A Model of Artificial Emotions for Behavior-Modulation and Implicit Coordination in Multi-robot Systems

Jérôme Guzzi
Dalle Molle Institute for Artificial Intelligence (IDSIA), USI-SUPSI
Lugano, Switzerland
jerome@idsia.ch

Alessandro Giusti
Dalle Molle Institute for Artificial Intelligence (IDSIA), USI-SUPSI
Lugano, Switzerland
alessandrog@idsia.ch

Luca M. Gambardella
Dalle Molle Institute for Artificial Intelligence (IDSIA), USI-SUPSI
Lugano, Switzerland
luca@idsia.ch

Gianni A. Di Caro
Carnegie Mellon University in Qatar
Doha, Quatar
gdicaro@cmu.edu

ABSTRACT

We propose a model of artificial emotions for adaptation and implicit coordination in multi-robot systems. Artificial emotions play two roles, which resemble their function in animals and humans: modulators of individual behavior, and means of communication for social coordination. Emotions are modeled as compressed representations of the internal state, and are subject to a dynamics depending on internal and external conditions. Being a compressed representation, they can be efficiently exposed to nearby robots, allowing to achieve local group-level communication. The model is instantiated for a navigation task, with the aim of showing how coordination can effectively emerge by adding artificial emotions on top of an existing navigation framework. We show the positive effects of emotion-mediated group behaviors in a few challenging scenarios that would otherwise require ad hoc strategies: preventing deadlocks in crowded conditions; enabling efficient navigation of agents with time-critical tasks; assisting robots with faulty sensors. Two performance measures, throughput and number of collisions, are used to quantify the contribution of emotions for modulation and coordination.

CCS CONCEPTS

• **Computing methodologies → Robotic planning**; **Multi-agent systems**; **Bio-inspired approaches**;

KEYWORDS

artificial emotions; multi-robot system; control architecture

ACM Reference Format:
Jérôme Guzzi, Alessandro Giusti, Luca M. Gambardella, and Gianni A. Di Caro. 2018. A Model of Artificial Emotions for Behavior-Modulation and Implicit Coordination in Multi-robot Systems. In *GECCO '18: Genetic and Evolutionary Computation Conference, July 15–19, 2018, Kyoto, Japan.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3205455.3205650

GECCO '18, July 15–19, 2018, Kyoto, Japan

ACM ISBN 978-1-4503-5618-3/18/07…$15.00
https://doi.org/10.1145/3205455.3205650

SUPPLEMENTARY MATERIAL

Annotated videos, detailed models and code to replicate our results are available at https://github.com/jeguzzi/artificial-emotions.

1 INTRODUCTION

In animals, emotions play a dual role: they act as internal modulators of behavior and as mediators of social interaction. The former role is exemplified by an animal's increased responsiveness to external stimuli when experiencing fear. The latter role is made possible by the fact that the animal *exposes* its emotional state (e.g., by changing body posture), thus implementing a simple yet immediate form of communication, resulting in a social modulation effect.

We present a general, high-level robot control architecture that proposes similar mechanisms, encompassing *artificial emotion generation and dynamics, behavior modulation*, and *emotion-based information sharing*. Artificial emotions are used to represent a robot's internal state in a compressed yet informative form. The internal emotional state modulates the robot's behavior and, at the same time, can be effectively exposed to other members of the system with very low bandwidth requirements, resulting in system-level modulation and social coordination.

Many previous works on affective computing [30] have developed computational emotion models [15, 26], often with explicit biological inspiration [21, 32, 34]. Different roles of emotions for multi-agent and multi-robot [12, 31] systems have been investigated in detail, namely: behavior modulation [7, 19, 20]; emotion-based information sharing [2, 10, 36]; compressed state representation [18].

Up to our knowledge, our work is the first which encompasses the different roles and functions that can be related to emotions into a *unified, coherent, high-level architecture*. We exploit the general notion of artificial emotions to obtain, at the same time, *adaptive behaviors for the single robots* and the *emergence of coordinated behaviors at the system-level*. In general, to obtain such results, either ad hoc tuning of robots' and system's parameters are required,

and/or potentially complex and bandwidth-consuming techniques for information sharing and coordination are needed.

In short, we aim to improve the performance of a multi-robot system by providing it with intrinsic adaptivity and emerging coordination. Emotion modeling is a good candidate to achieve this goal: after all, animal emotions exist because they have a positive impact on the species, i.e., they allow individuals and communities to perform and adapt better. A similar motivation is shared by several related works [3, 23, 33], which also experimentally demonstrate beneficial effects of artificial emotions. In particular, emotion-modulated social interaction has been used in multi-agent systems to avoid deadlock-like situations without centralized planning [28], and to improve group performance in a task-sharing scenario [3]. The specific emotion of fear, as a natural reaction to a situation perceived as potentially harmful, was investigated in [6] as a behavior modulator learned by a robot through interactions in the real world.

Note that a related research direction uses artificial emotions to specifically facilitate interaction with humans [1, 5, 9], also in the context of multi-robot systems [17, 35]. Our work presents a general architecture and does not explicitly pursue this specific goal. However, as we discuss in Section 2.5, in principle robots can expose their artificial emotions also to neighboring humans, and vice versa, which makes the framework rather general.

To illustrate the use of the model in a practical scenario of interest, we instantiate it for representing and using artificial emotions in the context of *multi-robot navigation and collision avoidance*. We consider a, potentially dense, multi-robot system where each robot moves independently. In order to obtain effective mobility patterns at the system level (i.e., reduced collisions, fluid motion), robots need to mutually adapt their mobility and implement some form of coordination. Instead of handcrafting task-specific adaptation and coordination schemes, we show that we can achieve these goals using rather generic artificial emotions. To this aim, we define, and include in the general architecture, a set of artificial emotions that are significant for the problem scenario. We also define quantitative performance measures, which we compute for three settings: without emotions; with emotion-based behavior modulation; with emotion-based behavior modulation and information sharing. The results show an advantage of using our emotion-based architecture to achieve effective coordination, without the need for ad hoc tuning of parameters nor defining handcrafted behaviors.

We remark that the use of the navigation scenario is purely for illustration purposes. A number of related works exist in this respect, that however set more specific and limited goals than our framework. A research direction, which is close to ours in this context, regards the use of *static* personality traits that parametrize navigation behaviors [13] as a space where to search for optimal strategies in heterogeneous swarms [11], whereas we make use of more general *dynamic* emotions.

Our **main contribution** is the high-level architecture for robot control based on artificial emotions and its formalization in terms of emotions representation and dynamics, behavior modulation, and information sharing for coordination, which is described in Section 2; Section 3 presents the application to multi-robot navigation; Section 4 describes our setup for validation and Section 5 reports quantitative results. Section 6 concludes the paper.

2 ARTIFICIAL EMOTIONS MODEL AND ARCHITECTURE

In this Section, we first introduce basic definitions and ideas behind our modeling approach for implementing and using artificial emotions (Sections 2.1 and 2.2). Then we describe: the dynamics regulating emotional changes (Section 2.3); how an emotional state determines behavior modulation (Section 2.4); and, finally, how emotion-based information sharing is performed to obtain system-level coordination (Section 2.5). Figure 1 illustrates the architecture encompassing all these aspects.

2.1 Internal state

We consider a robotic agent that, at time t, is engaged in some *task* T (e.g., assigned by an external operator) and is part of a (potentially heterogeneous) group of cooperative robotic agents. An agent is characterized by: a set of *abilities* A, a *personality* P, and a time-dependent *energy level* $L(t)$. The abilities are defined as the (fixed) set of sensory-motor capabilities of the robot, which are used to realize the current *behavior*, as selected according to some decisional policy. The personality is defined as the collection of fixed (or slowing changing) parameters that affect the way abilities are used to realize behaviors.

By using its abilities to sense the external environment and neighboring agents, the robot acquires at time t some information that we denote as $\boldsymbol{x}(t) = (x_1, x_2, ..., x_M) \in \mathbb{R}^M$. Assuming an intermediate cognitive layer that processes the raw measures gathered through the sensory apparatus of the robot, the information in $\boldsymbol{x}$, complemented by the internal data (A, T, L), is used to update the *internal state* $\boldsymbol{\mu}$ of the robot. This is parametrized by means of a collection of *micro-states* $\boldsymbol{\mu}(t) = (\mu_1, \mu_2, ..., \mu_N) \in \mathbb{R}^N$. Micro-states are features that are relevant to represent the current situation of the environment (including other robots' status) and of the robot itself for the given task. Each μ_n is a real-valued variable quantifying the level of activation of the corresponding n-th micro-state. Clearly, the precise definition of the micro-states is task and robot dependent. For instance, for a navigation task, relevant features to account for might include usual features such as the robot's velocity, presence of obstacles, patterns of nearby robots, energy levels, as well as higher level notions, such as how predictable the motion of a neighbor robot is, and so on. For a collective transportation task, relevant features might be quite different. In any case, the precise definition of micro-states that are accessible through sensors and are useful for the purpose of taking high-level control decisions is expected to be rather straightforward.

The current internal state $\boldsymbol{\mu}$ affects the emotional state as described in the next Section. In turn, the emotional state modifies the behavior generation for the current task T, and is, at the same time, communicated to other robots, thus promoting swarm-level coordination as discussed in Section 2.5.

2.2 Emotions

The use of the term micro-states is loosely borrowed from statistical mechanics, intending that different combinations of values for the features give expression to different *affective macro-states*, and multiple combinations of values can be associated to the same macro-state. More precisely, we define as $\boldsymbol{\varepsilon}(t) = (\varepsilon_1, \varepsilon_2, ..., \varepsilon_K) \in [0, 1]^K$

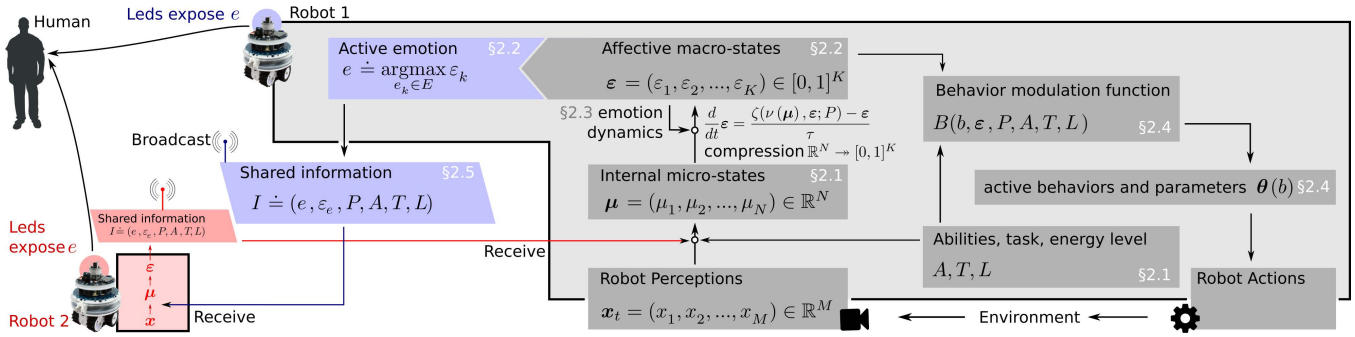

Figure 1: Information flow in the proposed modular architecture based on the use of affective states: exploded view for Robot 1 and summarized view for Robot 2. The central part illustrates the emotion dynamics of Robot 1, where affective macro-state $\boldsymbol{\varepsilon}$ depends on input observations from the environment $\boldsymbol{x}$, task T, personality P, abilities A, energy level L, and the emotional state of neighboring robots (red box). The current active emotion e (linked to the largest component of $\boldsymbol{\varepsilon}$) of each robot can be exposed through a very low-bandwidth link (like an RGB LED) and shared with humans.

the affective macro-states that result from the *compression* of the internal micro-states $\boldsymbol{\mu}(t)$ by a time independent surjective function $\nu : \mathbb{R}^N \twoheadrightarrow [0,1]^K$.

Affective macro-states should: (1) be part of a universally shared vocabulary (e.g., every human has an idea of what *fear* is and how it affects behaviors); (2) be relevant to the task and the characteristics of the robot; (3) be a compact representation (compression) of the internal state, i.e., $K \ll N$.

In this context, we define *emotions* as labels of affective macro-states. Namely, the affective macro-state $\varepsilon_k \in [0,1]$ is interpreted as the activation at time t of the k-th emotion e_k belonging to the set $E = \{e_1, \ldots, e_K\}$.

For example, given a collection of micro-states that describe "how well a sensor is working", "how often the robot's goal is changing", "how predictable a neighbor is", "how coherent the environment's representation is", …, the affective macro-state *confusion-level* aggregates this information and describes how high the activation of the emotion *confusion* is. In Section 3, we extend this example to a set of emotions that arises in a navigation scenario.

At any time, all emotions in E contribute to modulate the agent's behavior (see Section 2.4). Nonetheless, the macro-state that predominantly affects a robot at time t has a particular importance and is used to define the robot's *active emotion* as

$$e(t) \doteq \operatorname*{argmax}_{e_k \in E} \varepsilon_k(t). \tag{1}$$

The active emotion is the one piece of information that describes the agent's state best and that should be shared with neighbors (see Section 2.5). The adopted winner-takes-all strategy, other than finding some ground in the biological processes describing emotions' dynamics, brings an intrinsic *robustness for parameter setting*, since it's only the most prominent emotion that matters, rather than the precise relative weights between all the different emotions.

To make the overall response even more robust and *avoid oscillations*, we use a *hysteresis* mechanism with two dynamically defined thresholds $\epsilon_{low} < \epsilon_{high}$: the robot keeps its currently-active emotion until its activation decreases under ϵ_{low} or another emotion's level rises above ϵ_{high}.

2.3 Emotion dynamics

In humans, affective activations are temporally modulated [25], with psychologists distinguishing between sentiments (long lived), moods (medium lived), and emotions (short lived). Typical duration of emotions ranges from seconds up to minutes. We assume that emotion activations follow a form of *stimulus-response dynamic*[1]

$$\frac{d}{dt}\boldsymbol{\varepsilon}(t) = \frac{\nu(\boldsymbol{\mu}(t)) - \boldsymbol{\varepsilon}(t)}{\tau}, \tag{2}$$

where τ is the typical emotion life-time once the stimulus that caused the emotion is removed.

We generalize Eq. (2) to allow for an internal emotional dynamics ζ that stimulates the activation of an emotion from (other) emotions' activations too and is parametrized by the robot personality P:

$$\frac{d}{dt}\boldsymbol{\varepsilon}(t) = \frac{\zeta(\nu(\boldsymbol{\mu}(t)), \boldsymbol{\varepsilon}(t); P) - \boldsymbol{\varepsilon}(t)}{\tau}. \tag{3}$$

For example, a robot that is experiencing confusion (e.g., because of too many other robots erratically moving around) while executing its task, and has a personality P_1 — that results into an impulsive, quick-tempered emotion dynamics — would quickly become frustrated because unable to address the cause of confusion. In turn, the frustration affective state will induce an appropriate behavior modulation, that will hopefully facilitate the solution of the problem (e.g., decrease speed to avoid collisions with those robots). Instead, a robot with a different personality P_2 that is more resilient to changes, would not activate frustration and would respond differently to confusion.

2.3.1 Linear model for emotion dynamics. In the following, we present a *linear model* for emotion dynamics based on Eq. (3) that is parametrized by personality factors, which we consider as stable characteristics of the agents.

Although psychologists disagree on the best way to conceptualize individual personality differences, considerable evidence supports the *Five Factor Model* (FFM) [8], which we therefore use

[1] This is compatible with the concept of priming in psychology literature, i.e., how past stimuli influence the present behavior [29].

as an inspiration model. The FFM characterizes individual differences in terms of five personality factors $P_* \in [0,1]$: openness P_{ope}, conscientiousness P_{con}, extraversion P_{ext}, agreeableness P_{agr}, and neuroticism P_{neu}.

We assume a linear mapping ν between micro- and macro-states. Let $\Sigma_k^{\uparrow}, \Sigma_k^{\downarrow} \subseteq \{1, 2, \ldots, N\}$ be the disjoints subsets of micro-states that influence positively (↑) or negatively (↓) the k-th macro-state. We define each mapping ν_k as a weighted sum over $\Sigma_k^{\uparrow}$ and $\Sigma_k^{\downarrow}$

$$\nu_k(\boldsymbol{\mu}) \doteq \sum_{i \in \Sigma_k^{\uparrow}} \frac{\mu_i}{\lambda_{ik}} + \sum_{i \in \Sigma_k^{\downarrow}} \left(1 - \frac{\mu_i}{\lambda_{ik}}\right), \tag{4}$$

where λ_{ik} defines a normalization factor and corresponds to the value of the i-th internal state that stimulates a large activation of the k-th emotion. Besides the sensorial stimuli $\boldsymbol{x}$, which are taken into account by $\boldsymbol{\mu} = \boldsymbol{\mu}(\boldsymbol{x})$, the affective macro-state has its own dynamics ρ, where different emotions influence each other. This is modeled[2] through a linear relation defined by a square matrix (ξ_{ij}) of dimension K^2: $\rho_k(\boldsymbol{\varepsilon}) \doteq \sum_{i=1}^{K} \xi_{ik}\varepsilon_i$. The overall dynamics of the emotional activations results from the definition of ζ in Eq. (3) as the saturation over $[0,1]$, through an activation function $\phi : \mathbb{R} \to [0,1]$, of the weighted sum of two contributions:

$$\frac{d}{dt}\varepsilon_k(t) = \frac{\phi\left(\gamma_{k1}\left(\gamma_{k2}\nu_k(\boldsymbol{\mu}(t)) + \gamma_{k3}\rho_k(\boldsymbol{\varepsilon}(t))\right)\right) - \varepsilon_k(t)}{\tau_k}. \tag{5}$$

This sort of *non-linear saturation* is similar to psycho-physiological dynamics of human emotions [22].

Regarding the FFM model, we note that *neuroticism* stands for the sensitivity to input observations: $\gamma_{k2} \propto P_{\text{neu}}, \tau_k \propto 1/P_{\text{neu}}$. Similarly, openness (or spontaneousness) makes a robot more willful to transit through emotions: $\gamma_{k3} \propto P_{\text{ope}}$. All the personality traits contribute to defining the saturation level $\gamma_{k1} = \gamma_{k1}(P)$.

2.4 Behavior modulation

In the following, we describe the behavior update process, depicted in the right part of Figure 1.

We assume that the robot is provided with a set of Q *predefined behaviors* $\{b_1, \ldots, b_Q\}$. At time t, each behavior b_i can be activated and tuned by adjusting R_i parameters $\boldsymbol{\theta}(b_i; t) = \left(\theta_{i0}, \theta_{i1}, \ldots, \theta_{iR_i}\right) \in [0,1]^{R_i}$, where the first parameter signals the behavior's activation, i.e., $\theta_{i0} = 1$ if b_i is active and $\theta_{i0} = 0$ otherwise. At each control step, the robot uses the compressed information about its internal state (provided by the affective macro-state defined in Section 2.2), as well as information about its task, its abilities, energy level and personality, in order to select the set of active behaviors and set their parameters $\boldsymbol{\theta}$ through a behavior modulation function B

$$\boldsymbol{\theta}(b; t) = B(b, \boldsymbol{\varepsilon}(t), P, A, T, L(t)), \quad \forall b \in \{b_1, \ldots, b_Q\}. \tag{6}$$

Let us consider again the previous example, in which a robot was increasing its *confusion-level*. This would yield a number of behavior modulations, such as: decrease moving speed to avoid collisions; activate an help-seeking behavior to lower the reliance on its sensors; increase the number of raw data samples that are filtered to update $\boldsymbol{\mu}$. An extended example of behavior modulation in the context of multi-agent navigation is presented in Section 3.

[2] A concrete example of internal emotion dynamics and matrix ξ can be found in [24].

2.5 Emotion sharing for coordination

So far, we have considered behavior modulation for a single robot, but we are also interested in using artificial emotions to have *emerging coordination* in a multi-robot system, possibly with a minimal overhead. To this aim, *information sharing* is a key factor; but, which information should a robot share with neighbors? We maintain that a minimal yet effective approach is to expose information about the robot's *active emotion* $e(t)$ and the related *activation* $\varepsilon_{e(t)}(t)$.

This information can be interpreted as an extremely compressed, but very *meaningful*, representation of the robot's internal state: it encodes important clues about the robot's next actions, its surroundings, and its reactions to environmental changes. To improve its predictability, the robot should complement affective information with data that helps neighbors to interpret it, like personality P, abilities A, task T and energy level L.

Therefore, the robot publishes, at the end of a control step at time t, the message

$$I(t) \doteq (e(t), \varepsilon_{e(t)}(t), P, A, T, L(t)), \tag{7}$$

which represents another instance of the modulation mechanism described in Section 2.4 and defined by Eq. (6).

After an agent receives such message from a neighbor, it adds its content to its own internal state $\boldsymbol{\mu}$. This, in turn, closes the loop by modifying the recipient's affective macro-states, active emotion, and communication content. This yields an automatically-regulated shared *social emotional state* that, when correctly designed, can increase the performance of the system as a whole.

Neighbors' emotions provide a compact information content that is relevant to the task, the robot's abilities and the environmental status. They integrate the robot's own sensing information with data that is pre-filtered within the neighborhood. Ultimately, this enables a better predictability of the neighbors' behavior and promotes the emergence of coordination within the system.

Note that these same mechanisms can be used to let social coordination emerging in human multi-robot systems. On the one hand, humans share their emotions through nonverbal communication (e.g., facial expressions). If a robot is able to detect the emotion of a human neighbor, it can use this information just like the information received by neighboring robotic agents. On the other hand, if robots explicitly share their emotional status (e.g., by modulating their appearance through colored lighting), interacting humans can learn to interpret this appearance as an emotion, and better relate to the robots' behavior by means of *empathy*.

3 EXAMPLE: ARTIFICIAL EMOTIONS IN MULTI-ROBOT NAVIGATION

In this Section, we instantiate the proposed framework in the context of multi-robot navigation (see Figure 2) to provide an *illustrative example* of the relation between emotions, behavior modulation and local communication. We will then apply this model in *simulation* (Section 4) and report experimental results (Section 5). Although the navigation task by itself has no emotional content, we will show that introducing the emotions summarized in Table 1, has a positive impact on performance.

Three challenging situations are considered.

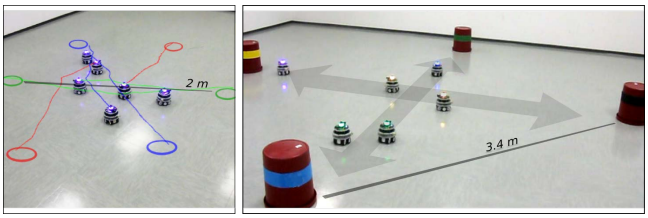

Figure 2: Multi-robot navigation scenarios: (left) CIRCLE scenario where the robots start on a circle and navigate to reach the diametrically opposite end; (right) CROSS scenario where robots navigate back-and-forth between opposite corners of a square.

Preventing and escaping deadlocks in crowds Reactive local navigation algorithms lack the planning and explicit coordination ability to handle very crowded (and/or cramped) environments with robots navigating to different targets; then, deadlocks may form, in which robots cluster are unable to proceed or escape.

Enabling effective activity of agents with critical tasks In a heterogeneous swarm, different agents sharing the same space may work on different tasks, some of which could be more time-critical than others; the swarm's behavior should account for this difference.

Assisting robots with sensing issues Sensing is a critical task for mobile robots; if a robot has a sensing malfunction, it may be assisted by a functioning robot which is heading to the same target and shares its sensing information.

The mappings ν and B between the agent's internal state, emotion activations, and behaviors are derived from few, legible, bio-inspired rules and common sense. For example, we impose that fearful agents should be more careful and slow down.

In the remainder of this section, we: describe the representation of the agent's internal state (Section 3.1); implement an instance of the linear dynamic model described by Eq. (5) (Section 3.2); present the agent's behaviors (Section 3.3) and their modulation through linear functions (Section 3.4). The whole model is formalized in detail in the Appendix provided as supplementary material.

3.1 Internal State

The relevant part of the agent's internal state $\boldsymbol{\mu} \in \mathbb{R}^5$ is given by:

$\mu_{\mathbf{eff}}$ **(efficacy)** ratio of speed towards target to optimal speed,

$\mu_{\mathbf{fs}}$ **(free space)** ratio of free space around the agent to agent's horizon,

$\mu_{\mathbf{nf}}$ **(nearby frustration)** fraction of neighbors in frustration,

$\mu_{\mathbf{td}}$ **(task delay)** relative duration of the travel towards the target with respect to the expected duration,

$\mu_{\mathbf{rot}}$ **(extra rotations)** how much the agent has to rotate to reach the desired direction.

3.2 Emotion dynamics

The emotions' activation functions $\nu_e : \mathbb{R}^5 \to [0, 1]$, for any $e \in E = \{\text{fear}, \text{frustation}, \text{urgency}, \text{confusion}\}$, are defined as following.

Table 1: Summary of emotions, micro-states that influence emotions dynamics, and (↑positive, ↓negative) behavior modulations in the context of multi-agent navigation.

Emotion	Stimuli (Section 3.2)	Behavior modulation (Section 3.3) Robot a	Neighbors
neutral	selected when no other emotion above threshold	default behavior and parameters	
fear	lack of free space, many frustrated neighbors	$\uparrow \psi, \downarrow v_{\text{opt}}$	
frustration	slow progress to target	$\uparrow$ trigger b^{escape}	
urgency	expected arrival close to deadline	$\downarrow m_*, \eta, \psi$	$\uparrow m_a$
confusion	too many changes of direction	$\downarrow v_{\text{opt}}$	$\uparrow m_a$ $\uparrow$ trigger b^{help}

Fear. Fear (of an incoming deadlock) increases when an agent is closely surrounded by other agents, and when the majority of the neighboring agents expose frustration:

$$\nu_{\text{fear}}(\boldsymbol{\mu}) = \frac{\mu_{\text{nf}}}{\lambda_{\text{nf}}} + \left(1 - \frac{\mu_{\text{fs}}}{\lambda_{\text{fs}}}\right). \tag{8}$$

Frustration. Agents get frustrated when they don't manage to advance quickly enough towards the target:

$$\nu_{\text{frustration}}(\boldsymbol{\mu}) = 1 - \frac{\mu_{\text{eff}}}{\lambda_{\text{eff}}}. \tag{9}$$

Urgency. Agents continuously monitor their progress towards the goal; when an agent detects slow progress compared to the time requirements of its task, urgency increases:

$$\nu_{\text{urgency}}(\boldsymbol{\mu}) = \frac{\mu_{\text{td}}}{\lambda_{\text{td}}}. \tag{10}$$

Confusion. When an agent notices to be following an unnecessarily complicated path to reach its target, confusion is increased because a sensing malfunction is suspected:[3]

$$\nu_{\text{confusion}}(\boldsymbol{\mu}) = \frac{\mu_{\text{rot}}}{\lambda_{\text{rot}}}. \tag{11}$$

Finally, we complete our example as a model

$$\frac{d}{dt}\varepsilon_k(t) = \frac{\phi\left(\gamma_k(P)\nu_k(\boldsymbol{\mu}(t))\right) - \varepsilon_k(t)}{\tau_k} \tag{12}$$

that ignores internal emotion dynamics ($\rho_k \equiv 0$ in Eq. (5)) and uses an activation function $\phi(x) = x|_0^1$, where we use the following notation for clamping $x \in \mathbb{R}$ in $[a, b] \subset \mathbb{R}$:

$$x|_a^b \doteq \begin{cases} a & \text{if } x \leq a \\ x & \text{if } a < x < b \\ b & \text{if } x \geq b \end{cases}. \tag{13}$$

In our example, personality P only affects urgency: agents with prioritized tasks have $\gamma_{\text{urgency}} = 100$, while all other agents have $\gamma_{\text{urgency}} = 1/4$; all others γ_k are set to one.

3.3 Behaviors

In our context, agents solve challenging navigation problems in crowded environments, where they need to reach one or more targets while avoiding other robots (see Figure 2). The primary

[3]This mechanism is related to motion sickness in humans, which is caused by conflicting visual and balance perceptions. Nausea and vomiting occur as a countermeasure of a suspected poisoning, which could be the root cause of such sensory issues.

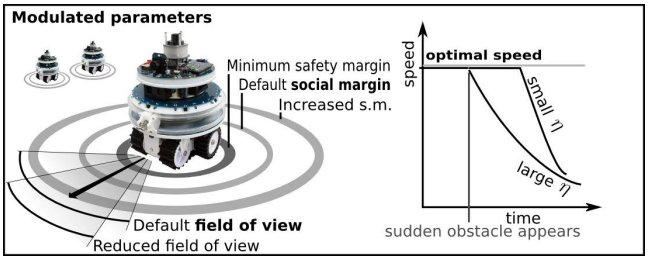

Figure 3: MarXbot robot and modulated parameters.

behavior is, therefore, a navigation behavior b^{nav}. Additionally, robots implement other two behaviors:

(1) b^{escape} is a behavior, alternative to b^{nav}, under which the target is temporarily ignored, and the robot only attempts to move towards the direction which allows for the longest free path. This behavior is triggered by agents with large frustration and contributes to avoid extreme crowding.
(2) b^{help} is a behavior, parallel b^{nav} and b^{escape}, under which a robot publishes its sensing information for others to use. Confused agents can exploit this information to safely navigate towards their target.

Many algorithms have been proposed for local navigation of robots while avoiding collisions with obstacles, other robots, or humans. We build b^{nav} upon the recently-proposed human-like heuristic [27] for robot navigation [14], which is based on a simple mechanism defined by the following few behavior parameters (Figure 3), which, in this work, we modulate through emotions:

Field of view ψ, is the width of the sensing cone within which a robot can perceive (with some uncertainty) non-occluded obstacles. The maximum value is dictated by robot hardware; it can be artificially decreased to obtain a less accommodating, more self-centered behavior where a robot tends to follow straighter paths and disregard neighbors which are not directly in front of him.

Optimal speed v_{opt}, is the normal moving speed that a robot keeps when there are no obstacles in the vicinity. Slow speed enables easier maneuvering in tight areas; moreover, when a slow robot frequently changes its direction, others will require fewer adjustments to their own trajectories.

Social margin m_a w.r.t. agent a. Robots plan their paths in order to keep a minimal distance from neighbors that is larger than the minimum safe distance. A reduced social margin leads to straighter paths and reduces the tendency to navigate around others, thus forcing them to deviate more from their path. A robot using a larger social margin tends to follow longer (thus less efficient) paths but is less affected by the unexpected motion of neighbors because of a larger headroom to smoothly navigate around them.

Caution (η), when confronted with an obstacle at a given distance, robots do not immediately stop; instead, they smoothly modulate their speed following the same rules observed in pedestrians and drivers [16]. In particular, the predicted time-to-collision with the obstacle should not exceed a given fixed time η. A large η results in a very cautious behavior in which robots tend to "brake early"; while robots with small η tend to brake at the last moment.

3.4 Behavior modulation

The behavior parameters are dynamically modulated by affective macro-states, and emotions trigger different behaviors in order to deal with specific issues; the model is summarized in Table 1.

$$\eta = \eta_0 \left(1 - 4\varepsilon_{\text{urgency}}\right)\Big|_{\eta_{\min}}^{\eta_0} \tag{14}$$

$$\psi = \psi_0 + (\psi_{\min} - \psi_0)\varepsilon_{\text{urgency}} + (\psi_{\max} - \psi_0)\varepsilon_{\text{fear}} \tag{15}$$

$$v_{\text{opt}} = v_0 \left(1 - \varepsilon_{\text{fear}} - \varepsilon_{\text{confusion}}\right)\Big|_{v_{\min}}^{v_0} \tag{16}$$

$$m_a = r_a + (m_0 - r_a)\left(1 - \varepsilon_{\text{urgency}}\right) + \begin{cases} m_{\text{conf}} & \text{if } e_t(a) = \text{confusion} \\ m_{\text{urg}} & \text{if } e_t(a) = \text{urgency} \\ 0 & \text{otherwise} \end{cases} \tag{17}$$

$$\theta_0^{\text{escape}} = 1 - \theta_0^{\text{nav}} = \begin{cases} 1 & \text{if } e_t = \text{frustration} \\ 0 & \text{otherwise} \end{cases} \tag{18}$$

$$\theta_0^{\text{help}} = \begin{cases} 1 & \text{if } e_t(a) = \text{confusion for a neighbor } a \\ 0 & \text{otherwise} \end{cases} \tag{19}$$

Fear. Fearful agents move slowly [Eq. (16)], and carefully monitor their environment using the largest field of view allowed by their sensors [Eq. (15)].

Frustration. Frustration triggers the escape behavior: frustrated agents move in a free direction, until frustration decreases and b^{nav} is resumed [Eq. (18)].

Urgency. Agents that feel urgency reduce their field of view [Eq. (15)]; therefore they tend to follow straighter paths because they ignore robots at the periphery of their sensing cone; similarly, they disregard the social margin around other agents [Eq. (17)] and reduce their caution [Eq. (14)], which results in faster (but less respectful) navigation; the navigation behavior remains safe as long as at least one of two agents in an on-collision pair has a large enough field of view. Neighbors artificially increase the social margin from robots with urgency [Eq. (17)], in order to watchfully keep a larger distance, avoid obstructing their path, and enforce safety — like drivers do when hearing a siren.

Confusion. Confused agents move slowly [Eq. (16)] and make neighbors increase their social margin [Eq. (17)] to reduce disturbance caused by the confused agents' erratic direction changes. Confused agent can exploit sensing information from neighbors once these activate the parallel behavior b^{help} [Eq. (19)].

4 EXPERIMENTAL SETUP

The navigation tasks are performed using a *realistic simulator* for large scale navigation experiments [14]. The robotics agents simulate the *marXbot* robots (see Figures 2 and 3), a small mobile platform [4], specifically designed for swarm robotics.

Simulation results are compared in three settings: *baseline*, in which no emotion-based modulation is implemented; *modulation*, where emotions are implemented and modulate the agents' behavior, but are not perceived by others; *modulation & communication*, where the neighbor's emotional state is fully taken into account.

To highlight the contribution of the emotions dynamics, we restrict our analysis to a model where the other factors (energy level,

abilities, personality and task) are uniform and constant. Moreover, because the set of emotions is small and easily interpreted, we could choose sensible values (listed in Table 2) for the free parameters introduced in Section 3. For example, to fix λ_{rot}, we estimate the value of the maximal amount that a robot, with normal sensing, would normally have to turn. We did not tune parameters to optimize performance.

Table 2: Top row: constant behavior parameters. Bottom row: emotion dynamic parameters.

η_{min}	η_0	ψ_{min}	ψ_0	ψ_{max}	v_{min}	v_0	m_0	m_{conf}	m_{urg}
0.1 s	0.5 s	$\pi/4$	$\pi/2$	π	0.03 m/s	0.3 m/s	0.2 m	0.2 m	0.4 m
$\tau_{\text{frustration}}$	τ_{fear}	τ_{urgency}	$\tau_{\text{confusion}}$	λ_{eff}	λ_{fs}	λ_{nf}	λ_{td}	λ_{rot}	
3 s	5 s	5 s	10 s	1	1/16	1/3	1	0.8 rad	

Performance measures. To compare the navigation performance of the swarm, we measure: (a) the *efficiency* $\in [0, 1]$ as the number of targets that the robots were able to reach, relative to the number of targets that the robots could reach in the same time while traveling in straight lines (i.e., the relative throughput); (b) the *collision hazard*, as the number of collisions per time unit and per agent.

5 EXPERIMENTAL RESULTS

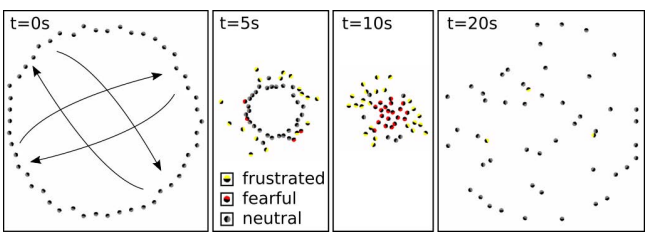

Figure 4: Time-lapse of the CIRCLE scenario for 50 robots: ($t = 0s$) robots start along a circumference and need to exchange positions with the agent at the opposite end; ($t = 5s$) robots converge at the center and ($t = 10s$) tend to form a cluster, which in this case is quickly unraveled as agents navigate towards their target ($t = 20s$). Frustrated robots are depicted in yellow; fearful robots in red. Animated videos of all scenarios are available in the supplementary material.

We report results for three experiments, in which robots attain a different high-level goal by means of the proposed framework.

Preventing and escaping deadlocks in crowds. We considered the CIRCLE task (Figure 4), in which agents initially placed along a circle need to exchange position with the agent at the opposite side. If many agents are considered, they tend to form a cluster in the center of the circle, which is difficult to escape and in which several collisions may occur.

Results. Figures 5(a, b) show that in the baseline case, efficiency drops dramatically as soon as more than 30 agents are considered; when emotions are implemented, relative efficiency decreases linearly with the number of robots even with swarms as large as 100 robots, and deadlocks never occur. The role of emotion sharing, in this case, is to promote the onset of fear ahead of time, i.e. *when a proper cluster has not yet formed*, but many frustrated robots can be perceived in the vicinity, which normally preludes to the formation of a deadlock. Fearful agents then slowly negotiate the difficult environment, which results in a much-reduced probability of collisions; this behavior additionally leads to a small (albeit statistically significant for $50 < N < 80$) efficiency improvement.

Enabling efficient activity of agents with time-critical tasks. We considered the CROSS task (Figure 2 Right), in which agents need to travel back and forth between two targets at opposite vertices of a square; this creates a crossroad in the middle where robots frequently need to navigate around each other.

Results. Figure 5(c) shows that the proposed modulation mechanism leads to a large increase in efficiency for agents with time-critical tasks (+36% for $N = 50$) that are more prone to experience urgency. When robots can also perceive the emotions of others, efficiency is further boosted (+54% for $N = 50$), as fearful/watchful agents tend to move out of the way of agents in a hurry.

Assisting robots with sensing issues. Considering the CROSS task as in the previous case with $N = 13$, temporary sensing malfunctions are simulated for 3 agents by stochastically adding white noise to their sensor measurements. An agent with a sensing malfunction will perceive the direction of its current target with a varying angular error.

Results. We observe that agents affected by a sensing malfunction become confused within a short time, as they perceive to be following excessively convoluted paths; the resulting information-sharing mechanism leads to a very limited efficiency loss on the whole swarm (-6% in the worst case, compared to -15% when the proposed mechanism is not implemented), see Figure 5(d). Note that, in this case, emotions are used to limit and optimize communication among agents by sharing sensing information only when useful. Without emotion sharing, the performance is slightly lower than the baseline, because the normal functioning agents become overly cautious and take longer paths around confused agents.

6 CONCLUSIONS

We presented a control architecture for robotic agents that is inspired by the dual nature of animal emotions, which act both as an internal behavior modulator, and as an implicit communication mechanism that allows for emerging coordination. We showed an implementation in the context of multi-robot navigation and collision avoidance, that was useful to solve three different issues and measurably increased the performance of the swarm behavior. Even though one could obtain similar results by handcrafting problem-specific behavior modulations, the approach we propose yields a simpler, more modular, more interpretable and reusable architecture. Future work will be aimed at validating the universality of the proposed emotion vocabulary and experimenting its application to real world scenarios in which multi-robot coordination is important.

REFERENCES

[1] H. S. Ahn and J. Y. Choi. 2012. Can we teach what emotions a robot should express?. In *Proc. of IEEE/RSJ Int. Conf. on Int. Robots and Sys. (IROS)*. 1407–1412.

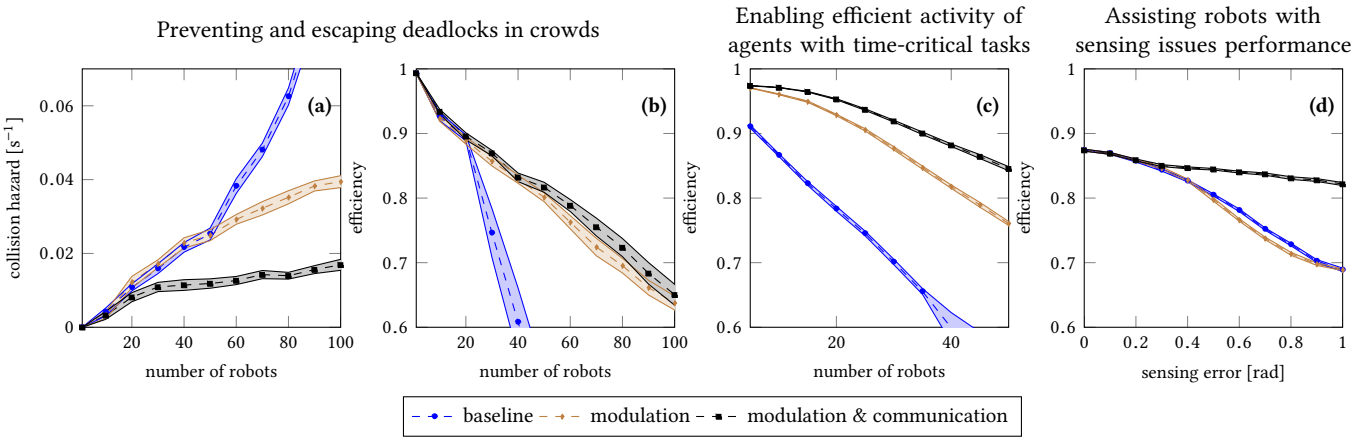

Figure 5: Quantitative performance measures for efficiency and collision hazard (a) and mean efficiency (b, c, d) are reported for the three considered cases: baseline (blue dots), modulation without communications (brown diamonds), modulation+communication (black squares). Areas between the lines denote 95% confidence intervals for the average over 25 randomly-initialized replicas.

[2] S. C. Banik, K. Watanabe, M. K. Habib, et al. 2008. An emotion-based task sharing approach for a cooperative multiagent robotic system. In *Proc. of the IEEE Int. Conf. on Mechatronics and Automation (ICMA)*. 77–82.

[3] S. C. Banik, K. Watanabe, and K. Izumi. 2008. Improvement of group performance of job distributed mobile robots by an emotionally biased control system. *Artificial Life and Robotics* 12 (2008), 245–249.

[4] M. Bonani et al. 2010. The marXbot, a miniature mobile robot opening new perspectives for the collective-robotic research. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. 4187–4193.

[5] C. L. Breazeal. 2004. *Designing Sociable Robots*. The MIT Press.

[6] A. Castro-Gonzalez, M. Malfaz, and M.A. Salichs. 2013. An Autonomous Social Robot in Fear. *IEEE Trans. on Autonomous Mental Develop.* 5, 2 (2013), 135–151.

[7] B. Cox and J.L. Krichmar. 2009. Neuromodulation as a robot controller. *IEEE Robotics & Automation Magazine* 16, 3 (2009), 72–80.

[8] J. M. Digman. 1990. Personality Structure: Emergence of the Five-Factor Model. *Ann. Rev. of Psych.* 41, 1 (1990), 417–440.

[9] J-M. Fellous and M. A. Arbib. 2004. *Who Needs Emotions: The Brain Meets the Robot*. Oxford University Press.

[10] A. Gage and R. Murphy. 2004. Affective recruitment of distributed heterogeneous agents. In *Proc. of Nat. Conf. on Artifical Intelligence (AAAI)*. 14–19.

[11] J. Godoy, I. Karamouzas, S. J. Guy, and M. L. Gini. 2016. Moving in a Crowd: Safe and Efficient Navigation among Heterogeneous Agents.. In *Proc. of the Int. Joint Conf. on Artificial Intelligence(IJCAI)*. 294–300.

[12] H. Gunes, B. Schuller, M. Pantic, and R. Cowie. 2011. Emotion representation, analysis and synthesis in continuous space: A survey. In *Proc. of the IEEE Int. Conf. on Automatic Face & Gesture Recognition and Workshops (FG)*. 827–834.

[13] S. J. Guy, S. Kim, M. C. Lin, and D. Manocha. 2011. Simulating heterogeneous crowd behaviors using personality trait theory. In *Proc. of the 2011 ACM SIGGRAPH/Eurographics symposium on computer animation*. ACM, 43–52.

[14] J. Guzzi, A. Giusti, L. M. Gambardella, G. Theraulaz, and G. A. Di Caro. 2013. Human-friendly robot navigation in dynamic environments. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*. 423–430.

[15] E. Hudlicka. 2011. Guidelines for Designing Computational Models of Emotions. *Int. Journal of Synthetic Emotions (IJSE)* 2, 1 (2011), 26–79.

[16] A Johansson. 2009. Constant-net-time headway as key mechanism behind pedestrian flow dynamics. *Physical review E* (2009).

[17] R. Khosla, N. Bianchi-Berthouze, et al. 2013. *Context-Aware Emotion-Based Multi-Agent Systems*. Springer.

[18] H. R. Kim, S. Y. Koo, and D. S. Kwon. 2010. Designing reactive emotion generation model for interactive robots. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. 2270–2275.

[19] T. Kuremoto et al. 2010. Autonomic Behaviors of Swarm Robots Driven by Emotion and Curiosity. In *Proc. of Int. Conf. on Life System Modeling and Sim. and Intell. Comp., and Int. Conf. on Intell. Comp. for Sustainable Energy and Env.: Part III*. LNCS, Vol. 6330. 541–547.

[20] C.P. Lee-Johnson and D.A. Carnegie. 2010. Mobile Robot Navigation Modulated by Artificial Emotions. *IEEE Transactions on Systems, Man, and Cybernetics – Part B* 40, 2 (2010), 469–480.

[21] J. Lin, M. Spraragen, and M. Zyda. 2012. Computational Models of Emotion and Cognition. *Advances in Cognitive Systems* 2 (2012), 59–76.

[22] R. Lowe. 2011. The feeling of action tendencies: on the emotional regulation of goal-directed behavior. *Frontiers in psychology* 2 (2011), 346.

[23] M. Malfaz, A. Castro-Gonzalez, R. Barber, and M.A. Salichs. 2011. A Biologically Inspired Architecture for an Autonomous and Social Robot. *IEEE Trans. on Autonomous Mental Development* 3, 3 (2011), 232–246.

[24] X. Mao, H. Bao, and Z. Li. 2008. Intelligent Emotion Decision System for Autonomous Agents. In *Proc. of Int. Conf. on Intell. Systems Design and Appl. (ISDA)*, Vol. 1. 189–194.

[25] G. Marreiros, R. Santos, C. Ramos, and J. Neves. 2010. Context-Aware Emotion-Based Model for Group Decision Making. *IEEE Intell. Sys.* 25, 2 (2010), 31–39.

[26] S. Marsella, J. Gratch, and P. Petta. 2010. Computational Models of Emotion. In *A blueprint for a affective computing: A sourcebook and manual*, K.R. Scherer et al. (Eds.). Oxford University Press.

[27] M. Moussaïd, D. Helbing, and G. Theraulaz. 2011. How simple rules determine pedestrian behavior and crowd disasters. *Proc. of the National Academy of Sciences of the United States of America* 17 (2011), 6884–6888.

[28] R. R. Murphy, C. L. Lisetti, R. Tardif, L. Irish, and A. Gage. 2002. Emotion-based control of cooperating heterogeneous mobile robots. *IEEE Transactions on Robotics and Automation* 18, 5 (2002), 744–757.

[29] S. T. Murphy and R. B. Zajonc. 1993. Affect, cognition, and awareness: affective priming with optimal and suboptimal stimulus exposures. *Journal of personality and social psychology* 64, 5 (1993), 723.

[30] R. W. Picard. 2000. *Affective Computing*. The MIT Press.

[31] L-F. Rodriguez and F. Ramos. 2012. Computational models of emotions for autonomous agents: major challenges. *Artificial Intelligence Review* (2012).

[32] L.-F. Rodriguez, F. Ramos, and Y. Wang. 2012. Cognitive Computational Models of Emotions and Affective Behaviors. *Int. Journal of Software Science and Computational Intelligence (IJSSCI)* 4, 2 (2012), 41–63.

[33] M.A. Salichs and M. Malfaz. 2012. A New Approach to Modeling Emotions and Their Use on a Decision-Making System for Artificial Agents. *IEEE Transactions on Affective Computing* 3, 1 (2012), 56–68.

[34] M. Sellers. 2013. Toward a comprehensive theory of emotion for biological and artificial agents. *Biologically Inspired Cognitive Architectures* 4 (2013), 3–26.

[35] R. Sun. 2008. *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation*. Cambridge University Press.

[36] W. Yuan and B. Zeng. 2008. Multi-Robot Task Allocation Using Abandoned-Undertaking Algorithm. In *Proc. of the Int. Conf. on Natural Computation (ICNC)*.