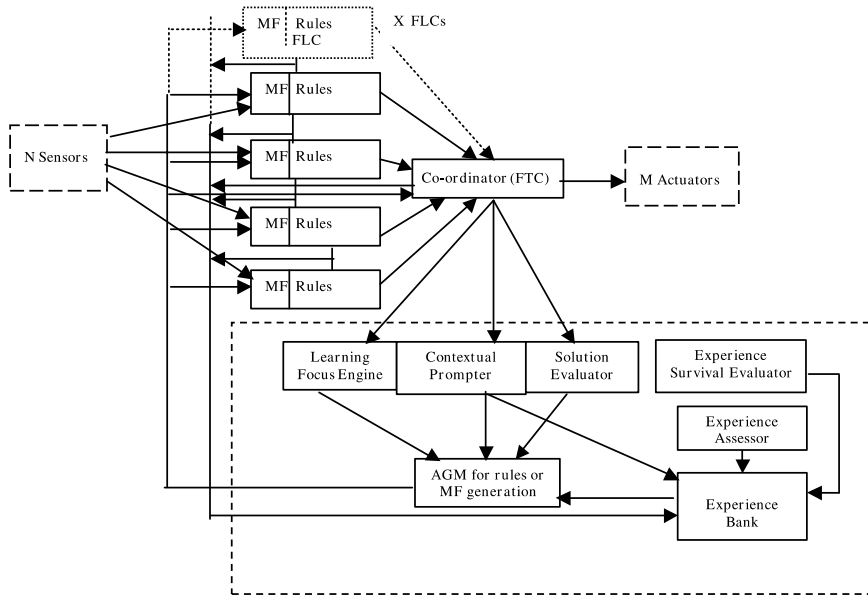


Fig. 6. Architectural overview of associative experience learning engine (UK patent no. 99-10539.7).

represented by parallel fuzzy logic controllers (FLC). Each FLC has two parameters that can be modified which are the *rule base* (RB) of each behaviour and its *membership functions* (MF), however in this paper we will concentrate on learning the RB. The behaviours receive their inputs from sensors. The output of each FLC is then fed to the actuators via the *co-ordinator* that weights its effect. When the system response fails to have a desired a response, the learning cycle begins.

The learning depends on the *learning focus* which is supplied by the *co-ordinator* (the fuzzy engine which weights contributions to the outputs). When the *learning focus* is learning an individual rule base for a behaviour, then each rule base for each behaviour is learnt alone, which is the case when learning the comfort behaviour rule base in the IB agent.

After the initial process of interactively learning the occupant's preferred behaviours, if the occupants behaviour indicates that there are changes in attitude or that a new situation has arisen, the system can add or delete rules to satisfy the occupant by re-entering the interactive mode. In the case of a new occupant in the room the system interactively monitors his actions during the first few minutes and after that the system takes control and fires the most suitable rule base from the *Experience Bank*. If this rule base is

not totally appropriate, the interactive system learning is started from the previously found best point in the search space rather than starting from random, this results in the speeding up of learning and adaptation. The experience assessor assigns each stored rule base in the *Experience Bank* a fitness value. When the *Experience Bank* is full, we have to delete some experiences. To assist with this the *experience survival evaluator* determines which rules are removed according to their importance (as set by the experience assessor).

When the past experiences do not satisfy the occupant's needs we use the best-fit experiences to reduce the search space by pointing to a better starting point, which is the experience solution with the largest fitness. We then fire an *adaptive genetic mechanism (AGM)* using adaptive learning parameters to speed the search for new solutions. The AGM is constrained to produce new solutions in certain range defined by the *contextual prompter* which is supplied by human constrains to avoid the AGM searching options where solutions are unlikely to be found. By using these mechanisms we narrow the AGM search space massively thus improving its efficiency. After generating new solutions the system tests the new solution and gives it fitness through the *solution evaluator*. The AGM provides new options until a satisfactory solution is achieved.

From a users viewpoint the system functions interactively as follows. A user is asked to select his preference for any given programmable setting. The system then tries to adapt its rules to achieve this setting. The user is prompted to confirm or deny his satisfaction with the result. The system then either tries to re-adjust the rules if the occupant is dissatisfied or, if the user is satisfied, the current rule set is accepted. Experiments to date show the *experience engine* achieves a satisfactory solution in a small number of iterations (typically 22) which takes only 3 min of the agent time. As we mentioned earlier this process would probably have to undertaken by a care assistant for some populations. In the next section we will explain the techniques in more detail.

We have compared our method, which is an online proactive method, with other offline supervised learning system like the Mendel–Wang approach [18], offline GA [15], ANFIS. We found, from the experimental results, that our system gives comparable results to that of the offline approaches while our system learns online through interaction with the user and has the ability to adapt to new users or situations. Offline methods have to repeat the learning cycle from the beginning and require that the initial training set plus any newly acquired data be used. Our system works by cause-effect actions in the form of fuzzy rules, based on the occupant's actions. The advantage of this is that the system responds and adapts to the users needs in an interactive manner. To gather all the given results we used an AEE specially adapted for IB control, shown in Fig. 7.

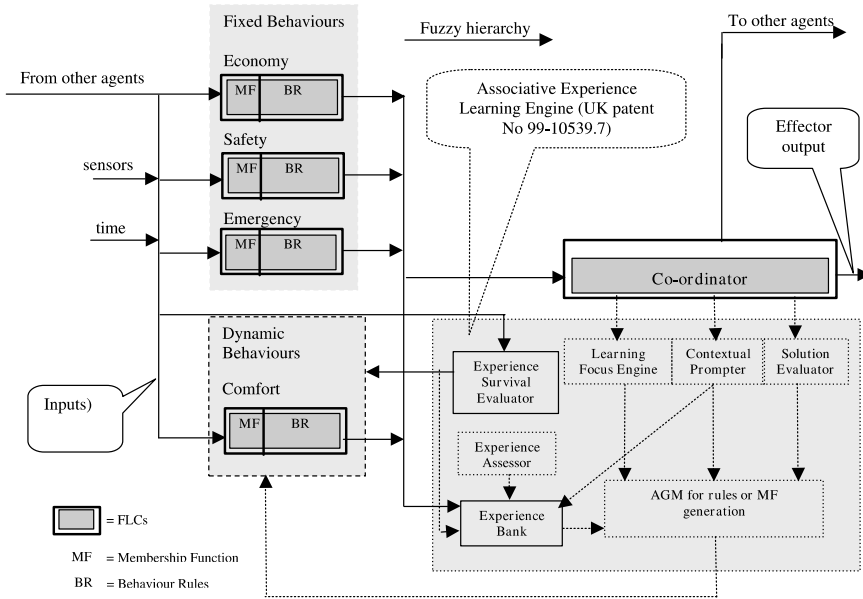


Fig. 7. Modified AEE embedded-agent architecture.

## 5. The associative experience learning engine applied in IB in detail

### 5.1. Identifying poorly performing rules

The rule-base is initialised to have all the outputs switched off. The GA population consists of all the rules consequents contributing to an action which is usually a small number of rules. As in the case with classifier systems, in order to preserve the system performance the GA is allowed to replace a subset of the classifiers (the rules in this case). The worst  $m$  classifiers are replaced by the  $m$  new classifiers created by the application of the GA on the population [9]. The new rules are tested by the combined action of the performance and apportionment of credit mechanisms. We will replace the actions of the worst four rules that participated in bad action at a given input.

In the learning phase the agent is introduced to different situations, such as having low temperature inside and outside the room and low illumination level, and the agent, guided by the occupants needs attempts to discover the rules needed in each situation. The learning system consists of learning different situations or episodes; in each situation only small number of rules will be fired. The model to be learnt is small and so is the search space. The accent on local models implies the possibility of learning by focusing at each step on a small part of the search space only, thus reducing interaction among partial

solutions. The interaction among local models, due to the intersection of neighbouring fuzzy sets means local learning reflects on global performance [3]. Moreover, the smooth transition among the different models implemented by fuzzy rules implies robustness with respect to data noise. So we can have global results coming from the combination of local models, and smooth transition between close models. Also dividing the learning into local situations can reduce the number of learnt rules. For example if we started learning with 81 rules, the agent discovered that during its interactive training with the occupant, it needed only 49 rules.

### 5.2. Fitness determination and credit assignment

The system fitness is determined by the *solution evaluator* and is evaluated by how much the system satisfies the room occupant desired input value or normal value (such as desired temperature) in a specific situation and how it reduces the normalised absolute deviation ( $d$ ) from the normal value. This is given by:

$$d = \frac{|\text{normal value} - \text{deviated value}|}{\text{max deviation}} \quad (3)$$

where the normal value will correspond to the value that gives the value desired by the human. The deviated value is any value deviating from the normal value. The maximum deviation corresponds to the maximum deviation that can occur. So the fitness of the solution is given by  $d1 - d2$  where  $d2$  is the normalised absolute deviation before introducing a new solution and  $d1$  is the normalised absolute deviation following the new solution. The deviation is measured using the agent's physical sensors, which gives the agent the ability to adapt to the imprecision and the noise found in the real sensors rather than relying on estimates from previous simulations.

The fitness of each rule at a given situation is calculated as follows. The crisp output  $Y_i$  can be written as in (1). If the agent has two output variables, then we have  $Y_{i1}$  and  $Y_{i2}$ . The normalised contribution of each rule  $p$  output ( $Y_{p1}, Y_{p2}$ ) to the total output  $Y_{i1}$  and  $Y_{i2}$  can be denoted by  $S_{r1}, S_{r2}$  where  $S_{r1}$  and  $S_{r2}$  is given by:

$$S_{r1} = \frac{Y_{p1} \prod_{i=1}^G \alpha A_{ip}}{\sum_{p=1}^M \prod_{i=1}^G \alpha A_{ip}}, \quad S_{r2} = \frac{Y_{p2} \prod_{i=1}^G \alpha A_{ip}}{\sum_{p=1}^M \prod_{i=1}^G \alpha A_{ip}} \quad (4)$$

We then calculate each rule's contribution to the final action  $S_c = (S_{r1} + S_{r2})/2$ . Then the most effective rules are those that have the greatest values of  $S_c$ . The fitness of the rule in a given solution is supplied by the *solution evaluator* and is given by:

$$S_{r_i} = \text{Constant} + (d_1 - d_2) \cdot S_c \quad (5)$$

$d_1 - d_2$  is the deviation improvement or degradation caused by the adjusted rule-base produced by the algorithm. If there is improvement in the deviation, then the rules that have contributed most will be given more fitness to boost their actions. If there is degradation then the rules that contributed more must be punished by reducing their fitness w.r.t other rules so we can give the chance to other useful actions and go away from the current bad actions.

### 5.3. Memory application

Zhou [24] presented the classifier system with memory (CSM) system that addresses the problem of long versus short term memory, i.e. how to use past experiences to ease the problem solving activity in novel situations. Zhou's approach is to build a system in which a short and long term memory are simultaneously present. The short term memory is just the standard set of rules found in every learning classifier system (the fuzzy rule base in our case). The long term memory is a set of rule clusters, where every rule cluster represents a generalised version of problem solving expertise acquired in previous problem solving activity. Each time the agent is presented a problem it starts the learning procedures trying to use long term experience by means of an appropriate initialisation mechanism. Thereafter, the system works as a standard classifier system (except for some minor changes) until an acceptable level of performance has been achieved. It is at this point that a generalizer process takes control and compresses the acquired knowledge into a cluster of rules that are memorised for later use in the long term memory.

In our system, as the agent begins learning, it has no previous experience and the *Experience Bank* is empty. But as it begins learning by GA, it begins filling the memory with different rule bases of different users. Each stored rule base is consisting of the rules that were learnt and the actions (consequences) that were learnt by the GA.

After monitoring the user's action for two minutes the agent then matches the fired rules during this time to sets of rules bases for different users stored in the *Experience Bank*. The system tries to identify which rule base is appropriate to the user based on the basis of actions taken by him during the first two minutes, and the rule base that contains the most similar actions to the occupant is chosen as a starting point for learning and adaptation.

Each time the agent is presented with a situation to solve, it begins checking if the consequences of firing the rules from a rule base extracted from the *Experience Bank* suits the new user or not. If these rules are suitable for the user then they are used for the comfort behaviours. If some actions are not suitable for the user, the system begins identifying the poorly performing rules as described in Section 5.1, then it fires the AGM to change these rules. This action helps to speed up the genetic search as it starts the search from the best found point in the search space instead of starting from a random point.

By doing this, our system does not need the matcher calculations used by [24]. This is because our system does not use the binary message coding, or “do not care”, conditions but uses perfect matches, hence we do not need the generalizer. The clusters are arranged in a queue starting from the most recent experiences.

Problems occur as the system begins accumulating experience that exceeds the physical memory limits. This implies that we must get rid of some of the stored information as the acquired experience increases. However we do not favour this, as this means that some of the experiences acquired by the agent will be lost. So for every rule base cluster we attach a *difficulty counter* to count the number of iterations taken by the agent to find a suitable rule base for a given user, we also attach a *frequency counter* to count how much this rule base have been retrieved. The *degree of importance* of each rule base cluster is calculated by the *experience survival valuer* and is given by the product of the *frequency counter* and the *difficulty counter*. This approach tries to keep the rules that have required a lot of effort to learn (due to the difficulty of the situation) and also the rules that are used frequently. When there is no more room in the *Experience Bank*, the rule base cluster that had the least *degree of importance* is selected for removal. If two rule base clusters share the same importance degree, tie-breaking is resolved by a least-recently-used strategy; derived from a “life invocation” flag, that is updated each time a rule is activated.

#### 5.4. Using GA to produce new solutions

If the rule base extracted from the *Experience Bank* is not suitable for the user, the GA starts its search for new consequences for the poorly performing rules. The fitness of each rule in the population is proportional to its contribution in the final action. If the proposed action by the new solution results in an improvement in performance then the rules that have contributed most will have their fitness increased more than the rules that have contributed less in this situation. If the result was a decrease in performance then the rules that have contributed most to this action will have their fitness less than the rules that have contributed less to this action. This allows us to go away from those points in the search space that cause no improvement or degradation in the performance. The parents for the new solution are chosen proportional to their fitness using the roulette-wheel selection process and the genetic operations of crossover and mutation are applied.

The proposed system can be viewed as a double hierarchy system in which the fuzzy behaviours are organised in a hierarchical form and the online learning algorithm is also a hierarchy. In the higher level we have a population of solutions stored in the *Experience Bank*. If the stored experiences leads to a solution then the search ends, if none of these stored experiences leads to a



solution then each of these experiences acquires a fitness assigned by the *Experience Assessor* by finding how many rule in the stored rule base are similar to the user's action in the first two minutes. The highest fitness experience is used as a starting position to the lower level GA that is used to produce new solution to the current situation.

The AGM is the rule discovery component for our system (as in the classifier system). We used Srinivas method [22] to adapt the control parameters (mutation and crossover probabilities). The strategy used for adapting the control parameters depends on the definition of the performance of the GA. In a non-stationary environment (which is our case), where the optimal solution changes with time, the GA should possess the capacity to track optimal solutions, too. The adaptation strategy needs to vary the control parameters appropriately whenever the GA is not able to track the located optimum. It is essential to have two characteristics in GA for optimisation. The first characteristic is the capacity to converge to an optimum (local or global) after locating the region containing the optimum. The second characteristic is the capacity to explore new regions of the solution space in search of the global optimum. In order to vary  $P_c$  (crossover probability) and  $P_m$  (mutation probability) adaptively, for preventing premature convergence of the GA, it is essential to be able to identify whether the GA is converging to an optimum. One possible way of detecting convergence is to observe the average fitness value  $f'$  of the population in relation to the maximum fitness value  $f_{\max}$  of the population.  $f_{\max} - f'$  is likely to be less for a population that has converged to an optimum solution than that for a population scattered in the solution space. The equations that determines  $P_c$ ,  $P_m$  are given by:

$$P_c = (f_{\max} - f') / (f_{\max} - f') \quad f'' \geq f' \quad (6)$$

$$P_c = 1 \quad f'' < f'$$

$$P_m = (f_{\max} - f) / 2(f_{\max} - f') \quad f \geq f' \quad (7)$$

$$P_m = 0.5$$

where  $f''$  is the larger of the fitness values of the solutions to be crossed,  $f$  is the fitness of the individual solutions. The method means that we have  $p_c$  and  $p_m$  for each chromosome. The type of crossover was chosen to be one point crossover for computational simplicity and real time performance.

In [22] this method was superior to the simple GA and gave faster convergence of rate 8:1. We use this adaptive method for finding the values of crossover and mutation probabilities. This way leads to fast convergent solutions, and adapts the GA for non-stationary environments and relieves the designer from determining these values, heuristically.

We use constrained GA search, which is constrained by the *contextual prompter* according to the occupant's needs so that if, for example, a temperature is too high for the room occupier then the AGM would not suggest

solutions increasing the temperature. In this way we can minimise the search space of the GA and achieve faster conversion.

In order to justify these techniques we have conducted various experiments with this adaptive GA (AGA) with open range and AGA with constrained range and simple GA (SGA) with constrained range for the problem of learning the comfort behaviour of different human attitudes as shown in Fig. 8. The SGA was tried with different parameters in the range [0.5 1.0] for  $p_c$  and [0.001 0.1] for  $p_m$ . And the best performing one was found to be  $p_c = 0.7$  and  $p_m = 0.002$  and is plotted in Fig. 8. It was found that the one with constrained AGA converges to a solution in average after *only 22 iterations and 3 min of the real agent time* (most of the time is consumed in interacting with the human). The AGA with open range converges also but after larger number of iterations (33 iterations in average) as it needs more time to explore the search space and determine its limits this takes about 11 min of the agent time to converge to a solution. The SGA with defined limits and  $p_c = 0.7$  and  $p_m = 0.002$  converge to a solution after in average 48 iterations and 24 min of the agent time. These experiments justify that our proposed algorithm converges in a very short time interval thanks to the AGA and the constrained GA.

We use binary coding in the GA. For each rule there are two actions which are the room heating and illumination. As we have seven output membership function, we decode each action by three bits as follows: *very very low* is 000, *very low* is 001, *low* 010, *normal* is 011, *high* is 100, *very high* 101, *very very high* 110. By doing this we have a chromosome length of 6 bits.

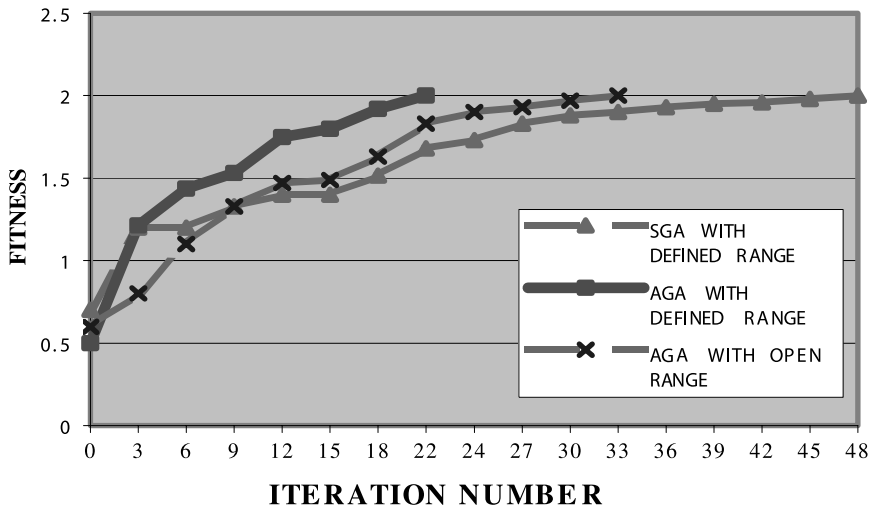


Fig. 8. The best fitness plotted against the number of iterations for different learning.

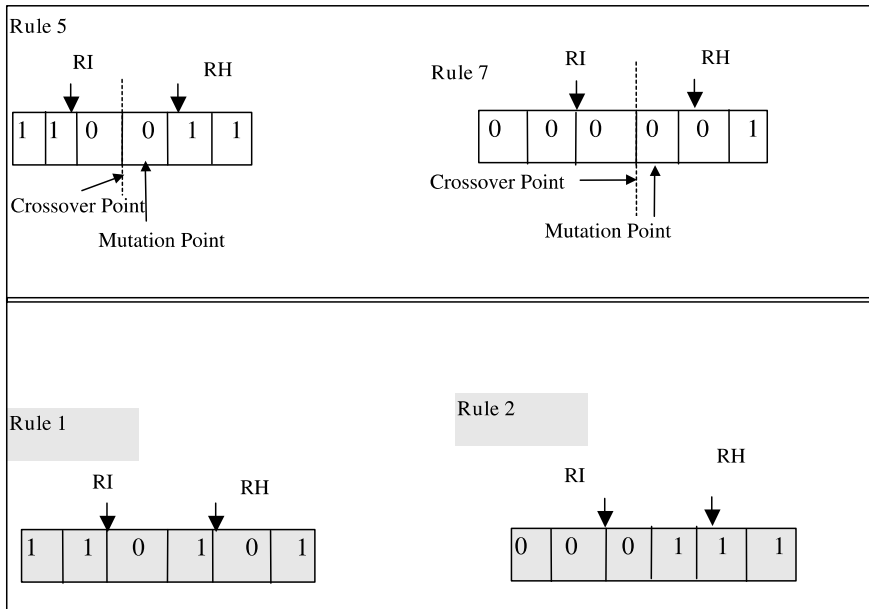


Fig. 9. An example of GA processes in our proposed classifier system in which rule 5 and rule 7 (which were selected due to their higher fitness values) are generating new consequent for rules 1, 2.

Fig. 9 shows a description of the GA operation in which the actions of rule number 5 and rule number 2 of the comfort behaviour are chosen for reproduction by roulette wheel selection due their high fitness. They have contributed more with their actions to the final action which caused improvement, or contributed less with their actions to final action which caused degradation. The adaptive crossover and mutation probabilities have been applied to both chromosomes. The resultant offspring were used to replace the consequences of rules 1 and 3 which were mostly blamed for the unsatisfactory responses. It is worth saying that at each situation the population consists only from that fired and caused the user dissatisfaction.

### 6. Experimental results

In our experiments we used an IB agent based on 68000 Motorola processor, the agent is equipped with light and heat sensors and actuators in the form of a heater and a light source; the IB agent is shown in Fig. 10. This agent is tested in a room with various conditions such as multiple occupancy, different levels of natural light and temperature and different times of the day and different




Fig. 10. The IB agent.

human desires. There is a built in economy behaviour that should switch the heat low and ventilation off after the room is vacated. There is also a safety behaviour that prevents the heat going below a minimum safe level (e.g. zero degrees which would result in pipes freezing). Centrally, there is the comfort behaviour for the occupant themselves which will be learnt using our patented fuzzy–genetic techniques. The agent deals in a proactive way with the occupant interacting with him to see if the action was completed satisfactory or not, and whether it is required to decrease or increase levels in question.

The agent was tried, in different environments, under different conditions such as hot sunny days and cold and dark days. The agent using our patented techniques was able to find a satisfactory rule base for different users in an average of 22 trials taking 3 min of the agent time. Our method also optimised the number of rules by using only the rules that are important to the room occupier. The system also has the ability to adapt later and to add or delete rules and modify the actions of the existing rules interactively to the room occupant's wishes.

We have tried different offline learning methods such as the Mendel–Wang method, offline GA and ANFIS, it was found that the Mendel–Wang approach

Table 1  
The rule base found by our method and by Mendel–Wang method

RTEMP	ONTEMP	RILLUM	ONILLUM	Our method		Mendel–Wang method	
				RH	RL	RH	RL
Low	Low	Low	Low	Very high	High	Very very high	Very high
Low	Low	Low	Norm	Very very high	Very very high	Very high	Very high
Low	Low	Norm	Low	Very high	Norm	Very high	Norm
Low	Low	Norm	Norm	Norm	Norm	Norm	Norm
Low	Low	Norm	High	Very high	Very low	Very high	Very low
Low	Low	High	Low	High	Very low	Very high	Very low
Low	Low	High	Norm	High	Very very low	Very high	Very very low
Low	Low	High	High	High	Very very low	Very high	Low
Low	Norm	Low	Low	Norm	Very very high	Norm	Very high
Low	Norm	Low	Norm	Norm	Norm	Norm	High
Low	Norm	Norm	Low	Very very high	Very low	Very very high	Very low
Low	Norm	Norm	Norm	High	Very low	Very high	Very very low
Low	Norm	Norm	High	High	Very very low	Very high	Very low
Low	Norm	High	Low	Very high	Low	Very very high	Very low
Low	Norm	High	Norm	Norm	Very very low	High	Very low
Low	Norm	High	High	Very high	Very very low	Very very high	Low
Low	High	Low	Low	Norm	Very high	Very very high	Very high
Low	High	Low	Norm	Very very high	High	Very very high	High
Low	High	Norm	Low	High	Norm	Very very high	High
Low	High	Norm	Norm	Very very low	Very low	Very very low	Very low
Low	High	Norm	High	Very very low	Very low	Very very low	Very low
Low	High	High	Norm	Very very low	Very low	Very very low	Very low
Low	High	High	High	Very very low	Very low	Very very low	Very low
Norm	Low	Low	Low	Norm	High	Norm	High
Norm	Low	Low	Norm	Norm	Very very high	High	High
Norm	Low	Norm	Low	Norm	Norm	Norm	Norm
Norm	Low	Norm	Norm	Norm	Norm	High	High
Norm	Low	Norm	High	Very very low	Very low	Very very low	Low

Table 1 (continued)

RTEMP	ONTEMP	RILLUM	ONILLUM	Our method		Mendel–Wang method	
				RH	RL	RH	RL
Norm	Low	High	Norm	Low	Norm	Very very low	Low
Norm	Low	High	High	Very very low	Very low	Very low	Norm
Norm	Norm	Low	Low	Very low	Very high	Very very low	Very very high
Norm	Norm	Norm	Low	Low	Norm	Very low	High
Norm	Norm	Norm	Norm	Very low	Very very low	Very very low	Norm
Norm	Norm	Norm	High	Very very low	Low	Very low	Very very low
Norm	Norm	High	Norm	Low	Very low	Very very low	Low
Norm	Norm	High	High	Very low	Low	Very very low	Low
Norm	High	Low	Low	Very very low	High	Very very low	Norm
Norm	High	High	High	Very very low	Very low	Very very low	Very low
High	Low	Low	Low	Very low	High	Very low	Very high
Norm	High	High	High	Very very low	Very low	Very very low	Very low
High	Low	Low	Low	Very low	High	Very low	Very high
High	Low	Norm	Low	Very low	Very very high	Very low	Very high
High	Low	Norm	Norm	Very very low	Low	Very low	Very low
High	Low	Norm	High	Norm	Very very low	Norm	Very low
High	Low	High	Norm	Norm	Low	Very very low	Low
High	Low	High	High	Very low	Low	Very low	Low
High	Norm	Low	Low	Norm	High	Very low	High
High	Norm	Norm	Low	Low	Very high	Very low	Very high
High	Norm	Norm	Norm	Very very low	Low	Low	Very low
High	Norm	Norm	High	Very very low	Very low	Very very low	Low
High	Norm	High	High	Very very low	Very low	Low	Very very low

had produced the best results. Table 1 shows the rules obtained by our method and the Mendel–Wang method, which is an offline rule extraction method. Note that our method had optimised the number of rules from an expected  $3^4 = 81$  rule base to only 49 rules. Also a problem with supervised learning techniques, such as Mendel–Wang, are that they need a person to supply a set of desired values. Which would be difficult since the values are not intuitive. If the person wants the heating turned to what might be the equivalent of 37% and the illumination to 40% (taken as percentages of their maximum output) then finding or deriving an appropriate value from the device setting for the purposes of generating a training rule would be difficult. This is why, although the rules extracted by our method are similar to the Mendel–Wang a rule, the difference is that in our case we interact with the person until he is satisfied. Another advantage our method is that our system can quickly adapt to any change, starting from the closest previous solution, while the supervised learning techniques are off-line and repeat the learning cycle from the beginning and requiring a full set of training data.

## **7. Conclusion and future work**

In this paper we have presented a novel online learning, adaptation and control algorithm based on a double hierarchical fuzzy–genetic system. In this prototype IB agent this system is composed of three fixed behaviours—the safety, emergency and economy behaviours and an adaptable rule set of comfort behaviours that are adapted according to the occupants actual behaviour. The system interactively learnt an optimised rule base for the comfort behaviour in a small set of interactions. The system produced similar rules to the rule base learnt by off-line supervised techniques but our system has many advantages such as the ability to adapt to changes in environmental conditions or the room occupant. It adapted in real time, on-line accommodating new data and rules, as they become known to the system. The system is memory based, possesses data, exemplar storage and retrieval capacities; leaning and improving through active interaction with the user and the environment. It includes parameters to represent short and long term memory, age, forgetting, and analyses itself in terms of behaviour, error and success. Thus we would argue that our approach substantially meets the criteria for intelligent systems set out by researchers such as Kasabov and Steels and described earlier in this paper which has not achieved by other methods in IB.

For our future work, we plan to move from an experimental bench based rig to IB equipped rooms with our intelligent, autonomous, agents and thereby extend the investigation to a richer set of environmental inputs which include a wider set of occupant behaviours. Also, we believe that an ideal IB agent

should require little or no involvement from the occupant and thus another goal of our future work is further minimise the interactive aspects of the agent.

### **Acknowledgements**

We are pleased to acknowledge the contribution from Malcolm Lear (Essex University) who built the agent hardware, sensors and test rig. We would also like to thank, Anthony Pounds–Cornish, Sue Sharples, Gillian Kearney, Robin Doling and Filiz Cayci with whom we have had many stimulating discussions on embedded-agent architectures.

### **References**

- [1] P. Angelov, R. Buswell, V. Hanby, Automatic generation of fuzzy rule-based models from data by genetic algorithms, in: *Proceedings of International Conference on Recent Advances on Soft Computing*, Leicester, UK, 2000.
- [2] A. Bonarini, F. Basso, Learning behaviours implemented as fuzzy logic and reinforcement learning, *Second Online Workshop on Evolutionary Computation*, 1996.
- [3] A. Bonarini, Comparing reinforcement learning algorithms applied to crisp and fuzzy learning classifier systems, in: *Proceedings of the Genetic and Evolutionary Computation Conference*, Orlando, Florida, 1999, pp. 52–60.
- [4] R. Brooks, Intelligence without representation, *Artificial Intelligence* 47 (1991) 139–159.
- [5] R. Brooks, Intelligent room project, in: *Proceedings of the 2nd International Cognitive Technology Conference (CT'97)*, Japan, 1997.
- [6] V. Callaghan, G. Clarke, Buildings as intelligent autonomous systems: a model for integrating personal and building agents, in: *The 6th International Conference on Intelligent Autonomous Systems (IAS-6)*, Venice, Italy, 2000.
- [7] M.H. Coen, Building brains for rooms: designing distributed software agents, in: *Proceedings of the Ninth Innovative Applications of AI Conference*, AAAI Press, 1997.
- [8] P. Davidsson, Energy Saving and Value Added Services; Controlling Intelligent-Buildings Using a Multi-Agent System Approach in *DA/DSM Europe DistribuTECH*, PennWell, 1998.
- [9] M. Dorigo, Genetics-based machine learning and behaviour based robotics: a new synthesis, *IEEE Transactions on Systems, Man, Cybernetics* (1993) 141–154.
- [10] T. Fukuda, N. Kubota, An intelligent robotic system based on fuzzy approach, *Proceedings of the IEEE* 87 (1999) 1448–1470.
- [11] A. Gegov, D. Azzi, G. Virk, Application of soft-computing techniques in modelling of buildings, in: *Proceedings of International Conference on Recent Advances on Soft Computing*, Leicester, UK, 2000.
- [12] D. Goldberg, *Genetic Algorithms in Search, Optimisation and Machine Learning*, Addison-Wesley, Reading, MA, 1989.
- [13] H. Hagrás, V. Callaghan, M. Colley, A fuzzy–genetic based embedded-agent approach to learning and control in agricultural autonomous vehicles, *1999 IEEE International Conference on Robotics and Automation*, Detroit, USA, 1999, pp. 1005–1010.
- [14] H. Hagrás, V. Callaghan, M. Colley, Online learning of fuzzy behaviours using genetic algorithms and real-time interaction with the environment, *1999 IEEE International Conference on Fuzzy Systems*, Seoul, Korea, 1999, pp. 668–672.



- [15] F. Hoffmann, Incremental tuning of fuzzy controllers by means of evolution strategy, in: Proceedings of the GP-98 Conference, Madison, Wisconsin, 1998, pp. 550–556.
- [16] N. Kasabov, Introduction: Hybrid intelligent adaptive systems, *International Journal of Intelligent Systems* 16 (1998) 453–454.
- [17] G. Linkens, O. Nyongeso, Genetic algorithms for fuzzy control, Part II: online system development and application, *IEE Proceedings Control Theory Applications* 142 (1995) 177–185.
- [18] J. Mendel, L. Wang, Generating fuzzy rules by learning through examples, *IEEE Transactions on Systems, Man and Cybernetics* 22 (1992) 1414–1427.
- [19] M. Mozer, The neural network house: an environment that adapts to its inhabitants, in: Proceedings of American Association for Artificial Intelligence Spring Symposium on Intelligent Environments, AAAI Press, 1998, pp. 110–114.
- [20] P. Robathan, *Intelligent Buildings Guide*, Intelligent Buildings Group and IBC Technical Services Limited, 1989.
- [21] A. Saffiotti, Fuzzy logic in autonomous robotics: behaviour co-ordination, in: Proceedings of the 6th IEEE International Conference on Fuzzy Systems, Barcelona, Spain, 1997, pp. 573–578.
- [22] M. Srinivas, L. Patnaik, Adaptation in Genetic Algorithms, *Genetic Algorithms for Pattern Recognition*, in: S. Pal, P. Wang, CRC press, 1996, pp. 45–64.
- [23] L. Steels, When are robots intelligent autonomous agents, *Journal of Robotics and Autonomous Systems* 15 (1995) 3–9.
- [24] H. Zhou, A computational model of cumulative learning, *Machine Learning Journal* (1990) 383–406.