Modeling Emotion Understanding in Stories: Insights from Traumatic Brain Injured Patients

Eleonora Ceccaldi, Cristina Battaglino[1],
Rossana Damiano[1], Valentina Galetto[3], and Marina Zettin[23]

[1] Dipartimento di Informatica, Università di Torino
[2] Dipartimento di Psicologia, Università di Torino
[3] Centro Puzzle, Torino

Abstract. In this paper, we describe a case study in emotion understanding in stories that leverages the contribution of traumatic brain injured people. In particular, we focus on moral emotions, leveraging the differences in moral functioning that characterizes traumatic brain injured patients.
By comparing the understanding of the moral and emotional aspects of characters' behavior in traumatic brain injured patients and in a control group of neurologically healthy subjects, we observed slight, yet meaningful differences in the two groups. We describe the test methodology and results, discussing their implications for the design of rehabilitation applications that leverage virtual characters.

Keywords: emotions; narrative understanding; rehabilitation

1 Introduction

Emotions play a fundamental role in stories, as acknowledged by scholars for centuries, with contributions that range from philosophy [2] and psychology [8] to narratology [15] and drama studies [24]. Emotions in stories have been studied also from a media-specific perspective, by taking into account the expressive means enabled by the media through which the story is delivered, from film [27] and videogame [28] to literary studies [15].

By creating an emotional bond with the characters, through mechanism such as sympathy, the audience experiences the process of identification that contemporary aesthetics sees as the core of the narrative experience [23, 13]. The study of how characters' emotions are understood by the audience is relevant to understanding how this bond is constructed. In this paper, we try to shed light on the understanding of characters' emotions through the insight provided by a group of people with traumatic brain injuries (TBI). Our attempt leverages the discrepancy observed in moral and emotional functioning in TBI patients to assess the role of the rational and emotional components in the understanding of stories. Traumatic brain injured patients are characterized by impairments of social cognition, as described by Milders et al. [21]. Although it is very difficult to find clear indications on the characteristics of such impairments, it is

clear that they involve moral functioning. According to Milders et al. [21], moral behavioural consequences of TBI tend to be more difficult than physical impairments for caregivers to deal with. Although there is no theoretical account on these moral impairments, the main differences on moral functioning for brain-damaged patients seem to pertain to tasks involving moral emotions [14, 22].

Computational models of characters, delivered by the AI research in virtual agents [3, 25], provide a formal model of the characters' behavior and emotions that can be compared with the expectations of the audience about what the characters will do and feel. Moreover, computational models of characters provide a formal framework for designing and implementing experiments about story understanding. Once validated, these models can be straightforwardly employed to drive the behavior of virtual characters in applications that range from entertainment to education and rehabilitation.

In previous work, Battaglino and Damiano [5] demonstrated through experiments with users that a computational model of agent (the Moral Emotional Agent, or MEA [6]) that encompasses moral values and emotions actually matches the expectations of the audience about the behavior and the emotions of characters in stories that concern moral dilemmas. Our methodology is aimed at comparing the differences in the understanding of emotions, and of moral emotions in particular, in traumatic brain injured patients and control subjects with no traumatic brain impairment. The approach we propose adapts to our case study the methodology described by Battaglino and Damiano [5] to evaluate the MEA model: the characters' emotions generated by the agent model were compared with the emotions attributed to the character by the human users, while here we compare them with the test and control group to study the differences between the two groups.

This paper is organized as follows. In Section 2 we describe the agent model we adopt for modelling story characters and its validations through experiments conducted with human subjects. The scenarios employed for the original experiments, re-used here for investigating the understanding of emotions by TBI patients, are described in Section 3. The experiment is described in 4 and its results are discussed in Section 5. Conclusion ends the paper.

2 Modelling story characters with the MEA agent

Moral values are especially relevant to the understanding of stories. According to Bruner, who stressed the importance of stories in cognitive and social psychology, one of the major functions of stories is the transmission of a culture's values [8]. In Bruner's view, narrative characters are taken by the audience to have not only beliefs and desires, but also values. In drama, the importance of values is accrued by the fact, clearly stated by Elder Olson, that "Our emotions are evidently functions of a system of values, and are regulated by that system" [24]. The Moral Emotional Agent (MEA) [6] accounts for the relevance of values through an explicit moral component: in the MEA model, the emotional state of the agent in response to a given situation is determined not only by its

congruence with the agent's goals but also by its congruence with the agent's values. The experiments described by Battaglino and Damiano in [5] show that the predictions of the model about the story characters' actions and emotions actually match the expectations of the users.

2.1 The MEA model

Designed to implement virtual characters for linear and interactive storytelling, the MEA model integrates into the well known Belief Desire Intention (BDI) model [7] an emotional component that relies on the cognitive theory of emotions by Ortony, Clore and Collins [1] (often referred to as OCC model). The BDI model, inspired by the pioneering work of Bratman [7] on bounded rationality, provides a theoretically sound and formally specified basis for the implementation of virtual agents [4, 18]. In the BDI model, the agent forms practical intentions (or plans) based on her/his beliefs in order to achieve her/his desires (or goals).

Here, we provide a sketch of the MEA model, focusing on the role of emotions in its deliberation component; details can be found in [6]. Formally, the agent features a set of *goals*, which form the motivational component of the agent. The agent's deliberation (Deliberation Process in Fig. 1) depends on the importance of its goals and on the possibility for the agent to achieve them. Each goal is associated with an importance of success and an importance of failure, and with three different set of conditions: adoption conditions, success conditions and failure conditions. When the agent believes that one of the adoption conditions of a goal is verified in the world, the goal becomes an *active intention* and can compete for being selected to become the *active focused intention*; if this is the case, the agent tries to find sequences of actions (i.e. plans) to achieve it. For each plan, the agent computes its cost and probability of success.

Following Van Fraassen's suggestion that moral values should be prioritized [12], the moral dimension of the agent is formed by a set of values organized in a scale. Each value is associated with a *priority* that specifies the importance of the value for the agent, and with a set violation conditions. When one of the violation conditions holds in the state of the world, the value is *at stake*. If the violation condition does not hold anymore (and no other violation conditions have become true), the value is brought back to balance.
The values are the moral drive of the agent: they constrain the behavior of the agent to its moral dimension and allow the agent to appraise the behavior of self and others in moral terms.

After the deliberation phase, the agent trades off its values against its goals to choose a plan for execution (see Planning and Anticipatory Appraisal Process in Fig. 1). The role of values in the agent's deliberation is mediated by the emotional appraisal: for each plan, the agent appraises the consequences of the plan in terms of the emotional states it would generate (Anticipatory Emotional Appraisal). To do so, the agent calculates the *Expected Emotional Reward* utility for each plan, based on the emotions generated by the expected consequences of

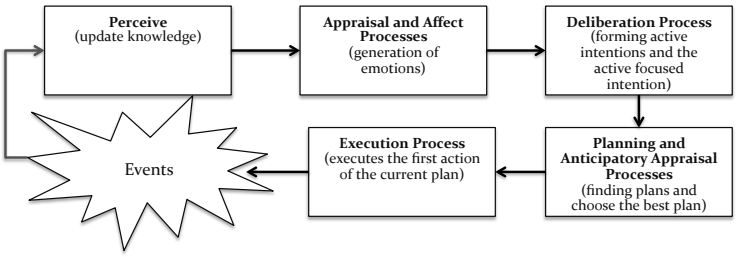

Fig. 1. A high level view of the agent's reasoning.

plans on the agent's goals and values. The agent will choose the plan that has the best trade-off between positive emotions and negative emotions.

The model was implemented into FAtiMa [11], a modular architecture designed to develop emotional agents.

2.2 Goals and values in emotional appraisal

Following an established practice in virtual agents, the emotional component of the MEA model relies on the paradigm of appraisal theories of emotions [25, 19]. In particular, the MEA model translates the appraisal process described by OCC [1] in terms of goal and value processing. In practice, the appraisal process generates a *desirability or (undesirability*) variable when an agent's goal is achieved (or not achieved); it generates a *probability* variable depending on the probability that an agent's plan succeeds; finally, it generates a *praiseworthiness* (or *blameworthiness*) variable when an agent's value is balanced (or put at stake)by the execution of some action. Given the appraisal variables, then, the affect generation process generates emotions according to the following domain-independent rules:

- **Joy (or Distress)** if a desirable (undesirable) appraisal variable is generated;
- **Hope (or Fear)** if a likelihood appraisal variable is high (or low);
- **Pride (or Shame)** if praiseworthy (blameworthy) appraisal variable is generated and the responsibility is self-caused;
- **Admiration (or Reproach)** if praiseworthy (blameworthy) appraisal variable is generated and the responsibility is other-caused.

The intensity of the emotions is based on multiplicative relationship between the importance of values and goals, the effort made (i.e. the cost of the plan) and the probability of success of the plan.

When both appraisal variables for values and goals are generated, the following compound emotions arise: **Gratification** (Joy and Pride), **Gratitude** (Joy and Admiration), **Remorse** (Distress and Self-Reproach), **Anger** (Distress and Reproach).

Thanks to the anticipatory emotional appraisal, the behavior of the agent is compelled by its moral values: for example, consider an agent who desires to eat a chocolate candy, but should steal in order to obtain one. In this case, the action that achieves the agent's goal threatens the moral dimension of the agent putting at stake the value "honesty". The emotional state governs the conflict that may arise from the agent's desires and moral values: the agent will opt for stealing only if the positive emotions generated by the perspective achievement of the goal of eating the chocolate candy outbalance the negative, moral emotions generated by the blameworthy action of stealing. The same rules are employed in the agent also for the appraisal of the events and the actions executed by self or others (Appraisal and Affect Processes in Fig. 1).

3 Narrative Scenarios

The experiment described by Battaglino and Damiano in [5] is based on a testing procedure that compares the predictions of the MEA model about the narrative characters with the expectations of human subjects. The experiment included three narrative scenarios, created with the help of a drama expert, and inspired by well known literary works. An implementation of the MEA agent model [10] was employed to model the behavior of a character in each narrative situation, in order to guarantee that the described characters' behavior and emotions were aligned with the choices and the emotional states predicted by the model. For each scenario, two different courses of action (namely, two *plans*) were generated by manipulating the priority of the values of the agent, but only one matched the actual character's behavior in the narrative situation that inspired the scenario (for simplicity, it was termed the "right" plan).

For each scenario, the subjects were asked to act and feel according to the characters' beliefs and values (as if they were doing practice in an acting school), choosing between the two courses of action generated by the system and selecting the appropriate emotions for the chosen option. The actions and emotions selected by the testers were compared with the actions and the emotional states generated by the model, in order to assess its validity and coverage. The experiment revealed that the predictions of the agent model substantially matched the expectations of the testers about the characters. Also, a comparison with the scale of values declared by the testers revealed that it has not influence on their expectations about the character's behavior and emotions [5].

Scenario 1. Wallace and uncle George's roses: *Wallace and Charlie are cousins. They live in the country where uncle George has a nursery of precious roses he brings to gardening contests. Uncle George is very jealous of his roses. Charlie wants to make a gift to his girlfriend and asks Wallace to give him the key of the nursery to get one.*
In this scenario, Wallace has to choose whether to be loyal to his cousin Charlie or to uncle George. The system finds two plans: the plan `giving_key` contains the action of giving the key to Charlie, thus deceiving uncle George, and puts at stake the value *Honesty* (priority 7.0), so Wallace will feel Shame emotion for

Scenario	Plan	Values at stake	Values balanced	Emotions
1. Wallace and uncle George's roses	giving_key	Honesty	-	Shame
	refusing_key	Loyalty	-	Shame
2. At school!	revenging	Pity	Justice	Shame, Joy, Pride, Gratification (Joy + Pride)*
	letting_go	-	-	Anger (continued)
3. A difficult choice	staying	Family	-	Shame, Joy, Distress, Remorse (Distress + Shame)*
	leaving	Happiness	-	Shame,Shame, Joy, Distress, Remorse (Distress + Shame)*
* as explained in the text, in this case the involved goals have also a role in the appraisal of emotions.				

Fig. 2. Overview of the narrative scenarios used in the experiment: for each alternative option (plan) of the character, we report the values it engages and the emotions it generate in the anticipatory emotional appraisal.

putting at stake this value; the plan `refusing_key` contains the action of refusing to give the key to Charlie and it puts at stake the value *Loyalty* to Charlie (priority 8.5), so Wallace will again feel Shame.
Assuming that the two plans have the same probability of success, Wallace's anticipatory appraisal leads him to choose the plan with the highest *EER*: in any case, Wallace will feel Shame, but will choose the course of actions that brings him to a state in which the Shame intensity is lower, `giving_key`: he gives the key to Charlie and feels Shame for putting at stake the value *Honesty*.

Scenario 2. At school!: *Tom is bullied by his classmate Pier. Pier has taken from Tom the role of director of the school newspaper, putting around lies about him. A few days later, Tom sees that Pier has forgotten his backpack with all his stuff in the locker room. Tom digs in Pier's backpack and finds evidence that Pier copied the class test. Tom is now uncertain about what to do, whether to take revenge against Pier or to pass through this situation.*
In this scenario, Tom has to choose if he wants to take vengeance or not. The system finds two plans: the plan `revenging` contains the action of humiliating Pier and puts at stake the value *Pity* (priority 7.5), so Tom will feel Shame, but it also brings back to balance the value *Justice* (priority 8.5) put at stake by

Pier, and satisfies the goal of being the director again, so Tom will feel Joy (for satisfying his goal) and Pride (for restoring his value *Justice*) as well (which, combined, give the Gratification emotion); the plan `letting_go` contains the action of letting it and has no effects on Tom's goals and values: If Tom performs this plan, the situation doesn't change and the value *Pity* is not put at stake: since the value *Justice* was put at stake by Pier before, Tom may still feel residual Anger towards Pier. Assuming that the two plans have the same probability of success, Tom chooses to execute the plan `revenging` and takes revenge on Pier: Tom will feel Shame for putting at stake the value *Pity*, but will also feel Shame and Gratification (Joy and Pride).

Scenario 3. A difficult choice: *New York, 2003. Mark and Lucy are married and they have a beautiful baby. Lucy has agreed to go a couple of years to Italy for the job of her dreams: working as a curator of a famous art gallery in Rome. Mark, however, has always wanted to be a judge in New York. Just when Lucy has officially accepted her job in Italy, Mark gets the seat as a judge in New York. Mark's desires are of being with his family and having the work of his dreams as well. Now, he has to choose whether to have the job or to stay with his family.* The system finds two plans: the plan `staying` contains the action of staying in New York without his family, which satisfies the goal of being a judge and threatens the goal of being with the family, putting at stake the value *Family* (8.0 priority), so he will feel Joy for satisfying his goal of being a judge, but also Shame (for putting at stake a value) and Distress (for threatening the goal of being with the family), which combine into the emotion of Remorse. The plan `leaving` contains the action of going to Italy, which satisfies the goal of being with the family but threatens his goal of being a judge and puts at stake the value *Happiness* (priority 8.5), so, again, Mark will feel Shame for putting at stake a value, Distress for threatening his goal of being a judge, but Joy for satisfying his goal of being with the family. The system chooses the plan `staying` since it puts at stake the value with a lower priority (Happiness), due to the equal importance of success of both goals: in any case Mark will feel Shame, but the anticipatory appraisal would make him choose the course of actions that brings him to a state of affairs in which the Remorse is lower.

4 Testing moral emotion understanding in TBI patients

The methodology described in Section 3 was adapted to a new sample made up of a group of TBI patients and a group of neurologically healthy participants. In the experiment, both groups (the test and control groups) were asked to execute the test; then, we compared the answers of the two groups with the predictions of the model and observed the differences between them. The goal was to gain insight on the moral and emotional functioning of the TBI patients and on the differences with the group of neurologically healthy subjects (which was known to correspond to the predictions of the agent model from the previous experiment). The MEA agent functioned here as a formally represented model of

the understanding of the character's behavior and emotions, validated through the experiment described in Section 3.

Our hypothesis was that the impairment of moral emotion reasoning that characterizes TBI patients would let differences emerge concerning moral emotions, without necessarily affecting their capability to reason about the moral implications of actions. Finding such pattern would also confirm the two-fold nature of the utilitarian and emotional nature embedded in the design of the MEA agent.

4.1 Description of the test

Test procedure and material. The experiment was conducted online, via a text-based web interface. For each scenario, a short text introduced the character and her/his values, then a narrative situation was described that put a stake the character's values. The scale of values was presented to participant not in a numerical format but with a figurative scale, in order to make the values priorities apparent at first glance (see Fig. 3). The task of identifying the expected course of action and emotions for the character was introduced to the participants as a game: the participants were asked to pretend they were exercising identification in an acting class, in order to leverage their capability to take the point of view of the character and behave "as if" they were the in the character's shoes. For each scenario, a pair of alternative actions were submitted to the participants, who also had to select a set of appropriate accompanying emotions. Each narrative situation was encoded in formal terms as described in Section 2 and the resulting formalization was employed to generate the behavior of the main character.

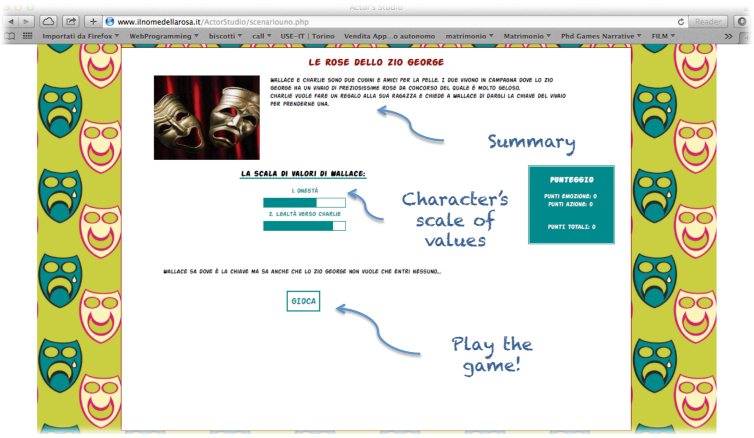

Fig. 3. A screenshot of the interface of the online testing system (in Italian).

First, two TBI patients, whose cognitive profile was similar to that of other TBI participants, were asked to participate in the study in order to test if their

cognitive impairments (i.e. reading difficulties, attention deficits) would cause them to find the tasks difficult or impossible to complete. Therefore, the methodologies were slightly modified in their verbal content in order to be simpler and shorter (as patients were very slow in reading, taking more than 60 min. for each test).

Participants. The clinical sample was made up of TBI patients recruited from the Puzzle rehabilitation center in Turin (Italy); the healthy participants were volunteers. 13 TBI patients and 13 control subjects participated in the experiment. TBI patients received the tests from the experimenter, while healthy participants completed the tasks at home, via an online form. The experimenter's help was needed to avoid noise from reading/comprehension/movement impairments often occurring in this kind of patients, as emerged in the aforementioned trials.

Measures. The participants' answers were transformed into scores in order to allow comparisons between groups. For this test, the following variables were taken into account, and a score was computed for each:

- **Action total score** (score=1 each time participants guessed the agent's 'right' action according to the MEA model);
- **Emotion total score** (score=1 each time participants chose the emotion predicted according to the MEA model);
- **Moral emotion score** (score=1 each time participants chose the emotion predicted according to the MEA model; only for moral emotions, being them admiration, contempt, gratification, gratitude, pride, shame, anger, blame);
- **Non moral emotion score** (score=1 each time participants chose the emotion predicted according to the MEA model; only for non moral emotions, being them joy, sadness, hope and fear);
- **Error total score** (score=1 each time participants both chose an emotion in disagreement with the MEA model and missed those predicted by the model);
- **Moral emotion error score** (score=1 each time participants both chose an emotion in disagreement with the MEA model and missed those predicted by the model; only for moral emotions);
- **Non moral emotion score** (score=1 each time participants both chose an emotion in disagreement with the MEA model and missed those predicted by the model, only for non moral emotions).

4.2 Results

We performed statistical tests to compare the performance of the two groups. We compared the means using Mann-Withney test (suitable for different-sized samples), using SPSS Statistics (IBM). Sample size was n=13 for TBI patients and n=11 for healthy participants (due to participants' drop-out).

Table 4 summarizes the results we found. No statistically significant difference was found between groups. Nevertheless, the mean difference for the moral emotion total score is close to significance ($p. < .063$). Mean difference for the

Measure	TBI M (SD)	CONTROL M (SD)	p-value (p<.05)
Action score	*2.15 (.081)*	*2.00 (.632)*	*.494*
Emotion total score	*3.38 (.768)*	*2.45 (2.067)*	*.106*
Moral emotion score	*2.15 (.081)*	*1.36 (1.567)*	*.063*
Non moral emotion score	*1.23 (.832)*	*1.09 (.831)*	*.820*
Error total score	*11.31 (2.394)*	*10.27 (2.832)*	*.303*
Error moral emotion score	*8.00 (2.160)*	*7.64 (2.111)*	*.776*
Error non moral emotion score	*3.31 (1.109)*	*2.64 (1.027)*	*.207*

Fig. 4. Overview of the experiment results.

Action total score is also worth mentioning; TBI patients don't seem to differ from healthy subjects for this kind of task, thus suggesting it to be more similar to a reasoning than to an emotional task. Percentages for action choice did not differ among the two groups (see Table 5); the more noticeable difference for the first scenario ("At school"), with TBI patients performing slightly better than the control group, should not surprise as controls had to become used to the methodologies by themselves.

As the table shows, scores for emotions (both moral and non-moral) are higher for TBI participants: although apparently surprising (according to our hypothesis), this can be explained by observing the error scores. For both moral and non-moral emotions, in fact, healthy participants show lower scores, thus suggesting that the TBI patients' better performance could be ascribed to their tendency to overselect emotions (as it can be seen in Table 6). For example, in the scenario "At school", TBI participants selected gratitude, pride and admiration as consequent emotions for the plans, while none of the controls did. However, we believe larger samples are needed to further assess this difference.

5 Discussion

Our work moves from the fact that, according to current literature, people that suffered from traumatic-brain injury tend to exhibit impaired moral functioning, even in the chronic phase of the injury. By "impaired moral functioning" we mean the meta-cognitive and interpersonal aspects of human behavior (e.g. being emphatic). With respect to the original experiment [5], our experiment was performed with TBI patients and with a control group of neurologically intact subjects. The goal of our experiment was to compare the expectations of TBI individuals about the narrative characters with the behavior and emo-

Scenario	Plan	TBI	Controls
1. Wallace and uncle George's roses	**giving_key**	46,2%	18,2%
	refusing_key	53,8%	81,8%
2. At school!	**revenging**	84,6%	81,8%
	letting_go	15,4%	18,2%
3. A difficult choice	**staying**	84,6%	81,8%
	leaving	15,4%	18,2%

Fig. 5. Overview of the plans selected by the TBI and control group for each scenario ('right' choice, predicted by the model, in bold).

Emotion	1. Wallace and uncle George's roses			2. At school!			3. Difficult Choice		
	M	**TBI**	C	M	**TBI**	C	M	**TBI**	C
Distress	-	**38,5%**	9,1%	-	**23,1%**	-	X	**23,1%**	18,2%
Joy	-	**23,1 %**	-	X	**23,1%**	36,4%	X	**76,9%**	54,5%
Fear	-	**-**	-	-	**46,2%**	54,5%	-	**23,1%**	9,1%
Hope	-	**15,4%**	18,2%	-	**-**	-	-	**30,8%**	9,1%
Pride	-	**23,1%**	9,1%	X	**23,1%**	27,3%	-	**38,5%**	36,4%
Shame	X	**23,1%**	27,3%	X	**23,1%**	18,2%	X	**23,1%**	18,2%
Admiration	-	**23,1%**	27,3%	-	**7,7%**	-	-	**30,8%**	18,2%
Reproach	-	**7,7%**	27,3%	-	**53,8%**	36,4%	-	**7,7%**	9,1%
Gratification	-	**7,7%**	18,2%	X	**53,8%**	-	-	**53,8%**	27,3%
Gratitude	-	**-**	-	-	**7,7%**	-	-	**23,1%**	9,1%
Remorse	-	**38,5%**	27,3%	-	**15,4%**	27,3%	X	**69,2%**	45,5%
Anger	-	**23,1%**	9,1%	-	**46,2%**	9,1%	-	**7,7%**	9,1%

Fig. 6. Overview of the emotions selected by the TBI and control group for each scenario (the first column of each scenario shows the emotion type predicted by the system). From top to bottom, the first group is non-moral emotions, the second group (light grey) is moral emotions, the third group (dark grey) is compound emotions.

tions generated by the MEA agent model, and to compare both with the control group, in search for insight about the nature of the impairment of TBI patients and the moral-emotional reasoning in general. To the best of our knowledge, no contributions thoroughly describe or help assessing these impairments.

Nonetheless, many authors have focused on the differences for moral judgment that can be found in brain-damaged individuals. Greene [14] and Martins et al. [20] illustrate how brain damaged individuals differ from healthy participants for moral personal dilemmas. According to Greene's perspective [14], personal moral dilemmas are those engaging emotional responses while evaluating the dilemma, whereas non-personal and non moral dilemmas elicit no emotional response. When facing a moral personal dilemma, TBI participants differ from neurologically intact individuals whilst they don't seem to differ when the dilemma elicits no emotion. According to Greene [14], personal moral dilemmas activate medial front gyrus, superior frontal gyrus while non-personal moral dilemmas engage dorsolateral and prefrontal areas; the same pattern can be observed when evaluating non-moral dilemmas. According to Moretti, Dragone and DiPellegrino [22], TBI individuals whose brain damage includes areas such as the ventromedial prefrontal cortex are able to perform moral judgements (in terms of right-wrong), but seem compromised in the emotional counterpart. Authors describe this as a "moral emotions selective impairment". Also, Hutcherson and colleagues [17] described the interplay between utilitarian and emotional appraisal during moral judgment involving to moral values. According to their work, utilitarian and emotional moral appraisals are computed independently and then integrated in a moral value response. We believe the emotional appraisal is what makes TBI patients' performance impaired, thus causing nonstandard moral judgments. Thus our results seem in line with current literature on emotional processing in brain damaged individuals; taken altogether these findings add further evidence to the suitability of the MEA model previously evaluated [5]. The difference we found is, in fact, what could be expected according to current studies on TBI.

We also had the secondary aim of gaining preliminary data for the development of a rehabilitative tool where TBI patients face moral situations, training their ability to perform moral reasoning and reasoning on their moral emotions. In this perspective, the results obtained were encouraging. On the one side, the experiment confirms the feasibility of this type of rehabilitation, accrued by the interest showed by the comments of the TBI patients: *A* said: "this test allowed me to think about the values that shape my moral judgment, helping me understand how every action comes after deep and elaborate reasoning"; *R.* believes this test "is important, as it makes my head start". On the other side, the indirect validation of the MEA model it provides opens the way to the exploitation of virtual agents in narrative-based rehabilitation application. The literature about serious games for rehabilitation mainly focuses on the design and implementation of empathetic agents who respond to the emotional states of the patients by providing emotional cues through verbal and non verbal behavior [16]. Our work goes in a complementary direction: here, we hypothesize the use of virtual

agents for narrative games that train the recognition of emotions and empathy in patients who have some impairment in this area. Following an established practice in rehabilitation [26], virtual characters may be exploited to train patients to appraise moral emotions in self and others.

6 Conclusions

Although it was not possible for us to find the key features of the impaired moral functioning in people who suffered from TBI described in literature, it is clear how, in order to shed more light on such differences, research must focus on emotion processing (rather than on reasoning). This follows what is stated by the mental models theory perspective on moral reasoning [9]: according to this perspective, moral reasoning is just regular reasoning that happens to concern moral issues. The TBI participants included in our study were in a late stage of their rehabilitation, thus their reasoning could be compared to those of neurologically intact individuals; nonetheless they differ from controls more significantly in their performance for those tasks requiring (moral or non-moral) emotional processing. We believe this indicates emotions as the key feature to investigate in order to analyze moral impairments following TBI.

We also believe that the interaction with virtual characters whose emotional state includes a moral emotional component could be useful to help TBI patients to train their emotional moral functioning and their understanding of moral (but also non-moral) emotions. So, our future work includes further investigations on the moral and emotional functioning of TBI patients with the help of virtual character by using larger groups of patients and more focused experiments.

References

1. A. Ortony, A.C., Clore, G.: The cognitive structure of emotions. Cambridge University Press Cambridge (1988)
2. Aristotle: Poetics. Oxford University Press (2013)
3. Aylett, R., Louchart, S., Dias, J., Paiva, A., Vala, M., Woods, S., Hall, L.: Unscripted narrative for affectively driven characters. IEEE Computer Graphics and Applications 26, 42–52 (2006)
4. Bates, J., Loyall, A.B., Reilly, W.S.: Integrating reactivity, goals, and emotion in a broad agent. In: In Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society. Citeseer (1992)
5. Battaglino, C., Damiano, R.: A character model with moral emotions: Preliminary evaluation. In: 5th Workshop on Computational Models of Narrative. p. 24 (2014)
6. Battaglino, C., Damiano, R., Lesmo, L.: Emotional range in value-sensitive deliberation. In: AAMAS. pp. 769–776 (2013)
7. Bratman, M.: Intention, Plans, and Practical Reason. Harvard University Press, Cambridge (MA) (1987)
8. Bruner, J.: The narrative construction of reality. Critical Inquiry 18(1), 1–21 (1991)
9. Bucciarelli, M., Khemlani, S., Johnson-Laird, P.N.: The psychology of moral reasoning. Judgment and Decision making 3(2), 121 (2008)

10. Cristina Battaglino, Rossana Damiano, J.D.: Telling the difference between asking and stealing: Moral emotions in value-based narrative characters. In: Seventh Intelligent Narrative Technologies Workshop (INT7)
11. Dias, J., Paiva, A.: Fatima modular: Towards an agent architecture with a generic appraisal framework. Proceedings of the International Workshop on Standards for Emotion Modeling (2011)
12. Fraassen, B.C.V.: Values and the heart's command. Journal of Philosophy 70(1), 5–19 (1973)
13. Giovannelli, A.: In Sympathy with Narrative Character, Journal of Aesthetics and Art Criticism, vol. I, pp. 83–95. Wiley Blackwell (2009)
14. Greene, J., Haidt, J.: How (and where) does moral judgment work? Trends in cognitive sciences 6(12), 517–523 (2002)
15. Hogan, P.C.: Affective narratology: The emotional structure of stories. U of Nebraska Press (2011)
16. Hudlicka, E.: Virtual affective agents and therapeutic games. Artificial Intelligence in Behavioral and Mental Health Care p. 81 (2015)
17. Hutcherson, C.A., Montaser-Kouhsari, L., Woodward, J., Rangel, A.: Emotional and utilitarian appraisals of moral dilemmas are encoded in separate areas and integrated in ventromedial prefrontal cortex. The Journal of Neuroscience 35(36), 12593–12605 (2015)
18. Lim, M.Y., Dias, J., Aylett, R., Paiva, A.: Creating adaptive affective autonomous npcs. Autonomous Agents and Multi-Agent Systems 24(2), 287–311 (2012)
19. Marsella, S.C., Gratch, J.: Ema: A process model of appraisal dynamics. Cognitive Systems Research 10(1), 70–90 (2009)
20. Martins, A.T., Faísca, L.M., Esteves, F., Muresan, A., Reis, A.: Atypical moral judgment following traumatic brain injury. Judgment and Decision Making 7(4), 478 (2012)
21. Milders, M., Ietswaart, M., Crawford, J.R., Currie, D.: Social behavior following traumatic brain injury and its association with emotion recognition, understanding of intentions, and cognitive flexibility. Journal of the International Neuropsychological Society 14(02), 318–326 (2008)
22. Moretti, L., Dragone, D., Di Pellegrino, G.: Reward and social valuation deficits following ventromedial prefrontal damage. Journal of Cognitive Neuroscience 21(1), 128–140 (2009)
23. Oatley, K.: A taxonomy of the emotions of literary response and a theory of identification in fictional narrative. Poetics 23(1), 53–74 (1995)
24. Olson, E.: Tragedy and theory of drama, vol. 22. Detroit: Wayne State University Press (1961)
25. Peinado, F., Cavazza, M., Pizzi, D.: Revisiting character-based affective storytelling under a narrative bdi framework. In: Interactive Storytelling, pp. 83–88. Springer (2008)
26. Riva, G.: Virtual reality in psychotherapy: review. Cyberpsychology & behavior 8(3), 220–230 (2005)
27. Smith, G.M.: Film structure and the emotion system. Cambridge University Press (2003)
28. Zagalo, N., Torres, A., Branco, V.: Emotional spectrum developed by virtual storytelling. In: Virtual Storytelling. Using Virtual Reality Technologies for Storytelling, pp. 105–114. Springer (2005)