

This document is published in:

IEEE Transactions on Autonomous Mental Development, June 2013, Vol. 5, Issue 2. 135-151.

DOI: 10.1109/TAMD.2012.2234120

© 2012 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

The funds have provided by the Spanish Government through the project called "A new approach to social robotics" (AROS), of MICINN (Ministry of Science and Innovation) and through the RoboCity2030-II-CM project (S2009/DPI-1559), funded by Programas de Actividades I+D en la Comunidad de Madrid and cofunded by Structural Funds of the EU.

An autonomous social robot in fear

Álvaro Castro-González, María Malfaz, Miguel A. Salichs

Abstract—Currently artificial emotions are being extensively used in robots. Most of these implementations are employed to display affective states. Nevertheless, their use to drive the robot’s behavior is not so common. This is the approach followed by the authors in this work. In this research, emotions are not treated in general but individually. Several emotions have been implemented in a real robot, but in this paper, authors focus on the use of the emotion of fear as an adaptive mechanism to avoid dangerous situations. In fact, fear is used as a motivation which guides the behavior during specific circumstances.

Appraisal of fear is one of the cornerstones of this work. A novel mechanism learns to identify the harmful circumstances which cause damage to the robot. Hence, these circumstances elicit the fear emotion and are known as fear releasers.

In order to prove the advantages of considering fear in our decision making system, the robot’s performance with and without fear are compared and the behaviors are analyzed.

The robot’s behaviors exhibited in relation to fear are natural, i.e. the same kind of behaviors can be observed on animals. Moreover, they have not been preprogrammed, but learned by real interactions in the real world.

All these ideas have been implemented in a real robot living in a laboratory and interacting with several items and people.

Index Terms—fear, emotions, decision making, social robot, autonomy.

I. INTRODUCTION

DUE to the recent interest on robotic applications for making easier our daily life, one of the most popular current research areas in robotics is Social Robotics. According to [1], a social robot can be defined as: “An autonomous or semi-autonomous robot that interacts and communicates with humans by following the behavioral norms expected by the people with whom the robot is intended to interact”. Therefore, it is expected that these robots have some kind of “social intelligence” which helps them to behave in an appropriate way. This kind of intelligence is related to emotions [2], so, an emotion-based control architecture seems to be an important element for socially intelligent robots [3].

The role of emotions in robotics has been extensively defended by many researchers [4][5][6][7][8] [9][3][10]. Most of them think that since emotions in humans and animals are essential for survival, autonomy, learning, decision-making, social interactions, etc., then, they are necessary for robots which are intended to live among us.

Nevertheless, according to Ziemke and Lowe [11], the main question is not if the robots need emotions, but if the human designers and users of robots need or want robots to have or at least express emotions. In fact, they stated that from an engineering perspective, the question rather is: Does building models of emotional/affective mechanisms into our robots

make them more natural, more useful, or more efficient? This question is very important, and in fact, some authors, such as Cañamero in [12] and Scheutz in [10], share the idea that emotions per se does not make robots more intelligent, but we must be able to prove how emotions improve somehow our robot’s performance. **In those referred works**, both authors propose a method **to evaluate** the usefulness of emotions in robots, that is, to compare their performance during the execution of the same task, with and without emotions.

The work presented in this paper is a continuation of previous research which is focused on autonomous and social robots. In previous works, authors have already applied emotions on virtual agents [13]. Moreover, they have presented the decision making system of the social robot Maggie [14] [15] which is based on drives, motivations, and emotions. **This robot learns, using reinforcement learning, the right action to select in order to maintain an internal equilibrium. The implemented emotions are happiness, sadness, and fear. The first two emotions, happiness and sadness, have been already studied in the cited works.** The work presented in this paper is centered on the use of fear and its application to robots.

The role of emotions in this system, as in [11], is focused on the behavioral organization of individual agents, rather than on their expression and recognition involved in social interactions. According to Arbib and Fellous [16], both approaches are referred as the “internal” and “external” aspects of emotions respectively. Others consider these aspects as emotional experience and emotional expression respectively [17].

Many psychologists nowadays agree that emotion is the result of an evaluation (appraisal) of the current situation [18] [19]. Moreover, the appraisal theory says that each emotion is elicited by a distinctive pattern of appraisal. Based on this idea, each emotion is treated by separate, that is, each of them has different appraisal mechanisms that release them.

In addition, each emotion has different roles. As already said, in this work, the emotions implemented in the decision making system of the robot are happiness, sadness, and fear. **In previous works, [14] [15] [13], is explained that happiness and sadness are generated from the evaluation of the wellbeing of the robot, and they are used as the reinforcement function in the learning process.**

This paper is focused on the emotion of fear, its generation, its role in the decision making process of the robot, and its influence on the robot’s performance.

In relation to its generation, it has been already stated that emotion is the result of an appraisal of the current situation. In the case of fear, as shown in the next section, this appraisal consists mainly on detecting a dangerous situation (emotional releaser). In nature, some of these emotional releasers are innate, or natural releasers, but others are learned by the

individual through its own experience. In this paper authors present an appraisal mechanism which allows the robot to learn to release fear in front of a new emerging dangerous situation, even if this event was not identified by the designer in advance.

On the other hand, the role of fear in humans and animals is basically to motivate the right behaviors to avoid those dangerous situations. Fear is used as an adaptive mechanism intended to prevent our robot of being harmed. This is, in fact, the unique role of fear according to Breazeal and Brooks [20]. Although in [21] is said that the fear responses (freezing or escape responses) are not learned and they are innate, in our approach, the robot does not have any predefined behavior. Thanks to the proposed decision making system and learning algorithms, the emotional behaviors are also learned by the robot and compared with the one expected from a living-being in order to show the validity of our approach.

At the end of the paper we evaluate, following the proposal of Cañamero and Scheutz, the performance of the robot using the emotion of fear in the decision making process. In order to do that, the robot lives in a real environment where dangerous situations arise. **Those results are compared with the ones obtained in the same environment and conditions, but without the fear emotion.** One of the main goals of this experiment is to study how the robot decides that a certain situation is dangerous (appraisal of fear).

The approach to fear presented in this paper is related to exogenous actions. These actions are not executed by the robot but by other objects causing asynchronous changes, which affect the robot. Harmful exogenous actions are the cornerstone to identify the dangerous situations.

When the exogenous actions executed in a state are always advantageous, the expected utility of this situation is high. If, from time to time, an exogenous action in a state is harmful to the robot, the expected utility of the state is still high. This is because, on average, the outcome is high. Both expected utilities are high, however, in the latter case, the situation is considered risky because an exogenous action could cause damage to the robot. By means of *fear*, these risky states are detected, and the right reactions are generated. The method proposed in this paper presents an approach to deal with these risky or dangerous situations.

This paper is organized as follows. Section II presents a brief review about the appraisal theory of emotions. Next, a review of some works related to this research topic is given in Section III. Then, Section IV exposes general ideas about the fear emotion and defines the scope of this work. In Section V, a general description of the robotic platform (hardware and software) is explained. Later, Section VI, presents a summary of the decision making system of the robot. Next, all parameters are configured to define the desired robot's personality for testing fear (Section VII). Section VIII details the novel mechanism intended to appraise dangerous situations (appraisal of fear). The next section, Section IX, presents the results of the experiments made in order to prove the usefulness of the fear emotion. Coming, some reflections about the observed results are expressed in Section X. Finally, the main conclusions of this paper are summarized in Section XI.

II. THE APPRAISAL THEORY

The appraisal theory is the claim that emotions are elicited by evaluations of situations [19]. According to this theory, it is the interpretations of situations, rather than the situations themselves, that cause emotions. Because appraisals intervene between situations and emotions, different individuals who appraise the same situation in significantly different ways will experience different emotions; and even a given individual who appraises the same situation in significantly different ways at different times will have different emotions. A good example would be a football match, the same situation (the result of the game) will produce different emotions depending on your team [8]. Another example can be observed on a student doing an exam: if he has studied hard all the semester, in general, he feels confident and relaxed; in contrast, if he has not studied enough, he experiences fear and gets nervous.

Following this theory, a situation cannot be tagged with an emotional value in advance, it is the interpretation each person makes of that situation which gives that individual evaluation.

In order to understand the appraisal mechanisms of emotions, that is, how they emerge in our brain, authors propose, as LeDoux in [22], that emotions must be studied one by one. According to him, the different emotions are produced by different brain networks. In fact, in [23] it is also explained how other theorists, inspired by the prototypical work of Darwin, have proposed that a small set of discrete emotions are underpinned by relatively separable neural system in the brain [24] [25]. Currently, it seems that different emotions involve different brain circuits, despite of same brain areas could be common.

Therefore, as stated in Section I, based on these ideas, emotions are studied separately, and more specifically, in this work we have focused our research on just one emotion: fear.

LeDoux [22] states that the function of this emotion is to detect danger and to produce reactions which increase the probabilities of survival in a dangerous situation. In other words, it is a defense mechanism. Therefore, the appraisal mechanism of fear is related to the evaluation of a situation as dangerous.

According to Rosis [26], there is a systematic confusion between two kinds of "evaluation":

- *Appraisal*: the automatic, implicit, intuitive orientation toward what is good and what is bad for the organism.
- *Evaluation*: the cognitive judgments relative to what is good or bad for someone (and why).

In fact, LeDoux relates this unconscious appraisal to emotion, and conscious evaluation to *feelings* [22].

On the other hand, Sloman [27] differentiates between primary emotions which have a reactive or automatic basis and secondary emotions that require a deliberative process to initiate them.

It is important to note that, in the approach presented in this paper, the appraisal of a dangerous situation will be based on an automatic process using associative learning. As will be shown later, the robot, using reinforcement learning, is able to identify dangerous situations without using any deliberative mechanism.

III. RELATED WORK

There are many researchers that have implemented emotions, or at least some emotionally inspired mechanisms, in their artificial systems. Some authors such as Breazeal [6] [20], Hirth et al [2], and Hollinger et al [28], focus their works on the external expression of emotions proving how they improve the human-robot interaction. Others implement emotions in their control architectures based on motivations in order to influence their decision making, see the works of Cañamero [29] [7] [12], Gadanho [4], and Parisi and Petrosino [30].

More specifically, Parsi and Petrosino, and Cañamero propose that emotions are a submechanism that exists in order to help the motivational decision mechanism to function more properly, by for example, modifying the intensity of the motivations, etc.. On the other hand, Gadanho uses emotions as reinforcement functions and for interrupting behavioural processes in order to deal with new and unexpected situations.

Most of the researchers consider emotions as a whole, that is, they define their role or functionality in their systems as a unique entity. Many of these researchers defend the position of views emotions as a limited set of dimensions. This set of dimensions defines an affective space which is used to appraise the emotion at each instant. This point of view minimizes the importance of distinctions among emotions [31].

One of the main differences of those approaches with the one presented in this paper is that the role of emotions is not define as a unique entity. On the contrary, it is proposed that each emotion has a particular role in the decision making process of the robot. Moreover, each emotion has its own appraisal process and is not related to a pre-defined event or situation (releaser event) as those authors do.

The work presented in this paper is implemented on a real social robot and, instead of studying the three cited emotions, authors focus on the fear emotion: its role, its appraisal process, and the learned emotional reactions.

This approach is not common in this research area, in fact, in relation to the treatment of the emotion fear, there are few works which study it as an independent emotion.

Herrera and Moffat [32] present an investigation of the role of fear as the anticipation of harm. According to them, dangerous is the situation that is appraised as potentially incompatible with some of our concerns. Nothing is dangerous per se, but in interaction. In this interaction we appraise this danger, and that is much dependent on each of us.

In their work, they use a simulated khepera robot in a prey/predator scenario. The goal is to maintain the robot around the target avoiding the predator. At first, the robot must be able to discriminate between the target object and the predator (to appraise the situation of being next to the target or next to the predator). In order to do this, they give valence to each object by observing the average motion of the robot. When the robot is being chased by the predator, which should result in a negative value, the average speed is high, while next to the target (positive value) is low.

Another interesting approach to the appraisal of fear is the one presented by Velásquez [33]. In his approach he considers natural (or innate) releasers and also, he included the capacity

of acquiring learned releasers. The natural releasers are for example, in the case of fear, situations in which the sensory systems would not work properly (dark environments), and the detection of archetypal predators. The learnt releasers correspond to the stimuli that they tend to be associated with and predictive natural releasers. For example, if a person *A* is naturally experiencing fear because of an unpleasant noise and, concurrently, *A* observes other person *B*, then, after several repetitions of this situation, person *A* will experience fear when he perceives person *B*.

In a later work, [34], Velasquez does not include any pre-wired cognitive elicitors, but rather allows them to be learned through emotional experiences, as the agent interacts with its environment. Actually, as introduced in section I, this is the same approach followed by the authors in this work: the releaser situation of the emotion fear (the dangerous situation) is learned by the robot using a novel mechanism to appraise that kind of situations.

IV. GENERAL ASPECTS OF FEAR

The present work is focused on the role of fear applied to an autonomous social robot. This research focuses on the same kind of fear that can be observed on people who have phobia about flying: they are terrified when they take a plane (even before) although the probability of a real plain crash is very low. Their fear cannot be controlled and their reactions are not intentional. Once they are on board, this hypothetical crash is not under their control, so they cannot make any decision to avoid it.

In general, in this work, two processes related to fear are covered: how the fear emotion emerges (the appraisal of fear already introduced in section II) and how it influences the robot's behavior (fear reactions).

Dual-process theories distinguish between intuition (fast and emotional) and reasoning (slow and controlled) as a basis for human decision-making. Related to the generation of fear, both approaches correspond with the appraisal and the evaluation of fear already presented in section II. Furthermore, reactions to fear can also be automatic or deliberative processes. Then, this dualistic approach is easily observed in fear. In some cases, fear is automatically elicited (mice are afraid of cats), but in others fear emerges due to a reasoning process (e.g. *due to the actual economic circumstances, I am afraid of loosing my job*). Moreover, this deliberative process affecting fear works as feedback to the intensity of fear (e.g. *if I loose my job, I will not get money, and then I will not be able to feed my family, and finally we will all die*). Besides, some responses to fear are automatically executed (a mouse runs away from a cat) but others are consciously carried out (e.g. *I have to work harder or I can be dismissed*).

Automatic processes are also referred as reactive processes by some researchers. Both terms, without distinction, can be used but authors prefer the automatic term.

Automatic processes, both for appraisal as well as for reactions, can be observed in animals. These are required for survival purposes. However, deliberative processes are specific of humans beings and this is one of our main characteristics.

The system proposed on this paper focuses on automatic aspects of fear: fear is automatically appraised once dangerous situations are identified; and the responses to fear are automatically selected because our system just considers the available information at that instant, there is not a model used to predict the effects in the future. Then, there is not reasoning behind our decision making system. The behavior is formed by selecting the most appropriate actions at each moment. Therefore, the whole process is an automatic process where deliberation is not involved.

Another classification of fear is related to the origin of the processes. Some fears are innate or inherited. This implies that these fears are more species specific than those acquired during life by experience. The latests correspond to learned fears which are considered in this work. Then, the robot learns both the releasers of fear and the proper reactions. Innate fears are considered by some researches as instincts which provide a key survival mechanism. Actually, animals without these instincts should have difficulties to reach adulthood.

In order to clarify the ideas presented in this section, several examples are presented in two tables: table I shows examples about how fear is generated; and table II presents examples of different reactions to fear. Tables I and II summarize the classifications of fear considered in this section. Each cell contains an example considering how the process has been acquired and how it is performed. Yellow cells correspond with the kind of fear the authors have implemented in this work. Red cells are those combinations which are impossible: something innate has been inherited so it is a species feature; in contrast, deliberation is a particular process of each individual; in consequence, deliberative-innate processes are not possible.

TABLE I
EXAMPLES OF DIFFERENT KINDS OF THE GENERATION OF FEAR

		How is it acquired?	
		Innate	Learned
How is it performed?	Appraisal (automatic)	Mice experience fear when they perceive the presence of a cat	Cats have fear when they perceive the presence of the veterinarian
	Evaluation (deliberative)		The global economic crisis

TABLE II
EXAMPLES OF DIFFERENT REACTIONS TO FEAR

		How is it acquired?	
		Innate	Learned
How is it performed?	Automatic	Mice escape from cats	Cats run away when they see the veterinarian
	Deliberative		I must work harder in order to keep the job position

In this work, fear is related to dangerous states where the robot can be harmed due to other's action, i.e. the robot will be afraid if it is in a situation where other agent can potentially damage it. Nevertheless, the action performed by the individual itself can also be harmful (imaging you walk a tightrope). In this case, these are risky actions and fear also comes up because of them (e.g. you are afraid of walking a tightrope). Risky actions have already been studied in virtual agents [35] and they will be considered in our robot in future works.

V. EXPERIMENTAL PLATFORM: THE SOCIAL ROBOT MAGGIE

The present work has been implemented and tested in the social robot Maggie. This robot and its control architecture are summarized in this section. For more detailed information, please refer to [36] [37] [14] [38].

Maggie (Figure 1) is a social and personal robot intended to perform research on human-robot interaction and to improve autonomy. It was conceived for personal assistance, entertainment, to help handicapped people, to keep people accompanied, etc. Its external friendly look facilitates its social robot task.



Fig. 1. Our social robot Maggie interacting with children

In relation to its hardware, Maggie is a computer-controlled system with a wheeled base which allows the robot to move through the environment. Its arms, neck, and eyelids movements provide Maggie with a life-like appearance. The vision system uses a camera in the head and, thanks to it, Maggie can recognize people and play several games. The laser telemeter and the ultrasound sensors are used by the navigation system to maneuver in the environment. By means of an infrared emitter/receiver, Maggie also operates different home appliances such as televisions or music players. Touch sensors on the surface of the body and a touch screen situated in the chest are used for a direct interaction with people. Moreover, two RFID antennas are used for identifying objects. In order to provide verbal interaction, our robot is equipped with a text-to-speech module and an automatic speech recognition system. In addition, verbal interaction is improved by a natural dialog system which provides Maggie with a natural oral communication.

Maggie was conceived as a fully autonomous robot that makes its own decisions. In order to achieve it, Maggie is controlled by an architecture developed by the authors' research group: the Automatic-Deliberative (AD) architecture. It considers two levels, the automatic and the deliberative levels.

In the AD architecture [39], both levels are formed by skills, which endow the robot with different sensory and motor capacities, and process information [40]. These skills are coordinated by the decision making system.

The proposed decision making system has a bidirectional communication with the rest of the control architecture. On

one hand, the decision making system selects the proper action which satisfies the most urgent need. This action is translated into a skill (deliberative or automatic one). On the other hand, the decision making system needs information from the environment in order to update the state of the robot and to assess the suitability of the skills activated. This information is provided by the sensors of the robot.

The proposed decision making system works in an automatic way. Deliberative processes look into the future to assess the effects of actions made in the present. Our system does not predict the future, but it considers the current state and the learned values from past experiences. This means that our decision making system is not a deliberative process but an automatic one which considers the knowledge acquired in previous experiences.

VI. THEORETICAL APPROACH TO OUR DECISION MAKING SYSTEM

The social and autonomous robot Maggie is endowed with a decision making system based on drives, motivations, emotions, and self-learning.

This system has been extensively described in previous works [13] [14] [15]. **Nevertheless, we will try to explain the system and some basic concepts: homeostasis, drives, and motivations.**

Homeostasis means maintaining a stable internal state [41]. According to the homeostatic approach, the human behavior is oriented to the maintenance of the internal equilibrium. The internal state can be parameterized by several variables, which must be at an ideal level. When the value of these variables differs from the ideal one, an error signal occurs: the drive [29]. In this approach, the drives are considered as the internal needs.

According to Hull [42], the drives increase the general excitation level of an animal and they are considered as properties of deficit states which motivate behavior.

The word *motivation* derives from the Latin word *motus* and indicates the dynamic root of behaviour, which means those internal factors, rather than external ones, that urge the organism to take action [43]. In other words, the motivational state is a tendency to correct the error, i.e., the drive, through the execution of behaviours.

The aim of the presented decision making system is to achieve a fully autonomous robot which learns to make decisions in order to maintain its needs within an acceptable range (**homeostasis**). For this purpose, it uses a reinforcement learning algorithm in order to decide which action is the most suitable in every state. In this work, the well-known Q-Learning algorithm [44] is employed, so the $Q(s, a)$ values must be computed. The proposed reinforcement function is the variation of the wellbeing of the robot. This variation is produced **by the passing of time and by the effects of the execution of an action.** The wellbeing is defined as a function of its drives, and it measures the degree of satisfaction of its internal needs. That is, the ideal wellbeing is defined as the lack of needs or drives, i.e. the maximum wellbeing is reached when all drives are satisfied (the value is 0). Therefore,

when drives increase, the wellbeing falls. Mathematically, it is expressed as:

$$Wb = Wb_{ideal} - \sum_i \alpha_i \cdot D_i, \quad (1)$$

where α_i are the ponder factor that weight the importance of each drive D_i in the wellbeing of the robot; and Wb_{ideal} is the ideal value of the wellbeing of the robot which has been set to 100. In this work, all drives have the same importance, then, α_i is 1 for all drives.

In order to determine the state of the robot, the internal and external states are required. The latest is defined by the relation of Maggie to the objects of the environment. The former, the internal state, is established by the *strongest* motivation, which is called the dominant motivation.

In order to calculate the intensity of each motivation, we have considered the idea of Hull [42] with states that motivation is determined by two factors: the associated drive and the incentive, that is, the presence of an external stimulus that predicts the future reduction of the need, as shown in (2),

$$\begin{aligned} \text{If } D_i < L_d \text{ then } M_i &= 0 \\ \text{If } D_i \geq L_d \text{ then } M_i &= D_i + w_i \end{aligned} \quad (2)$$

where M_i is a motivation, D_i is the related drive, w_i is the related external stimulus, and L_d is called the activation level. It has a constant value of 10. If a drive is under this value, the related motivation will not be considered for the dominant motivation.

According to this model, the general idea is that we are motivated to eat when we are hungry, and also, when we have food in front of us, although we do not really need it.

It is important to note that in this decision making system, there are no predefined motivational actions. This means that the robot does not necessary know in advance which actions, or skills, to select in order to satisfy the drive related to the dominant motivation. There is a repertory of actions and they can be executed depending on the relation of the robot with its environment, i.e. the external state. For example, the robot is able to interact with people as long as it is accompanied by someone.

Again, the main goal of the robot is to learn the right relationship between states and actions. That is, to learn the best action to execute in every state in order to maximize its wellbeing (by satisfying its drives).

In summary, the decision making loop follows the next steps: first, it determines the state (internal and external); then, an action is selected (randomly selected during learning or the best one when learning is over); and finally, the state-action evaluation ($Q(s, a)$), using the wellbeing variation as the reinforcement, is accomplished.

VII. SETUP OF THE DECISION MAKING SYSTEM

In this section, the configuration and parameters of the decision making system of the robot are presented.

The proposed system can be fine-tuned in order to adapt it to different requirements. This paper presents a configuration suitable to test the effectiveness of fear. However, all values and parameters can be changed and different policies of behavior will be obtained in future experiments. This can be seen as the personality of the robot.

A. Internal state: drives and motivations

As expressed by equation (2), the intensity of each motivation is affected by two factors: internal needs and external stimuli. The internal needs are the drives, and their values depend on inner parameters. The external stimuli are the objects situated in the environment altering the robot motivations.

The selected drives in this work are:

- **Energy:** this drive is necessary for survival.
- **Boredom:** the need of fun or entertainment.
- **Calm:** the need of peace.
- **Loneliness:** this is the lack of social interaction.

All these drives represent the deviation from the ideal state. This ideal state corresponds to the value zero for all drives.

Since we want Maggie to be an autonomous social robot, based on past works [45] and experiences, and considering the drives defined (each motivation is connected to a drive), the motivations that have been considered are:

- **Survival:** it refers to the energy dependence. This motivation is connected to the need of *energy*. Then, the *survival* motivation is the most critical one. This is the major requirement to be achieved by an autonomous robot.
- **Fun:** this motivation is related to entertainment purposes and its associated drive is *boredom*. The *fun* motivation refers to the need of entertainment of the robot itself. This means that this drive can be satisfied when Maggie is having fun and this is achieved when it is dancing.
- **Relax:** it is linked to a peaceful environment and it is related to the drive *calm*. In contrast with *fun*, *relax* is its counterpoint: it searches for noiseless conditions.
- **Social:** it corresponds to the need of human-robot interaction. It is associated to the *loneliness* drive. As presented in section V, Maggie is a social robot so one of its main goals is to stablish relationships with people. This attitude is enforced by this motivation.
- **Fear:** this motivation arises in dangerous situations and it guides the robot toward a secure state. More details are presented on section VIII.

All motivations have been defined considering that Maggie is a social robot designed to interact with and move among people. Then, its behaviors have to be as natural as possible, i.e. they have to be comprehensible by humans sharing the environment with the robot.

The use of fear as a motivation in our robot is one of the novel ideas presented in this paper. As you can see, *fear* is treated in a different way than the other motivations. Fear is considered a motivation but there is not a drive related to it because fear does not represent a deficiency in any need. However, it is able to lead the robot's behavior.

1) *Dynamics of drives and motivations:* In a similar way to any need on humans or animals, drives fluctuate. A person is not hungry just after having lunch, but hunger increases as time passes. After we eat and the digestive process has begun, the need of energy is inhibited due to satiety signals. These satiety signals slowly dissipate until the hunger again takes over. Then, drives vary according to several signals and parameters [46]. Drives in our robot evolve in an analogous way. The evolution functions of drives are set by the designer and they affect the behavior of the robot. Since drives temporally evolve from scratch, motivations do as well.

Figure 2 shows the dynamics of all drives. The evolution functions for all drives do not have to be all equal. In fact, each drive fluctuates according to different functions.

Drives evolution is determined by three factors: the increasing function, the satisfaction time, and the saturation level. Following, each of this components is explained for each drive.




Fig. 2. Comparison of drives progression.

a) *Increasing Functions:* In the implementation proposed in this paper, *boredom*, *loneliness*, and *calm* drives linearly increase but with different slopes. It means that, as time goes by, these drives become bigger and bigger, and so do the corresponding motivations.

Considering that being social is one of the main characteristics of our robot, interaction with people is one of the most relevant aims. Therefore, the *loneliness* drive is the fastest one. This means that the motivation associated to this drive, *social*, more frequently compete to be the dominant one. Consequently, the behaviors learnt for this motivation are exhibited more often. It ends up with a robot whose main behavior is the one related to human-robot interaction.

The *boredom* drive goes after. This is because authors conceived Maggie as a nice robot for people and, a robot having fun is more attractive than a passive one. The *fun* motivation leads the robot to perform enjoyable reactions.

Finally, *calm* evolves slighter so it is the slowest drive. This implies that it is harder to exceed its activation level in order to struggle for being the dominant motivation. In addition, this drive just evolve when music is been played: Maggie needs to relax after it has been listening music for a while. All in all, this provokes that the *relax* motivation scarcely becomes the dominant motivation.

In order to achieve a fully autonomous robot, power autonomy is the first step. Therefore, the most relevant inner need, due to the necessity of survival, is the *energy* drive.

Therefore, this drive evolves as battery level varies. Then, its value matches the battery level.

As said before, the *fear* motivation is different. Theoretically, there is not drive linked to this motivation. However, from a computational point of view, a drive needs to be linked to the *fear* motivation. *Fear* drive value will be risen to the maximum at once when a dangerous situation is detected. This can be seen on the right of Figure 2. These dangerous states are not predefined but they are learned by the robot itself through interaction with the environment. The appraisal of the fear emotion is detailed on section VIII. This is another novel idea presented in this work.

b) *Satisfaction times*: After a drive is satisfied, it does not immediately start evolving, there is a *satisfaction time* before it evolves again. The same idea occurs to human beings: once we have eaten, we do not feel hungry again but it takes some time before we need to eat again. The satisfaction time has been empirically set and they are summarized on table III. At the very beginning of Figure 2, satisfaction times can be observed.

TABLE III
SATISFACTION TIMES FOR ALL DRIVES

Drive	Satisfaction time
energy	-
boredom	30s
calm	30s
loneliness	60s
fear	-

Since *energy* mirrors battery level, it does not make sense to consider its satisfaction time.

Besides, considering the previous definition of the *fear* motivation, satisfaction time does not make sense in relation with fear.

c) *Saturation levels*: In order to avoid an unstopped increase in the value of the drives, a saturation level is defined for each one: once a drive has reached its saturation value, it does not exceed this value. Different drives have different saturation values which affect the dominant motivation in case of a never-ending expansion of the drives. This can be seen as an emergency control mechanism in case that several drives are saturated and their motivations compete to be the dominant one. In this situation, there are predefined priorities that determine the dominant motivation in those exceptional situations. These priorities can be seen as inherited knowledge or instincts in living beings which allow them to face extreme situations. Table IV presents the sorted list of saturation levels.

TABLE IV
SATURATION LEVEL FOR ALL DRIVES

Drive	Saturation level
energy	20
fear	19.9
loneliness	17.8
boredom	17.7
calm	17.6

In our implementation, *energy* has the highest saturation level because it is the most urgent since it is related to survival: if the energy drive is saturated it means that the battery level is really low and it is critical to get the batteries recharged.

Fear is the second one and it is over the rest of drives. As explained before, when a dangerous situation is perceived, the *fear* value is set to its maximum, which corresponds to the saturation value. This value is over the others because *fear* represents a really dangerous situation which must be avoided somehow as soon as possible. Just survival can be more urgent than *fear*.

The rest of the saturation values were fixed considering the same reasons used for the evolution functions of the drives.

2) *External Stimuli*: According to equation (2), the motivation value is the summation of the drive and the external stimulus. Just like human beings or animals can feel thirsty when they see water, the motivations are influenced by objects in the world. These are called the external stimuli for motivations. These stimuli have more or less influence: their values depend on the states related to the objects.

In the implementation presented here, the external stimuli are listed on table V. When music is been played, the robot perceives it and the motivation to have fun increases. If Maggie perceives the docking station, survival motivation is promoted. Lastly, the presence of a person close to the robot strengthens its social motivation. **All the external stimuli have been empirically set to the value of 2.**

TABLE V
ALL EXTERNAL STIMULI USED IN THIS WORK

Motivation	External stimuli	State related to ext.stim.
fun	music	listening
survival	docking station	plugged
social	any person	close

B. The robot in the world: sensing and acting

The world is sensed by the robot in terms of objects and its related states to these objects. Objects are not limited to physical objects but to abstract objects too.

The world where Maggie is living in is limited to the laboratory and the objects the robot perceives and interacts with are: a music player, the music in the lab, the docking station for supplying energy, and the people around the robot. The states related to all these items must also be defined and the transitions from one state to another are detected by several skills running on Maggie.

Maggie interacts with the world through the objects and their potential actions. These actions are implemented as skills in the AD architecture. Following, each item and its related actions are introduced.

1) *Music player*: In order to operate the music player, the robot has to be close enough to it. Three different states have been considered: *far*, if the robot is in a position where it is not able to operate the player; when the robot is close to the player and is already working, the state is *near-on*; and, finally, when the robot is close but the player is off the state is *near-off*.

These two states, *near-on* and *near-off*, are required to avoid sending the same command twice to the player. **The data required to determine the position related to the player is provided by the geometrical navigation system. Controlling remote infrared home appliances (such as the music player) is achieved by an infrared emitter/transmitter located at Maggie’s belly. All details can be read at [38].**

Then, the possible actions with the *music player* item are:

- Go to player: Maggie approaches the music player to operate it.
- Play music: music is played at the player when it is off. This action produces a change of state in other object: the music, from *non-listening* to *listening*.
- Stop music: music is stopped when it is being played. This action produces a change of state in other object: the music, from *listening* to *non-listening*. This transition keeps a peaceful atmosphere.
- Idle: it just represents the possibility to remain next to the player for a while.

2) *Music*: The robot’s environment is the lab, and *music* can be playing there. Then, the robot can be *listening*, or *not*, to *music*. Just when the robot is *listening* to *music*, it is able to *dance*. If *music* is mute, it cannot *dance*. **An infrared emitter is used to play/stop the music when Maggie is close to the player.**

About the music, there is just one possible action:

- Dance: the robot moves its body with the music. This action is just executed when Maggie is *listening* to music. This action can be executed at every place inside the lab because the music is perceived from anywhere in the room.

3) *Docking station*: The *docking station* is the source of energy. If the robot is *plugged*, the battery is charging; otherwise, the robot is *unplugged* and the battery level decreases. **In order to find the docking station, the robot relies on the navigation system and the information from the laser telemeter. Eventually, to determine if it is plugged or not, a data acquisition device is in charge of reading the battery data.**

The attainable actions with the docking station are:

- Charge: Maggie approaches the docking station, plugs to it, and stays there until the battery is full. At the end of this action the robot is *plugged* and the battery is recharged.
- Remain: it holds plugged for a while.

4) *Person*: Our robot is intended to interact with people. Hence, people are considered as “objects” of the environment. Regarding interaction, a person has to be close enough to touch, speak or recognize. For that reason there are two states in relation to a person: *present* and *absent*. **These states are determined by merging two signals: bluetooth form the mobile phones and personal RFID tags.**

The *person* item offers just an available action:

- Interact: with this action the robot is not executing any particular ability but it is evaluating the action that the person is carrying out. This means that the value of the *interact* action captures the effect of other agents’ actions

over the robot’s wellbeing. These are exogenous actions because they are executed by other agents.

An overview of the robot’s world is displayed in Figure 3. It provides a good perspective of the scenario and the objects the robot interacts with during the experiments.




Fig. 3. Overview of our lab and objects the robot interacts with

Once an action is selected and executed, it causes effects over the drives. When an action has successfully ended, its effects are applied. If an error occurs during the execution of a skill or its result is not satisfactory, this situation is notified and its effect is not applied. In our experiments, the effects of actions can be positive or negative in terms of robot’s drives. A positive effect reduces the value of a drive, this is, it reduces a robot’s need (this **generally** implies an increase on robot’s wellbeing). Actually, when the drive is set to zero, it is said that the action satisfies the drive. Some actions can also “damage” some drives of the robot increasing their values (**so in most cases** the robot’s wellbeing drops). All effects are shown in table VI.

TABLE VI
EFFECTS OF ACTIONS

Action	Object	Drive	Effect
stop cd	cd player	calm	set to 0
dance	music	boredom	set to 0
positive interaction	person	social	set to 0
negative interaction	person	social	+10

In the current implementation the effect of an action is applied only to one drive. Nevertheless, in future works it could happen that, for example, dancing affects the Energy drive or interacting with people satisfies also Boredom.

VIII. HOW FEAR EMERGES. THE APPRAISAL

The fear emotion can be considered as an adaptive response to threatening situations [17]. As commented in section IV, some of these threatening situations are innate, but others are learned. In this work, authors have aimed their attention at

learned releasers of fear. Therefore, this section exposes how a dangerous state is detected (appraisal of fear) following a learning process.

The idea is similar to what happens when a person A beats person B . Since this fact causes an intense emotional experience, even if A has just sporadically hit B , B will remember this situation and its dramatic effects for long time. Therefore, whenever A is close to B , B relives this situation and evaluates its possible consequences. The final result is that B is afraid of A .

Threatening situations, or dangerous states, are those where significant harm can be caused.

In our approach, this harm is inflicted by actions performed by active objects, which are capable of executing actions themselves. According to our implementation, these active objects are the people.

As already introduced in Section I, these actions, which are not executed by the robot but by other object from the environment, are called exogenous actions. Exogenous actions leads to complex domains. These domains are quite hard to model because exogenous actions are difficult to foresee.

Besides, the exogenous actions (and their effects) are mixed with the actions executed by the robot and their effects. For example, considering the example presented at the beginning of this section, if person B is walking and A hits B , B suffers pain but, using Reinforcement Learning, B does not know if this negative reward, the pain, is because of the walking or because A hit him. Therefore, in a reinforcement learning framework, the reward of an action executed by the robot could be altered by an exogenous action. A key issue is to undoubtedly identify the effects caused by the actions of the robot and the effects of the exogenous actions.

In this first implementation of *fear* in our robot, in order to distinguish the effects of robot's actions from the effects of exogenous actions, our implementation considers the exogenous actions when the robot is executing the *interact* action. This action does not have any effect, so the resulting effects during its execution, can be certainly assigned to an exogenous action. This is, the robot estimates how good the other agent's actions (exogenous actions) are because all the effects and transitions are due to these actions. The performance of the exogenous actions is also measured using the variation of the robot's wellbeing. Therefore, the exogenous action, somehow, implies an agent-robot interaction.

Three different processes are involved in the generation of fear:

- Storing the worst experiences
- Detecting new dangerous states
- Updating the fear motivation

Some of these processes can occur in parallel.

A. Storing the worst experiences

In short, in our approach, *fear* is appraised considering the states where the exogenous actions, few times, have damaged the robot. Considering the definition of the robot's wellbeing (equation (1)), dangerous states are identified as the states where exogenous actions, which caused a considerable decay

on the wellbeing, have been executed few times with negative consequences. Usually, these states correspond to situations where the robot, in general, is not damaged but, sporadically, some adverse exogenous actions cause harm to the robot. As previously stated, using the standard Q-learning, in this situation the utility value, the Q value, is still high. This causes that the robot will not learn to avoid this situation. For this reason, the worst Q values must be stored in order to remember the worst experiences at each state. This is similar to animals which remember their worst experiences and relive them when they are facing the same situation. Then, the dangerous states are determined using the worst value of all available exogenous actions. For this reason, in addition to the Q values, the Q_{worst} values have been computed according to the next equation for each iteration:

$$Q_{worst}^{obj_i}(s, a_{exog}) = \min(Q_{worst}^{obj_i}(s, a_{exog}), r + \gamma \cdot V_{worst}^{obj_i}(s')) \quad (3)$$

where obj_i is an active object which is able to carry out actions, a_{exog} is the exogenous action executed by obj_i from state s , the resulting state is s' , r is the reward corresponding to the variation of the robot's wellbeing, and γ is the discount factor. $V_{worst}^{obj_i}(s')$ means the best $Q_{worst}^{obj_i}$ value from the new state and it corresponds to:

$$V_{worst}^{obj_i}(s') = \max_{a \in A_{obj_i}^{exog}}(Q_{worst}^{obj_i}(s', a)) \quad (4)$$

$V_{worst}^{obj_i}(s')$ computes the best possible action among the $Q_{worst}^{obj_i}$ values from the state s' . In other words, it stores the value of the least harmful action from the new state.

The states considered for the appraisal of fear correspond to the external state of the robot. This is, the state related to objects in the world. This is because it is considered that a state is dangerous independently of the internal state. For example, in humans, if you are afraid of spiders, you will experience fear if you see a spider, independently of any internal need; i.e. it does not matter if a person suffering arachnophobia is hungry or thirsty, he is terrified when he sees a spider. Likewise, the states during the appraisal of fear are just related to the objects in the world.

Olteanu [47] states that the evaluation of internal and external situation is a crucial process for the appraisal of emotion. In this work, the variation of robot's wellbeing and external state of the robot are involved in the appraisal of fear.

B. Detecting new dangerous states

Once the Q_{worst} values are computed, they are used to identify the dangerous states. These dangerous states are recognized by the robot itself, so they are not pre-programmed in advance.

A state is considered as a dangerous situation when its $Q_{worst}^{obj_i}$ value is below a certain threshold: L_{danger} . On the contrary, this is considered as a safe state. Mathematically, it is expressed like:

$$\begin{aligned}
&\text{If } Q_{worst}^{obj_i}(s, a_{exog}) < L_{danger} \Rightarrow \\
&\quad s \text{ is a dangerous state; } \forall s \in S_{obj_i} \\
&\text{If } Q_{worst}^{obj_i}(s, a_{exog}) \geq L_{danger} \Rightarrow \\
&\quad s \text{ is a safe state; } \forall s \in S_{obj_i}
\end{aligned} \tag{5}$$

where S_{obj_i} is the set of all states related to object i .

C. Updating the fear motivation

As explained before, in this work, fear is considered as a motivation which is able to govern the robot's behavior. Once the dangerous states are identified, the *fear* motivation is able to be the dominant one and to lead the robot's behavior. Whenever the robot transits to a dangerous state, *fear* emerges. In a formal way, if s is the current robot's state, the fear value is updated according to the next equation.

$$\begin{aligned}
&\text{If } s \text{ is a dangerous state} \Rightarrow \text{Fear} = \textit{high} \\
&\text{If } s \text{ is a safe state} \Rightarrow \text{Fear} = \textit{low}
\end{aligned} \tag{6}$$

High and low values of fear correspond to the presence and to the absent of fear respectively. Their corresponding numerical values are 19.9 and 0 and the comparison of *fear* with other drives was shown in Figure 2.

Behaviors to avoid states that harm is recurrently provoked from can be directly learned by reinforcement learning, since their Q values are low. However, sporadic danger from a particular state cannot be managed by those algorithms. The proposed mechanism for the appraisal of fear has been specifically designed to consider these dangerous situations. All in all, our appraisal of fear perfectly deals with both circumstances.

IX. EXPERIMENTAL RESULTS

This section validates and analyzes the use of fear in our social robot. As previously exposed, fear has been considered as a motivation which incites the robot to behave. The experiment consists of comparing the performance of our robot with and without fear as a motivation. For both cases, there are two phases: first, the exploring session where Maggie tries every action in order to learn the right policy to act; then, the exploiting session where the learned policy is employed to decide how to behave. During exploitation, learning is frozen and the best action is always selected.

All the experiments have been achieved by real robot-environment interaction and this interaction has been evaluated according to its effect over the robot's wellbeing. Previous knowledge has not been given to the robot in advance, so it has learned from the ground up. The learning algorithm applied in this work is a variation of the Q-Learning algorithm: the Object-Q-Learning algorithm, which has been already presented in previous works [15]. **The fact that previous knowledge is not given in advance to the robot implies that all the Q-values have the same initial value. In this experiment is set to 1.**

In our experiments, an iteration corresponds to the execution of an action by the robot. The robot decides at each iteration the action to be selected. The probability of an action to be

selected is determined by the Q value associated to this action for the current state, and by the level of exploration. During the exploring sessions, all actions have the same probability of execution. This is required in order to guarantee that all actions are tried many times from all the possible states. At exploitation, the best action is always selected. The best action is the most convenient in terms of the robot's wellbeing. That is, action a is the best action when the robot is at state s if the $Q(s, a)$ value is the highest one among all available actions.

The interactions between the robot and the environment, where experiments are accomplished, take a considerable amount of time. Hence, the learning phase has been established around 700 iterations, which lasts more than seven hours **split in a couple of days**.

As justified in [48], at some point, exploring must stop and the learned values must be exploited. Considering this approach, after 500 iterations, learning starts to progressively decrease until it is totally suspended. After this point, the Q values do not change anymore and the best policy is always selected.

Since this work has been implemented in a social robot intended to interact with people, the object *person* has been considered as the active object which shares the environment with Maggie and interacts with it. Then, the exogenous actions are those actions executed by the people around Maggie. The exogenous actions affect the external state as well as the internal state of the robot. For example, when a person approaches Maggie, the state related to this person (the external state) has changed, and it is not due to the robot's actions. Moreover, the actions accomplished by a person may affect some robot's drives (the internal state): e.g. if a person hits the robot, the *social* drive soars, i.e. the need of a positive social interaction increases. Again, all these consequences are not caused by the robot but by the people's action (the active objects' action).

In these experiments, two people interact with the robot: *Alvaro* and *Perico*. Both alternatively approach Maggie, one by one. *Perico* always interacts with positive actions: he strokes the robot or he says compliments to Maggie. This results on the satisfaction of the social drive, which is set to 0. *Alvaro* generally acts in a positive way too. However, sometimes, he hits or offends Maggie. This is reflected in the robot's wellbeing through an increment of ten units in the *social* drive (equation (7)). In general, both users benefit the robot but *Alvaro* occasionally causes harm to it.

$$\text{If the robot is harmed} \Rightarrow D_{social} = D_{social} + 10 \tag{7}$$

The hits and strokes to the robot are perceived by the tactile sensors in its body surface. The compliments and insults are processed and interpreted by the robot's dialog system [49].

Next, several results are presented in order to show the goodness of fear in an autonomous social robot. First, the results of the appraisal of fear are analyzed. Then, how fear has influenced the robot's behavior is commented. Later, the adaptability of our method is demonstrated by comparing the different learned reactions to fear depending on the user's

behavior. Finally, the usefulness of fear and its advantages are proved.

A. Results on the appraisal of fear

During the experiments, considering that the maximum “punishment” of a negative exogenous action corresponds to a penalty of ten points to the *social* drive, and based on observations during trials, L_{danger} has been set to -10 points. As a consequence, whenever the robot is in a state where there is a $Q_{worst}^{obj_i}$ value below this threshold, this is considered as a dangerous state. Therefore, the *fear* motivation suffers a drastically increment in its value as shown in equation (6).

As already said, the consequences of the actions executed by both users (*Alvaro* and *Perico*) over the robot’s wellbeing are perceived by Maggie. In order to do it, Maggie is endowed with an action called *interact*. This action makes Maggie to do nothing, therefore, it is possible to evaluate how the other’s actions affect it. Thus, translating equation (5) into our experiment, it results on equation (8).

$$\begin{aligned} \text{If } Q_{worst}^{Alvaro}(s, interact) < -10 &\Rightarrow \\ s \text{ is a dangerous state; } \forall s \in S_{Alvaro} & \\ \text{If } Q_{worst}^{Perico}(s, interact) < -10 &\Rightarrow \\ s \text{ is a dangerous state; } \forall s \in S_{Perico} & \end{aligned} \quad (8)$$


Figure 4 depicts the evolution of the worst Q values associated to the exogenous actions. As can be seen in Figure 4(a), since all *Perico*-Maggie interactions are favorable from a robot’s point of view, its Q_{worst}^{Perico} value slightly decreases from its initial value 1, and it remains stable around value 0. In contrast, the Q_{worst}^{Alvaro} value associated to *Alvaro*’s interaction is significantly reduced (Figure 4(b)). This is due to the number of interactions where *Alvaro* has hit or offended Maggie. This number is low in comparison to the total amount of interactions: during the learning phase, *Alvaro* harmed Maggie five times of thirty-seven interactions (13’5%).

Note that the exogenous actions have been executed when a person is *present*. Therefore, considering equation (8), the worst Q values are associated to the state s when $s = Alvaro\ is\ present$ or $s = Perico\ is\ present$. Naturally, if a person is absent, his actions do not interfere on the robot’s “life”. Therefore, potentially dangerous states are *Alvaro is present* and *Perico is present* because our robot can be damaged from them.


Looking into Figure 4, the robot does not know anything about dangerous states, or what to be afraid of, until iteration 182. At this point, *Alvaro* hits the robot one more time, and $Q_{worst}^{Alvaro}(present, interact)$ reaches the value -11.2097 . This value is under the selected threshold (-10) and, therefore, the robot determines that being next to *Alvaro* can be harmful. From this iteration on, if *Alvaro* is close to the robot, this is identified like a dangerous state and, as a result, the *fear* motivation is rocketed. Consequently, *fear* potentially becomes the dominant motivation, so it guides the robot’s behavior.

B. How fear influences the social motivation

As presented in section VII-A, the social motivation is related with the need of positive human-robot interaction.



(a) Q_{worst} value for action executed by Perico



(b) Q_{worst} value for action executed by Alvaro

Fig. 4. Q_{worst} values of exogenous actions.

Interaction with people has a great positive effect over this motivation and this is achieved through the interaction with *Alvaro* and *Perico*. Then, these actions are the most suitable skills to be executed: this is the reason because the highest Q values among all actions, when *social* is the dominant motivation, correspond to *interact with Alvaro* and *interact with Perico* (see Figure 5).

The interaction with *Alvaro* must be detailed. *Alvaro*’s actions are, most of the times, favorable. Despite of the small percentage of hurting actions, the final interaction-with-*Alvaro*, its Q value, is quite high. However, the small number of hurting actions are enough to scare Maggie. Maggie is afraid of *Alvaro* because of the few negative interactions, which cause a 10 points penalization on its *social* drive. The upper plot in Figure 5 depicts the evolution of the $Q^{Alvaro}(Alvaro\ is\ present, interact\ with\ Alvaro)$ value when *social* is the dominant motivation. Around the iteration 100 and 180, this value decreases because there has been an important decrement on the robot’s wellbeing due to negative interactions. This is enough for Maggie to detect and remember the dangerous situation: hereafter, whenever *Alvaro* is close to the robot, this situation is appraised as a dangerous state, and the *fear* motivation value exceeds the *social* motivation value. Therefore, whenever *Alvaro* is present, the *social* motivation is not the dominant one again, and this Q value is not updated anymore.

C. Learned fear reactions: escaping

As previously shown, the proposed system is able to identify the dangerous situations which have not been previously defined. Using the learning mechanism, the robot determines what action must be selected to avoid these situations. In this




Fig. 5. Learnt Q-values associated to exogenous actions when **social** is the dominant motivation

experiment, the users (*Alvaro* and *Perico*) approach Maggie, one by one, and stay there. At that point, since Maggie is accompanied, it must decide to interact or to execute another action. In this work, dangerous states are associated to the presence of *Alvaro*. Then, the robot learns how to “escape” from *Alvaro*. This can be explained by observing Figure 6 where the results are shown. The actions which imply a displacement on the geometrical position of the robot are *go to player* and *charge*. The former moves the robot towards the *cd player* and the last gets the robot plugged to the docking station. Both actions make the robot to move away from *Alvaro*. As can be observed, these two actions are the most appropriated ones when *fear* is the dominant motivation since they have the highest values. Therefore, when the robot is scared (*Alvaro* is beside Maggie), it moves to the *docking station* if it is close to the *cd player*, or to the *cd player* if it is plugged. This can be seen as a run-away behavior learned by the robot itself and it is similar to what animals do when they are afraid.

D. Learned fear reactions: freezing

Since humans are unpredictable autonomous agents, different reactions to fear can be observed depending on the person involved in the situation.

In the results presented on Figure 6, both users alternatively approach Maggie with the intention of achieving some human-robot interaction. Recalling, *Perico* always achieves positive human-robot interactions, and *Alvaro*, once in a while, causes harm to Maggie. As a consequence, Maggie is afraid of *Alvaro* and, as exposed in the previous section, it learns to escape from him.

However, our system is flexible enough to learn different behaviors according to diverse people’s attitude. In this experiment, users have been trained to behave in a different way: now, *Alvaro* and *Perico* separately approach Maggie and they chase the robot. Users will leave when they get bored due to the robot’s inactivity. Again, *Alvaro* occasionally damages Maggie. Considering these damages, *fear* comes out on Maggie when *Alvaro* is present.

A new learning session has been conducted, similar to the previous ones but with *Alvaro*’s new behavior. Results can be

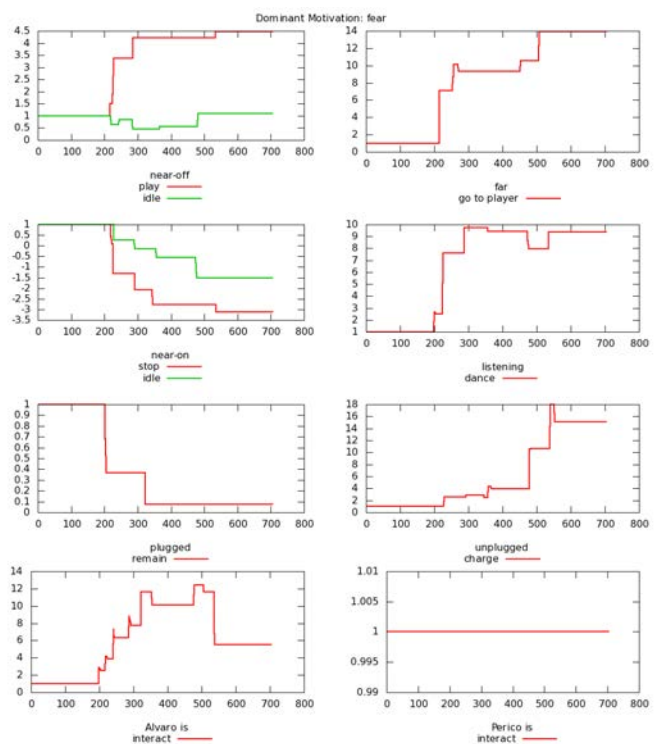


Fig. 6. Learned Q values when **fear** is the dominant motivation.

observed on Figure 7. In this case, the behavior learned when *fear* is the dominant motivation is related to the *idle* action, when Maggie is close to the music player (both, with music on and off), and to the *remain* action, when it is plugged. This is because the Q values associated to these actions are the highest ones (upper three plots on the left column of Figure 7). These actions share that they cannot be externally perceived because they do not make any expression or movement, they give the impression of inactivity. Therefore, the robot bores *Alvaro* and he moves away from Maggie. After this happens, *fear* ceases resulting on the following benefit for the robot.

Summarizing, in this experiment the cause of fear has not been changed (the presence of *Alvaro*) and it has been perfectly identified again. However, the reaction to fear is totally different. Our method nicely works with users conducting in diverse manners and the proper avoidance behavior is learned in each situation.

The new learned behavior dealing with fear can be justified considering that some animals paralyze when facing a dangerous situation. It seems that they are “frozen” by fear.

E. Does Maggie need fear?

This last section of the experiments tries to justify the use of *fear* as a motivation. Here, the performance of the robot, **during the exploiting phase**, is measured and compared with the results obtained from experiments where fear does not exist. In this section, the same motivations considered in previous experiments are employed (all motivations introduced on section VII-A).

Two different learning sessions have been realized, both using reinforcement learning algorithms. First, the robot learns




Fig. 7. Learned Q values when **fear** is the dominant motivation. Alvaro chases the robot until getting bored or interacting with Maggie.

to behave without considering *fear* as a motivation. In consequence, the motivations present in this session are: *survival*, *fun*, *relax*, and *social*. In the second session, the same four motivations are considered plus the *fear* motivation.

During both sessions, the robot learns the right policy to satisfy its needs. However, the session considering *fear* learns an additional behavior in relation with this motivation (escape or freeze). Each learned policy is used during an exploiting session. These exploiting sessions last around 80 minutes and the best action is always selected at each iteration.

The learning and the exploiting sessions are performed in the same environment considering our two well-known users: *Alvaro* and *Perico*. In this case, the users approach Maggie and remain there until **the user or Maggie moves away**.

In order to evaluate both configurations, the results obtained during both exploiting sessions are compared. The performance indicators considered in previous works [50] have been employed: the average wellbeing and the percentage of permanence in a certain security zone. Besides, in this paper, authors present also the results about the percentage of time the robot is without a dominant motivation (all drives are below the activation level L_d presented on section VI) that gives an idea about how *comfortable* the robot is. Finally, the number of times the robot is harmed is also compared.

1) *Average wellbeing*: Since the variation of the wellbeing was used as the reward during the learning phase, the robot tends to maximize it. Table VII presents the values corresponding to the average wellbeing with and without *fear* during the exploiting sessions. The average wellbeing when *fear* does not exist is slightly higher. This can be seen as a disadvantage of

using *fear*. However, this is understandable considering that, when *fear* is included as motivation, the number of drives used to compute the wellbeing is bigger, so the wellbeing value is lower.

This drawback can be observed in nature too: a fearful person is not in a pleasant situation, his wellbeing decreases due to the anxiety suffered because of the fear. As a consequence, the person is distressed while he is afraid. However, other benefits can be obtained from *fear*.

TABLE VII
AVERAGE WELLBEING DURING EXPLOITING SESSIONS

without fear	with fear
87.77	86.72

2) *Permanence in the secure area*: These benefits are related to other reliable performance rate: the percentage of time the robot’s wellbeing remains in a security zone. If the robot’s wellbeing is within this area, it can be said that the robot is “fine” because its wellbeing is high. Thus, the percentage of time the wellbeing remains in this area gives an idea about how well the robot is performing.

In order to establish the limits of the secure area, the ideal wellbeing value (100) and the activation levels for motivations (10) are considered. Since all drives simultaneously evolve and several motivations can compete for the dominance, the security area width was set to 15. Consequently, it is considered that when the robot’s wellbeing is between 100 and 85, it is within the secure area.

Table VIII shows the percentage of permanence within the secure area **during the exploiting sessions**. As can be seen, when *fear* is included as a motivation, the wellbeing is almost the 70% of iterations within the secure area, which represents a 5% more than when *fear* is not used. This is coherent because *fear* is used to avoid dangerous states where the robot can be damaged. Once the robot is harmed, the wellbeing decreases enough to move out the secure area.

TABLE VIII
PERMANENCE WITHIN THE SECURE AREA DURING EXPLOITING SESSIONS

without fear	with fear
65%	69.5%

3) *Non dominant motivation*: Moreover, if there is not a dominant motivation, it means that all the internal needs and external stimuli are not strong enough to induce a behavior. Hence, it can be considered that the robot is in a comfortable situation. The percentage of time that a dominant motivation does not exist, **during the exploiting session**, proves how pleasant the robot’s “life” is. Table IX shows that considering *fear*, the 78% of the time there is not dominant motivation. On the other hand, when the robot lives without *fear*, the percentage is reduced to 72%. Once again, these numbers show how *fear* provides a better quality of “life”.

4) *Number of times the robot has been damaged during the exploiting session*: The most relevant result of using *fear* is related to the damage caused by *Alvaro* to the robot **when it**

TABLE IX
PERCENTAGE WITHOUT A DOMINANT MOTIVATION DURING EXPLOITING SESSIONS

without fear	with fear
72.22%	78%

“lives” according to the learnt policy of behavior. When *fear* is not implemented, the robot tries to interact with both users in order to satisfy its social need **since it has not learnt to identify that being next to Alvaro is dangerous**. This action leads Maggie to, some times, be harmed by *Alvaro*. As depicted in table X, this happens six times of twenty-three interactions between Maggie and *Alvaro*.

Now, considering *fear* as a motivation in our system, the robot does not interact with *Alvaro* at all, **so he could not hurt it. This is because, as shown in previous sections, the robot learnt to avoid the interaction with Alvaro**. Focusing again on table X, by means of *fear*, the dangerous situations are totally averted. In fact, the robot has not been damaged when *fear* is implemented (0/0). Therefore, **we can conclude that *fear* improves the performance of the robot since** provides a safety mechanism to avoid situations where the robot can be damaged.

TABLE X
HARM / INTERACTIONS WITH ALVARO DURING EXPLOITING SESSIONS

without fear	with fear
6/23	0/0

In conclusion, despite of the fact that the average wellbeing is hardly worse, *fear* provides significant benefits. Specially the fact that harm is totally avoided.

X. DISCUSSION

Since social robots move and interact with humans sharing the same areas, one of the main requirements for social robotics is a natural robot’s behavior in terms of similarity to humans’ behaviors, or at least animals’ behaviors (these are perfectly understandable and accepted by people). One of the advantages of using emotions in robots is that emotions allow animal-like responses to certain situations. Fear has been successfully implemented in our robot Maggie in order to provide a natural mechanism of avoiding dangerous situations.

Considering the experiments achieved by Klüver and Bucy [51], monkeys’ behavior were studied in relation with fear. Normal monkeys are afraid of people: when a person tries to approach a monkey in a cage, it escapes running to other corner and remains there. In contrast, monkeys with bilateral temporal lobectomies (suppression of the area of the brain where the amygdala, thought to be the brain structure in charge of the fear regulation and other emotions, is located) experienced some kind of fearlessness: people approached them, touched them, and even stroke them and picked them up. Studies with rats and lynxes reflect the same results. Therefore, fear provides monkeys, and animals in general, with the escaping behavior required at certain situations to survive (humans can be dangerous for monkeys). This same kind of

behavior has been shown by our robot. Maggie has learned that when a certain situation is dangerous, it moves to other place far from where the danger is. When fear is not included as a motivation, Maggie’s behavior corresponds to the same exhibited by an animal suffering a temporal lobectomy, similar to Klüver and Bucy’s monkeys: it is not able to identify the dangerous situations.

In fact, Maggie learns the proper behavior to avoid dangers. As presented on the experiments (sections IX-C and IX-D), depending on different people attitudes, the danger-avoidance behavior could differ: as exposed in the previous paragraph, one behavior is to run away from where the danger is, but the other is to remain still until the threatening person gets bored and goes. This is also a common human behavior observed in terrified people: some people are stunned when they face a great danger. Other example can be observed in some chickens: after a chicken is frighten, it crouches down and trembles with fear.

However, the origin of this behavior differs: in animals, this is an unconscious bodily reaction which makes muscles tensed. In our robot, the reaction is provoked because the learned values indicate that the danger will disappear after. Nevertheless, both responses, in animals and in our robot, are automatic because the exhibited fear behavior is formed without any perspective into the future, just by executing the best action at each moment. The selection of these actions considers the current available information. Then, there is not any planing looking into the future, thus, there is not deliberation.

In this work, reactions to fear have been learned by our robot through interaction with its world. In animals, some reactions to fear are inherited, this is, they are instinctive. Instincts are innate behaviors that are not highly dependent on specific learning experiences performed by the individual [17]. In fact, instinctive behaviors have been learned by the species through evolution. Our experiments have shown that the results obtained from evolution and from our proposed mechanism are similar: escaping or freezing reactions are observed in both. This can be seen as another proof of the good performance of our system because the behaviors exhibited as consequence of fear are analogous: the reactions to fear learned by our robot are comparable to those innate reactions exhibited by animals.

Fear on animals is related to anxiety. Anxiety and its bodily reactions are proportional to the intensity of danger and, by extension, to the intensity of fear. One does not feel the same level of fear when one takes a ride on the roller coaster than when a criminal points you with a gun. However, in our work this is not considered, and fear is a binary variable: it is afraid or it is not. Therefore, the level of fear perceived by Maggie is constant for all the circumstances that evoke fear.

Moreover, in our system, once a danger state is identified, it is not forgotten ever. This is based on the theory that memories associated with fear are quickly formed and long-lasting [46]. However, this situation could lead a robot to suffer some kind of *anxiety disorders* typical of humans beings. Imagine a long-lasting experiment which takes several days. At the beginning, during the first hour, the robot identifies the presence of person *A* as a danger because *A* has hit the robot few times. Despite

the fact that all the rest of actions carried on by *A* during the rest of the days were always positive, our system remembers always the painful initial interactions between Maggie and *A*. Consequently, if *A* is present, then fear emerges on Maggie during the rest of the experiment.

From a psychological perspective, this can be seen as an inappropriate experience of fear which is related to anxiety disorders. There are some points in common with Post-traumatic Stress Disorder (PTSD). Approximately, PTSD is related to intense or unrealistic worries suffered when the stimuli related to a past trauma are present. Even if person *A* damaged Maggie at the very beginning, and he has not done it again in several days, which suggests that this behavior hardly will be repeated, fear arises in the presence of *A*. Also, similarities with a phobia provoked by exposure to situations leading to avoidance behaviors can be found. In particular, a sort of social phobia can be identified because any social interaction with a specific person is avoided, even if it seems that he will not induce any damage. In future works, fear will be enhanced with mechanisms to take into account these features and make it more flexible.

At this point, it seems that traumas on humans are very hard to re-program. This is exactly what happens to our robot as well.

XI. CONCLUSION

The goal of our autonomous social robot is to learn what to do in every situation in order to survive and to maintain its needs satisfied.

In this work, authors have presented a novel approach of using fear in a social robot to guide its behavior. Moreover, an original appraisal mechanism of fear has been implemented and it allows to identify non-predefined dangerous situations. Fear is treated as a motivation which moves to behave.

The fear motivation is elicited when a dangerous situation is detected. These circumstances are not predefined, but they are appraised by the robot through interaction. Therefore, the robot is able to identify by itself the conditions which cause fear.

Permanent harmful exogenous actions can be easily avoided by reinforcement learning algorithms. However, when few negative experiences in relation to exogenous actions have been suffered in a specific situation, it is not easily identified as a potential dangerous situation. Nevertheless, the presented method is able to assess them as a dangerous situations too.

Once the dangerous states are recognized, the robot is able to learn what to do for avoiding them. This is achieved when fear becomes the dominant motivation. It has been show how the robot naturally learns different behaviours depending on the user's attitude, and those are similar to what animals do: it runs away or freezes.

As proved on section IX-E, the average wellbeing does not improve when fear is considered. People in fear live distressed, and this fact is shown in our robot as well. However, some other benefits justify the use of fear. First of all, by means of fear the robot has avoided all harmful exogenous actions: Maggie has not been hurt anymore by a user. Moreover, the

permanence on comfortable levels of wellbeing is better when fear is present since it is not hit anymore. Additionally, the quality of life can be also measured as the amount of time that a particular behavior is not required, i.e. there is not dominant motivation. Also in this case, fear outperforms.

From a human-robot interaction point of view, as already said, the behaviors displayed by the robot are rather animal-like. This helps to improve the interaction when robot is *living* with people.

ACKNOWLEDGMENT

The authors gratefully acknowledge the funds provided by the Spanish Government through the project called "A new approach to social robotics" (AROS), of MICINN (Ministry of Science and Innovation) and through the RoboCity2030-II-CM project (S2009/DPI-1559), funded by Programas de Actividades I+D en la Comunidad de Madrid and cofunded by Structural Funds of the EU.

REFERENCES

- [1] C. Bartneck and J. Forlizzi, "A design-centered framework for social human-robot interaction," in *2004 IEEE International Workshop on Robot and Human Interactive Communication*. Kurashiki, Okayama, Japan, 2004, pp. 31–33.
- [2] J. Hirth, N. Schmitz, and K. Berns, "Towards social robots: Designing an emotion-based architecture," *International Journal of Social Robotics*, vol. 3, no. 3, pp. 273–290, 2011.
- [3] A. Kelley, *Who Needs Emotions? The Brain Meets the Robot*. Oxford University Press, 2005, ch. Neurochemical networks encoding emotion and motivation: an evolutionary perspective.
- [4] S. Gadanho, "Reinforcement learning in autonomous robots: An empirical investigation of the role of emotions," Ph.D. dissertation, University of Edinburgh, 1999.
- [5] F. Michaud, P. Pirjanian, J. Audet, and D. Ltourneau, "Artificial emotion and social robotics," *Distributed Autonomous Robotic Systems*. Springer Verlag, Hiedelberg, Germany, pp. 121–130, 2000.
- [6] C. Breazeal, *Designing Sociable Robots*. The MIT Press, 2002.
- [7] L. Cañamero, *Emotions in Humans and Artifacts*. MIT Press, 2003, ch. Designing emotions for activity selection in autonomous agents.
- [8] A. Ortony, *Emotions in Humans and Artifacts*. MIT Press, 2003, ch. On making Believable Emotional Agents Believable, pp. 188–211.
- [9] R. C. Arkin, *Who needs emotions? The brain meets the robots*. Oxford University Press, 2004, ch. Moving up the food chain: Motivation and Emotion in behavior-based robots.
- [10] M. Scheutz, "Useful roles of emotions in artificial agents: a case study from artificial life," in *AAAI 2004*. AAAI press, Menlo Park, 2004, pp. 42–48.
- [11] T. Ziemke and R. Lowe, "On the role of emotion in embodied cognitive architectures: From organisms to robots," *Cognitive Computation*, vol. 1, no. 1, pp. 104–117, 2009.
- [12] L. Cañamero, "Emotion understanding from the perspective of autonomous robots research," *Neural Networks*, vol. 18, pp. 445–455, 2005.
- [13] M. A. Salichs and M. Malfaz, "A new approach to modeling emotions and their use on a decision-making system for artificial agents," *IEEE Transactions on affective computing*, vol. 3, no. 1, pp. 56–68, 2012.
- [14] M. Malfaz, A. Castro-Gonzalez, R. Barber, and M. Salichs, "A biologically inspired architecture for an autonomous and social robot," *Autonomous Mental Development, IEEE Transactions on*, vol. 3, no. 3, pp. 232–246, sept. 2011.
- [15] A. Castro-González, M. Malfaz, and M. A. Salichs, "Learning the selection of actions for an autonomous social robot by reinforcement learning based on motivations," *International Journal of Social Robotics*, vol. 3, no. 4, pp. 427–441, 2011.
- [16] M. A. Arbib and J. M. Fellows, "Emotions: from brain to robot," *Trends in Cognitive Sciences*, vol. 8 (12), pp. 554–561, 2004.
- [17] E. Kandel, J. Schwartz, and T. Jessell, *Principles of Neural Science*. Elsevier, 1991.

- [18] K. Scherer, *Handbook of cognition and emotion*. Wiley, New York, 1999, ch. Appraisal theory, pp. 637–663.
- [19] I. Roseman and C. Smith, “Appraisal Theory: Overview, Assumptions, Varieties, Controversies,” in *Appraisal Processes in Emotion Theory Methods Research*, K. R. Scherer, A. Schorr, and T. Johnstone, Eds. Oxford University Press, 2001, ch. 1, pp. 3–19.
- [20] C. Breazeal and R. Brooks, *Who Needs Emotions: The Brain Meets the Robot*. MIT Press, 2004, ch. Robot Emotion: A Functional Perspective.
- [21] J. LeDoux, *Synaptic Self: How brains become who we are*. Penguin Books, 2002.
- [22] J. Ledoux, *The Emotional Brain: The Mysterious Underpinnings of Emotional Life*, ser. A Touchstone book. Simon & Schuster, 1998. [Online]. Available: <http://books.google.es/books?id=7EJN5I8sk2wC>
- [23] T. Dalgleish, “The emotional brain,” *Nature Neuroscience Reviews*, vol. 5, no. 7, pp. 583–589, 2004.
- [24] P. Ekman, “An argument for basic emotions,” *Cognition and Emotion*, vol. 6(3/4), pp. 169–200, 1992.
- [25] A. Damasio, T. Grabowski, A. Bechara, H. Damasio, L. Ponto, J. Parvizi, and R. Hichwa, “Subcortical and cortical brain activity during the feeling of self-generated emotions,” *Nature Neurosci*, vol. 3, pp. 1049–1056, 2000.
- [26] F. Rosis, C. Castelfranchi, P. Goldie, and V. Carofiglio, *Emotion-Oriented Systems*, ser. Cognitive Technologies. Springer Berlin Heidelberg, 2011, ch. Cognitive evaluations and intuitive appraisals: can emotion models handle them both?, pp. 459–481.
- [27] A. Sloman, “Architectural Requirements for Human-like Agents Both Natural and Artificial. (What sorts of machines can love?),” in *Science*, K. Dautenhahn, Ed. John Benjamins, 2000, ch. 7, pp. 163–195. [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.46.9493>
- [28] G. Hollinger, Y. Georgiev, A. Manfredi, B. Maxwell, Z. Pezzementi, and B. Mitchell, “Design of a social mobile robot using emotion-based decision mechanisms,” in *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*, 2007, pp. 3093–3098.
- [29] L. Cañamero, “Modeling motivations and emotions as a basis for intelligent behavior,” in *First International Symposium on Autonomous Agents (Agents’97)*, 148-155. New York, NY: The ACM Press., 1997.
- [30] D. Parisi and G. Petrosino, “Robots that have emotions,” *Adaptive Behavior*, vol. 18, no. 6, pp. 453–469, 2010.
- [31] R. S. Lazarus, *Appraisal processes in emotion: Theory, methods, research*. New York: Oxford University Press, 2001, ch. Relational meaning and discrete emotions, pp. 37–67.
- [32] C. Herrera and D. Moffat, “Fear: Appraisal of danger or anticipation of harm,” in *AAAI 2005 Fall Symposium on Anticipatory Cognitive Embodied Systems*, 2005.
- [33] J. Velásquez, “Modeling emotion-based decision-making,” *Emotional and intelligent: The tangled knot of cognition*, pp. 164–169, 1998.
- [34] —, “When robots weep: Emotional memories and decision making,” in *Proceedings of AAI-98*, 1998.
- [35] M. Malfaz and M. Salichs, “Learning to avoid risky actions,” *Cybernetics and Systems*, vol. 42, no. 8, pp. 636–658, 2011.
- [36] M. A. Salichs, R.Barber, A. M.Khamis, M.Malfaz, J. F.Gorostiza, R.Pacheco, R.Rivas, A.Corrales, and E.Delgado, “Maggie: A robotic platform for human-robot social interaction,” in *IEEE International Conference on Robotics, Automation and Mechatronics (RAM 2006)*. Bangkok, Thailand, 2006, pp. 1–7.
- [37] R. Barber and M. A. Salichs, “Mobile robot navigation based on event maps,” in *International Conference on Field and Service Robotics*, June 2001, pp. 61–66.
- [38] J. Salichs, Álvaro. Castro-González, and M. A. Salichs, “Infrared remote control with a social robot,” in *FIRA RoboWorld Congress 2009*, Springer, Ed. Incheon, Korea.: Springer, August 2009, pp. 86–95.
- [39] M. A. Salichs and R. Barber, “A new human based architecture for intelligent autonomous robots,” in *4th IFAC Symposium on Intelligent Autonomous Vehicles*, Sapporo, Japan, 2001, pp. 85–90.
- [40] R. Rivas, A. Corrales, R. Barber, and M. A. Salichs, “Robot skill abstraction for ad architecture,” in *6th IFAC Symposium on Intelligent Autonomous Vehicles*, Toulouse, 2007.
- [41] K. C. Berridge, “Motivation concepts in behavioural neuroscience,” *Physiology and Behaviour*, no. 81, pp. 179–209, 2004.
- [42] C. L. Hull, *Principles of Behavior*. New York: Appleton Century Crofts, 1943.
- [43] J. Santa-Cruz, J. M. Tobal, A. C. Vindel, and E. G. Fernández, “Introducción a la psicología,” 1989, facultad de Psicología. Universidad Complutense de Madrid.
- [44] C. J. Watkins and P. D. Dayan, “Technical note: Q-learning,” *Machine Learning*, vol. 8(3), pp. 279–292, 1992.
- [45] M. Malfaz, A. Castro-González, R. Barber, and M. A. Salichs, “A biologically inspired architecture for an autonomous and social robot,” *IEEE Transactions on Autonomous Mental Development*, vol. 3, no. 3, pp. 232–246, 2011.
- [46] M. Bear, B. Connors, and M. Paradiso, *Neuroscience: Exploring the brain*. Lippincott Williams & Wilkins, 2001.
- [47] A. Olteanu, I. Simion, A. Purcăreanu, and N. Bizdoacă, “Robotic Architecture for Experiments on Emotional Behavior.”
- [48] M. Malfaz, “Decision making system for autonomous social agents based on emotions and self-learning.” Ph.D. dissertation, Carlos III University of Madrid, 2007.
- [49] F. J. Fernández de Gorostiza Luengo, “Programación natural de un robot social mediante diálogos,” Ph.D. dissertation, Carlos III University of Madrid, 2010. [Online]. Available: <http://hdl.handle.net/10016/9359>
- [50] M. Malfaz and M. Salichs, “Using muds as an experimental platform for testing a decision making system for self-motivated autonomous agents,” *Artificial Intelligence and Simulation of Behaviour Journal*, vol. 2 (1), pp. 21–44, 2010.
- [51] H. Kluver and P. Bucy, “Preliminary analysis of functions of the temporal lobes in monkeys,” *Archives of Neurology & Psychiatry*, vol. 42, no. 6, p. 979, 1939.