

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Automatic Annotation of Characters' Emotions in Stories

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1537906> since 2016-11-30T10:34:51Z

Publisher:

Springer International Publishing

Published version:

DOI:10.1007/978-3-319-27036-4_11

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Automatic annotation of characters' emotions in stories

Vincenzo Lombardo¹, Cristina Battaglino¹, Rossana Damiano¹, and Antonio Pizzo²

¹ Department of Computer Science and Cirma, University of Torino**

² Department of Humanities and Cirma, University of Torino

Abstract. The emotional states of the characters allow the audience to understand their motivations and feel empathy for their reactions to the story incidents. Consequently, the annotation of characters' emotions in narratives is highly relevant for story indexing and retrieval but also editing and analysis. In this paper, we address the construction of tools for the annotation of characters' emotions in stories, opening the way to the construction of a corpus of narratives annotated with emotions.

Keywords: emotion annotation · narrative corpora · emotion appraisal

1 Introduction

Despite the strong interest for affect in social media, tools for the annotation of emotions of narrative resources still lack. Stories, a pervasive feature of today's media - with narrative contents generated daily by amateurs and industry and made available on several platforms and devices - are a domain where affect is recognizably relevant, being intrinsic to the very notion of story (see for example [1]). In this paper, we address the annotation of characters' emotions in stories, with the goal of providing a resource that can support the development of emotional modules in story generators as well as in interactive narrative systems. With the paradigm of crowdsourcing annotation in mind, we propose a workflow for the annotation of characters' emotions that enforces a coherent emotion model in annotation through the use of rules.

In order to devise a set of annotation tags, we resort to cognitive theories of emotions. The core of such cognitive theories is the notion of appraisal [17]: emotions stem from how a character appraises a given situation with respect to its personal perspective. Basically, if the character appraises some event as beneficial, she/he feels positive emotions, such as happiness; if she/he appraises some event as deleterious, she/he is worried or disgusted.

The notion of emotional appraisal relies on the detection, and possibly annotation, of story incidents and how characters are engaged in those incidents, in terms of motivations. This requirement has brought to the integration of the cognitive theories of emotions into the computational systems that implement the

** The affiliations of the authors refer to the preparation of this work.

characters through the well known BDI (Belief-Desire-Intention) model of intentional agency [4, 5]. Virtual characters implemented through this model 1) rely on their knowledge (or Beliefs) about the world to achieve their goals (Desires) through plans of actions (Intentions), and 2) appraise the emotions triggered by the primitives of the BDI model.

The approach to the annotation of emotion pursued in this paper is semi-automatic: the annotators introduce schematic descriptions of the story incidents and of the goals of the characters that motivate the incidents, and a set of rules compute the characters' emotions, actually the emotion tags that describe the way characters react to the story incidents. The purpose of this rule-based approach is to enforce coherence in the implementation of the appraisal model and limit the intervention of the human annotation to the identification of the story incidents and characters' goals. Following the OCC theory of emotion appraisal [19], the rule set assumes that characters are driven by the achievement/failure of their goals and the respect/violation of their values, engaging in conflicts that are the input to their emotional states.

This paper is structured as follows. After reviewing the related work in story annotation in Section 2, we describe how the reference model of emotional appraisal has been translated in terms of a story representation formalism (Section 3) and discuss its use in story annotation. Section 4 contains the description of the rules for annotating the characters' emotions, illustrated through well known examples. Conclusion ends the paper.

2 The annotation of stories

The first attempts at describing the story content in a systematic way date back to the first half of the last century, prompted by availability of the large collections of narratives gathered by folklore studies. These attempts were mainly aimed at classifying stories into categories (see, e.g., Thompson's "motif index" of folk literature [27]) and distilling the basic structures that emerge from corpora (see, e.g., Propp's detailed account of the structure of Russian fairy tales [22]). However, such models were only partially encoded with the use of formal languages, a limitation they share with some notable today's resources for narrative content, such as, e.g., TvTropes³.

Some 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, thanks to the formal nature of linguistic structuralism. The adoption of the Proppian model has led to the creation of repositories for tasks that range from the creation of fictional story worlds [9] to narrative generation [11]. On the linguistic side, the DramaBank Project [8] is a repository of semantically annotated narratives, oriented to the surface generation of different stories from shared nuclei [25]. Being concerned more with the encoding of plots rather than characters, the DramaBank annotation language has specific

³ tvtropes.org

operators for causality and intentionality in stories, such as *Attempt to cause*. All these projects, however, do not account for the annotation of characters' emotions.

The Narrative Knowledge Representation Language (NKRL) proposed by [28] also provides highly developed tools for the annotation of the narrative content that account for the linguistic expression of stories. Mostly focused on the relation between the semantics of stories and their textual presentation, it opens to the application of sentiment analysis to the narrative text [29], a related topic of the affect annotation, but it does not acknowledge the role of the characters and the annotation of the emotions they feel.

The computational account (and, consequently, the annotation) of characters' emotions strongly relies on an explicit account of the characters' mental states in terms of what they believe and intend at any point of the story development. These character-centered description issues, not encompassed by the annotation languages above, are accounted for by the ontology-based story annotation in [14] and [15]. The ontology of drama called Drammar grounds the representation of characters upon the notion of agents' intention (realized through the notion of plan). 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, longer term plans consisting of lower-level, shorter term plans (called subplans). Goals originate from the values of the agents, i.e., put at stake or balanced through the plan actions. The representation of dramatic characters is formalized through the rational agent paradigm, or Belief Desire Intention (BDI) paradigm [4] (which has already seen some applications in the computational storytelling community [18] [20]). The dramatic scenes of the story 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 dramatic qualities *par excellence*. A relevant feature of this ontology is that it was designed with the goal of being independent from specific genres of media, making it suitable for the annotation of heterogeneous narrative resources.

3 Rule-based annotation of emotions

Given the representation of the story in terms of characters' goals, values and action plans, the emotion annotation rules add to the characters a set of emotional labels. The rules we propose are based on the cognitive theory of emotions proposed by Ortony, Clore and Collins in 1988 [19] (OCC). A number of computational models of emotions, since the pioneering work by [7], rely on this theory [24, 10, 6] to model the emotional states of virtual agents, or on appraisal theories in general [23, 16].

3.1 Background: the OCC model of emotions

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), mapped 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), mapped onto the compliance (or conflict) with moral values (e.g., Pride or Shame), following the computational model of moral emotions in [3]; the *affection* for an entity involved in the situation. Notice that the utilitarian dimension can be also appraised by the agent from the point of view of another agent, thus generating emotions oriented to other agents (e.g. Happy-for 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): the affection towards the target cannot be computed, being intrinsic to the appraising agent. Notice that the same situation can be simultaneously appraised as an event, an action or an entity: so, for example, another agent's action may be appraised as an intentional act, as a mere event, or both, giving rise to different emotion types.

Finally, when a situation is appraised as event or act, its temporal dynamics becomes relevant: if the appraised situation is still ongoing, a *prospect-based* emotion will be generated 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 they 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.

3.2 Mapping the OCC model onto the rules

In order to be compliant with the Semantic Web and the Linked Open Data initiatives, we have adopted the Semantic Web Rule Language (*SWRL*) to encode

the emotional appraisal process. This choice permits the possible use of the emotion appraisal process as a web service in a future of distributed resources in the web.

The SWRL rule language[12] augments an OWL ontology with a rule layer, adding the possibility to declare arbitrary Horn clauses expressed as *IF THEN* rules. A *SWRL* based system is therefore composed of ordinary *OWL* axioms plus a set of *SWRL* rules. The antecedent and consequent of the rules consist of lists of atoms, which may be *OWL* class expressions, property definitions, or built-ins. Most of the current available *DL* reasoners, such as *Pellet* or *Hermit*, support inferences based on *SWRL*. Encoding emotion generation using *SWRL* rules enables the automatic generation of the emotions of the characters in a scene that has been previously annotated in the Drammar ontology. However, notice that the rules may be easily translated into a different rule formalism in order to apply them to an annotation carried out with an equivalent formalism.

Translating the computational model of emotions into the rules for the generation of emotion labels involves a mapping of the elements of the appraisal process (appraising agent, situation and dimension of appraisal) onto the primitives of the Drammar ontology. In order to make the OCC theory compliant with the BDI agent approach incorporated in Drammar, we have adopted the operationalization of the theory provided by [2], which relies on the introduction of the notion of value for the appraisal of moral emotions. Basically, the rule antecedent represents a character's appraisal of a situation, and is based on the storyworld states determined by the character's goals achievement (e.g., killed an enemy of the country), the character's values put at stake or balanced (e.g., the freedom of the country), and the plans the character is committed to (e.g., ambushing the enemy during a street parade). The rule consequent asserts what emotions the character feels as a consequence of the appraisal and what is the target of the emotion. When the antecedent of the rule is true in the annotated story, the rule fires and, as a result, an emotion of the type prescribed by the rule is added to some character. For example, if a character's goal is not achieved in a given scene due to the occurrence of an (unpredictable) event, the rule derives that what happened in the scene is appraised by the character as undesirable and thus the consequent generates a Distress emotion for the character. Notice that, following the *componential view* of a computational model of emotions presented in [17], the *appraisal derivation* is encoded in the antecedent of the rules, while the *affect derivation* is encoded in the rule consequent.

The appraised situation is mapped onto a scene of the drama and the appraising agent is mapped onto a character featured in the scene. 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. The support and conflict relations over goals determine the desirability/undesirability of the plans executed by the agents in the scene, seen as ongoing or accomplished processes; the way the plans affect the agents' values (putting them at stake or bringing them back to balance) determines the praiseworthiness/blameworthiness of the agent who executes them

(self or other); finally, the affection toward other agents determine the emotions oriented to others.

3.3 Toward a corpus of emotionally annotated stories

The rules for assigning emotions to characters are deployed in a semi-automatic story annotation pipeline. The advantage of using a formally encoded set of rules for emotion assignment is that they allow the human annotator to check her/his assessment of the characters' emotional state against the emotion types returned by the rules. Here, we sketch a workflow that uses the rules set to construct a corpus of narratives annotated with characters' emotions, with the goal of fostering the creation of narrative corpora annotated with with characters' emotions.

1. **Corpus selection.** A corpus of narrative works is selected by experts in drama and narrative theory, with the goal of providing a comprehensive corpus spanning through ages, genres and media.
2. **Segmentation.** Works are segmented by an expert before being fed to the parallel human and semi-automated annotation procedures, so as to provide a basic alignment. Notice that the segmentation is often explicitly marked in most media, from classic drama (with explicit acts, scenes, etc.) to Hollywood movies (with acts, sequence, scenes, etc.). However, in case the segmentation is not marked in the original work, the segmentation made by the annotators can still be considered reliable, as argued by [13] based on an empirical basis.
3. **Human annotation.** Each story is annotated by at least two different annotators, trained in dramatic narration and selected based on their familiarity with the work. The annotation follows the schema implied by the appraisal model, namely the set of incidents occurring in each story scene, together with the links to the characters' plans, goals and values and the conflict/support relations over them.
4. **Rule based computation of emotions.** A reliable version of the annotation is generated by accounting for the agreement between the annotators: this step requires the annotated elements (plans, goals and values) to be consistent not only from a logical perspective, but also from a narratological perspective, so the supervision of an expert may be required. The annotation is then fed to a reasoner⁴ for the application of the *SWRL* emotion rules presented in Section 3.

Notice that the annotators may provide useful insights on the correctness of the rules, leading to minor changes and/or refinements of the rules.

4 Rules for the automatic annotation of emotions

The rules for the automatic annotation of emotions can be classified according to the relations that holds over the primitives of the agent model.

⁴ Such as Pellet, www.clarkparsia.com/pellet

The appraisal of an event as *desirable* (or *undesirable*) depends on the relation between a goal of the appraising agent and another agent’s goal, achieved by the plan of the other agent in the scene. This relation is expressed through the properties `inConflictWith` or `inSupportOf`, respectively: an event is desirable if the goal it achieves is `inSupportOf` of the agent’s goal, undesirable otherwise.

The appraisal of an action as *praiseworthy* (or *blameworthy*) depends on the relation between a character’s value and a plan another agent is committed to (or by the agent itself) as a way to achieve some goal. This relation 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 (in Drammar, this equates to saying that the value is a `ValueEngaged` in the effects of a plan), the plan is (or else, the actions contained in it are) 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 execution of a plan 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. Notice that this is in line with the observations about prospect-based emotions made in [26].

Figure 1 illustrates the rules for emotion generation. The table is divided into three main sections, that correspond, respectively, to Well-being (and Prospect-based) emotions, Fortune-of-others emotions, and Attribution emotions, respectively. Compound emotions complete the table.

Well-being emotions, such as Distress and Joy, depend on the relation between a `Goal ?G` and a `Goal ?GSA` owned by an `Agent`. An event is desirable if it encompasses a plan that achieves a goal `?G inSupportOf` of the agent’s goal `?GSA`, undesirable if the goal `?G` is `inConflictWith` the agent’s goal (notice that `?G` and `?GSA` can be the same goal).

Fortune-of-others emotions, such as Happy-for another agent, depend 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 ?GOA` (see Fig. 1). For example, if the `Agent ?SA` loves another `Agent ?OA` and the `Goal ?G` is `inSupportOf` the `Goal ?GOA` 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 ?P`. 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) (see Fig. 1).

Event	A goal g_i is achieved by a Plan p_i in the scene	
	Appraisal Variables	Appraising Agent feels
Well-being and Prospect-based emotions		
Undesirable		
Goal g_i is in conflict with appraising agent's goal	Prospect of the event	
	No prospect	
	The appraising agent doesn't have a belief about the relation between the goal g_i and its goals	DISTRESS
	Uncertain	
	The appraising agent believes the relation between the goal g_i and one of its goals; the plan p_i is not still accomplished	FEAR
	Confirmed	
The appraising agent believes the relation between the goal g_i and one of its goals; the plan p_i is accomplished	FEAR-CONFIRMED	
Disconfirmed		
The appraising agent believes the relation between the goal g_i and one of its goals; the plan p_i is not accomplished	RELIEF	
Prospect of the event		
Goal g_i is in support of appraising agent's goal	No prospect	
	The appraising agent doesn't have a belief about the prospect of the event and the Plan p_i is accomplished	JOY
	Uncertain	
	The appraising agent believes the relation between the goal g_i and one of its goals; the plan is not still accomplished	HOPE
	Confirmed	
	The appraising agent believes the relation between the goal g_i and one of its goal; the plan is accomplished	SATISFACTION
Disconfirmed		
The appraising agent believes the relation between the goal g_i and one of its goals; the plan is not accomplished	DISAPPOINTMENT	
Fortune-of-others Emotions		
Undesirable		
Goal g_i is in conflict with another agent's goal	Relations	
	Love	
	The appraising agent loves the other agent	PITY
	Hate	
The appraising agent hates the other agent	GLOATING	
Desirable		
Goal g_i is in support of another agent's goal	Relations	
	Love	
	The appraising agent loves the other agent	HAPPY-FOR
	Hate	
The appraising agent hates the other agent	RESENTMENT	
Attribution Emotions		
Praiseworthy		
The plan p_i rebalances one of the values of the appraising agent	Agency	
	Self	
	The appraising agent intends the plan p_i	PRIDE
	Other	
Another agent intends the plan p_i	SHAME	
Blameworthy		
The plan p_i puts at stake one of the values of the appraising agent	Agency	
	Self	
	The appraising agent intends the plan p_i	ADMIRATION
	Other	
Another agent intends the plan p_i	REPROACH	
Compound Emotions		
Desirable and Praiseworthy		
Goal g_i is in support of another agent's goal and the plan p_i rebalances one of the values of the appraising agent	Agency	
	Self	
	The appraising agent intends the plan p_i	GRATIFICATION
	Other	
Another agent intends the plan p_i	GRATITUDE	
Undesirable and Blameworthy		
Goal g_i is in conflict with another agent's goal and the plan p_i puts at stake one of the values of the appraising agent	Agency	
	Self	
	The appraising agent intends the plan p_i	REMORSE
	Other	
Another agent intends the plan p_i	ANGER	

Fig. 1: SWRL Rules encoded in the Drammar Ontology.

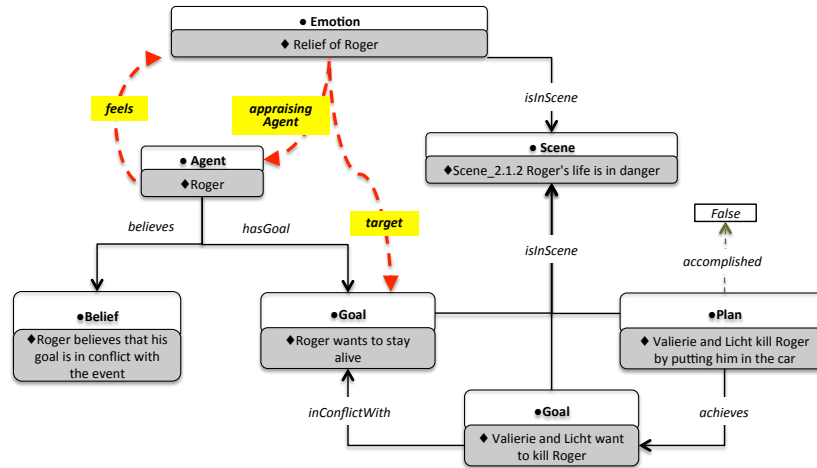


Fig. 2: The annotation of the scene used by the emotion rule module for the Agent “Roger”. The property `target`, `feels` and `appraisingAgent` are inferred by the rule for Relief emotion.

Attribution emotions can be self- or other-directed: if the **Plan** $?P$ that the **Agent** $?SA$ considers praiseworthy (or blameworthy) is a plan intended by the agent $?SA$ itself, the **Agent** $?SA$ feels Pride (or Shame). Otherwise, if the **Plan** $?P$ plan is intended by another **Agent** $?OA$ in the scene, $?SA$ feels Admiration or Reproach (see Fig. 1).

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

4.1 Examples

In order to illustrate how the rules work on a story scene annotated with the narrative and dramatic elements, we resort to two examples taken from the movie “North by Northwest” by Alfred Hitchcock (1959), a tale of mistaken identity where the main character Roger tries to prove that he is not the ‘double’ George Kaplan. As a Hollywood classic, the characters’ emotions in this movie can be identified unambiguously enough to provide straightforward examples for emotion assignment [21].

Fig. 2 describes the activation of the SWRL rule for Relief of the agent Roger in 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 putting him into a moving car. Roger eventually manages to exit from the car before it falls off a cliff. We only report the salient elements that are needed to illustrate

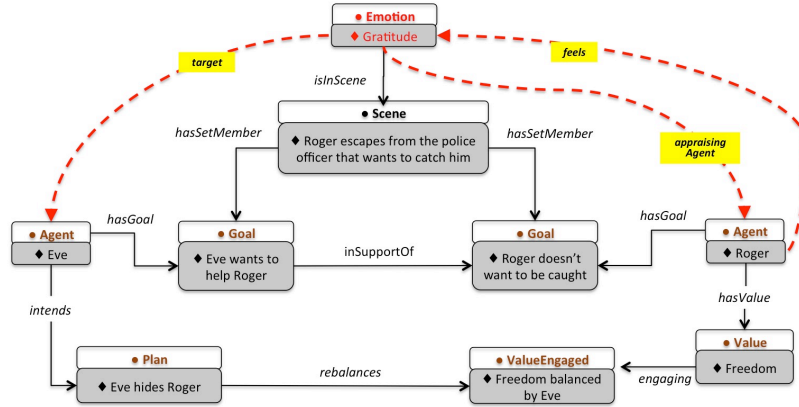


Fig. 3: The representation of the example: Eve helps Roger to hide from the police officers. The dotted lines indicate the emotion felt by Roger.

the activation of the Relief SWRL rule. 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* by focusing on the conflict between the two goals: Valerie and Licht want to kill Roger; Roger wants to stay alive. The application of the SWRL rules correctly outputs Roger’s Relief as the emotion triggered in the scene. In Fig. 2, the appraised 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 property `accomplished` set to `false`, meaning that that the event is *disconfirmed*. 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 (see Fig. 1). 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`).

Fig. 3 shows the annotation of 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. As a spy of the USA government, Eve knows that Roger is not an assassin, so she helps him. The incidents described above are contained in the **Scene** “*Roger escapes from the police officer that wants to catch him*”. The occurrence of the event (Roger’s escape) featured in the scene is motivated by the following goals: Roger has the **Goal** of not being caught by the police officers (*Roger doesn’t want to be caught*); Roger’s goal is supported by Eve’s **Goal** of helping Roger (*Eve wants to help Roger*). Roger **feels** an **Emotion** of Gratitude toward Eve, because her goal is `inSupportOf` of Roger’s goal of not being caught

and her plan brings back to balance Roger's Value of *Freedom* (rebalances arrow from Eve's Plan *Eve hides Roger* to Roger's value). The Gratitude rule asserts that the Agent *Roger* is the appraisingAgent who feels the Emotion *Gratitude*, with the Agent *Eve* as target (property target).

5 Conclusion

Since emotions are relevant for the comprehension and the fruition of stories by their public, they can be useful to implement innovative tools for search and editing of stories. Differently from sentiment detection in linguistic corpora, which mainly relies on lexical and semantic resources, the annotation of characters' emotions requires an understanding of how the characters' conflicting motivations in a story determine the generation of their emotions.

In this paper, we have described a system for the automatic generation of characters' emotions in stories, with the appraisal process 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. Based on the representation of story characters as BDI agents augmented with moral values provided by the Drammar ontology, the rules encode the OCC model of emotions, integrated with the notion of values into the BDI model.

In this paper, we have illustrated our approach through examples taken from a classic Hollywood movie. As a future work, we plan to use the rule set for a narrative annotation campaign, aimed at adding emotional labels to a comprehensive set of narrative works, spanning through media, ages and genres.

References

1. Quesenberry Keith A and Coolson Michael K. What makes a super bowl ad super? five-act dramatic form affects consumer super bowl advertising ratings. *Journal of Marketing Theory and Practice*, 22(4):437–454, 2014.
2. Cristina Battaglino and Rossana Damiano. Emotional appraisal of moral dilemma in characters. In *ICIDS*, pages 150–161, 2012.
3. Cristina Battaglino, Rossana Damiano, and Leonardo Lesmo. Emotional range in value-sensitive deliberation. In *AAMAS*, pages 769–776, 2013.
4. M.E. Bratman. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge (MA), 1987.
5. P. R. Cohen and H. J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213–261, 1990.
6. 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.
7. Clark Davidson Elliott. *The Affective Reasoner: A Process Model of Emotions in a Multi-agent System*. PhD thesis, Northwestern University, Evanston, IL, USA, 1992. UMI Order No. GAX92-29901.
8. David Elson. Dramabank: Annotating agency in narrative discourse. In *LREC*, pages 2813–2819, 2012.

9. Chris R Fairclough and Pádraig Cunningham. A multiplayer opiate. *International Journal of Intelligent Games & Simulation*, 3(2), 2004.
10. Patrick Gebhard. Alma: a layered model of affect. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 29–36. ACM, 2005.
11. Pablo Gervás. Propp’s morphology of the folk tale as a grammar for generation. In *CMN*, pages 106–122, 2013.
12. Ian Horrocks, Peter F Patel-Schneider, Harold Boley, Said Tabet, Benjamin Grosf, Mike Dean, et al. Swrl: A semantic web rule language combining owl and ruleml. *W3C Member submission*, 21:79, 2004.
13. Vincenzo Lombardo and Rossana Damiano. Commonsense knowledge for the collection of ground truth data on semantic descriptors. In *Proceedings of the 2012 IEEE International Symposium on Multimedia (ISM 2012)*, pages 78–83. IEEE Computer Society, 2012.
14. Vincenzo Lombardo and Antonio Pizzo. Multimedia tool suite for the visualization of drama heritage metadata. *Multimedia Tools and Applications*, pages 1–32, 2014.
15. Vincenzo Lombardo and Antonio Pizzo. Ontology–based visualization of characters intentions. In *Interactive Storytelling*, pages 176–187. Springer, 2014.
16. Stacy C Marsella and Jonathan Gratch. Ema: A process model of appraisal dynamics. *Cognitive Systems Research*, 10(1):70–90, 2009.
17. Stacy C. Marsella, Jonathan Gratch, and Paola Petta. Computational models of emotion. In K. R. Scherer, T. Bänziger, and Roesch, editors, *A blueprint for an affectively competent agent: Cross-fertilization between Emotion Psychology, Affective Neuroscience, and Affective Computing*. Oxford University Press, Oxford, 2010.
18. E. Norling and L. Sonenberg. Creating Interactive Characters with BDI Agents. In *Proceedings of the Australian Workshop on Interactive Entertainment IE2004*, 2004.
19. A. Ortony, G. Clore, and A. Collins. *The Cognitive Structure of Emotions*. Cambridge University Press, 1988.
20. F. Peinado, M. Cavazza, and D. Pizzi. Revisiting Character-based Affective Storytelling under a Narrative BDI Framework. In *Proc. of ICIDIS08*, Erfurt, Germany, 2008.
21. Carl Plantinga. *Moving viewers: American film and the spectator’s experience*. Univ of California Press, 2009.
22. V. Propp. *Morphology of the Folktale*. University of Texas Press, 1968.
23. Stefan Rank and Paolo Petta. Appraisal for a character-based story-world. In *IVA*, pages 495–496, 2005.
24. W Scott Reilly. Believable social and emotional agents. Technical report, DTIC Document, 1996.
25. Elena Rishes, Stephanie M Lukin, David K Elson, and Marilyn A Walker. Generating different story tellings from semantic representations of narrative. In *Interactive Storytelling*, pages 192–204. Springer, 2013.
26. Bas R Steunebrink, Mehdi Dastani, and John-Jules Ch Meyer. The occ model revisited. In *Proc. of the 4th Workshop on Emotion and Computing*, 2009.
27. Stith Thompson. Myths and folktales. *The Journal of American Folklore*, 68(270):482–488, 1955.
28. Gian Piero Zarri. Conceptual and content-based annotation of (multimedia) documents. *Multimedia Tools and Applications*, 72(3):2359–2391, 2014.
29. Gian Piero Zarri. Sentiments analysis at conceptual level making use of the narrative knowledge representation language. *Neural Networks*, 58:82–97, 2014.