FAtiMA Toolkit: Toward an Accessible Tool for the Development of Socio-emotional Agents

SAMUEL MASCARENHAS, [MANUEL GUIMARÃES](#), RUI PRADA, and
PEDRO A. SANTOS, INESC-ID and Instituto Superior Técnico, Universidade de Lisboa, Portugal
JOÃO DIAS, Faculty of Science and Technology, University of Algarve, CCMAR and INESC-ID, Portugal
ANA PAIVA, INESC-ID and Instituto Superior Técnico, Universidade de Lisboa, Portugal

More than a decade has passed since the development of FearNot!, an application designed to help children deal with bullying through role-playing with virtual characters. It was also the application that led to the creation of FAtiMA, an affective agent architecture for creating autonomous characters that can evoke empathic responses. In this article, we describe the FAtiMA Toolkit, a collection of open-source tools that is designed to help researchers, game developers, and roboticists incorporate a computational model of emotion and decision-making in their work. The toolkit was developed with the goal of making FAtiMA more accessible, easier to incorporate into different projects, and more flexible in its capabilities for human-agent interaction, based upon the experience gathered over the years across different virtual environments and human-robot interaction scenarios. As a result, this work makes several different contributions to the field of Agent-Based Architectures. More precisely, the FAtiMA Toolkit's library-based design allows developers to easily integrate it with other frameworks, its meta-cognitive model affords different internal reasoners and affective components, and its explicit dialogue structure gives control to the author even within highly complex scenarios. To demonstrate the use of the FAtiMA Toolkit, several different use cases where the toolkit was successfully applied are described and discussed.

CCS Concepts: • **Computing methodologies** → **Intelligent agents**;

Additional Key Words and Phrases: Embodied agents, affective computing, cognitive architecture, emotions, social robots

ACM Reference format:

Samuel Mascarenhas, Manuel Guimarães, Rui Prada, Pedro A. Santos, João Dias, and Ana Paiva. 2022. FAtiMA Toolkit: Toward an Accessible Tool for the Development of Socio-emotional Agents. *ACM Trans. Interact. Intell. Syst.* 12, 1, Article 8 (March 2022), 30 pages.
<https://doi.org/10.1145/3510822>

The reviewing of this article was managed by associate editor Duarte, Carlos.

This study received Portuguese national funds from FCT - Foundation for Science and Technology through project: UIDB/50021/2020 and UIDB/04326/2020 and SLICE PTDC/CCI-COM/30787/2017 and by the European Commission Horizon 2020 project Rage (<http://www.rageproject.eu/>) Grant agreement No 644187.

Authors' addresses: S. Mascarenhas, M. Guimarães (corresponding author), R. Prada, P. A. Santos, and A. Paiva, INESC-ID and Instituto Superior Técnico, Universidade de Lisboa, R. Alves Redol 9, Lisbon, Portugal, 1000-029; emails: samuel.mascarenhas@gaips.inesc-id.pt, {manuel.m.guimaraes, rui.prada, pedro.santos}@tecnico.ulisboa.pt, ana.paiva@inesc-id.pt; J. Dias, Faculty of Science and Technology, University of Algarve, CCMAR and INESC-ID, Estr. da Penha 139, Faro, Portugal; email: joao.dias@gaips.inesc-id.pt.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

2160-6455/2022/03-ART8 \$15.00

<https://doi.org/10.1145/3510822>

1 INTRODUCTION

The FATiMA Toolkit is a collection of tools designed to be used in tandem for the creation of cognitive agents with socio-emotional skills. The toolkit was developed in the context of the RAGE¹ and Slice² Projects with the goal of improving the FATiMA Modular agent architecture [18]. FATiMA was initially developed in 2005 for the purpose of driving the behavior of autonomous 3D characters in a serious game about bullying [19] (the acronym stands for Fearnot AffecTive Mind Architecture). At the time, its distinguishing feature was a computational model of an emotion theory named the **Ortony, Clore, and Collins Model of Emotions (OCC)** [50] that affected the character's planning and behavioral reactions. Since then, many other computational models of social and psychological constructs were added to the architecture. Some examples include empathy [55], psychological drives [41], and culture [43]. The addition of all these constructs had a significant impact on the complexity of using the architecture. To facilitate this problem, a modular version was developed [18] that allowed users to turn off the cognitive modules they were not interested in using for their scenarios. Still, other usability issues remained in the modular version, which can also be found in other similar agent-based tools.

In a 20-year retrospective article about pedagogical agents [34], Johnson and Lester recognize that, despite the fact that there is plenty of empirical evidence of the effectiveness in the use of agents to promote learning, the lack of better tools for creating such agents is severely diminishing a more widespread adoption. Naturally, researchers have limited resources to spend on documenting, marketing, and supporting the tools they develop for doing research, and this affects adoption by outsiders. However, there are more relevant and recurring problems in the design and implementation of agent-based tools that further aggravate this issue. Some of the problems are more of a technical nature, such as the use of monolithic systems that do not favor heterogeneity of components, the framework paradigm, and how complex systems handle real-time interactions, but others can be more conceptual, such as how to explicitly structure the dialogue system (central vs. distributed approach) or the lack of tools that facilitate the process of authoring characters and situations. These problems require a more flexible approach than the ones often used in cognitive agents.

In this article we discuss some of these problems, such as the ones mentioned before, and our proposed solutions that guided the creation of the FATiMA Toolkit. In particular, we will discuss the use of meta-cognition, which, broadly speaking, refers to the human capacity of "thinking about thinking," as a solution. This capacity was represented by introducing meta-beliefs that are associated to dynamic reasoning processes, thus allowing for non-declarative knowledge to be interwoven together with declarative knowledge within a unified rule-based system. The added flexibility from this approach allowed us to tackle the identified problems, as described later in the article. Our goal is that these solutions can serve as guidelines to other researchers who are developing similar systems.

To showcase the added flexibility and ease of application to different interactive environments and settings, we will describe different case studies where the FATiMA Toolkit has been successfully applied. We will start with two commercial video games that were developed by an external company using the toolkit: Space Modules, where players have to solve problems of customers who can be more or less angry, and Sports Team Manager, where players must assemble the most optimally performing sailing team by resolving conflicts and managing the interactions between teammates. We also describe a Virtual Reality experience, a Game Jam involving FATiMA, and how a search algorithm such as Monte Carlo Tree Search was implemented in the Toolkit and applied

¹<http://rageproject.eu/>.

²<https://gaips.inesc-id.pt/slice/project.html>.

to a survival game. Finally, we will present some the results of using FAtiMA in a course on Game AI at IST, University of Lisbon, for the past 3 years and how we used the toolkit to investigate the impact of group-based emotions in a robotic partner.

In the remainder of this article we will start with a brief analysis of related work on computational models, frameworks, and tools for the development of socio-emotional agents. Next we will start by providing an overview of how the FAtiMA Toolkit works along with its core components in Section 3. In the Sections 4 and 5 we will describe the main innovations of the FAtiMA Toolkit. Agent-based frameworks have been found to create problems for their users and developers. For instance, having a monolithic structure for agent-based development is quite inflexible and restrictive; hence, we decided to use a library-based design where each component is independent of the others. In these two sections (Sections 4 and 5) we discuss how we addressed these issues and the thought process behind our solutions. The first section describes the framework paradigm and our library-based design and how the introduction of meta-beliefs and reasoners allowed us to seamlessly integrate multiple cognitive processes. The latter sections delved into how and why we moved toward an explicit dialogue structure that combines advantages of dialogue tree structures with agent-based decision-making. In the following section, Section 6, the developed authoring tool is described along with its main objectives. At last, the article concludes with the description and analysis of several different case studies that have been successfully developed using FAtiMA.

2 RELATED WORK

As a cognitive agent architecture, the cornerstone of FAtiMA is its ability to synthesize emotions in a way that is grounded in human psychology with a particular focus on social interaction. Adopting a more physiological approach, the PSI model [6] establishes a cognitive motivation-based theory that aims at modeling the human psyche by formalizing its functioning in a computer simulation that regulates its needs autonomously. In PSI, all actions and behavior ultimately have the purpose of achieving a finite set of primary, predefined drives. Instead of modeling explicit emotion categories, PSI sees emotions as emerging from specific modulations of cognitive and motivational processes. These modulations are realized by emotional parameters. Depending on the cognitive resources and the motivational state of the organism in a given situation, these parameters are adjusted, resulting in more or less careful or forceful ways of acting, as well as more or less deliberate cognitive processing. For instance, in the case of an event with a high arousal level, it is likely that the system needs to react urgently to the situation at hand, which requires a fast decision process. PSI is an example of a physiologically inspired model, of which the Cathexis model [71] and the work of Cañamero [10] are additional examples. Another type of a cognitive architecture that explores the interplay between several cognitive processes (e.g., reactive, deliberative) and sees emotions as mechanisms to regulate these processes is Sloman's CogAff framework [62]. Although these type of models are very useful to understand and explain cognitive, behavioral, and emotional processes, such complex and sometimes low-level descriptions of behavior are not helpful to design socio-emotional characters for storytelling-based scenarios, where simpler models are easier to understand and to use.

EMA [23–25] is an appraisal-based computational model that follows Smith and Lazarus's cognitive-motivational-emotive psychological theory [64] and extends SOAR [49] by incorporating emotions. In EMA, appraisal and coping are built on the causal representations developed for decision-theoretic planning augmented by explicit cognitive representations of desires, intentions, beliefs, the state of current plans, and corresponding probabilities. Considering the example where a flying bird suddenly enters the room, EMA analyzes the impact of this event in its current plans and goals. It discovers that there is a plausible hypothesis where the bird attacks and injures the

agent. This possible future event creates a threat to an existing goal (being healthy), which will be considered undesirable and will then generate an emotion. Coping strategies are proposed to maintain desirable or overthrow undesirable in-focus emotion instances.

The WASABI architecture [7, 8] follows a dimensional approach where it simulates the whole agent's emotional state as a point in a continuous three-dimensional space, PAD (Pleasure, Arousal, and Dominance). Pleasure determines the intrinsic contribution of an event to the agent's well-being. Arousal represents the agent's level of being awake and reactive to stimuli. Finally, Dominance represents the agent's perceived capability to change the environment or the cause of the event causing a given emotion. Whenever an emotionally relevant event is perceived, the pleasantness of the event is appraised, generating an emotional impulse that will change the current level of pleasure and arousal in the PAD space, and which will slowly decay toward a neutral level. The level of Dominance is determined from the situational context from the cognitive layer of architecture. Primary emotions are used to trigger the corresponding facial expressions, while secondary emotions are sent to the deliberation component, which can use this information to activate particular behaviors. We believe that this mechanism of making decisions based on emotional state information is important for an emotional agent architecture. However, while in WASABI the actual decision-making rules are tightly coupled to the implementation, in the FAtiMA Toolkit we adopted a declarative approach that allows authors to specify their own rules on how emotions affect behavior, by using meta-beliefs about the agent's emotional state.

The Virtual Human Toolkit [31] is a well-known architecture that is also designed to facilitate the creation of autonomous conversational characters. Moreover, the architecture is also highly modular. However, unlike the FAtiMA Toolkit, the modules that are provided are focused in handling aspects that are more related to the embodiment of a character rather than its cognitive abilities. For example, the toolkit includes a multimodal sensing framework [58] and a nonverbal behavior generator [39] that is based on the BML standard [36]. Another distinguishing aspect is that the Virtual Human Toolkit adopts a framework approach that enforces the use of a message broker middleware for communicating with the available modules.

GAMYGDALA [51] is a computational model of emotions that is based on the OCC theory [50]. Similar to the FAtiMA Toolkit, this emotional appraisal engine was designed to be more accessible to game developers. In the perspective of the authors, one of the issues with affective models that are tightly coupled to a specific reasoner is that their use is limited to games where that specific type of reasoning is desirable. With that in mind, GAMYGDALA implements an appraisal mechanism that is independent from the how the characters reason. Essentially, developers need only to provide a list of goals for each character and then specify which events will block or facilitate each goal. Based on that information, the engine will determine the changes made to the character's emotional state. Similar to GAMYGDALA, the appraisal engine in the FAtiMA Toolkit is an independent component that is not tied to a particular type of reasoning. However, this is achieved in a different manner. Namely, developers are able to incorporate their game-specific AI as meta-beliefs, which can then be used by the appraisal process as regular beliefs.

The concept of meta-cognition, which refers to "cognition about cognitive phenomena" [30], is one of the inspirations behind CLARION. CLARION [66] is an example of a modularly structured cognitive architecture that explicitly represents meta-cognitive mechanisms to enable a dual-process model of explicit and implicit cognition. In [67] the author states that "without meta-cognitive control, a model agent may be blindly single minded." In order to promote a greater flexibility in its core cognitive algorithms, the notion of meta-cognition was also integrated into the FAtiMA Toolkit.

In the past, when authoring complex intelligent agent scenarios, using frameworks such as the one's in Table 1, authors would rely on a team of engineers or the developers of the architecture

Table 1. Table Comparing Different Agent-modeling Frameworks in Terms of Their Base Cognitive Theory and Emotional Models Used

Name - Year	Base Cognitive Theory	Emotion Theory
ALEC - 2003	CLARION	Appraisal: Sloman [61]
FAtiMA - 2005	BDI	Appraisal: OCC Theory [50]
EMA - 2009	BDI, Planning	Appraisal: Smith and Lazarus [63]
SOAR-EMOTE - 2009	SOAR	Appraisal: Scherer [57], Gross and John [26]
WASABI - 2009	BDI	Dimensional: PAD [47]
GAMYGDALA - 2014	-	Appraisal: OCC Theory [50]

to perform the manual task themselves. However, recent developments in the **Natural Language Processing (NLP)** and **Machine Learning (ML)** fields have led to the automatization, or at least partial automatization, of this task. For example, in the Human-Robot Interaction field, data-driven methods for learning non-verbal behavior, hand gestures, eye gaze, and animations from human behavior have been extensively applied [16, 37, 56, 73].

The generation of social behavior, on the other hand, is a broad and ill-defined problem, and thus hard to learn. Research has focused on how to extract information from existing datasets that contain social examples such as stories using NLP. In the AI Planning field this approach has been quite used [60]: Story Framer uses Natural Language descriptions, written by users, as input and is able to learn planning domain models [32, 42]. In a similar vein, Janghorbani et al. [33] introduced an authoring assistant tool to automate the process of domain generation from natural language description of virtual characters with the objective of bridging the authoring and the planning tasks.

Stories provide quite a good base for models to learn social behavior from. As a consequence, researchers have focused on creating automated tools for story comprehension [17] and to generate narrative plots themselves [68]. For instance Ammanabrolu et al. [2] developed an automated storytelling tool to order the plot of stories with causal and commonsense knowledge [3] using COMET [9], a transformer-based language model designed for commonsense inference.

Other approaches to easing the authoring burden associated with Social Agent architectures is to automatically generate content for interactive scenarios, the most common one being dialogue, in particular, taking into account the high importance of emotions in these contexts, affective dialogue. Language models have been used to automatically create affective dialogue with promising results [5, 21]. For example, Antunes et al. [4] was able to use the GPT-2 model [52], which has been quite successful in open-domain, affective, and data-to-text generation, and condition it with emotionally connotated event structures to create a model that supports the social sharing of emotions in human-agent communication [4]. Goswamy et al.'s work [22], for instance, the authors were able to incorporate affective information within the GPT-2 model without compromising on the grammatical correctness of the outputted result, outperforming existing affective text generation models. We will discuss some of the ongoing work of integrating these techniques in the FAtiMA Toolkit at the end of this article.

3 FATIMA TOOLKIT AND ITS CORE ELEMENTS

The FAtiMA Toolkit is a collection of tools and assets designed for the creation of characters (virtual or robotic) with social and emotional intelligence. The main focus of the framework is not on mimicking the behavior of a real person; rather, inspired in Social Science and Emotional Theories, its goal is to create socio-emotionally agents that can interact with humans and other agents in

Fig. 1. FAtiMA Toolkit components.

interactive scenarios, scenarios such as Social Skills Training Environments and serious and even commercial games, as we will describe in Section 4.

Figure 1 shows a diagram of all the existing components within the FAtiMA Toolkit. Each is the result of continued iterations and experience accumulated while working with developers since the beginning of the RAGE Project, now transitioning into the Slice Project. The entire toolkit is open source, available on GitHub.³ Moreover, the FAtiMA Toolkit's official website⁴ gathers all available information regarding the tool including Training Materials, Publications, Demonstrations, Tutorials, and Starter Kits.

An affective agent created with the FAtiMA Toolkit will be initialized by the Role Play Character component. The name of the component was chosen based on the notion that FAtiMA agents are typically designed to be able to act according to a specific social role in a coherent manner. This component is responsible for managing the structures that constitute the internal state of the agent: their beliefs, memory, and emotional state. Authors can also add more complex and believable behavior to the agents, which we call Role Play Characters, by using other components. For instance, in order for agents to be able to perceive events in the world around them and have emotional reactions, the Emotional Appraisal component can be used. To define the agents' ability to make decisions based on their beliefs and goals, one can use the Emotional Decision Making asset. All of these different components were developed as libraries, which the developers or authors can directly import and explore more easily without having to worry about compatibility issues. This is one of the main contributions of this work and we will further discuss it in Section 4. We will now describe the core elements that compose the FAtiMA Toolkit.

³<https://github.com/GAIPS-INESC-ID/FAtiMA-Toolkit>.

⁴<https://fatima-toolkit.eu/>.

3.1 Naming Convention

The FAtiMA Toolkit uses a structured naming scheme called **Well-Formed Names (WFNs)**. These names are strings that follow a set of syntactic rules, which are based on the notions of symbols and variables inspired in first-order logic:

- **Symbols** represent constant entities (actions, objects, name of properties, name of relations); e.g.,: |Attack|, |John|, |Table|, |Brown|, |Loves|, |A1|, etc.
- **Variables** are written between square brackets and represent an entity or value that is not specified yet, which can be replaced by symbols; e.g.,: |[x]|, |[target]|.
- **Composed names** represent properties or relations between two or more Symbols; e.g.,: |Color(Sky, Blue)|, |Has(John, [x])|.

3.2 Beliefs

Beliefs represent the internal state of the agent, its beliefs about the surrounding world including itself and other agents. In the FAtiMA Toolkit they are represented by Composed Names, which have a value attributed to them. There is no rigid convention for how beliefs are declared as long as they are consistent across the scenario. For instance, there are multiple ways of representing a belief in which the agent has five apples:

```
Has(Apple)=5
Has(Apple, 5)=True
Food(Apple)=5
Apple(5) = True
```

3.3 Events

As typical of an agent model, the main input are events. These can refer either to properties that change or to actions that occur. Events are represented as WFNs. More precisely, they are composed names, a tuple, that has five different elements:

```
<type , subject , action , target , timestamp >
```

The |type| represents the type of event that occurred. It can be an “Action-Start,” “Action-End,” or “Property-Change” event; the type is related to its effect on the decision making or knowledge base of the agents. The |subject| is the one responsible for the |action| toward |target|. Finally, the |timestamp| represents the timeframe in which the action was executed; if it is not specified that agent stores the event with the timeframe it is currently in. For example:

```
Event(Action-End, Sam, Turn(On), TV)
Event(Property-Change, Sam, State(TV), On)
```

The first aforementioned example corresponds to an event where agent |Sam| finishes the action of turning on the TV and the follow-up example is an event indicating that the property |TV(Is)| changed its value to |On| due to agent |Sam|.

3.4 Role Play Character

The *Role Play Character* handles the events in the following manner. If the event is a property change, the corresponding belief in the **Knowledge Base (KB)** is updated accordingly. If the event concerns the performance of an action, the agent checks if it recognizes the agent performing it. If not, the performer agent is added to the list of known agents. Subsequently, the event is processed by the *Emotional Appraisal* component, which updates the emotional state of the agent

with new emotions that the event might trigger. As will be detailed later, this process uses a rule-based approach. Finally, the event is stored in the *Autobiographical Memory*, registering the event's logical name, the current time, and the emotions it caused.

The output of the *Role Play Character* are actions or, more precisely, high-level decisions about which actions the agent should perform next. These actions can be read by external components upon request. The execution of these actions is left in charge of modules that are more closely linked to the agent's embodiment, such as a behavior scheduler or a speech synthesizer. This type of module falls outside the scope of the FATiMA Toolkit. To reach a decision, the *Role Play Character* executes the algorithm defined in the *Emotional Decision Making* and returns the result. Similar to the *Emotional Appraisal*, this algorithm also uses a rule-based approach. More precisely, it goes through the list of existing decision rules and applies, in each of them, a general unification algorithm between the list of the rule's activation conditions and the agent's set of beliefs. The algorithm will find all the possible sets of valid substitutions that satisfy the rule's conditions. Each of these substitution sets will originate a concrete decision to execute a certain action. To illustrate this process, consider the following example of a decision rule:

```
Action: Eat
Target: [ object ]
Priority: 1
Conditions:
  Has([ object ]) = True
  Edible([ object ]) = True
```

A single variable is present: `[object]`. As such, for this action to be executed a valid substitution must be found for it. Consider the following set of beliefs: "`|Has(Bag)=True|, |Has(Apple)=True|, |Edible(Apple)=True|`". In this situation the unification algorithm would unify `[object]` with `|Apple|` but not with `|Bag|` as the latter fails to pass the last conditional check.

By expressing the agent's decisions using conditions with logical variables, the action space of the agent grows and adapts according to its beliefs. This means that if an agent is able to find another object that it believes to be edible, the previous rule would enable the action of eating it automatically. However, this type of rule-based reasoning by itself can be quite limiting when trying to handle scenarios that require a more complex decision-making strategy. For example, consider an agent that is intended to be able to play a game of go or poker in a proficient manner. In that case, it would be quite difficult to come up with decision rules based on logical belief checks. Because of this limitation, it is quite common to find other algorithms governing the decision making of an agent, such as heuristic search, planning, reinforcement learning, and deep learning. Any of these algorithms will be more advantageous in a particular domain and less so in others. For instance, deep learning has become immensely popular recently, but it typically requires a large corpus of data to be useful. Also, the original version of FATiMA included a classical continuous planner, although it was difficult to explore its full capabilities in some conversational domains.

3.5 Knowledge Base

Although the notion of beliefs and meta-beliefs has been already discussed, there are two important aspects that have not yet been mentioned. The first aspect is that the *Knowledge Base* is capable of storing beliefs about the beliefs of other agents. In humans, this ability is referred to as a Theory of Mind. It was implemented here by specifying an additional field for each belief named `[perspective]`. The value of this field corresponds to the name of the agent that holds the

belief. By default, the unification algorithm uses the agent's own perspective when searching for valid substitutions. However, this can be changed by using the following meta-belief in a condition: $|\text{ToM}([\text{agent}], [\text{belief}])|$. This meta-belief will retrieve the value of $[\text{belief}]$ from the perspective of $[\text{agent}]$. Other meta-beliefs can also use this form of perspective taking within their algorithm as is the case of the *Emotional Appraisal* component when it is inferring the emotional state of other agents. The final important aspect about beliefs is that they also have a certainty value associated to them that can range from 0 to 1. The purpose of this value is to affect the perceived utility of each decision as will be explained later when describing the *Emotional Decision Making* component.

3.6 Emotional State

The *Emotional State* is the structure that holds the set of active emotions the agent is currently experiencing. FAtiMA's emotional model is based on the OCC theory of emotions [13], which is an appraisal theory that follows a categorical approach. Formally, an emotion is defined as a tuple with five elements:

$\langle \text{type}, \text{valence}, \text{intensity}, \text{cause}, \text{target} \rangle$

The $|\text{type}|$ of the emotion represents its categorical name. In total, the OCC theory proposes 22 emotion types that are psychologically distinct. Examples include Joy, Fear, Distress, Pride, and Reproach, among others. The $|\text{valence}|$ of an emotion indicates if the emotion is positive or negative, which depends on its type. For instance, Joy is considered to be positively valenced, while Distress is viewed as negative. The $|\text{intensity}|$ of an emotion is a number greater than zero that indicates how strongly is the agent experiencing the emotion. As times goes on, this value becomes smaller until it finally reaches zero and the emotion is removed from the emotional state. Finally, the $|\text{cause}|$ corresponds to the event that triggered the emotion and the $|\text{target}|$ is the name of the agent or object that the emotion is directed at. An emotional state can have multiple emotions with their causes and targets, which can be either manually added by the author by creating a tuple for each emotion or automatically generated by the Emotional Appraisal component. Below is an example where the agent's internal state has an emotion of the type "Love" with an intensity of 2 towards "Peter." The cause id corresponds to an automatically generated identifier that is attributed to an event in the agent's autobiographic memory.

Emotion Type: Love

Intensity: 2

Cause ID: 0

Target: Peter

One advantage that FAtiMA has in using a categorical approach to emotions, in comparison to a dimensional model, is that it is possible for agents to experience multiple emotions at the same time with different intensities. For example, the agent might experience Distress after losing its queen on a chess game while simultaneously feeling Admiration for its opponent for doing a very good play. Another advantage is that it is also possible to specify different thresholds and decay rates for different emotion types. This can be used to simulate certain personality traits through the agent's emotional biases. For instance, using the previous example, an agent that is meant to be perceived as highly competitive in the chess example should be set with a lower threshold for Distress and a higher threshold for Admiration.

Besides the list of active emotions, the emotional state also implements the $|\text{Mood}([\text{agent}])|$ meta-belief that was previously mentioned. The concept of mood is an affective state that is more broad in the sense that it is not typically provoked by a specific event. Similar to emotions, mood is also

characterized by a positive or negative valence, but its duration is longer. Based on these notions, mood is represented in FATiMA as a meta-belief whose value ranges from -10 to 10. This value is updated each time a new emotion is added. In the case of positive emotions, the mood value will increase based on the emotion's intensity. Conversely, negative emotions decrease the mood of the agent, also based on their intensity.

3.7 Emotional Appraisal

The purpose of the *Emotional Appraisal* library is to generate new emotions by doing a subjective assessment of the events received as input. The notion that emotions should be generated in this manner stems from the ideas put forward by appraisal theories of emotion [48]. In a broad sense, these theories propose that the reason people have different emotional responses when facing the same situation is because they conceptualize it in a different manner. More specifically, people will make a series of judgments about a situation and the result of those judgments will determine the emotional response.

The appraisal model that is implemented inside the *Emotional Appraisal* library uses appraisal rules to simulate the aforementioned psychological process of judging events. Formally, an appraisal rule has the following components:

`<event , target , appVariables , conditions >`

The parameter `|event|` is a well-formed name that is used to match the name of the events received. The `|target|` is the name of the agent or object that the rule applies to. The third element, named `|appVariables|`, consists of a list of appraisal variables and their respective numerical values. Finally, `|conditions|` is a list of logical conditions that need to be verified according to the agent's beliefs in order for the rule to be activated.

For the purpose of illustrating how the specified algorithm uses these rules to generate emotions, consider the following simple example:

```
Event : Event ( Action-End , [ x ] , Smile , SELF )
Target : SELF
AppVariables :
    Desirability = [ d ]
Conditions :
    RapportLevel ( SELF , [ x ] ) = [ d ]
```

Considering this rule, what would happen when an agent, named Sam, perceives the agent John smiling, which is represented as `|Event(Action-End,Sam,Smile,John)|`? First, the algorithm tries to unify the rule's `|event|` with the event name. In this case, the unification is possible with the following substitution set: $\{[x]/\text{John}\}$. It is important to mention that `|SELF|` is a special symbol that can be used in a rule and is always replaced by the name of the agent, in this case `|Sam|`. Because at least one substitution set was found, the next step is to unify the rule's `|conditions|` with the agents' beliefs, taking into account the already existing substitutions from the previous step. If we assume that agent Sam believes $\text{RapportLevel}(\text{Sam}, \text{John}) = 5$, then the rule's `|conditions|` are made true by applying the following substitutions: $\{[x]/\text{John}, [d]/5\}$. The final step is now to generate an emotion based on the type and value of the appraisal variables specified in the rule. Following the OCC theory [50], a positive desirability results in the generation of a Joy emotion whose intensity is linearly proportional to the desirability value. Assuming this intensity is higher than Sam's threshold value of Joy, the emotion is added to the *Emotional State* and becomes active until its intensity decays below a minimum value.

As we mentioned, the framework follows the model proposed by Ortony, Clore, and Collins [50]. In order to support the full range of OCC emotions, FAtiMA uses five different appraisal variables:

- **Desirability:** Are the consequences of an event desirable for me?
- **Desirability for others:** Are the consequences of an event desirable for other(s)? Combined with the “Desirability” appraisal variable, it describes emotions related to the Fortune of others and how desired that fortune is for the agent.
- **Praiseworthiness:** Is the action praiseworthy within the agent’s social norms?
- **Goal Probability:** How far is the agent from reaching a particular goal?
- **Like:** Is the action/object/agent appealing?

The relation between these variables, their valence, and how each emotion is computed is briefly described in the OCC Model. During the implementation of those descriptions, we’ve made some adaptations as there was no exact equation on how, for instance, the intensity is calculated. Nevertheless, FAtiMA is able to generate all 22 different emotions described in the OCC Model of Emotions.

3.8 Emotional Decision Making

As previously mentioned, the purpose of the *Emotional Decision Making* is to deliberate on the action that the agent should perform next. This is done using a rule-based mechanism that is able to consider both the beliefs (regular and meta) of the character and its emotional state. Decision rules are defined as the following:

`<action , target , conditions , priority , layer >`

The attribute `|action|` is a well-formed name that can be either a symbol (e.g., Smile, Eat) or a composed name (e.g., Offer(Gift)) in which the inner symbols are treated as the action’s parameters. The `|target|` of the action is the name of the agent or object that the action is aimed at. The parameter `|conditions|` is a list of logical conditions used to check if the action can be executed. The `|priority|` is a numerical value that represents how important it is for the agent to execute the action. Finally, the `|layer|` parameter is a symbol that indicates in which layer is the action considered in the decision-making process. It allows the agent to consider different groups of actions at different time scales. Typically, there is a need for at least two layers when it comes to decision making, one that is reactive and allows for fast decisions and another that is slower and focused on more strategic decisions. Note that the combination of these two layers is usually referred to as a hybrid agent architecture [11].

When describing the *Role Play Character* and the notion of meta-beliefs, we already used some examples of decision rules and how they work concerning the logical unification process. As such, we will now focus on the purpose of the last two attributes of a decision rule, namely `|priority|` and `|layer|`. Typically, there will be multiple decision rules in the chosen configuration of the *Emotional Decision Making* component. Even if there is a single decision rule, it can generate multiple actions as a result of finding more than one set of valid substitutions. In that case, the agent needs to deliberate on which concrete action it should execute. This is what the `|priority|` of the decision rule is used for. More precisely, the *Emotional Decision Making* will return the list of all actions found ordered based on their `|priority|` score, from highest to lowest. Note that the value of `|priority|` can be a logical variable that is determined by using both regular and meta-beliefs, similar to the appraisal variables in an appraisal rule. If that is the case, a numerical substitution must be found in the rule’s `|conditions|` or the action is not selected. Alternatively, the value of `|priority|` can also be defined as a constant value. Finally, the `|priority|` score of a decision rule is affected negatively if

the substitution set contains uncertain beliefs. This is done by multiplying the priority score with the certainty value of the beliefs that were used in the rule's substitutions.

The following example captures an "Offer Gift" action toward a [target]. The variable's value will be calculated by the unification algorithm using the agent's knowledge base; for example, in this case, if the agent was to have "CloseFriend(SELF, Peter) = True" belief in its KB, the variable [target] would be substituted by "Peter." The final condition pertains to the Emotional State of the agent; if the agent is in a "good" mood, then the condition will be true. It uses a meta-belief to calculate the mood, which we will address in Section 4.2.

```
Action: Offer ( Gift )
Target: [ target ]
Priority: 1
Conditions:
    CloseFriend ( SELF , [ target ] ) = True
    Mood ( SELF ) > 2
```

One of the typical challenges associated with creating a decision-making solution for a virtual agent or a robot is being able to handle real-time interaction. This means that some of the agent's decisions must be made rather quickly in response to what is happening in the environment and what the other agents are doing. The purpose of the |layer| attribute is to solve this issue. Essentially, this attribute is used to group decision rules by their complexity and purpose. For instance, decision rules related to facial expressions can be labeled as "Expressive," whereas rules to decide upon which move to play in a given game can be labeled as "Deliberative." Then, when making a decision request, it is possible to specify a specific |layer| as a parameter and the algorithm will only process decision rules of that |layer|. Without this mechanism, every decision rule would have to be analyzed at each request, which would quickly become too inefficient for handling real-time interaction.

4 AN ACCESSIBLE INTELLIGENT AGENT'S TOOLKIT

The FATiMA Toolkit was specifically developed with the intent of making agent-based solutions more accessible and reliable for its users. Taking this into account, we worked with several different authors, developers, and students in order to identify the main problems when using this type of framework. In this section we will describe these issues and how we addressed them; we hope that our approach and solutions can serve as a guideline to other researchers that are developing similar systems.

4.1 Library-Based Design

Many agent-based tools are developed by adopting a framework paradigm. More precisely, they define a rigid monolithic structure that acts as the skeleton of the whole system. For instance, in the case of FATiMA Modular, the architecture dictates a core algorithm that implements the perception-action cycle of an agent. Within this core algorithm it is possible to add or remove components of different types. Due to this problem, the design of the core algorithm is faced with the highly difficult challenge of having to be generic enough to encompass all possible forms of a perception-action cycle in an agent.

Another very important issue in the use of a framework paradigm is that, generally speaking, frameworks are hard to combine together due to their nature. Because of this issue, the risk in adopting a particular framework for a given project is substantially higher. In the context of the

RAGE Project where we got to work with several different developers and received feedback from people within the industry, we feel that this type of concern is quite important, especially in a large project. Additionally, it also plays a major role in explaining why, for instance, only a small number of game engines are actually used in the video game industry and many companies still decide to develop an in-house solution.

During the development of the FAtiMA Toolkit, we transitioned from a framework paradigm to a library-based approach. Library-based design means that the components work as individual libraries that can be used on their own. From a software standpoint, this equates to while having a complex and robust architecture, the pieces that make part of it can function independently and can be picked apart to be used in another puzzle.

In the FAtiMA Toolkit context it means that, if developers are only interested in having a solution to generate emotions, they are able to use the *Emotional Appraisal* library in isolation. If a developer is only interested in having a character that has an Autobiographic Memory, then it can simply use the *Role Play Character* library alone.

As we've mentioned, the components can be used autonomously; however, they work best together. The configuration shown in Figure 1 depicts how the existing components are combined together to create an agent with emotions and a complete and rich perception-action cycle.

4.2 Meta-beliefs

In Section 3.4 we discussed the main limitation of rule-based approaches: they become difficult to handle once scenarios require more complex decision-making strategies. The solution we adopted to overcome this problem was to move toward meta-cognition. The term, which was coined by Flavell [30], refers in a broad sense to the human capacity of “thinking about thinking.” Humans are capable, to a certain extent, to regulate and become conscious of the ways they make decisions in different situations. Due to its broad definition, several terms and phenomena have been associated with metacognition by different researchers [70]. For the purpose of this work, we focused on modeling aspects that pertain to an individual's procedural knowledge related to decision making and emotional awareness. More precisely, we defined a formal notion of meta-beliefs that was added to the *Knowledge Base* in addition to the notion of regular beliefs. On a conceptual level, the key distinction to be made between a meta-belief and a regular one is that the latter refers to descriptive knowledge, whereas the former refers to knowledge produced by affective or reasoning processes. For example, in a game of chess, a person will have beliefs about the names of the different pieces that are on the board as well as their positions and color (e.g., “I believe the piece on position b4 is a black queen”). These would fall into the category of descriptive beliefs. However, when playing the game, the person will also form beliefs about their best moves and act according to those beliefs, assuming they want to win (e.g., “I believe my best move now is to move the black queen from b4 to f4”). Similarly, a person might also have beliefs of their own emotional state or the emotional state of the adversary (e.g., “I believe my opponent is angry due to losing her queen”).

On a practical level, the notion of meta-belief was implemented in the following manner. Similar to a regular belief, its syntactic structure is a composed name with a root symbol and a sequence of literals within parentheses that are separated by commas. The difference is that meta-beliefs are registered in the *Knowledge Base* as procedures, with their name becoming a reserved keyword. As such, the unification algorithm is able to identify them when parsing the composed names within a logical condition. Whenever it encounters a meta-belief it will dynamically execute its associated code to retrieve its value. The code can be either very simple or quite complex, depending on the nature of the meta-belief itself. For example, one of the meta-beliefs that the *Role Play Character*

adds is $|Mood([a])|$, which simply retrieves the mood value of agent $[a]$. Comparatively, the meta-beliefs added by the *Social Importance* and the *CiF-CK* modules are more complex. These two modules are labeled as *Reasoners* in Figure 1 precisely because their main purpose is to augment the reasoning capabilities of the agent's decision making and appraisal through the addition of meta-beliefs. The "(...)" displayed in the diagram is to indicate that external developers can also extend the toolkit by adding additional reasoner modules with their own set of meta-beliefs without changing the toolkit's source code. Going back to the chess example, one possibility would be to add a chess engine as a reasoner module that would add the meta-belief " $BestChessMove([board-id],[player-color],[depth])$ ". That meta-belief could then be used in a decision rule like this:

```
Action : PlayChess
Target : [m]
Priority : 1
Conditions :
    Playing ( Chess ) = [ id ]
    PlayerColor ( SELF , Chess ) = [ c ]
    BestChessMove ( [ id ] , [ c ] , Mood ( SELF ) ) = [ m ]
```

As this example illustrates, one of the advantages of meta-beliefs is that they can be used inside logical conditions in an identical manner as regular beliefs. Moreover, they can be nested inside each other. The nesting is possible because each meta-belief has full access to the contents of the *Knowledge Base*, which includes both regular and other meta-beliefs. In this particular case, $|Mood(SELF)|$ is nested inside $|BestChessMove|$ with the aim of making the $[depth]$ parameter of the chess search depend on the agent's current mood. The result would be that the agent, similar to a human, would come up with better moves in the game if it was in a good mood. To summarize, the key contribution of meta-beliefs, as defined in this manner, is that they allow for multiple reasoning and affective processes to be interwoven together under a unified rule-based system. We will now discuss in further detail some of the components that have previously been mentioned.

4.3 Reasoner Components

As stated previously, the role of reasoner components is to augment the capabilities of the *Emotional Appraisal* and the *Emotional Decision Making* by adding meta-beliefs that can be used in logical conditions. Therefore, their use is optional, which is important as they may require additional authorial effort. With that said, the two reasoners that are shown in Figure 1 work in conjunction to augment the ability of the agent to perceive and act appropriately in different socio-cultural contexts. More precisely, the first one, named *Social Importance*, was originally added to FAtiMA in order to be able to create groups of agents that would act and feel according to different cultural values [43]. The second one, named *CiF-CK*, is based on a model [29] that was used to add social behavior to Non-Player Characters in the popular game Skyrim [65]. The model is an adaption of the "Comme il Faut" architecture [46].

4.3.1 Social Importance. This component adds the meta-belief $|SI([target])|$, which calculates the amount of **social importance (SI)** the $|[target]|$ agent has in the agent's perspective. This is a numerical value, ranging from 1 to 100, that signifies the extent to which the agent is willing to act in the interest of the $|[target]|$ agent. The concept is based on a sociological theory of human motivation that conceptualizes human action in accordance to two fundamental relational dimensions: status and power [35]. According to this theory, we are motivated to confer status to those who deserve it and refrain from claiming more status than what we perceive to have in the eyes of others. The amount claimed and conferred by each action depends a great deal on cultural

conventions, but generally, the higher the request, the more status it claims. For a FAtiMA agent, the purpose of SI is to implement this notion of status so it can better navigate and understand different relational contexts. For instance, a stranger asking a personal question is inappropriate but not a friend. This is because, all else being equal, the SI of a friend should be higher than the SI of a stranger. These conventions are implemented in the component as attribution rules, which are defined as the following tuple:

```
<target , conditions , siValue >
```

These rules work in a similar manner to the appraisal and decision rules described previously. Essentially, the unification algorithm will process each rule individually and try to find valid substitutions for the rule's conditions. If so, the `|siValue|` defined in the rule is added to the `|target|`'s total SI. Note that these rules can refer to beliefs about properties of other agents, such as whether or not they are a family member, or they can refer to their past actions in the environment, such as the amount of times they were rude toward the agent, for example. Also, by using the meta-belief `|ToM([target], SI(SELF))|`, the agent is also capable of inferring how much SI it has in the perspective of the `|target|` agent. This is done by using the beliefs associated to the `|target|` agent perspective when going through the attribution rules. The following example describes a simple Social Importance Rule where the agent has an SI value (20) attributed to their close friends:

```
Action: Close Friends
Target Variable: [ t ]
Value: 20
Conditions:
    CloseFriends([ t ]) = True
```

4.3.2 CiF-CK. In most scenarios that are focused on social interaction it is likely there will be various social conventions that agents should follow such as introducing themselves in the beginning of a conversation or saying goodbye at the end. The purpose of the *CiF-CK* component is precisely to make agents aware of these existing social exchanges as well as capable of executing them appropriately. The main motivation is to be able to automatically add a layer of social behavior that can be reused across several domains.

The primary knowledge structures of *CiF-CK* are *Social Exchanges*. These are meant to represent an established sequence of interaction steps between two agents. Formally, they are defined as the following tuple:

```
<name, target , steps , conditions >
```

The parameter `|name|` is used to identify the type of social exchange (e.g., Flirt, Greet) and `|target|` is the name of the agent with whom the social exchange is to be performed. The parameter `|steps|` is a list of symbols that represents the sequence of actions that both agents perform in a back-and-forth manner (e.g., `{|Initiate, Answer, Finalize|}`). Finally, `|conditions|` are logical conditions that can be of two types: (1) `|starting-conditions|` and (2) `|mode-conditions|`. The first type is used to define when a Social Exchange can start, whereas conditions of the second type are used to compute the desire an agent has to perform a specific social exchange in a certain mode. As such, each `|mode-condition|` has the name of the mode it refers to (e.g., Positive, Negative, Sarcastic) and it also has a value associated with it. The desire for the agent to perform a step of the social exchange using a particular mode is given by the sum of the values of all its `|mode-conditions|` that are verified. The following is an example that captures what we've just discussed; the target only positively responds if it is in a good mood.

Social Exchange Name: Flirt
 Description: Teases [t]
 Target: [t]
 Steps: Initiate , Response
 Starting Conditions:
 SI([t]) > 10
 IsAttractive([t]) = True

Mode Condition 1:
 Value: 5
 Mode: Positive
 Condition:
 Mood(SELF) > 5

Mode Condition 2:
 Value: 3
 Mode: Negative
 Condition:
 Mood(SELF) < 0

The social exchanges that are defined in the *CiF-CK* component can be used in decision rules through the following meta-belief: $SE([se-name],[target],[step],[mode])$. This meta-belief will automatically find valid substitutions for the defined variables and return a numerical value representing the desire of the agent to perform a specific step of a social exchange in a certain mode. For example, the value of $SE(Flirt, John, Initiate, Sarcastic)$ is determined by first checking if the starting conditions of the “Flirt” social exchange are verified and the current step in the execution order is “Initiate.” If all this applies, the algorithm then sums and returns the value of all the “Sarcastic” conditions that are verified. The returned value can then be directly linked to the priority of the decision rule that uses this meta-belief.

Note that there is a synergy between the *CiF-CK* and *Social Importance* component as it is possible to use conditions based on SI in the conditions of a social exchange, as shown in the example above. In the following section, we will now explain how a decision that involves dialogue such as the previous example of flirting is handled by the toolkit.

5 EXPLICIT DIALOGUE STRUCTURE

As previously mentioned, the first use case of FATiMA was to drive the behavior of non-player characters in FearNot!, a serious game about bullying. Rather than offering a linear story for the player to experience, the game chose to use an educational role-play approach, in which the story emerges from the interactions of characters playing their designated role (e.g., bully, victim, victim’s friend, etc.). Like the player, each autonomous agent is controlling their designated character according to their own knowledge and perspective about the world and the other characters. Given that the dialogue in the story emerges from the individual decisions of the different agents, this approach to interactive storytelling is commonly referred to as being distributed or character centric [12]. Alternatively, the other common approach for designing an interactive storytelling system is the centralized or plan-based approach. In it there is a central AI component, typically referred to as a drama manager. This component is in charge of planning and directing all the dialogue actions

Fig. 2. Simple dialogue tree with each circle representing a dialogue state and each arrow a dialogue action.

that are made by each character and then adapting the story plan in accordance to how the player acts. The game *Faade* [45] is a well-known interactive drama where this approach was used with success. There are both upsides and downsides in each approach. On one hand, the central planning approach provides a greater deal of control and predictability to the author, but its rigidity makes it difficult to scale up. On the other hand, the distributed approach allows for much more flexibility in the way characters respond but loses in terms of predictability of the outcome, which also makes it hard to scale up when trying to ensure a particular authorial or educational goal from the narrative.

Neither of the two approaches mentioned have been largely adopted in the video games industry so far due to its more risk-averse nature. For the most part, developer studios still rely heavily on building scripted cut scenes mixed with multiple-choice options for the player. This usually leads to a branching narrative format. Consequentially, most game developers become accustomed to creating and managing dialogue trees, preferring to find ways by which their practical limitations can be compensated rather than switching to a distributed or centralized planning approach. Motivated by this preference, most commercial tools for building dialogue systems for video games also use dialogue trees as their principal component.

Because dialogue trees are so prevalent in game development, we decided to incorporate them in the FAtiMA Toolkit but maintaining a character-centric paradigm. This led to a hybrid solution that is meant to be more familiar to a game developer while still maintaining the flexibility and richness of having multiple autonomous agents controlling the dialogues of their characters. To explain how the proposed dialogue system works, consider the tree structure shown in Figure 2. The same structure could be represented in the following tabular format as demonstrated by Table 2.

Creating this simple dialogue between a player and a character would be quite straightforward in a system that is based on dialogue trees. However, if one were to use an approach based on autonomous agents, then the tree itself would no longer be explicitly represented anywhere. Instead, it would have to be implicitly defined in the agents' goals, beliefs, internal state, and actions. As such, to achieve the same effect, the author would have to configure an agent whose cognitive decisions would lead it to perform the dialogue action *d3* after the player performed *d1* and *d4* after the player performed *d2*.

For such a simple dialogue, the agent-based solution is less practical. However, dialogues are usually quite less constrained than the example given. As one may notice, both in *s1* and in *s2* there is only one possible response for the character to select. As such, there is no need for a decision-making process. But what if we want to add multiple possible replies for the character to select from? For instance, consider that in state *s3* we would like to add the dialogue shown in Table 3, when the character is in a negative mood.

Table 2. Detailed Dialogue States Described in Figure 2

Id	Current State	Next State	Utterance
d1	s1	s2	What are you doing?
d2	s1	s3	How are you feeling?
d3	s2	s4	Nothing special.
d4	s3	s4	I am feeling great.

Table 3. A New Dialogue State

Id	Current State	Next State	Utterance
d5	s3	s4	None of your business.

In a system with dialogue trees, adding this dialogue for a character requires the existence of a conditional mechanism that will select it in the right context. In this case, an example would be a rule associated to *s3* that would check the mood of the character and select *d5* in case it was negative and *d4* if it was positive. The scalability issue of this solution becomes apparent when we want a range of characters that exhibit diverse dialogue strategies in the same contexts. At this point, the rules encoded in each state must account for all the desirable character variations. This is the point where an agent-based solution becomes more valuable as each agent will have its own separate logic to select dialogue actions. However, due to this separation, it becomes difficult to grasp how the narrative flows.

The solution adopted in the FATiMA Toolkit was to have an explicit dialogue tree structure but with a crucial distinction. The representation of the tree structure does not contain any logic that decides what the characters say. Instead, it merely informs agents controlling characters of how many options they can choose from in a given dialogue state. Essentially, it is as if the agents were other human players that can select freely among a finite set of dialogue options. In order for the agents to reason about their choice, dialogue options can be associated to different semantic tags pertaining to the meaning and the style of the utterance.

This information is accessible in the reasoning process of the agent through the following meta-belief:

`ValidDialogue ([ cs ], [ ns ], [ m ], [ s ])`

This meta-belief will go through the list of defined dialogues and will build a set of substitutions for each dialogue that it finds, also taking into account any existing constraints from previous conditions. The four different arguments are *[cs]*, current state; *[ns]*, next state; *[m]*, meaning; and *[s]*, style. For every substitution found, the value returned is “True.” To demonstrate the usefulness of this meta-belief consider the following decision rule for a “Speak” action:

Action : `Speak ([ cs ], [ ns ], [ m ], [ s ])`

Target : `[ x ]`

Priority : 1

Conditions :

`DialogueState ([ x ]) = [ cs ]`

`ValidDialogue ([ cs ], [ ns ], [ m ], [ s ]) = True`

This single decision rule is capable of handling any basic dialogue tree like the one shown in Figure 2 where the characters’ dialogues are determined just by the current state of the conversation. Additionally, for the character to choose to speak using a “Rude” style when its mood is negative, adding the following decision rule with a higher priority would be sufficient:

```

Action : Speak ([ cs ], [ ns ], [ m ], Rude )
Target : [ x ]
Priority : 2
Conditions :
    DialogueState ([ x ]) = [ cs ]
    ValidDialogue ([ cs ], [ ns ], [ m ], Rude ) = True
    Mood (SELF) < 0

```

The purpose of these examples is to demonstrate that, ultimately, the decisions are still being made autonomously according to the logic of each agent. As such, the flexibility that comes with an agent-based solution is not sacrificed. Instead, this flexibility is combined with a centralized representation of the possible dialogues, which is the paradigm that most game developers are already accustomed to. As a result, it becomes much less cumbersome for them to adopt this agent-based solution, even for scenarios with a simple dialogue structure. This centralized perspective for the definition of dialogues, meaning that they are not restricted in the mind of a single agent, is also aligned with the idea that dialogues are often social practices, of common knowledge, that can be used by different agents in similar situations.

6 INTEGRATED AUTHORING TOOL

The accessibility of an agent-based framework like the FAtiMA Toolkit is determined in large part by its authoring experience. Ideally, a person would be able to come up with different ideas for how a character behaves in one or more contexts and then easily configure different FAtiMA agents that would match those ideas without having to write code.

With the previous goal in mind, the FAtiMA Toolkit provides an integrated authoring tool that supports the creation and configuration of scenarios by non-coders. As shown in Figure 3, this is a GUI tool with a hierarchical structure that contains several editors, a simulator, and a Web API. At the top, there is a main editor where general aspects of the scenario are configured. These include the scenario's name, a short description, and a list of character profiles. On the second level of the hierarchy, there are six different editors with their own specific purpose.

The purpose of the Dialogue Editor is to define all the dialogue actions that will be available for the agents and the players to select from. Additionally, there is also the possibility to export and import dialogues from other formats, including MS Excel. The Dialogue Editor also provides a validation mechanism that runs a depth-first search through all the dialogue states and tells the author which states are not reachable and how many end states exist. This information is used to quickly detect any possible errors in state transitions. Finally, it is also possible to create and view a State Machine Graph of all possible dialogues and their corresponding states and transitions.

The Character Editor is where the author configures the different agent profiles that are part of the scenario. For each agent, the author is able to specify their initial emotional state, the initial beliefs of their knowledge base, the previous events the agent will be able to recall in their autobiographic memory, and the agent's goals.

The Cognitive Rules section is responsible for configuring different cognitive components of the agents. In the Emotional Appraisal Editor, the author defines all the appraisal rules and parameters that drive the generation process of emotions. In the Emotional Decision Making Editor, the author specifies the rules that will guide the decision-making process of the agent. In the Social Importance Editor, the author specifies the rules that will tell the agent how it should perceive and act toward others from a relational perspective. Finally, in the CiF-CK Editor, the author can specify different patterns of micro-interactions that are common in human interaction.

Fig. 3. Integrated authoring tool: diagram of existing components.

To help the authoring process, the author is capable of defining what are the consequences of an agent performing a specific action in the *World Editor*. For instance, the author can specify that after each dialogue action, the property “Has(Floor)” will assume the value corresponding to the name of the agent to whom the speak action was directed. This is particularly relevant in a multi-party conversation with multiple agents talking to one another. An additional rule could then be added that after a certain amount of time has passed, if the agent who has the floor remains silent, then the property “Has(Floor)” will assume the name of another random agent. The effects of actions might also be programmed directly in the game itself and communicated to the agents via events. The main benefit of using the *World Editor* is that these effects are configurable without having to recompile the game. Also, the action effects defined in the *World Editor* will be visible to the *Simulator*.

The *Simulator* component allows the author to quickly run and test the defined scenario. As such, it becomes easier for the author to be able to do iterative testing and to reuse previous configurations from other scenarios and adapt them accordingly. Testers can play the role of any of the loaded characters or just become a spectator as the scenario unfolds. Additionally, it also features several debug tools such as a real-time agent inspector capable of querying the agent’s Knowledge Base and display their Emotional State, Memory, and Goals along with showing information regarding which rules are being triggered and its consequent effects.

The *FATiMA Toolkit’s Web API* module allows users to run a server on their own machine and through GET and POST functions interact with *FATiMA* and its engine. Because they are built around HTTP protocols, nearly any programming language can be used to access them. As a result, *FATiMA* is now able to be used by a wide array of different technologies and engines.

One final consideration about the accessibility of the authoring process is that all of the configuration files are stored in a JSON format that is easy to read and change using a text editor.

7 USE CASES

To illustrate the potential and versatility of the *FATiMA Toolkit*, we will now present several different recent case studies where it has been successfully applied. We will start by enumerating each in Table 4 and proceed to describe each more in-depth within the following sections.

Table 4. FAtiMA Use Cases

Name	Project	Publication	URL
Space Modules	RAGE Project	[27, 44, 72]	5
Sports Team Manager	RAGE Project	[27, 44, 72]	6
Virtual Reality Police Interview	RAGE Project	[28]	7
FAtiMA's Game Jam	RAGE Project	[27]	8
Don't Starve Together	Master Thesis	[1]	9
Game AI Course	University of Lisbon	[27]	10
Sueca Card Game with Robotic Partners	AMIGOS Project	[15]	11

7.1 Space Modules Inc.

Space Modules Inc (see Figure 4) is a serious game that was developed in the RAGE Project by the company PlayGen¹² for an educational organization in the Netherlands named Stichting Praktijkleren.¹³ In total, the game was used and evaluated across six participating schools with the total number of students around 1,100. It is a single-player game where the players must provide customer service on behalf of a spaceship part manufacturer. The challenge lies in the fact that each customer has a unique emotional profile. For instance, some customers get angrier more quickly, whereas other get angrier if you do not solve their problem promptly. The score that is attributed to the player is derived from the emotional state of the character at the end of the conversation. Therefore, players must learn how to manage the emotions of their customers in addition to providing the right solutions to their problems. Concerning the interaction, a limited amount of dialogue choices is given to the player at each turn until the conversation ends and a new customer call arrives. During the conversation, the NPC provides nonverbal feedback that is indicative of its current emotional state. It is up to the player to pay close attention to these signals and the conversation unfolds.

It is important to highlight that the integration of the FAtiMA Toolkit in the code of Space Modules Inc. was handled directly by the game developers at PlayGen. By following an available example, they were able to use the toolkit to drive the emotional dynamics of the characters as well as their responses. Moreover, the pedagogical experts were able to code the many dialogues that occur within the game using the explicit dialogue structure that was described in Section 5. This was achieved by first defining a list of symbols that would be used to categorize the meaning and the style of the different dialogue alternatives. We then configured a small set of appraisal and decision rules that were sensitive to those symbols and to the game's pedagogical goals. After this was complete, the pedagogical experts were then able to use the provided authoring tool to code all the dialogues of the game, taking into account the several different spaceship devices and their associated problems. This is a very encouraging result for the accessibility of the toolkit given that our involvement in the authorial process of previous serious games that used FAtiMA was much higher.

⁵<https://play.google.com/store/apps/details?id=com.PlayGen.SpaceModulesInc>.

⁶<http://playgen.com/play/sport-team-manager/>.

⁷<https://fatima-toolkit.eu/showcase/>.

⁸<https://itch.io/jam/fatima-jam>.

⁹<https://fatima-toolkit.eu/fatima-toolkit-mcts-and-dont-starve-together/>.

¹⁰<https://fenix.tecnico.ulisboa.pt/cursos/meic-t/disciplina-curricular/845953938489420>.

¹¹<https://gaips.inesc-id.pt/amigos/>.

¹²<http://playgen.com/>.

¹³<http://stichtingpraktijkleren.nl/>.

Fig. 4. SpaceModules: conversation screen (left); score screen (right).

7.2 Sports Team Manager

Sports Team Manager is another game developed by Playgen during the RAGE Project that makes use of the FAtiMA Toolkit. Here, the objective is to assemble the most optimally performing sailing team by resolving conflicts and managing the team's interactions.

In order to better identify the character's skills and personalities, players must interview each member. The team has a set of predetermined roles, each with overlapping skill requirements [44]. Additionally, players must take into account the team's morale and work on the social relationships within the team. Each customer/sailor can have different emotional profiles; thus, providing different challenges and conflict situations, which the player must solve, will arise seamlessly.

Once again, the framework's integration with the game engine was handled directly by the developers. FAtiMA is responsible for managing the agent's beliefs, reactions, and actions. In this case, the beliefs are related to information such as their last position in the team, skill ratings, opinion ratings, and event states. Sports Team Manager is also able to take full advantage of the Autobiographic Memory component of FAtiMA. Because each of the player's decisions is saved into the character's memory, a history of events can be preserved. As such, it can also be reloaded in further play sessions, allowing for the possibility of a persistent game emotional state and decision making.

7.3 FAtiMA Virtual Reality Demo: Police Interview

During the final years of the RAGE Project, we had the opportunity to develop a Virtual Reality experience that makes full use of the FAtiMA Toolkit's abilities. It also makes use of other RAGE-developed assets, such as L2F's "Speech I/O"¹⁴ and **Behaviour Mark-up Language (BML)** Realizer,¹⁵ showing players, developers, and other stakeholders the synergy between the tools developed throughout the project.

The demo is set in an interrogation room where users play the part of a Police Interviewer and interact with a Suspect, who is a Role Play Character. The player's objective is to obtain the most information about the suspect they can without losing control of the conversation. In order to interact with the subject, players can select questions from a wide array of possible choices

¹⁴<https://www.gamecomponents.eu/content/226/speechio>.

¹⁵<https://www.gamecomponents.eu/content/231/bml-realizer>.

Fig. 5. FAtiMA Virtual Reality demo.

that are displayed in a virtual computer screen as seen in Figure. After selecting a question, the virtual suspect will answer it, using Text to Speech, according to its mood and its emotions. The questions are sorted through various different topics such as “Marriage,” “Children,” and “Work,” among others. These questions and answers are managed through the Integrated Authoring Tool Asset’s Dialogue Manager. This experience, also called FAtiMA VR Demo, was developed using the Unity Game Engine and is compatible with most Virtual Reality headsets. Figure 5 shows the Player perspective within the game. The screen shown on the left is where the Players can take a look and select one of the possible questions to ask the Virtual Suspect.

This Social Skills Training scenario was later used in a Social Presence [40] study. It was found that interacting with an intelligent virtual character in Virtual Reality had a stronger impact compared to the same interaction in a traditional non-immersive platform, in terms of both presence and believability [28].

7.4 FAtiMA’s Game Jam

To promote the use of the FAtiMA Toolkit, an online “hackathon” was created where participants were encouraged to create video games with engaging characters powered by the FAtiMA Toolkit. The hackathon was hosted on the itch.io website¹⁶ and lasted around 2 months.

The winner was a game called “The Princess and the Shapeshifter”¹⁷ that was created for the competition with original art and design. The developers took full advantage of FAtiMA’s character-driven approach to allow players to take the role of the characters they interacted with. Throughout the game, players must interact with six different characters/agents, with different personalities and abilities, and must persuade them into letting players take something from them so that they can “shape-shift” into them.

7.5 Don’t Starve Together

Don’t Starve Together is an open-world survival video game where, as in any survival video game, the objective is to survive for as long as possible [20]. These games, typically, have several

¹⁶<https://itch.io/jam/fatima-jam>.

¹⁷<https://petzi.itch.io/the-princess-and-the-shapeshifter>.

different non-playable characters that players can interact with during their play-through. The few autonomous characters of Don't Starve Together present limited behavior and are not subject to the same rules as the players, in terms of survivability. The lack of AI for NPCs capable of creating believable characters affects the player experience. This is the primary motivation behind Almeida and Rui's Master Thesis [1].

In order to improve the quality of the agents within the game, Almeida and Rui proposed the development of a platform that enables the creation of NPCs controlled by agency-based models for survival games. Using the FATiMA Toolkit, as the model for agency, and Don't Starve Together, they've implemented and published the framework with an example character, Walter. Walter can be authored using the FATiMA Toolkit and its actions are reflected within the gameplay of a Don't Starve Together session.

7.6 Game AI Course

The toolkit has been put to test in a course on Game AI at IST, University of Lisbon, for the past 3 years by over 100 students. During the semester, students work on four different projects; the first three cover a wide range of fields within Artificial Intelligence: Movement and Collision Avoidance, Pathfinding and Monte Carlo Tree Search, and Decision Making algorithms, correspondingly. For the fourth and final project students are given the option to choose the topic they want to delve into, within AI and Games. Here, they can propose their own assignment or they can choose one of the four different options we suggest, two of which use FATiMA. One combines the FATiMA's planning capabilities with Monte Carlo Tree Search applied to a "survival" game, and the other makes full use of the Toolkit's ability to create interesting conversational scenarios.

7.6.1 FATiMA, Monte Carlo Tree Search, and Don't Starve Together. Since Almeida and Rui's Thesis [1], we've challenged Games AI students to improve the character (Walter), create their own model, and try different approaches. For example, we've found that instead of creating a model with, let's say, 30 different decision-making rules, a **Monte Carlo Tree Search (MCTS)** approach has better results in terms of survivability. In this case, students designed the agent's action space and their own interpretation of the environment around the agent and, using FATiMA's Meta-Beliefs, created their own MCTS algorithm within the framework. When agents used this meta-belief, it would run an MCTS to decide which of the available Don't Starve Together-based actions was the best according to their survivability goals.

In terms of results, the standard "survival" games' metric is number of days survived. The original "Walter" survived for 1.8 days with 19 different decision rules. Two years later students have been able to create an agent capable of surviving 5 days, on average.

Over the past 2 years we've had over 20 different students in 10 groups working on their own version of this project. Their results prove both the planning capabilities of FATiMA and its adaptability to different contexts and environments.

7.6.2 FATiMA Toolkit and Conversation Scenarios. By opting into this project, students were given the task of creating at least two different conversation scenarios, one with a single character interacting with the player and another with two characters engaging in conversation with the player at the same time. While the project did not require students to work with all of FATiMA's components, many of the groups chose to use most of them. Many of the students started by using the Emotional Appraisal and the Emotional Decision Making modules, but in order to enrich their scenarios, they resorted to other reasoner components. Social Importance, for instance, was used to manage the friendship or romance between agents, while the CIF-CK component was often used to accommodate social conventions such as agents greeting each other in the beginning of a conversation or saying goodbye at the end. In order to debug and test their scenarios each group

Fig. 6. Mixed human-robot teams playing Sueca.

used the Simulator and the Dialogue Editor quite a bit. One particularly useful option often shown in their reports was the option to create a graph tree with the authored dialogue states.

In order to get acquainted with the tool, students had 1.5-hour workshops on the FAtiMA Toolkit. Additionally, we also shared with the students a version of the toolkit that is integrated with the Unity game engine and uses components to realize the body and expression of the characters developed by other members of the RAGE Project.

Over the last 3 years we've had over 30 groups of students (approximately 90 different people) creating and designing conversational scenarios using the FAtiMA Toolkit for the course of AI on Games. Students were given no limitation; as long as the non-player characters had believable emotional responses they were free to use any kind of topic for the conversation. The result was a wide number of all kinds of surprisingly interesting scenarios.

Scenarios included a police officer pulling over the player, after-death experiences featuring ghosts, and Obama and Trump debating about building a wall, and a scenario where players were at the gates of heaven and had to convince the gatekeeper to let them in. In the latter case, to be successful, players had to avoid upsetting the gatekeeper too much. Other groups opted for a more serious theme such as a job interview or a shopping scene with a father, his son, and a shopkeeper.

7.7 Sueca Card Game with Robotic Partners

Our final case study involves the use of the FAtiMA Toolkit to control the decisions of social robots rather than virtual characters. There is also another important difference. Rather than focusing exclusively on dialogue actions, the case study revolves around a team-based card game. More precisely, two robots play the game of *Sueca* with two human partners, as shown in Figure 6. This is a Portuguese trick-taking card game that is played by four players divided in two teams against one another.

From a research point of view, the main purpose of this game scenario was to investigate how robots can become more effective teammates. With that goal in mind, we decided to explore a particular type of affective phenomenon that has not yet received a lot of attention in neither virtual agents nor social robots. The phenomenon in question is group-based emotions, which is the term that is used to describe emotions that people experience due to their sense of membership to a given group [38]. A common scenario where these emotions occur is when playing or watching competitive sports.

One of the interesting aspects of group-based emotions is that it combines appraisal theories of emotion with self-categorization theories of social identity [69]. Based on the ideas of the model

proposed by Goldenberg et al., the *Emotional Appraisal* component of FAtiMA was used to generate group-based emotions by having the robot appraise both its actions and the actions of its partner as if both have the same subject, namely, their team. As an example, the robot felt and expressed “group shame” whenever it or its partner performed a bad play during the game. A user study was conducted where one of the robots in the game had group-based emotions and the other one felt individual-based emotions instead [15]. The results show that the robot with group-based emotions was perceived as both more likeable and more trustworthy.

From a technical perspective, the two robots are fully autonomous in their behaviors. As detailed in Correia et al. [14], a **Perfect Information Monte Carlo (PIMC)** algorithm was used for the robot’s ability to play the game in a proficient manner. The algorithm was combined with FAtiMA’s appraisal mechanism so that emotions are generated based on how well the agent is doing in the game and are expressed through both utterances and non-verbal animations. In total, about 200 utterances were created for this scenario using the explicit dialogue structure that was defined previously. As for the non-verbal expressions, these are driven by Nutty Tracks, one of the components within the SERA ecosystem [54]. The FAtiMA Toolkit’s library design allowed the Emotional Appraisal component to be adapted to fit other frameworks and use different algorithms. Together with its follow-up study, over 100 participants interacted with the robots through FAtiMA-designed scenarios [14, 15].

8 CONCLUSIONS

The importance of affective behavior in characters (and agents) is widely acknowledged, in particular, in the gaming industry and in social robotics. Emotions capture our attention and make our interaction with artificial entities more natural and engaging. Consequentially, a lot of effort has been put into creating tools that facilitate the ability for these artificial entities to be able to portray different emotions and social behavior with an increased degree of realism and believability. The use of affective models based on appraisal theories has the potential to automate the process of generating emotions for both virtual agents and social robots in a manner that is congruent with human psychology. Their application is particularly relevant when the social dynamics of the interaction are not trivial. This was the case of FearNot! [19], the serious game that led to the creation of the original FAtiMA architecture.

Despite their potential, the use of more complete affective models outside research has still been quite limited as there are both technical and conceptual issues that are hindering a more widespread adoption. One such issue that is on the technical side is the typical framework approach to agent models as it imposes limitations in how they can be applied and combined with other tools. On the conceptual side, cognitive architectures are designed with the intention of being able to function in many different domains. As such, they need to be able to encompass different types of algorithms for decision making and emotional appraisal. Finally, interaction scenarios with virtual agents and social robots tend to be heavily focused on dialogue, but it is often the case that dialogue actions are defined as regular actions in an agent model without any further representation.

In this article we described how these important issues were addressed in the development of the FAtiMA Toolkit. Rather than using a framework approach, the toolkit uses libraries that are combined together to create affective agents. Moreover, through the use of meta-beliefs, it is possible to use various specialized algorithms under a unified rule-based approach that governs both the agent’s emotional appraisal and decision making. The toolkit also has a unique approach to model dialogue that uses a centralized state-based structure for defining all the dialogue possibilities but leaving the decision logic in the agent model. Finally, a GUI-based authoring tool is also provided to facilitate the process of creating agent-based scenarios and testing them iteratively.

The advantages brought by the toolkit have been explored by several different users in a number of different applications. In this article we presented a few of them. PlayGen, a UK-based game studio, used the FAtiMA Toolkit in both of their games: Space Modules Inc. and Sports Team Manager. Players must deal with social dilemmas to accomplish their objectives and reach a better score. FAtiMA was used to manage these social interaction scenarios and to handle each of the character's behaviors. The FAtiMA Virtual Reality Demo - Police Interrogation is an experience whose focus is to display the high level of immersion a FAtiMA-modeled character could afford in a Virtual Reality setting. We have also illustrated the results of an Online Game Jam powered by FAtiMA and how a Master Thesis was able to implement a search algorithm through FAtiMA in a survival game environment. Finally, we also presented the outcomes of FAtiMA in a course on Game AI at IST, University of Lisbon, for the past 3 years and how we used the toolkit to investigate the impact of group-based emotions in a robotic partner.

Regarding future work, we are planning to provide additional tools to the toolkit with the goal of further alleviating the authoring effort, and improving control and verifiability of the socio-emotional behavior of the designed characters. For instance, the hybrid dialogue system design allows us to easily replace the user's role by an autonomous agent, which can be built from a particular user model, and then automatically test the dialogue structure and decision-making rules, determining the probability of reaching each particular state or expressing a particular behavior. We believe this to be an important mechanism to provide game designers so that they can more easily detect potential problems or unreachable states and have some degree of authorial control over the interaction.

Finally, the potential of data-driven approaches we've discussed in Section 2 is quite promising. Nonetheless, it is important to highlight some of the drawbacks of these approaches. Data-heavy systems are prone to echoing the biases present in the datasets [59], limiting their ability to produce generic and independent results. Moreover, the complexity of deep learning algorithms, extensively used in data-centric approaches, makes it so that the resulting model is inherently uninterpretable to human users [53]. Taking these issues into consideration, we have been working on the addition of some of these techniques in the FAtiMA Toolkit with the objective of partially automating its authoring task. We believe that using a combination of theory- and data-driven approaches, it is possible to mitigate each other's shortcomings and provide a clearer and more approachable authoring experience. This work is ongoing.

REFERENCES

- [1] Fabio Almeida and Prada Rui. 2018. *Creating an Agent-based Framework for Don't Starve Together*. Master's Thesis. IST, University of Lisbon, Portugal.
- [2] Prithviraj Ammanabrolu, Wesley Cheung, William Broniec, and Mark O. Riedl. 2020. Automated storytelling via causal, commonsense plot ordering. *arXiv preprint arXiv:2009.00829*.
- [3] Prithviraj Ammanabrolu, Wesley Cheung, Dan Tu, William Broniec, and Mark O. Riedl. 2020. Bringing stories alive: Generating interactive fiction worlds. *CoRR abs/2001.10161* (2020). <http://arxiv.org/abs/2001.10161>.
- [4] Ana Antunes, Joana Campos, João Dias, Pedro A. Santos, and Rui Prada. 2021. EEG model: Emotional episode generation for social sharing of emotions. In *Proceedings of the 21st ACM International Conference on Intelligent Virtual Agents*. 1–8.
- [5] Nabiha Asghar, Pascal Poupart, Jesse Hoey, Xin Jiang, and Lili Mou. 2018. Affective neural response generation. In *European Conference on Information Retrieval*. Springer, 154–166.
- [6] Joscha Bach, Dietrich Dörner, and Ronnie Vuine. 2006. Psi and microPsi: A novel approach to modeling emotion and cognition in a cognitive architecture. In *Proceedings of the 7th International Conference on Cognitive Modeling, Trieste*. 20–25.
- [7] Christian Becker-Asano. 2008. *WASABI: Affect Simulation for Agents with Believable Interactivity*. Ph.D. Dissertation. Faculty of Technology, University of Bielefeld. IOS Press (DISKI 319).

- [8] Christian Becker-Asano and Ipke Wachsmuth. 2010. Affective computing with primary and secondary emotions in a virtual human. *Autonomous Agents and Multi-Agent Systems* 20, 1 (2010), 32–49. <https://doi.org/10.1007/s10458-009-9094-9>
- [9] Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. Comet: Commonsense transformers for automatic knowledge graph construction. *arXiv preprint arXiv:1906.05317*.
- [10] Dolores Cañamero. 1997. Modeling motivations and emotions as a basis for intelligent behavior. In *Proceedings of the 1st International Conference on Autonomous Agents (AGENTS'97)*. ACM, New York, NY, 148–155. <https://doi.org/10.1145/267658.267688>
- [11] Carlos Carrascosa, Javier Bajo, Vicente Julián, Juan M. Corchado, and V. Botti. 2008. Hybrid multi-agent architecture as a real-time problem-solving model. *Expert Systems with Applications* 34, 1 (2008), 2–17.
- [12] Marc Cavazza, Fred Charles, and Steven J. Mead. 2002. Character-based interactive storytelling. *IEEE Intelligent Systems* 17, 4 (2002), 17–24.
- [13] Gerald L. Clore and Andrew Ortony. 2013. Psychological construction in the OCC model of emotion. *Emotion Review* 5, 4 (2013), 335–343.
- [14] Filipa Correia, Patrícia Alves-Oliveira, Tiago Ribeiro, Francisco Melo, and Ana Paiva. 2017. A Social Robot as a Card Game Player. In *Thirteenth Artificial Intelligence and Interactive Digital Entertainment Conference*.
- [15] Filipa Correia, Samuel Mascarenhas, Rui Prada, Francisco S. Melo, and Ana Paiva. 2018. Group-based emotions in teams of humans and robots. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 261–269.
- [16] Ferdinand De Coninck, Zerrin Yumak, Guntur Sandino, Remco C. Veltkamp, and B. CleVR. 2019. Non-verbal behavior generation for virtual characters in group conversations. In *IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR'19)*. IEEE.
- [17] Irene-Anna Diakidoy, Antonis Kakas, Loizos Michael, and Rob Miller. 2015. STAR: A system of argumentation for story comprehension and beyond. In *2015 AAAI Spring Symposium Series*.
- [18] Joao Dias, Samuel Mascarenhas, and Ana Paiva. 2014. FAtiMA modular: Towards an agent architecture with a generic appraisal framework. In *Emotion Modeling*. Springer, 44–56.
- [19] João Dias and Ana Paiva. 2005. Feeling and reasoning: A computational model for emotional characters. In *Portuguese Conference on Artificial Intelligence*. Springer, 127–140.
- [20] Klei Entertainment. 2016. Don't Starve Together. Klei Entertainment.
- [21] Sayan Ghosh, Mathieu Chollet, Eugene Laksana, Louis-Philippe Morency, and Stefan Scherer. 2017. Affect-LM: A neural language model for customizable affective text generation. *arXiv preprint arXiv:1704.06851*.
- [22] Tushar Goswamy, Ishika Singh, Ahsan Barkati, and Ashutosh Modi. 2020. Adapting a language model for controlled affective text generation. In *Proceedings of the 28th International Conference on Computational Linguistics*. 2787–2801.
- [23] Jonathan Gratch and Stacy Marsella. 2003. Modeling coping behavior in virtual humans: Dont worry, be happy. In *Proceedings of 2nd International Joint Conference on Autonomous Agents and Multiagent systems*.
- [24] Jonathan Gratch and Stacy Marsella. 2004. A domain-independent framework for modeling emotion. *Journal of Cognitive Systems Research* 5, 4 (2004), 269–306.
- [25] Jonathan Gratch and Stacy Marsella. 2006. EMA: A computational model of appraisal dynamics. In *European Meeting on Cybernetics and Systems Research*.
- [26] James J. Gross and Oliver P. John. 2002. Wise emotion regulation. In *The Wisdom in Feeling: Psychological Processes in Emotional Intelligence*, L. F. Barrett and P. Salovey (Eds.). The Guilford Press, 297–319.
- [27] Manuel Guimarães, Samuel Mascarenhas, Rui Prada, Pedro A. Santos, and João Dias. 2019. An accessible toolkit for the creation of socio-emotional agents. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2357–2359.
- [28] Manuel Guimarães, Rui Prada, Pedro A. Santos, João Dias, Arnav Jhala, and Samuel Mascarenhas. 2020. The impact of virtual reality in the social presence of a virtual agent. In *Proceedings of the 20th ACM International Conference on Intelligent Virtual Agents*. 1–8.
- [29] Manuel Guimarães, Pedro Santos, and Arnav Jhala. 2017. CiF-CK: An architecture for social NPCs in commercial games. In *2017 IEEE Conference on Computational Intelligence and Games (CIG'17)*. 126–133. <https://doi.org/10.1109/CIG.2017.8080425>
- [30] John H. Flavell. 1979. Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist* 34 (1979), 906–911.
- [31] Arno Hartholt, David Traum, Stacy C. Marsella, Ari Shapiro, Giota Stratou, Anton Leuski, Louis-Philippe Morency, and Jonathan Gratch. 2013. All together now: Introducing the virtual human toolkit. In *13th International Conference on Intelligent Virtual Agents*.
- [32] Thomas Hayton, Julie Porteous, Joao Ferreira, Alan Lindsay, and Jonathon Read. 2017. Storyframer: From input stories to output planning models. In *Workshop on Knowledge Engineering for Planning and Scheduling (KEPS'17). The 27th International Conference on Automated Planning and Scheduling (ICAPS'17)*.

- [33] Sepehr Janghorbani, Ashutosh Modi, Jakob Buhmann, and Mubbasir Kapadia. 2019. Domain authoring assistant for intelligent virtual agents. *arXiv preprint arXiv:1904.03266*.
- [34] W. Lewis Johnson and James C. Lester. 2016. Face-to-face interaction with pedagogical agents, twenty years later. *International Journal of Artificial Intelligence in Education* 26, 1 (2016), 25–36.
- [35] Theodore D. Kemper. 2011. *Status, Power and Ritual Interaction*. Ashgate. 316 pages.
- [36] Stefan Kopp, Brigitte Krenn, Stacy Marsella, Andrew N. Marshall, Catherine Pelachaud, Hannes Pirker, Kristinn R. Thórisson, and Hannes Vilhjálmsón. 2006. Towards a common framework for multimodal generation: The behavior markup language. In *International Workshop on Intelligent Virtual Agents*. Springer, 205–217.
- [37] Taras Kucherenko, Dai Hasegawa, Gustav Eje Henter, Naoshi Kaneko, and Hedvig Kjellström. 2019. Analyzing input and output representations for speech-driven gesture generation. In *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents*. 97–104.
- [38] Toon Kuppens and Vincent Y. Yzerbyt. 2012. Group-based emotions: The impact of social identity on appraisals, emotions, and behaviors. *Basic and Applied Social Psychology* 34, 1 (2012), 20–33.
- [39] Jina Lee and Stacy C. Marsella. 2006. Nonverbal behavior generator for embodied conversational agents. In *6th International Conference on Intelligent Virtual Agents*.
- [40] Kwan Min Lee, Wei Peng, Seung-A Jin, and Chang Yan. 2006. Can robots manifest personality? An empirical test of personality recognition, social responses, and social presence in human–robot interaction. *Journal of Communication* 56, 4 (2006), 754–772.
- [41] Mei Yii Lim, João Dias, Ruth Aylett, and Ana Paiva. 2012. Creating adaptive affective autonomous NPCs. *Autonomous Agents and Multi-Agent Systems* 24, 2 (2012), 287–311.
- [42] Alan Lindsay, Jonathon Read, Joao Ferreira, Thomas Hayton, Julie Porteous, and Peter Gregory. 2017. Framer: Planning models from natural language action descriptions. In *27th International Conference on Automated Planning and Scheduling*.
- [43] Samuel Mascarenhas, Nick Degens, Ana Paiva, Rui Prada, Gert Jan Hofstede, Adrie Beulens, and Ruth Aylett. 2016. Modeling culture in intelligent virtual agents. *Autonomous Agents and Multi-Agent Systems* 30, 5 (2016), 931–962.
- [44] Samuel Mascarenhas, Manuel Guimarães, Rui Prada, Joao Dias, Pedro Santos, Kam Star, Ben Hirsh, Ellis Spice, and Rob Kommeren. 2018. A virtual agent toolkit for serious games developers. In *Proceedings of the Computational Intelligence and Games Conference (CIG'18)*. IEEE.
- [45] Michael Mateas and Andrew Stern. 2003. Façade: An experiment in building a fully-realized interactive drama. In *Game Developers Conference*, Vol. 2. 4–8.
- [46] Josh McCoy, Mike Treanor, Ben Samuel, Brandon Tarse, Michael Mateas, and Noah Wardrip-Fruin. 2010. Authoring game-based interactive narrative using social games and comme il faut. In *Proceedings of the 4th International Conference & Festival of the Electronic Literature Organization: Archive & Innovate*. Citeseer.
- [47] Albert Mehrabian. 1996. Pleasure-arousal-dominance: A general framework for describing and measuring individual differences in temperament. *Current Psychology* 14, 4 (1996), 261–292.
- [48] Agnes Moors, Phoebe C. Ellsworth, Klaus R. Scherer, and Nico H. Frijda. 2013. Appraisal theories of emotion: State of the art and future development. *Emotion Review* 5, 2 (2013), 119–124.
- [49] Allen Newell. 1994. *Unified Theories of Cognition*. Harvard University Press.
- [50] Andrew Ortony, Gerald L. Clore, and Allan Collins. 1990. *The Cognitive Structure of Emotions*. Cambridge University Press.
- [51] Alexandru Popescu, Joost Broekens, and Maarten Van Someren. 2014. Gamygdala: An emotion engine for games. *IEEE Transactions on Affective Computing* 5, 1 (2014), 32–44.
- [52] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog* 1, 8 (2019), 9.
- [53] Arun Rai. 2020. Explainable AI: From black box to glass box. *Journal of the Academy of Marketing Science* 48, 1 (2020), 137–141.
- [54] Tiago Ribeiro, Andre Pereira, Eugenio Di Tullio, and Ana Paiva. 2016. The SERA Ecosystem: Socially Expressive Robotics Architecture for Autonomous Human-Robot Interaction. In *AAAI Spring Symposium Series*.
- [55] Sérgio Hortas Rodrigues, Samuel Mascarenhas, João Dias, and Ana Paiva. 2014. A process model of empathy for virtual agents. *Interacting with Computers* 27, 4 (2014), 371–391.
- [56] Najmeh Sadoughi, Yang Liu, and Carlos Busso. 2017. Meaningful head movements driven by emotional synthetic speech. *Speech Communication* 95 (2017), 87–99.
- [57] Klaus R. Scherer. 1999. Appraisal theory. In *Handbook of Cognition and Emotion*, T. Dalgleish and M. J. Power (Eds.). 637–663.
- [58] Stefan Scherer, Stacy C. Marsella, Giota Stratou, Yuyu Xu, Fabrizio Morbini, Alesia Egan, Albert Rizzo, and Louis-Philippe Morency. 2012. Perception markup language: Towards a standardized representation of perceived nonverbal behaviors. In *The 12th International Conference on Intelligent Virtual Agents (IVA'12)*.

- [59] Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. 2019. The woman worked as a babysitter: On biases in language generation. *arXiv preprint arXiv:1909.01326*.
- [60] Avirup Sil and Alexander Yates. 2011. Extracting strips representations of actions and events. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing 2011*. 1–8.
- [61] Aaron Sloman. 1987. Motives, mechanisms, and emotions. *Cognition and Emotion* 1, 3 (1987), 217–233.
- [62] Aaron Sloman. 2002. How many separately evolved emotional beasts live within us? In *Emotions in Humans and Artifacts*. MIT Press, 29–96.
- [63] Craig A. Smith and Richard S. Lazarus. 1993. Appraisal components, core relational themes, and the emotions. *Cognition & Emotion* 7, 3–4 (1993), 233–269.
- [64] Craig A. Smith, Richard S. Lazarus, et al. 1990. Emotion and adaptation. In *Handbook of Personality: Theory and Research*, 609–637.
- [65] Bethesda Game Studios. 2015. *The Elder Scrolls v: Skyrim*. Bethesda Game Studios.
- [66] Ron Sun. 2006. The CLARION cognitive architecture: Extending cognitive modeling to social simulation. *Cognition and Multi-agent Interaction* (2006), 79–99.
- [67] Ron Sun. 2007. The motivational and metacognitive control in CLARION. *Modeling Integrated Cognitive Systems* (2007), 63–75.
- [68] Pradyumna Tambwekar, Murtaza Dhuliawala, Lara J. Martin, Animesh Mehta, Brent Harrison, and Mark O. Riedl. 2018. Controllable neural story plot generation via reinforcement learning. *arXiv preprint arXiv:1809.10736*.
- [69] John C. Turner, Michael A. Hogg, Penelope J. Oakes, Stephen D. Reicher, and Margaret S. Wetherell. 1987. *Rediscovering the Social Group: A Self-categorization Theory*. Basil Blackwell.
- [70] Marcel V. J. Veenman, Bernadette H. A. M. Van Hout-Wolters, and Peter Afflerbach. 2006. Metacognition and learning: Conceptual and methodological considerations. *Metacognition and Learning* 1, 1 (2006), 3–14.
- [71] Juan D. Velásquez and Pattie Maes. 1997. Cathexis: A computational model of emotions. In *Proceedings of the 1st International Conference on Autonomous Agents (AGENTS '97)*. ACM, New York, NY, 518–519. <https://doi.org/10.1145/267658.267808>
- [72] Wim Westera, Rui Prada, Samuel Mascarenhas, Pedro A. Santos, João Dias, Manuel Guimarães, Konstantinos Georgiadis, Enkhbold Nyamsuren, Kiavash Bahreini, Zerrin Yumak, et al. 2020. Artificial intelligence moving serious gaming: Presenting reusable game AI components. *Education and Information Technologies* 25, 1 (2020), 351–380.
- [73] Yeyao Zhang, Eleftheria Tsipidi, Sasha Schriber, Mubbasir Kapadia, Markus Gross, and Ashutosh Modi. 2019. Generating animations from screenplays. *arXiv preprint arXiv:1904.05440*.

Received March 2020; revised December 2021; accepted January 2022