A Formalization of Emotional Planning for Strong-Story Systems

Alireza Shirvani, Stephen G. Ware
Narrative Intelligence Lab, University of Kentucky
Lexington, KY 40506
{ashirvani, sgware}@uky.edu

Abstract

Characters that are capable of expressing emotions seem more lifelike to a human audience. Interacting with emotional characters may cause a human user to feel empathy and to care for their fate. In turn, the audience may feel more engaged in the story. We take a step to address the lack of computational models of emotions for strong-story systems and propose to extend strong-story state-space narrative planners, already equipped with intentionality and belief, to reason about character emotions. Our proposed system is multi-agent, highly domain-independent, and focused on reasoning and decision making. In this paper, we evaluate to what extent our system can accurately model a set of emotions and in doing so, improve the believability of story characters.

Introduction

Reasoning about the emotions of virtual agents helps in designing them and reasoning about their behavior (Meyer 2006). Emotions offer guidance about the possible consequences of actions and, in turn, motivate other actions in order to face the resulting emotions (Alfonso Espinosa, Vivancos Rubio, and Botti Navarro 2014). In this paper, we focus on the incorporation of emotions into narrative planning.

Narrative planners can be considered as strong-story systems (Riedl and Bulitko 2013) because they prioritize author requirements while also finding explanations for character actions. In contrast, strong-autonomy systems focus on agent simulations within the constraints defined by the author in which agents are unaware of any overarching narrative. Simulation games or exploratory learning environments benefit from strong-autonomy systems, where it is not always necessary to guide the user's narrative experience.

Strong-story systems have a centralized coordinator, often referred to as the experience manager (Riedl and Bulitko 2013), that coordinates all actions of virtual characters to prevent deviating from the desired narrative. This makes them useful for educational and training applications with specific instructions and goals, which may require reasoning about every user interaction with virtual characters.

In particular, state-space narrative planners explore and prune the space of all possible character actions to only those that can be explained (Ware et al. 2019). This gives them the freedom to choose what character actions to include to generate a story and thus enables them to "achieve the highest degree of leverage over the virtual world to bring about the desired narrative experience for the user" (Riedl and Bulitko 2013).

Incorporating emotions into existing narrative planners has the following potential benefits. First, it allows the planner to infer each character's emotions at every state, e.g. a character hopes to achieve a certain goal and is disappointed when they realize they cannot. The planner can then choose to communicate those emotions to the reader via text or animations. Knowing what emotions are triggered for each character after each action, we can also better explain why those characters chose to take those actions, specifically how characters always take action to experience positive emotions. Moreover, reasoning about such character qualities will allow them to generate a larger set of possible stories that could not be created by previous narrative planners. For instance, a character may feel fear for one of their goals and proactively form plans to eliminate that fear.

Our model of emotion is inspired by Gratch's emotional planning (1999). According to the appraisal theory, emotions are appraised based on events and goals. For instance, if an events achieves an agent's goal, they feel joy. An important benefit of emotional planning is that it can determine the significance of events to goals based on domain-independent features of plans. In narrative planning, a plan is only valid for an agent if it achieves their goal, which, in terms of emotions, triggers their joy. Our goal is to expand that definition to also use appraisals of other emotions. More specifically, we say that a plan is valid if it triggers any positive emotion, e.g. a plan that eliminates an agent's fear of their goal being threatened, even if it does not achieve any of their goals.

This paper builds on our previous state-space narrative planning framework (Shirvani, Ware, and Farrell 2017) to define which stories make sense based on character emotions. We describe and evaluate the knowledge representation, not the search process. We strive for high domain independence and focus on agent reasoning and decision making, rather than fine-tuning physiological manifestations of emotions.

Related Work

Emotions have been mostly incorporated into strong-autonomy systems rather than strong-story systems. EMA (Marsella and Gratch 2009) is an appraisal-based computational model of emotion that describes the agent-environment relationship in terms of appraisal variables, and employs those variables to produce emotional responses. The emotional responses may then trigger coping mechanisms to cope with the situation. FLAME (El-Nasr, Yen, and Ioerger 2000) presents a computational model of emotion using fuzzy logic to map events and observations to emotional states, as well as several inductive learning algorithms to learn patterns of events, associations among objects, and expectations. ALMA (Gebhard 2005) considers a hierarchy of affect types with emotions as short-term, mood as mid-term, and personality as long-term affects, and utilizes the relationships between these affect types as defined by Mehrabian (1996). ALMA's AffectML, its XML-based affect modeling language, can then be used to describe appraisal rules and personality profiles for characters in order to compute real-time emotions and moods.

Cavazza et al. (2009) consider emotion as an essential part of the story domain. Focusing on speech and dialog, they implement emotional planning by including character feelings in domain actions (Pizzi, Cavazza, and Lugrin 2007; Cavazza et al. 2007). Pérez-Pinillos, Fernández, and Borrajo (2011) define a PDDL domain model that integrates agent drives, emotions, and personality. Agent actions can change their emotional state based on hand-authored values, and, in turn, the effects of an action can be different for an agent depending on its emotional state.

Many strong-autonomy systems are single-agent simulations of virtual characters that simulate agents in a specific scenario, which makes it difficult to adapt them to different scenarios (Kenny et al. 2007). It is possible for strong-autonomy systems to simultaneously run multiple agent simulations to replicate the multi-agent nature of strong-story systems. However, these systems are limited to the actions those agents decide to take, which may be arbitrary when multiple actions are reasonable. The centralized reasoning in strong story systems considers all reasonable actions to find one that best meets the author's intentions. On the other hand, many planning-based strong-story systems are either designed for specific story domains or introduce a domain model that increases author burden.

Narrative Planning with Intentionality and Belief

Example Story Domain

We use the following example domain throughout this paper. *Tom* is sick and needs medicine. He has two coins and his utility function is his number of coins plus two times his number of medicines. He could either buy it in town from a *Merchant* or make it using herbs in the forest. However, contrary to his beliefs, there are no herbs in the forest.

He also knows there is a *Bandit* in the forest that could steal his coins. Tom can buy a sword from the merchant that prevents the bandit from robbing him. The utility function of both the bandit and the merchant is the number of coins they have. Figure 1 presents an example of an expanded search space for this story domain.

Classical and Narrative Planning

Given some initial world state, a planning algorithm searches for a sequence of actions to achieve a goal (Russell and Norvig 2009). Actions have preconditions which must be true in the state immediately before they occur and effects which change the state. A state space is a graph whose nodes are states and whose directed edges represent actions. An edge $s \xrightarrow{a} s'$ exists if action a's precondition is satisfied in state s and applying a's effects would change the state to s'. A plan is a sequence of actions, or a path through this graph.

Narrative planners (Porteous, Cavazza, and Charles (2010), Young et al. (2013), and many others) extend this formalism to tell believable stories. They reason about the author's goal, as well as the beliefs and goals of each character (i.e. agent). Rather than expressing goals as propositions, we use utility functions:

Definition. A *utility function* $s \rightarrow \mathbb{R}$ maps a state s to a real number, with higher preferable to lower.

Definition. A *narrative planning problem* is $\langle s_0, U, A, C, U_c \rangle$,where s_0 is the initial state, U is the author's utility function, A is a set of actions that can occur, C is a set of agents, and U_C is a set of utility functions, one for each agent. $U(c, s)$ expresses the utility of agent c in state s.

Actions $a \in A$ have preconditions and effects, as in classical planning. Similar to Riedl and Young's model of intention (2010), actions also define which agents take the action.

Definition. For each action $a \in A$, let $Con(a)$ denote the set of *consenting characters* who must have a reason to take the action.

Actions also change the beliefs of characters. The mechanics of how beliefs change are defined by Shirvani, Ware, and Farrell (2017); we present only the definitions needed to define our model of emotion.

Definition. For some sequence of actions π in a state s, let $\alpha(\pi, s)$ denote the state *after* taking the actions in π. In Figure 1, $\alpha(\{Go(Tom, Forest)\}, s_1) = s_2$. α is only defined when, for every action $a \in \pi$, the precondition of a is satisfied in the state immediately before taking a.

Definition. For some character c in a state s, let $\beta(c, s)$ denote what c *believes* the state to be when it is actually s. In Figure 1, $\beta(Tom, s_1) = s_{10}$. In s_1 there are no herbs in the forest, but in s_{10} there are, so Tom wrongly believes there are herbs in the forest.

Various narrative planning frameworks differ in how they define explained actions.

Definition. In state s, an action a is *explained for* character $c \in Con(a)$ when there exists a sequence of actions π such that a is the first action in π and:

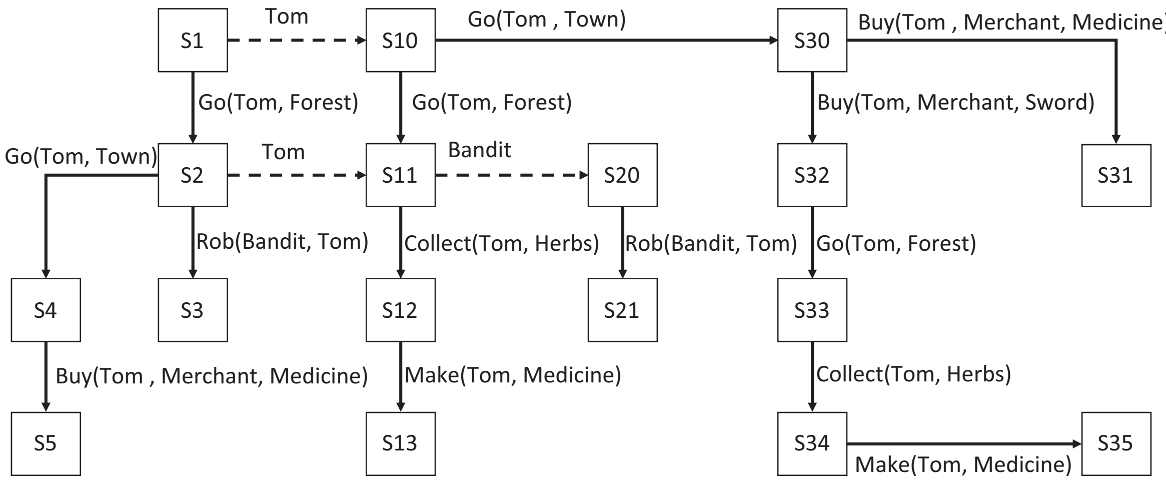

Figure 1: A part of the state space of the example story. The dashed edges represent epistemic edges.

1. $U(c, \alpha(\pi, \beta(c, s))) > U(c, s)$
2. Every action after a in π is explained.
3. π does not contain a strict subsequence that also meets these three requirements.

In other words, an action makes sense for a character when that character can imagine a plan starting with that action that (1) they believe will lead to a higher utility, (2) the plan makes sense for the other consenting characters, and (3) it doesn't contain unnecessary or redundant actions.

Definition. In state s, an action a is *explained* when, for all consenting characters $c \in Con(a)$, a is explained for c in s. In other words, an action is explained when it is explained for all the characters that need a reason to take it. Characters can have different reasons. Tom can buy the medicine because he wants it, and the merchant will sell it because she wants coins. The merchant has no reason to give away the medicine, so Tom cannot expect her to.

Definition. In state s, a sequence of actions π is *explained* when, for all actions $a \in \pi$, a is explained in the state before a occurs. In other words, a sequence is explained when all its actions are explained.

Definition. A *solution* to a narrative planning problem is an (1) explained sequence of actions that (2) increases the author's utility and (3) does not contain a strict subsequence that also meets these three requirements.

Note that one character can expect another character to act. We call this *anticipation* (Shirvani, Farrell, and Ware 2018), but a character can only anticipate actions that help them increase their utility. They should also be able to anticipate harmful sequences.

Definition. A sequence of actions π is *expected* for character c in state s when π is an explained sequence of actions in s.

This definition is similar to that of an explained sequence of actions, but it need not lead to a higher utility. An explained plan is one that a character would want to take because it helps achieve their goals, but characters can still expect plans that are bad for them, e.g. Tom can expect to be robbed, even though he does not want to be robbed. In keeping with the ideals of a strong story system, characters can expect many sequences, not just one. Characters do not commit to a single expectation (what a BDI system might call an *intention*), but can expect any sequence that meets these requirements. This enables the planner to choose from a wide variety of believable stories when trying to meet the author's requirements.

In the next section, we will expand these definitions to incorporate emotions. We will show how emotions are triggered as a consequence of actions and how characters can also expect actions to trigger certain emotions for them.

Narrative Planning with Emotions

The OCC Model of Emotion

The OCC model of emotion (Ortony, Clore, and Collins 1990) defines 22 different emotions which are triggered based on the consequences of events, actions of agents, and aspects of objects. In this paper, we will focus on the 12 emotions that are related to consequences of events. These 12 emotions are referred to as well-being emotions and are presented in Table 1. Table 1 also presents how each emotions is triggered based on an agent's utility value.

We also define a very simple character relationship model as follows. The relationship between every pair of characters is defined as a real number in [-1, 1]. Notation $R(c, c')$ represents the relationship value of characters c and c'. Two characters, c and c', are considered *friends* if $R(c, c') > 0$, and otherwise *non-friends*.

Positive and Negative Emotions

This section defines how emotions are triggered based on character utility functions. Table 1 also presents a concise description of these definitions.

Table 1: Well-being emotions and their corresponding appraisal variables

Emotion	Appraisal	Relationship to goals
Joy	The occurrence of a desirable event	Utility increases
Distress	The occurrence of an undesirable event	Utility decreases
Hope	An unconfirmed desirable event	Expects a higher utility value
Fear	An unconfirmed undesirable event	Expects a lower utility value
Satisfaction	A confirmed desirable event	Achieves the expected higher utility value
FearsConfirmed	A confirmed undesirable event	Achieves the expected lower utility value
Disappointment	A disconfirmed desirable event	No longer expects the higher utility value
Relief	A disconfirmed undesirable event	No longer expects the lower utility value
HappyFor	A desirable event that is desirable for another	Consents to an action or utility increases while increasing a friend's utility
Resentment	An undesirable event that is desirable for another	Utility decreases by an action that a friend consented to or increases a friend's utility
Gloating	A desirable event that is undesirable for another	Consents to an action or utility increases while decreasing a non-friend's utility
Pity	An undesirable event that is undesirable for another	Utility of self and a non-friend decreases

Positive Emotions

1. $\boldsymbol{Joy(c, a, s)}$ is triggered for character c at state s after taking/observing action a if $U(c, s) > U(c, s')$, such that $\alpha(a, s') = s$. The intensity of *Joy* is equal to how much c's utility increases after a or $U(c, s) - U(c, s')$.

   In Figure 1, Joy is triggered for Tom in state *s5* because his utility increases to 3.

2. $\boldsymbol{Hope(c, u, s)}$: Character c feels Hope to achieve utility u as long as there is at least one expected plan π starting from state s, such that $u = U(c, \alpha(\pi, \beta(c, s)))$ and $u > U(c, s)$. The intensity of *Hope* is equal to how much c's utility increases when it reaches u or $U(c, s) - u$.

   In state $s1$ of Figure 1, Tom hopes for utility values 4 (by making the medicine himself) or 3 (by buying the medicine).

3. $\boldsymbol{Satisfaction(c, u, s)}$ is triggered for character c at state s if $U(c, s) =$ hoped utility u. The intensity of *Satisfaction* is equal to that of the corresponding *Hope*.

   In Figure 1, Satisfaction could have been triggered for Tom in *s13* if there were herbs in the forest.

   Although Joy is always triggered when Satisfaction is triggered, they are triggered for different reasons and also the opposite is not true, since a character's utility may increase by surprise.

4. $\boldsymbol{Relief(c, u, s)}$ is triggered for character c at state s if c no longer fears utility u—Fear is defined later—and $U(c, s) > u$. The intensity of *Relief* is the reciprocal of the intensity of the corresponding *Fear*.

   In Figure 1, Relief is triggered for Tom at state *s32* because *Tom* buys a sword and no longer can be robbed.

5. $\boldsymbol{HappyFor(c, c', a, s)}$: Character c feels happy for character c' at state s after action a if for c, $c \in Con(a)$ or $U(c, s) > U(c, s')$, for c', $U(c', s) > U(c', s')$, and for c and c', $R(c, c') > 0$, such that $\alpha(a, s') = s$. As a result, $R(c, c')$ increases by a predefined value σ. The intensity of *HappyFor* is equal to how much c''s utility increases or $U(c', s) - U(c', s')$.

   Assuming $R(Tom, Merchant) > 0$, HappyFor is triggered for Tom at state *s5* because after buying the medicine, the merchant's utility is increased by 1.

6. $\boldsymbol{Gloating(c, c', a, s)}$: Character c feels gloating towards character c' at state s after action a if for c, $c \in Con(a)$ or $U(c, s) > U(c, s')$, for c', $U(c', s) < U(c', s')$, and for c and c', $R(c, c') \leqslant 0$, such that $\alpha(a, s') = s$. As a result, $R(c, c')$ decreases by a predefined value σ. The intensity of *Gloating* is equal to how much c''s utility decreases or $U(c', s') - U(c', s)$.

   Assuming $R(Tom, bandit) \leqslant 0$, Gloating is triggered for the bandit at state *s3* because the bandit's utility increases to 2 and Tom's utility decreases to 0.

Negative Emotions

1. $\boldsymbol{Distress(c, a, s)}$ is triggered for character c at state s after taking/observing action a if $U(c, s) < U(c, s')$, such that $\alpha(a, s') = s$. The intensity of *Distress* is equal to how much c's utility decreases after a or $U(c, s) - U(c, s')$.

   In Figure 1, Distress is triggered for Tom in state *s3* because his utility reduces to 0.

2. $\boldsymbol{Fear(c, u, s)}$: Character c feels the Fear of decreasing their utility to u as long as there is at least one expected plan π starting from state s, such that $u = U(c, \alpha(\pi, \beta(c, s)))$ and $U(c, \alpha(\pi, \beta(c, s))) < U(c, s)$. The intensity of *Fear* is equal to how much c's utility decreases when it reaches u or $u - U(c, s)$.

   In Figure 1, Tom fears his utility to decrease to 0 because he expects that the bandit could steals his coins.

3. $\boldsymbol{FearsConfirmed(c, u, s)}$ is triggered for character c at state s if $U(c, s) =$feared utility u. The intensity

of *FearsConfirmed* is equal to that of the corresponding *Fear*.

In Figure 1, FearsConfirmed is triggered at $s3$ when *Tom* is robbed as he feared he would be. Although Distress is always triggered when FearsConfirmed is triggered, they are triggered for different reasons and also the opposite is not true, since a character's utility may decrease by surprise.

4. $\boldsymbol{Disappointment(c, u, s)}$ is triggered for character c at state s if c no longer hopes for utility u and $U(c, s) < u$. The intensity of *Disappointment* is the reciprocal of that of the corresponding *Hope*.

   In Figure 1, Disappointment is triggered for Tom in state $s2$ because he realizes there are no herbs in the forest.

5. $\boldsymbol{Resentment(c, c', a, s)}$: Character c feels resentment for character c' at state s after action a if for c, $U(c, s) < U(c, s')$, for c', $c' \in Con(a)$ or $U(c', s) > U(c', s')$, and for c and c', $R(c, c') > 0$, such that $\alpha(a, s') = s$. As a result, $R(c, c')$ decreases by a predefined value σ. The intensity of *Resentment* is equal to how much c's utility decreases or $U(c, s) - U(c, s')$.

   Assuming $R(Tom, bandit) > 0$, Resentment is triggered for Tom at state $s3$ because the bandit's utility increases to 2 and Tom's utility decreases to 0.

6. $\boldsymbol{Pity(c, c', a, s)}$: Character c feels pity for character c' at state s after action a if $U(c, s) < U(c, s')$, $U(c', s) < U(c', s')$, and $R(c, c') \leqslant 0$, such that $\alpha(a, s') = s$. As a result, $R(c, c')$ increases by a predefined value σ. The intensity of *Pity* is equal to how much c''s utility decreases or $U(c', s) - U(c', s')$.

Emotional Planning

Based on characters' expected emotions, we redefine explained actions as follows.

Definition. In state s, an action a is *explained for* character $c \in Con(a)$ when there exists a sequence of actions π such that a is the first action in π and:

1. A positive emotion is triggered for c in $\alpha(\pi, \beta(c, s))$.
2. Every action after a in π is explained (based on our new definition) for all its consenting characters.
3. π does not contain a strict subsequence that also meets these three requirements.

Based on this definition, characters are no longer limited to acting on plans that increase their utility. Moreover, the best sequence of actions for a character is no longer the one that maximizes that character's utility, but one that maximizes/minimizes the intensity of their positive/negative emotions.

Currently, based on the 12 emotions we defined in this paper, an explained sequence of actions could trigger Joy, Relief, Satifaction, HappyFor, and Gloating. A sequence of actions triggers Joy and Satisfaction for a character when that character's utility increases after that sequence. These are the sequences that previous narrative planners that do not reason about emotions consider explained. In addition to those sequences, characters can now consent to actions based on their friendship or rivalry to feel HappyFor or Gloating. Characters can also act in response to their fears (*expected* sequences of actions that decrease their utility). An example of an emotional plan is when Tom decides to buy a sword (s10, s30, and s32 in Figure 1). Buying a sword is explained because, with the sword, he is relieved that the bandit can no longer rob him. His utility not only does not increase, but also decreases for using one of his coins. In short, the proposed model allows characters to act emotionally rather than rationally.

Evaluation

We show that (1) the set of stories generated by our system are a superset of those generated by previous narrative planners that do not reason about emotions. We also claim that (2) our model of emotion accurately determines character emotions similar to a human audience's expectations when reading a story. We also claim that (3) when participating in an interactive story generated by our system, human players find the characters that follow our model of emotions more believable than those without emotions.

First Hypothesis

The first hypothesis is that the set of stories generated by the proposed model (*A*) is a superset of the set of stories generated by narrative planners that do not reason about emotion (*B*). Due to space limitations, we present only the sketch of the formal proof.

For all stories in B, a character plan is explained if it increases that character's utility. This is equivalent to the character feeling Joy. Therefore, all the stories in B are also in A. However, A also includes stories in which character plans result in Relief, HappyFor, and Gloating. For instance, it is an explained plan for Tom to go to town and buy a sword. This plan is explained because it makes him feel Relief. This story is not in B because "buying a sword" is not an explained action. In other words, there is no sequence of actions that includes "buying a sword" and increases Tom's utility. Therefore, the proposed model can not only generate all stories that don't reason about emotion, it can also generate new stories that previous models could not.

Second Hypothesis

We evaluated our operationalization of a set of OCC well-being emotions by presenting the example story in Figure 1 to human readers. After reading the description of the story, Tom's actions were presented to the participants one step at a time. Each step was a translation of the corresponding action using simple natural language templates. After a set of steps, we asked what they think Tom may feel at that moment, and participants could choose from one of six emotions.

We only considered the six basic emotions, Joy, Distress (Sadness), Hope, Fear, Relief, and Disappointment, to avoid overwhelming the participant with all 12 well-being emotions. We did not include Satisfaction and FearsConfirmed since in this example story, they were always triggered respectively when Joy and Distress were triggered.

Results We collected the results of 70 participants from Amazon's Mechanical Turk (AMT). Only 2 participants chose not to answer all questions and their responses were removed. There were a total of 7 questions that presented 6 emotions to the participant to choose from. The results were first analyzed to investigate participant agreement. Using Krippendorff's α (Krippendorff 2012), the inter-rater reliability agreement was $\simeq 0.4$. We then used the binomial exact test (Howell 2012) to determine the correct answer to each question. We considered all answers that the participants agreed on as the correct answer. For 5 questions, the participants agreed on exactly one option ($p < 0.05$), and for 2 questions, the participants agreed on two options ($p < 0.05$—for both options) [1]. To calculate the accuracy of our model, the correct answer to each question was then compared to how our model answers that question. Accordingly, the accuracy of our model was 100% for the 6 considered emotions and the considered short story.

Third Hypothesis

To evaluate character believability, we generated a short interactive story in which the participant plays the role of the main character of the story. The rules of the story are similar to our example domain in Figure 1, where the player's goal is to have medicine. One difference between this interactive story and our example domain is that we mention the forest bandit to the player, and the player has the option to first go to town and buy a sword. However, the bandit will never rob the player regardless of the sword.

The story also includes two non-player characters (*John* and *William*), one who expresses emotions through text and facial expressions, and one who doesn't. The participant can view both characters' portraits[2] and thoughts, which may change after certain player choices. For instance, after the player comes back from the forest, both non-player characters state that the bandit stole their axe and they can no longer collect lumber and buy bread.

Figure 2 presents an example of these two characters. At different steps, the emotional character can express *Happy*, *Sad*, or *Scared* facial expressions, and express their thoughts using emotion keywords, e.g. *hope*, *fear*, and so on. For different participants, the emotional character is randomly chosen to be John or William, or to be shown on the left or the right. The player has several opportunities to help either character or none of them, e.g. they could give them a sword or a coin. We hoped the expressions of the emotional character would cause the player to feel empathy and thus, help that character.

The player's goal is to buy the medicine and go back to the cottage. After satisfying this goal, we asked them how believable they found each character and whether they would have helped them both if they could.

[1] For instance, participants agreed that Tom feels Relief and Sadness when he spends a coin to buy a sword (sad for losing a coin and relieved for having a sword.

[2] Character images are taken from the virtual environment *Camelot* (Samuel et al. 2018)

Figure 2: An example of the two characters.

Results We collected the results of 70 participants from AMT, from which 15 did not finish the experiment and their incomplete data was discarded. Using the binomial exact test and Bonferroni correction for testing multiple hypotheses (Holm 1979), the following results were obtained for the rest of the participants.

- The players chose to buy a sword before going to the forest ($p < 0.03$). This shows in an interactive story, human players would act on their fears in addition to their goals. We can use this finding to say that the new stories generated by our model that include plans to trigger Relief make sense to a human audience.
- The players chose to help the emotional character ($p < 0.01$) by giving them one of their coins. The players also stated that they would have helped both characters if they could ($p < 0.01$). This shows that presenting character emotions, only through text and a few facial expressions, made players more empathetic towards that character.
- The players agreed that the emotions and reactions of the emotional character were somewhat to very believable ($p < 0.01$) and more so than the character that expresses no emotions ($p < 0.01$) [3].

Conclusions and Future Work

We introduced a computational model of emotion for strong story state-space narrative planners that reason about intentionality and belief. Our model works independently of the story domain provided by the author and strives for minimum addition to author burden. Our work is limited in that it currently only includes 12 OCC emotions and only a few, i.e. Relief, HappyFor, and Gloating, directly affect the planning process. In this paper, we validated 6 emotions and provided initial results to show that characters generated by our model improve believability and empathy. We plan to expand our model of emotion to include the other OCC emotions and to combine it with computational models of personality, such as (Bahamón and Young 2013; Bahamón, Barot, and Young 2015; Shirvani and Ware 2019; Shirvani 2019) to further improve character believability.

[3] The latter refers to when players chose very believable for the emotional character and somewhat believable or not at all believable for the other character, or chose somewhat believable for the emotional character and not at all believable for the other.

References

Alfonso Espinosa, B.; Vivancos Rubio, E.; and Botti Navarro, V. J. 2014. Extending a BDI agents' architecture with open emotional components. Technical report, Department of Information Technology, Universitat Politeócnica de Valeóncia.

Bahamón, J. C., and Young, R. M. 2013. CB-POCL: a choice-based algorithm for character personality in planning-based narrative generation. In *Proceedings of the 2013 Workshop on Computational Models of Narrative*, 4–23.

Bahamón, J. C.; Barot, C.; and Young, R. M. 2015. A goal-based model of personality for planning-based narrative generation. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 4142–4143.

Cavazza, M.; Lugrin, J.-L.; Pizzi, D.; and Charles, F. 2007. Madame bovary on the holodeck: immersive interactive storytelling. In *Proceedings of the 15th ACM international conference on Multimedia*, 651–660.

Cavazza, M.; Pizzi, D.; Charles, F.; Vogt, T.; and André, E. 2009. Emotional input for character-based interactive storytelling. In *Proceedings of The 8th International Conference on Autonomous Agents and Multi-Agent Systems-Volume 1*, 313–320.

El-Nasr, M. S.; Yen, J.; and Ioerger, T. R. 2000. FLAME—fuzzy logic adaptive model of emotions. *Autonomous Agents and Multi-agent systems* 3(3):219–257.

Gebhard, P. 2005. ALMA: a layered model of affect. In *Proceedings of the fourth international joint conference on Autonomous Agents and Multi-Agent Systems*, 29–36.

Gratch, J. 1999. Why you should buy an emotional planner. In *Proceedings of the Autonomous Agents 1999 Workshop on Emotion-based Agent Architectures (EBAA'99)*, 99–465. Citeseer.

Holm, S. 1979. A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics* 65–70.

Howell, D. C. 2012. *Statistical methods for psychology*. Cengage Learning.

Kenny, P.; Hartholt, A.; Gratch, J.; Swartout, W.; Traum, D.; Marsella, S.; and Piepol, D. 2007. Building interactive virtual humans for training environments. In *Proceedings of Interservice/Industry Training, Simulation, and Education Conference*, 911–916.

Krippendorff, K. 2012. *Content analysis: An introduction to its methodology*. Sage.

Marsella, S. C., and Gratch, J. 2009. EMA: A process model of appraisal dynamics. *Cognitive Systems Research* 10(1):70–90.

Mehrabian, A. 1996. Analysis of the big-five personality factors in terms of the PAD temperament model. *Australian Journal of Psychology* 48(2):86–92.

Meyer, J.-J. C. 2006. Reasoning about emotional agents. *International Journal of Intelligent Systems* 21(6):601–619.

Ortony, A.; Clore, G. L.; and Collins, A. 1990. *The cognitive structure of emotions*. Cambridge university press.

Pérez-Pinillos, D.; Fernández, S.; and Borrajo, D. 2011. Modeling motivations, personality traits and emotional states in deliberative agents based on automated planning. In *International Conference on Agents and Artificial Intelligence*, 146–160. Springer.

Pizzi, D.; Cavazza, M.; and Lugrin, J.-L. 2007. Extending character-based storytelling with awareness and feelings. In *Proceedings of the 6th international joint conference on Autonomous Agents and Multi-Agent Systems*, 1–3.

Porteous, J.; Cavazza, M.; and Charles, F. 2010. Applying planning to interactive storytelling: Narrative control using state constraints. *ACM Transactions on Intelligent Systems and Technology* 1(2):1–21.

Riedl, M. O., and Bulitko, V. 2013. Interactive narrative: An intelligent systems approach. *AI Magazine* 34(1):67–67.

Riedl, M. O., and Young, R. M. 2010. Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research* 39:217–268.

Russell, S. J., and Norvig, P. 2009. *Artificial Intelligence: a modern approach*. Pearson, 3 edition.

Samuel, B.; Reed, A.; Short, E.; Heck, S.; Robison, B.; Wright, L.; Soule, T.; Treanor, M.; McCoy, J.; Sullivan, A.; et al. 2018. Playable experiences at AIIDE 2018. In *Proceedings of the Fourteenth Artificial Intelligence and Interactive Digital Entertainment Conference*, 275–280.

Shirvani, A., and Ware, S. G. 2019. A plan-based personality model for story characters. In *Proceedings of the 15th AAAI international conference on Artificial Intelligence and Interactive Digital Entertainment*, 188–194.

Shirvani, A.; Farrell, R.; and Ware, S. G. 2018. Combining intentionality and belief: Revisiting believable character plans. In *Fourteenth Artificial Intelligence and Interactive Digital Entertainment Conference*, 222–228.

Shirvani, A.; Ware, S. G.; and Farrell, R. 2017. A possible worlds model of belief for state-space narrative planning. In *Proceedings of the Thirteenth Artificial Intelligence and Interactive Digital Entertainment Conference*, 101–107.

Shirvani, A. 2019. Towards more believable characters using personality and emotion. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 230–232.

Ware, S. G.; Garcia, E.; Shirvani, A.; and Farrell, R. 2019. Multi-agent narrative experience management as story graph pruning. In *Proceedings of the Fifteenth Artificial Intelligence and Interactive Digital Entertainment Conference*, 87–93.

Young, R. M.; Ware, S. G.; Cassell, B. A.; and Robertson, J. 2013. Plans and planning in narrative generation: a review of plan-based approaches to the generation of story, discourse and interactivity in narratives. *Sprache und Datenverarbeitung, Special Issue on Formal and Computational Models of Narrative* 37(1-2):41–64.