Rule based appraisal of emotions in drama

V. Lombardo[1], C. Battaglino[1], R. Damiano[1], and A. Pizzo[2]

[1] Department of Computer Science and Cirma, University of Torino
[2] Department of Humanities and Cirma, University of Torino
vincenzo|rossana|battagli@di.unito.it,antonio.pizzo@unito.it

Abstract. In stories, the emotional charge of the characters plays an important role in engaging the audience. The emotional states of the characters allow the audience to understand their motivations and to perceive their reactions to a dramatic situation. In this paper, relying on a semantic representation of the drama features, we present and evaluate an emotional rule system that generates the characters' emotions based on their representation of their mental states.

Keywords: emotion annotation, drama ontology

1 Introduction

Computing characters' emotions is relevant for a number of tasks ranging from retrieval to editing. Consider, for example, the following scenarios: a system, conceived for the general public, that searches a (multimedia) story bank (such as, e.g., [6], see below), through an effective tool that goes beyond mere editorial metadata (title, author, etc.), able to answer queries of the type "the novel where a woman drowns her husband with the help of her lover but eventually goes insane from remorse"; an environment for assisted drama editing (such as, e.g., Dramatica[3]), where the writer can visualize the course of characters' emotions along the plot and assess their timing and coherence.

Cognitive theories of emotions can provide a systematic account of characters' emotions in stories. According to cognitive theories [13], emotions stem from how a character *appraises* a given situation with respect to its own goals and moral standards: if it appraises some event as beneficial, it is happy; if it appraises some event as deleterious, it is worried or disgusted; etc. Since the notion of appraisal advocates an intentional account of agency, cognitive theories of emotions have been integrated into virtual characters by using the well known BDI model, which provides the required primitives for the appraisal process [2, 3].

In this paper, we leverage a computational model of emotions [1], based on the OCC theory of emotion appraisal [15], to create a set of rules that compute characters' emotions based on a description of their goals and values. We assume a BDI based description of the characters [11, 12], where characters are driven by their goals and respond to the violations of their values, engaging in conflicts

[3] http://dramatica.com

that are the input to their emotions. suitable to develop functionalities such as the search and editing functions mentioned above.

This paper is structured as follows: after surveying the related works about how appraisal theories are encoded in intelligent agents (Section 2), we illustrate the basic encoding of the drama facts through the Drammar ontology (Section 3). Section 4 presents the emotional rules system, while Section 5 presents the experiment.

2 Related Work

A varieties of recent projects have investigated the creation of story repositories with formal tools. Propp's work, in particular, has been the object of formalization with AI tools in fields that range from the creation of fictional story worlds [7] to narrative generation [8]. The DramaBank Project [6] is a repository of semantically encoded narratives, based on a formal annotation, oriented at the surface generation of different stories from shared nuclei [18]. The DramaBank annotation language accounts for causality and intentionality in stories with specific operators, such as *Attempt to cause*, but does not account for an emotional level in characters, since they are mostly concerned with the encoding of plots rather than character structures. The Narrative Knowledge Representation Language (NKRL) proposed by [19] also provides tools for the annotation of the narrative content, but it does not acknowledge the role of the characters and their emotions.

The integration of emotions into virtual characters' architectures has seen its first, pioneering approach in [5], emotional states are explained as a consequence of specific configurations of mental states (e.g., beliefs and goals), that are the output of a person's appraisal of the environment she/he is situated in. For example, a situation may be *desirable* with respect to the person's goals, or it may be appraised as *immoral* because it contains some immoral action with respect to the moral beliefs of the person. A number of computational models of emotions, including [5], rely on the appraisal theory proposed by Ortony, Clore and Collins [15] (OCC). A relevant feature of computational models of emotions is that the emotion appraisal process is carried out in a domain–independent fashion. In [17], the independence of the appraisal process from the domain is limited to the desirability of events, which is based on goal processing; the appraisal of actions as praiseworthy and blameworthy, on the contrary, is reduced to the principles such as "help my goals to succeed" or "do not cause my goals to fail". In [4], the appraisal of events is independent from the domain, and is carried out by processing the syntactic information encoded in the representation of plans (e.g., the probability of success) and goals (e.g., the success or failure conditions). The system, however, does not contain the necessary information for generating the appraisal variables, which are necessary to Attribution emotions. In [1], Attribution emotions, such as Pride or Shame, are derived from the evaluation of actions in a domain–independent way, based on the notion of moral values (such as 'honesty', 'freedom').

3 Drammar Ontology

Drammar[4] is a computational ontology for the representation of the elements of the drama (for details about the encoding, see [10]). For a description of the theoretical foundations for dramatic elements see [10].

Drammar representation of characters centers upon the notion of agents' intention (realized through a plan) and the goal achieved (or tried to achieve). A plan consists of the actions that are to be carried out in order to achieve some goal; plans are organized hierarchically, with high–level behaviors formulated as lower–level plans (called subplans). Goals originate from the values of the agents that are engaged by the plans, i.e., put at stake or balanced through the plan actions, given the beliefs (i.e., the knowledge) of the agents. The representation of dramatic characters is formalized through the rational agent paradigm, or BDI (Belief, Desire, Intention) paradigm [2] (which has already seen some applications in the computational storytelling community [14] [16]).

The scenes are the places for the interplay of the actions that are carried out by the agents to achieve their goals. The scene is built in order to orchestrate the conflicts (or, alternatively, the support relations) over the goals and to induce into the agents the emotions sought after by the author of the drama. The emotions felt by the agents are the dramatic qualities *par excellence* and are computed through the appraisal operation. The appraisal operation, encoded through SWRL rules, will be addressed in detail in the next section. The repre-

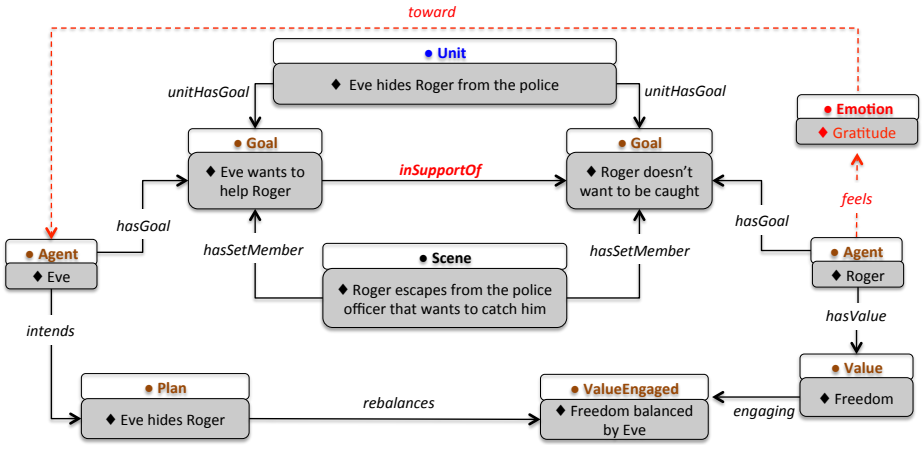

Fig. 1: The representation of the example: Eve helps Roger to hide from the police officers.The dotted lines indicate the annotation of emotion felt by Roger.

sentation example in Fig. 1 refers to a scene taken from the "North by Northwest"

[4] An available version of the ontology, encoded in the OWL 2 RL language, can be downloaded at http://www.di. unito.it/ vincenzo/FTP_SWJ/

movie by Alfred Hitchcock, a tale of mistaken identity where the main character Roger tries to prove that he is not the 'double' George Kaplan. In the example, we model the scene in which Eve helps Roger to hide from the police officers who want to catch him, because they believe that he is an assassin. Eve is a spy of the USA government and knows that Roger is not an assassin, so she helps him. Roger feels Gratitude toward Eve, because her goal is in support of Roger's goal of not being caught and her plan re-balances Roger's *Freedom* value.

In Fig. 1, the incident described above is encoded in the `Unit` *"Eve hides Roger from the police"* (top). What motivate this unit are the following goals: 1) Roger's goal to not be caught by the police officers; 2) Eve's goal to help Roger; with the first goal being supported by the second one. The plan devised by Eve to achieve her goal engages Roger's value of Freedom. Goals and engaged values are handled through a scene structure depicted in the figure. In the example, emotions are represented by the properties `feels` and `toward` instantiated by the rules, so that Roger feels gratitude toward Eve.

4 Rule-based emotion generation

The automatic annotation of emotions is conducted via a set of rules, informed on a computational model of the emotional agent, namely the Moral Emotional Agent described in [1].

As anticipated in Section 2, in OCC theory emotions are activated as a consequence of a person's (here, an agent's) subjective appraisal of a given situation. The *appraisal process* encompasses the following elements: the appraising *agent*, the appraised *situation*, the *dimension* of appraisal. Depending on the configuration of these elements, different emotion types are generated. The OCC theory acknowledges three main dimensions of appraisal: the utilitarian dimension of *desirability* (or undesirability), that [1] map onto the achievement (or failure) of goals, following an established tradition in computational models of emotions (e.g., Joy or Distress); the moral dimension of *praiseworthiness* (or blameworthiness), that [1] map onto the compliance (or conflict) with moral values (e.g., Pride or Shame); the *affection* for an entity involved in the situation. The utilitarian dimension can be also appraised by the agent from the point of view of another agent, thus generated other agent-oriented emotions (e.g. Pity or Reproach).

The *target* of the emotion, then, varies depending on the appraisal of the situation as a mere event or as an intentional act: in the former case, the target of the emotion is the event itself and the relevant dimension of appraisal is the desirability of the event; in the latter case, the target is the agent who intentionally performed the act and the relevant appraisal dimension is the praiseworthiness of the action. A third case is the appraisal of a specific entity (e.g., an object or a person) involved in the situation according to an affective, subjective inclination (e.g., Love and Hate): here, we do not consider this case since the affection towards the target is intrinsic to the appraising agent and cannot be computed. If the appraised situation is still ongoing, a *prospect-based* emotion will be gen-

erated based on the agent's expectation about its outcome (e.g., Hope or Fear). Otherwise, the generated emotion type depends on the actual outcome of the event with respect to the dimensions of desirability and praiseworthiness (e.g., Relief).

In OCC, emotions are grouped into emotion families depending on the appraisal dimensions. When the appraisal dimension is desirability, *Well–being* emotions are generated; these can be *Prospect–based* if the refer to the prospective accomplishment of events. The appraisal of actions according to the moral dimension gives rise to *Attribution* emotions. The appraisal of situations from the perspective of other agents gives rise to *Fortune–of–Others* emotions.

In previous work, we chose SWRL rule language [9] as the formal tool for encoding the emotion annotation rules [10]. The SWRL rules augment the OWL–based representation with a rule layer built on top of it, adding the possibility to declare arbitrary Horn clauses expressed as *IF THEN* rules. Encoding emotion generation using *SWRL* rules enables the automatic generation of the emotions of the characters in a scene annotated in Drammar.

Translating the computational model of emotions into the emotion generation rules involves a mapping of the elements of the appraisal process (appraising agent, situation and dimension of appraisal) onto the primitives of the Drammar ontology. Basically, the rule antecedent represents a character's appraisal of a situation, and is based on the character's goals, values and plans (e.g., a goal achieved or not, a value put at stake, a plan the character is committed to). The rule consequent asserts what emotions the character feels as a consequence of the appraisal and what is the target of the emotion.

The appraised situation is mapped onto a scene of the drama and the appraising agent is mapped onto a character featured in the scene. Modelling the appraisal dimension requires a more complex mapping. The content of the scene is represented as a set of variables that correspond to goals (`Goal` in Drammar), achieved by plans (`Plan` in Drammar), and values (`Value` in Drammar), engaged by the execution of plans. Appraisal dimensions are represented as relations over this set of variables. The appraisal of an event as *desirable* (or *undesirable*) depends on the relation between a goal of the appraising agent and another's agent goal, achieved by the plan of the other agent in the scene. The relation is expressed through the properties `inConflictWith` or `inSupportOf`: an event is desirable if the goal it achieves is `inSupportOf` of the agent's goal, undesirable otherwise. Notice that, in this case, a plan is construed as an (intentional) event, in line with the OCC theory.

The appraisal of an action as *praiseworthy* (or *blameworthy*) depends on the relation between a character's value and a plan committed by another agent (or by the agent itself) as a way to achieve some goal. The relation between a value and a plan is expressed by the property `atStake` concerning one of the values of the character: if the value is put `atStake` as a consequence of the execution of a plan in the scene, the plan is blameworthy; otherwise, if a value is not at stake anymore after the execution of a plan, the plan is praiseworthy.

The temporal dynamics of the appraised situation, relevant for Prospect–based emotions, is grasped by a property describing the status of the plan execution in the agent's expectations. The status of a prospect event is expressed by the property `accomplished` of a plan, whose value is a string. A plan accomplishment can be *uncertain* (i.e., "uncertain") if the agent expects the plan to achieve its goal, *successful* (i.e., "true") if the plan has been successfully executed and has achieved its goal as expected, *failed* (i.e., "false") if the plan has not achieved its goal, differently from what expected. The Fig. **??** illustrates the rules for emotion generation. *Well-being* emotions, such as Distress and Joy, depend on the relation between a `Goal` *?G* and a `Goal` $?G_{SA}$ owned by an `Agent`. An event is desirable if it encompasses a plan that achieves a goal *?G* `inSupportOf` of the agent's goal $?G_{SA}$, undesirable if the goal *?G* is `inConflict` with the agent's goal.

Fortune-of-others emotions, such as Happy-for another agent, depends on the agent's emotions *Love/Hate* for another agent encoded in the representation and on the (un)desirability of an event for an other agent's *Goal* $?G_{OA}$. For example, if the `Agent` *?SA* loves another `Agent` *?OA* and the `Goal` *?G* is `inSupportOf` the `Goal` $?G_{OA}$ of the other `Agent` *?OA*, *?SA* feels *Happy-for* for the other agent *?OA*. Otherwise, the agent feels *Gloating* toward the other agent.

Attribution emotions arise when the agent appraises the consequences of an action with respect to its values. This happens when an `Agent` *?SA* owns a `Value` *?V* that is a *ValueEngaged ?VE* in the effects of the `Plan`. The `Agent` *?SA* appraises the `Plan` *?P* as *praiseworthy* if the value *?VE* is re–balanced by the plan (i.e., the data property `atStake` of *?V* is *false* as a consequence of the plan); the `Plan` *?P* is blameworthy if *?VE* is put at stake by the plan (i.e., the data property `atStake` of *?V* is *true* as a consequence of the plan). Attribution emotions can be self– or other–directed: the `Agent` *?SA* feels Pride or Shame if it intends the `Plan` *?P* and the plan is, respectively, praiseworthy or blameworthy. Otherwise, if another `Agent` *?OA* in the scene intends the `Plan` *?P*, *?SA* feels Admiration or Reproach.

Compound emotions arise when the agent feels Well-being emotions and Attribution emotions at the same time. Gratification (Remorse) emotion rule fires if the `Agent` *?SA* feels *Joy* (*Distress*) and *Pride* (*Shame)* in the `Scene` *?S*, *SWRL Gratification*(*SWRL Remorse*) rule fires and *?SA* also feels *Gratification (Remorse)*. Gratitude (Anger) emotion rule fires if the `Agent` *?SA* feels *Joy* (*Distress*) and *Admiration* (*Reproach)* in the `Scene` *?S*, *SWRL Gratitude*(*SWRL Anger*) rule fires and *?SA* also feels *Gratitude (Anger)*.

In the following (Fig. 2), we describe the activation of the SWRL rule for Relief for the agent Roger in the running example taken from the "North by Northwest" movie by Alfred Hitchcock. In particular, we focus on the scene in which two foreign spies, Valerie and Licht, believing that Roger is George Kaplan, try to kill him by forcing him to drink bourbon and by putting him into a moving car. Roger manages to exit from the car before it falls off a cliff. The `Scene` *"Scene_2.1.2 Roger's life is in danger"* has one `Agent`: the main character *"Roger"*. The emotional charge of the scene is usually described in

the traditional *misè en scene* focusing on the conflict between the two goals: Valerie and Licht want to kill Roger; Roger wants to stay alive. Given the event represented by Valerie and Licht's goal and by their failed plan, the system succeeds in calculating the resulting characters' emotional charge. Following the SWRL rules, the system outputs Roger's relief as the emotions triggered in the scene that corresponds to the unit. In (Fig. 2), the event is represented by the `Plan` *"Valerie and Rick kill Roger by putting him in the car"* that achieves Valerie and Licht's `Goal` *"Valerie and Licht want to kill Roger"*. The plan has the data

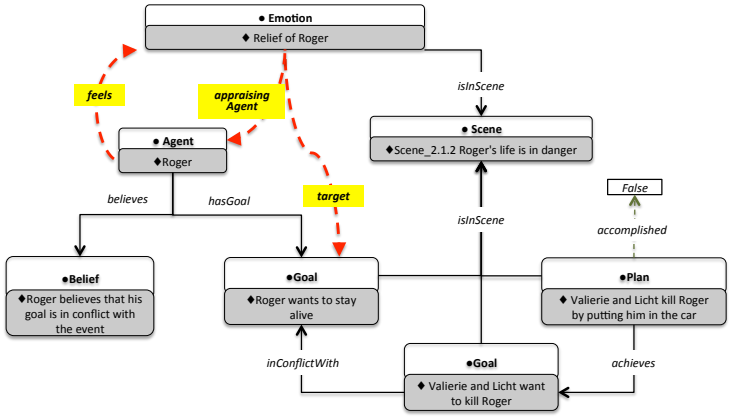

EVENT: the Plan ``*Valerie and Licht kill Roker by putting him in the car''* that achieves the Goal ``*Valerie and Licht want to kill Roger''*

Fig. 2: The annotation of the scene for the `Agent` *"Roger"*. The property `target`, `feels` and `appraisingAgent` are inferred by the rule for Relief emotion.

property `accomplished` set to *false*, this means that the event is *discofirmed*. The `Agent` *Roger* has the `Goal` *"Roger wants to stay alive"* that is `inConflictWith` Valerie and Licht's goal and the agent believes that his goal is in conflict with the event. Thus, the `Agent` *Roger* appraises the event as an *undesirable disconfirmed* event that leads to the activation of the Relief SWRL rule. The Relief rule consequent asserts that the `Agent` *Roger* is the `appraisingAgent` that feels the `Emotion` *Relief of Roger*, with the `Goal` *"Roger wants to stay alive"* as target (property `target`).

5 Evaluation & Discussion

In this section, we describe an experiment that aims at evaluating the application of the emotional rules presented in Section 4 on the data obtained by the manual annotation of stories by experts.

Experimental Protocol. The annotated corpus included two Hollywood movies, the historical romance *Casablanca* (by Michael Curtiz) and the unlikely thriller *North by northwest* by Alfred Hitchcock, respectively; an opera, *Carmen* (George

Bizet, libretto of Henri Meilhac and Ludovic Halévy), and the Greek tragedy *Oedipus the King* (Sophocles). The characters whose emotions are annotated are: *Roger* (North by northwest movie), *Rick*, *Ilsa* and *Laszlo* (Casablanca movie), *Carmen*, *Don Josè*, and *Michaela* (Carmen opera), and finally *Oedipus* (Oedipus Greek tragedy).

Each drama in the corpus was segmented into units and analyzed by an annotator who identified the segment's main incidents and then annotated the main actional elements of the units and the OCC–classified emotion types felt by the main characters. The annotators were students of the Media and Arts program, trained in dramatic narration; each work was annotated by a different annotator, selected based on her/his familiarity with the work. Subsequently, for each segment identified by the annotator, a drama scholar annotated the goals, plans, and values involved in the segment in the formal language of *Drammar*. Then, the annotation was fed to a reasoner[5] for the application of the *SWRL* emotion rules presented in Section 4.

We compared the improvement brought about by the rule with the results of a preliminary experiment, described in [10]. With respect to previous work [10], the rule set presented in Section 4 contains a monotonically more fine–grained encoding of the agent's expectations about prospect events, and of the agent's appraisal of the relation between its goals and the goal achieved in the appraised situation. This improvement allows us to discriminate between Prospect-based emotions and Well-being emotions, thus avoiding conflicts in rule activation. For the comparison, we availed ourselves of the following measures: *Human Annotated Emotion Types Detection* and *Tokens Accuracy*. The *Human Annotated Emotion Types Detection* represents the capability of a system of detecting the set of emotions types (i.e., the emotional range) annotated by humans for each character in the corpus. It is calculated by computing precision and recall of the generated emotion types on the emotion types annotated by the human annotators. This measure is not dependent on the number of tokens of a specific emotion types. The *Tokens Accuracy* represents the accuracy of a system in generating the number of emotions tokens annotated by humans for a given character in the corpus. It is calculated by computing precision and recall on the emotions type tokens (i.e., the single instances of each emotion type). This measure takes into consideration the number of times that the human annotators or the systems generate a specific emotion type.

Results. Regarding the *Human Annotated Emotion Types Detection*, the average precision is 0.88 and the average recall is 0.94 (see Table 1). With respect to previous results [10], we obtained an higher average precision (0.88 against 0.71) and an higher average recall (0.94 vs 0.89). In particular, the improvement regards the precision for characters who felt emotions types that belong to Prospect-based emotions such as Roger (0.73 vs 0.69), Rick (0.71 vs 0.62), Ilsa (1 vs 0.5), Laszlo (1 vs 0.8) and Oedipus (1 vs 0.79) (see Table 1).

Regarding the *Tokens Accuracy* measure, the average precision is 0.72 while the average recall is 0.93 (see Table 2). With respect to previous results, that

[5] Pellet, www.clarkparsia.com/pellet

show an average precision and recall equal to 0.51 and 0.85, respectively, the improvement is more apparent when we computed the *Tokens Accuracy* measure because it considers also the number of times that a certain emotions type is annotated by humans and generated by the systems (see Table 2 - Roger (0.62 vs 0.32), Rick (0.52 vs 0.43), Laszlo (0.83 vs 0.5) and Oedipus (0.91 vs 0.62)).

	NbN		**Casablanca**		**Carmen**			**OedipusAll**	
	Roger	Rick	Ilsa	Laszlo	Carmen	D. Juan	Micaela	Oedipus	
Precision	0.73	0.71	1	1	0.8	0.79	1	1	0.88
Previous Precision	0.69	0.62	0.5	0.8	0.5	0.79	1	0.79	0.71
Recall	1	0.83	1	1	0.8	1	1	0.89	0.94
Previous Recall	1	0.83	0.75	1	0.8	1	1	0.78	0.89

Table 1: Detection of the emotions types (*Human Annotated Emotion Types*).

These improvements are due to the fact that our emotional rules with respect to those presented in [10] do not include the appraisal of Hope (Fear) as part of the appraisal of Disappointment (Relief) and Satisfaction (Fear-confirmed) and the appraisal of Joy (Distress) as part of the appraisal of Relief (Disappointment). For example, 0 tokens of the Fear emotion type are annotated by

	NbN		**Casablanca**		**Carmen**			**OedipusAll**	
	Roger	Rick	Ilsa	Laszlo	Carmen	D. Juan	Micaela	Oedipus	
Precision	0.62	0.52	0.75	0.83	0.57	0.55	1	0.91	0.72
Previous Precision	0.32	0.43	0.5	0.5	0.27	0.47	1	0.62	0.51
Recall	1	0.90	1	0.83	0.8	1	1	0.91	0.93
Previous Recall	1	0.90	0.66	0.83	0.6	1	1	0.83	0.85

Table 2: Precision and recall on emotions tokens (*Tokens Accuracy*).

the human annotator for *Roger* in *North by Northwest*: while the previous rule system generated 6 emotion tokens of this emotion type, our rules discriminate the appraisal in a more efficient way and are in line with the human annotation.

6 Conclusion

In this paper, we described a system for the automatic generation of characters' emotions in stories, encoded in a set of SWRL rules. A rule based system alleviates the task of manual annotation of characters' emotions by providing a coherent and founded model for character emotion generation through a variety of media. We designed and ran an experiment where the emotions automatically generated by the rules were compared to the emotions assigned by human annotators to story characters on a corpus of stories ranging from traditional to new

media. The experiments showed a good performance of the model with respect to the annotation provided by the humans.

References

1. C. Battaglino, R. Damiano, and L. Lesmo. Emotional range in value-sensitive deliberation. In *Proc. of the 12th Int. Conf. on Autonomous Agents and Multi-Agent Systems*, (AAMAS'13), pages 769–776, 2013.
2. M.E. Bratman. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge (MA), 1987.
3. P. R. Cohen and H. J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213–261, 1990.
4. João Dias, Samuel Mascarenhas, and Ana Paiva. Fatima modular: Towards an agent architecture with a generic appraisal framework. In *Workshop on Standards in Emotion Modeling*, Leiden, 2011.
5. C. D. Elliott. *The Affective Reasoner: A process model of emotions in a multi-agent system*. PhD thesis, Northwestern University, 1992.
6. D. Elson. Dramabank: Annotating agency in narrative discourse. In *LREC*, pages 2813–2819, 2012.
7. C. R. Fairclough and P. Cunningham. A multiplayer opiate. *Int. Journal of Intelligent Games & Simulation*, 3(2), 2004.
8. P. Gervás. Propp's morphology of the folk tale as a grammar for generation. In *Proc. of Int. Workshop CMN*, pages 106–122, 2013.
9. I. Horrocks, P. F. Patel-Schneider, and H. et al. Boley. Swrl: A semantic web rule language combining owl and ruleml. *W3C Member submission*, 21:79, 2004.
10. V. Lombardo, C. Battaglino, R. Damiano, A. Pizzo, and A. Lieto. Coupling conceptual modeling and rules for the annotation of dramatic media. *Semantic Web*, 00:1–32, 2014.
11. V. Lombardo and R. Damiano. Semantic annotation of narrative media objects. *Multimedia Tools and Applications*, 59(2):407–439, 2012.
12. V. Lombardo and A. Pizzo. Multimedia tool suite for the visualization of drama heritage metadata. *Multimedia Tools and Applications*, pages 1–32, 2014.
13. S. C. Marsella, J. Gratch, and P. Petta. Computational models of emotion. In *A blueprint for an affectively competent agent*. Oxford University Press, Oxford, 2010.
14. E. Norling and L. Sonenberg. Creating Interactive Characters with BDI Agents. In *Proc. of the Australian Workshop on Interactive Entertainment IE2004*, 2004.
15. A. Ortony, G. Clore, and A. Collins. *The Cognitive Structure of Emotions*. Cambrigde University Press, 1988.
16. F. Peinado, M. Cavazza, and D. Pizzi. Revisiting Character-based Affective Storytelling under a Narrative BDI Framework. In *Proc. of ICIDIS08*, Erfurt, Germany, 2008.
17. W. S.Scott Reilly. Believable social and emotional agents. Technical report, DTIC Document, 1996.
18. E. Rishes, S. M. Lukin, D. K. Elson, and M. A. Walker. Generating different story tellings from semantic representations of narrative. In *Interactive Storytelling*, pages 192–204. Springer, 2013.
19. G. P. Zarri. Conceptual and content-based annotation of (multimedia) documents. *Multimedia Tools and Applications*, 72(3):2359–2391, 2014.