

Modeling Appraisal in Theory of Mind Reasoning

Mei Si, Stacy C. Marsella, David V. Pynadath

Information Sciences Institute
University of Southern California
Marina del Rey, CA 90292
meisi@isi.edu, marsella@isi.edu, pynadath@isi.edu

Abstract. Cognitive appraisal theories, which link human emotional experience to their interpretations of events happening in the environment, are leading approaches to model emotions. In this paper, we investigate the computational modeling of appraisal in a multi-agent decision-theoretic framework using POMDP based agents. We illustrate how five key appraisal dimensions (motivational relevance, motivation congruence, accountability, control and novelty) can be derived from the processes and information required for the agent’s decision-making and belief maintenance. Through this illustration, we not only provide a solution for computationally modeling emotion in POMDP based agents, but also demonstrate the tight relationship between emotion and cognition. Our model of appraisal is applied to three different scenarios to illustrate its usage. We also discuss how the modeling of theory of mind (recursive beliefs about self and others) is critical for simulating social emotions.

1 Introduction

In recent years virtual agents have been used in a wide range of domains to interact with people, such as tutoring systems (e.g. [1–3]), entertainment systems (e.g. [4–6]) and virtual salesmen (e.g. [7]). The computational modeling of emotion has become a key aspect of virtual agent designs. Incorporating models of emotion into virtual agents can make the agent behave more lifelike and expressive as well as give the agent the capacity to understand the emotion of others.

Computational models of emotion used in virtual agents have often been based on cognitive appraisal theories [8–12], which argue that a person’s subjective assessment of their relationship to the environment, the person-environment relation, determines the person’s emotional responses. This assessment occurs along several dimensions, called appraisal variables or checks, including motivational congruence, novelty, control, etc. Emotion is decided by the combination of results from these checks. For example, an unexpected event that is incongruent with the person’s motivations and is beyond the capacity of the person to cope with may lead to fear responses.

The work we report here investigates the computational modeling of appraisal within an existing multi-agent decision-theoretic framework – Thespian [3] for interactive narratives, in which a human user can play a role in a story and interact with virtual characters.

Our approach in modeling appraisal in Thespian has been to keep separate the processes by which an agent’s representation of the person-environment relation is constructed and belief maintenance processes already support appraisal processes. The construction of the person-environment representation is a product of a Thespian agent’s standard decision-making and belief maintenance processes. We illustrate how key appraisal variables can then be straightforwardly extracted from this representation of an agent’s decision making and belief maintenance processes. Through this illustration, we not only provide a solution for computationally modeling appraisal in POMDP based agents, but also demonstrate the tight relationship between emotion, decision-making and belief revision.

This work is in the spirit of, and closely related to, work on the EMA model of emotion [13] that argues that appraisal can leverage the representations formed by an agent’s cognitive decision-making. A key difference here is Thespian’s modeling of other agents and the role it plays in decision-making and belief revision. Agents in Thespian possess beliefs about other agents that constitute a fully specified, quantitative model of the other agents’ beliefs, policies and goals. In other words, the agents have a theory of mind capability with which they can simulate others. Compared to computational models that do not have an explicit theory of mind, Thespian’s explicit representation of agents’ subjective beliefs about each other enables the model to better reason about social emotions. For example, agents can reason about other agent’s cognitive and emotional processes both from the other agent’s and its own perspective.

In the work reported here, we focus on five appraisal variables: motivational relevance, motivational congruence, accountability, control and novelty. We demonstrate the application of our model of appraisal in three different scenarios, including the Little Red Riding Hood fairy tale, a small talk between two persons and a firing-squad scenario as described in [14]. The Little Red Riding Hood story will be used as an example to motivate the discussion throughout this paper.

2 Related Work

Cognitive appraisal theories have had an increasing impact on the design of virtual agents. In FLAME, El Nasr et al. [15] use domain-independent fuzzy logic rules to simulate appraisal. Moffat and Frijda [16] build an agent framework called WILL, in which concerns and relevance are simply evaluated as the discrepancies between the agent’s desired state and the current state. Similar to WILL, the Cathexis model [17] uses a threshold model to simulate basic variables, which they call “sensors”, related to emotion. The OCC model of appraisal [10] has inspired many computational systems. Elliott’s [18] Affective Reasoner

uses a set of domain-specific rules to appraise events based on the OCC theory. Both EM [19] and FearNot! [2] deployed the OCC model of emotion over plan based agents using domain-independent approaches.

The work on EMA [13] follows the Smith and Lazarus theoretical model of appraisal [11]. EMA defines appraisal as domain-independent processes over a plan-based representation of the person-environment relation, termed a causal interpretation. Cognitive processes maintain the causal interpretation while appraisal processes map appraisal relevant features of this representation to appraisal dimensions. Whereas the construction of the causal interpretation by cognitive processes is treated as distinct from appraisal, the form of the representation is designed to reduce appraisal to simple (and fast) pattern matching.

3 Thespian Agent

Thespian is a multi-agent framework for authoring and simulating interactive narratives, in which a human user can play a role in a story and interact with virtual characters. Thespian is built upon PsychSim [20], a multi-agent system for social simulation based on Partially Observable Markov Decision Problems (POMDPs).

Thespian’s basic architecture uses PsychSim’s POMDP based agents to control each virtual character, with the character’s personality and motivations encoded as the agent’s goals. This section introduces components in a Thespian agent that are relevant to our new cognitive appraisal model, including the agent’s state, dynamics, goals, beliefs, policies and relationships.

State A character’s state is defined by a set of state features, such as the name and age of the character, and the relation between that character and other characters, e.g. affinity. Values of state features are represented as real numbers.

Dynamics Dynamics define how actions affect agents’ states. For example, small talk among a group of agents can increase their affinity with each other by 0.1. The effects of actions can be defined with probabilities. For examples, the author may define that when the hunter kills the wolf, he can only succeed 60% of the time.

Goals We model a character’s personality profile as its various goals and their relative importance (weight). Goals are expressed as a reward function over the various state features an agent seeks to maximize or minimize. For example, a character can have a goal of maximizing its affinity with another character. The initial value of this state feature can be any value between 0.0 and 1.0; this goal is completely satisfied once the value reaches 1.0. An agent usually has multiple goals with different relative importance (weights). For example, the character may have another goal of knowing another character’s name, and this goal may be twice as important to the character as the goal of maximizing affinity.

Beliefs Thespian agents have a “Theory of Mind” that allows them to form mental models about other agents. The agent’s subjective view of the world includes its beliefs about itself and other agents and *their* subjective views of the world, a form of recursive agent modeling. An agent’s subjective view (mental model) of itself or another agent includes every component of that agent, such as state, beliefs, policy, etc.

An agent’s belief about its own or another agent’s state is represented as a set of real values with probability distributions. The probability distribution of the possible values of a state feature indicates the character’s beliefs about this value. For example, a character’s belief about its affinity with another character could be $\{.1 \text{ with probability of } 90\%, .9 \text{ with probability of } 10\%\}$. For the simplicity of demonstration, in this paper we only give examples using the expected values of the state features. The expected value of a state feature is simply calculated as $\sum_{i=0}^n \text{value}_i * P(\text{value}_i)$.

Each agent can have one mental model of itself and multiple mental models of other agents. The agent’s belief about another agent is a probability distribution over alternative mental models, e.g. in the Red Riding Hood story, Red can believe that there is a 40% chance the wolf has a goal to eat people and a 60% chance otherwise.

An agent’s beliefs get updated in two ways. One is through dynamics. Upon observation of an action, within each mental model the corresponding dynamics will be applied, and the related state features’ values are updated. The other way an agent changes its beliefs is through adjusting the relative probability of alternative mental models. Each observation serves as an evidence for plausibility of alternative mental models, i.e. how consistent the observation is with the predictions from the mental models. Using this information, the probabilities of the mental models are updated based on Bayes’ Theorem [21].

Policy In Thespian, all agents use a bounded lookahead policy. Each agent has a set of candidate actions to choose from when making decisions. When an agent selects its next action, it projects into the future to evaluate the effect of each option on the state and beliefs of other entities in the story. The agent considers not just the immediate effect, but also the expected responses of other characters and, in turn, the effects of those responses, and its reaction to those responses and so on. The agent evaluates the overall effect with respect to its goals and then chooses the action that has the highest expected value. For example, when Red decides her next action after being stopped by the wolf on her way to Granny’s house, the following reasoning happens in her “mind,” using her beliefs about the wolf and herself. For each of her action options, e.g. talking to the wolf or walking away, she anticipates how the action affects each character’s state and utility. Next, Red needs to predict the wolf’s responses to each of her potential movements by simulating the wolf’s lookahead process. Similarly, for each of the wolf’s possible action choices, e.g. asking Red a question or continuing small talk, Red calculates the immediate expected states and utilities of both the wolf and herself. Then, Red simulates the wolf anticipating her responses in

turn. The lookahead process is only boundedly rational – the recursive reasoning stops when the maximum number of steps for forward projection is reached. For example, if the number of lookahead steps is set to be one, the wolf will pick the action with highest utility after simulating one step of Red’s response rather than several rounds of interaction. Similarly based on the wolf’s potential responses in the next step, Red calculates the utilities of her action options and chooses the one with the highest utility¹.

When an agent has multiple mental models about other agents, currently by default the agent uses the most probable mental models for reasoning about others’ responses. The user (the author of the interactive narrative) can, however specify alternative rules for the agent, such as to consider worst/best case scenarios. In that case, the utilities of actions will be evaluated within each mental model during the agent’s decision-making process. In Section 4.2, an example is provided for utilizing this function in evaluating an agent’s appraisal of control.

Relationships PsychSim has a built-in capability of modeling static and dynamic social relationships between agents which in turn can influence the agent’s decision-making and belief update. Specifically, PsychSim agents maintain a measure of support (or affinity) for another agent. This capacity is relevant to our subsequent discussion of accountability so we briefly touch upon it here. Support is computed as a running history of their past interactions. An agent increases (decreases) its support for another, when the latter selects an action that has a high (low) reward, with respect to the preferences of the former.

4 Computational Model of Appraisal

In this section we illustrate how appraisal dimensions can be derived using processes involved and information gathered, in the agent’s decision-making processes. We first present the overall appraisal process, which specifies when appraisal happens and where the related information comes from, and then present our algorithms for evaluating the five appraisal dimensions.

4.1 Appraisal Process

Smith & Lazarus describe appraisal as a continuous process, that people constantly reevaluate their situations – the “appraisal-reappraisal” cycle [11]. In our computational model, we also try to capture this phenomenon. Upon observing a new event – an action performed by an agent or the human user, each agent appraises the situation and updates its beliefs. The calculation of motivational relevance, motivational congruence, novelty and accountability depends only on the agent’s beliefs about other agents’ and its own utilities in the current step and the previous steps, and therefore can always be derived immediately (see

¹ Note that at run time Thespian/Psychim agents do not need to perform the lookahead reasoning, rather they use compiled policies which are precompiled offline [22].

Section 4.2 for the details.) Depending on the extent of reasoning the agent performed in the former steps, the agent may or may not have information immediately regarding its control of the situation. However, when the agent makes its next decision, its control is automatically evaluated. These evaluations in turn will affect the agent’s emotion. In fact, the agent may reevaluate along every appraisal dimension as it obtains more updated information about other characters. In the examples given in this paper, we report appraisal produced using information gathered in the agent’s previous lookahead process. However, it could be based on either the expectation formed in previous steps, or the lookahead process being performed at the current step. The agent may also express both emotional responses in sequence.

Thespian agents have mental models of other agents; they can not only have emotional responses to the environment but also form expectations of other agents’ emotions. During the decision-making processes, the lookahead process calculates the agent’s belief about states and utilities of every possible action choice of each of the agents. This information is kept in the agent’s memory as its expectations. To simulate expectation of another agent’s emotional responses, the observing agent’s beliefs about the other agent is used for deriving appraisal dimensions. For instance, agent A can use its beliefs about agent B to evaluate the motivational relevance and novelty of an event to agent B. When evaluating appraisal dimensions relating to oneself, the agent will use its belief about itself. If the observing agent has multiple mental models about another agent, it uses the mental model with highest probability to simulating appraisal.

4.2 Appraisal Dimensions

In this section we provide pseudo-code for evaluating the five appraisal dimensions (motivational relevance, motivation congruence or incongruence, accountability, control and novelty) using states/utilities calculated during the agent’s decision-making process.

Motivational Relevance & Motivational Congruence or Incongruence

Motivational relevance evaluates the extent to which an encounter touches upon personal goals, and motivational congruence or incongruence measures the extent to which the encounter thwarts or facilitates the personal goals [11].

Algorithm 1 Motivational Relevance & Motivation Congruence

preUtility: utility before the event happens

curUtility: utility after the event happens

$$\text{Motivational Relevance} = \text{abs} \frac{\text{curUtility} - \text{preUtility}}{\text{preUtility}}$$

$$\text{Motivational Congruence} = \frac{\text{curUtility} - \text{preUtility}}{\text{abs}(\text{preUtility})}$$

We model these appraisal dimensions as a product of the agent’s utility calculations which are integral to the agent’s decision-theoretic reasoning. We use the

ratio of relative utility change and the direction of utility change to model these two appraisal dimensions. The rationale behind this is that the same amount of utility change will result in different subjective experiences depending on the agent’s current utility. For instance, if eating a person increases the wolf’s utility by 10, it will be 10 times more relevant and motivationally congruent if the wolf’s original utility is 1 (very hungry) compared to the original utility of 10 (less hungry). Algorithm 1 gives the equations for evaluating motivational relevance and motivational congruence or incongruence. When the calculated value of *Motivational Congruence* is negative, the event is motivationally incongruent to the agent, to the extent of $\text{abs}(\text{Motivational Congruence})$.

Accountability

Accountability characterizes which person deserves credit or blame for a given event [11]. Various theories have been proposed for assigning blame/credit, e.g. [23]. The reasoning usually considers factors such as who directly causes the event, does the person foresee the result, does the person intend to do so or is it coerced, etc.

Just as the appraisal of motivational relevance and motivation congruence can be performed as part of existing Thespian/PsychSim decision-making and belief update processes, we argue here that accountability can be treated as an improvement to PsychSim’s existing approach to model affinity support relationships between agents.

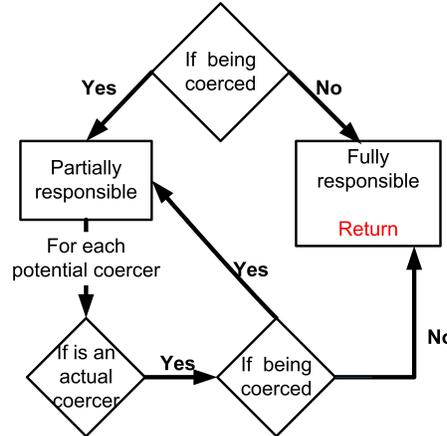


Fig. 1. Accountability

In Figure 1 we use a diagram to illustrate our algorithm for determining accountability. For the simplicity of the algorithm, we assume that the agent expects others to always foresee the effects of their actions. This assumption is reasonable for most of the virtual agents because normally a person would expect others to project into the future the same number of steps in their decision-making process as what the person will do themselves.

Our algorithm first looks at the agent which directly causes the harm/benefit, and judges if this agent is the one who should be fully responsible. The function *If_Coerced()* is called to determine if the agent was coerced to perform the action. If the agent was not coerced, it should be fully responsible and the reasoning stops there. Otherwise, each of the coercers will be judged on whether it was coerced by somebody else or volunteered to do the action. We will trace limited steps back in the history to find out all the responsible agents.

Algorithm 2 If_Coerced(*actor*, *pact*)

```

# actor: the agent being studied
# pact: the action performed by actor
# preUtility: actor's utility before doing pact

for action in actor.actionOptions() do
  if action  $\neq$  pact then
    #if there exists another action which does not hurt actor's own utility
    if utility(action)  $\geq$  preUtility then
      Return F
  if utility(action)  $<$  preUtility then
    Return F
Return T

```

Algorithm 2 gives the pseudo code for determining if an agent was coerced, and Algorithm 3 answers the question of who coerced it. Coercion is computationally defined in our model as: if other than the action chosen by the agent, all its action options lead to a drop in its utility (i.e. the agent will be punished if it chooses any other actions); however, if all of the agent's action options will result in a utility drop, the agent is regarded as not being coerced for picking that particular action which hurt the other agent's utility.

Algorithm 3 Is_Coercer_For(*agent*, *actor*, *agent_pact*, *actor_pact*)

```

# check if agent coerced actor
# agent_pact: the action performed by agent
# actor_pact: the action performed by actor

for action in agent.actionOptions() do
  if action  $\neq$  agent_pact then
    Simulate action agent_pact
    if If_Coerced(actor,actor_pact) == F then
      Return T
Return F

```

To decide who coerced an agent, we treat each agent that acted between the coerced agent's current and last actions as a potential coercer. For each potential coercer, if the coerced agent would not have been coerced in case the potential coercer had made a different choice, then the potential coercer is judged as actually being a coercer. This is shown in Algorithm 3.

Control

The appraisal of control evaluates the extent to which an event or its outcome can be influenced or controlled by people [12]. It captures not only the individual's own ability to control the situation but also the potential for seeking instrumental social support from other people.

In our computational model, if the agent is modeled as using the most probable mental models in reasoning about other agents' actions, the evaluation of control is straightforward. It is whether the agent anticipates an event will happen in the near future that will take it out of the unfavorable situation. If the agent is modeled as considering all alternative mental models of other agents in its decision making, we factor in the probabilities of the mental models.

Algorithm 4 Control(*preUtility*)

preUtility: the agent's utility before falling into the unfavorable situation

control \leftarrow 0

for *m1* in *mental_models_about_agent1* **do**

for *m2* in *mental_models_about_agent2* **do**

for *m3* in *mental_models_about_self* **do**

 #project limited steps into the future using this set of mental models

 lookahead(*m1,m2,m3*)

 #*curUtility*: the agent's current utility after the lookahead process

if *curUtility* \geq *preUtility* **then**

control \leftarrow *control* + *p(m1)* * *p(m2)* * *p(m3)*

Return *control*

Algorithm 4 gives the pseudo code for evaluating control when considering alternative mental models. This algorithm first looks at whether there is a solution within individual mental models set about self and other agents, then the probabilities of these mental models to be correct, and therefore the event, if being predicted, will actually happen in the future. For example, assume Granny has two mental models about the wolf, which are: a 60% chance the wolf will die after the hunter shoots at it and a 40% chance the wolf will not; also Granny has two mental models regarding the hunter's location, which are a 50% chance the hunter is close by and therefore has a chance to rescue her and a 50% chance the hunter is far away. After Granny is eaten by the wolf, the only event that can help her is that the wolf is killed by the hunter. Therefore, she would evaluate her control as: 60% * 50% = 30%. Note that when using information generated in the past reasoning process for deciding control, the reasoning process, though it happened in the past, must contain information regarding the current moment and the future. Algorithm 4 contains pseudo-code for the three-agent interaction case. It is straightforward to configure it to be applied when more or less agents are in the interaction.

Novelty

In this work, we adapt Scherer’s definition of “novelty at the conceptual level”, that is, novelty describes whether the event is expected from the agent’s past beliefs [12].

Novelty appraisals can be treated as a byproduct of the agent’s belief maintenance. Specifically, in a multi-agent context the novelty of another agent’s behavior can be viewed as the opposite side of the observing agent’s beliefs about the other agent’s motivational consistency, i.e. the more consistent the event is with the observing agent’s beliefs about the other agent’s motivation, the less novel. We define novelty as $1 - consistency$, where *consistency* is calculated using one of the methods proposed by Ito et al. [21].

$$Consistency(a_j) = \frac{e^{rank(a_j)}}{\sum_j e^{rank(a_j)}} \quad (1)$$

Novelty is calculated based on the most probable mental model that the observing agent has about the actor of the event. The algorithm first ranks the actor’s alternative actions’ utilities in reversed order ($rank(a_j)$). The higher the event’s utility ranks compared to other alternatives, the more novelty. For example, if from Red’s perspective the wolf did an action which has the second highest utility among all five alternatives, the novelty Red experiences is calculated as $1 - \frac{e^3}{\sum_{0-4} e^j} = 0.37$.

5 Sample Results

In this section we provide additional examples to illustrate the use of our computational model of appraisal in modeling social interactions. In particular, in Scenario 1 we demonstrate the tight relationship between emotion and cognitive decision-making by showing how appraisal is affected by the depth of reasoning in decision-making. In Scenario 2 we provide a complex situation for accountability reasoning and show that the result of our model is consistent with another validated computational model of social attribution.

5.1 Scenario 1: Small Talk

In this example, two persons (A and B) take turns talking to each other. Both of them have goals to be talkative and obey social norms. In fact, just the norm following behavior itself is an incentive to them – they will be rewarded whenever they do an action that is consistent with social norms. Table 1 contains the two persons’ appraisals of motivational relevance regarding each other’s actions. We did not include results of other appraisal dimensions as they are less interesting in this scenario.

We provide a comparison of appraisal results when the person’s previous reasoning process takes a different number of steps. It can be observed in Table 1 that a person appraises the other person’s initiatives as irrelevant when

it performs shallow reasoning (lookahead steps = 1). In this case, even though the person has predicted the other person’s action, the action does not bring him/her immediate reward. Once the person reasons one step further, he/she finds out that by opening up a topic the other person actually provides him/her a chance to engage in further conversation and perform a norm following action. The person will then appraise the other person’s behavior as relevant.

Table 1. Example 1: Small talk between two persons

Step	Action	Perspective	Lookahead Steps	Motivational Relevance
1	A greets B	B	1	0
		B	2	$3e^{10}$
2	B greets A	A	1	0
		A	2	0.99
3	A asks B a question	B	1	0
		B	2	0.99
4	B answers the question	A	1	0
		A	2	0.49

5.2 Scenario 2: Firing-squad

We implemented this scenario from [14] to illustrate accountability reasoning in which agents are coerced and have only partial responsibility. The scenario goes like this:

In a firing-squad, the commander orders the marksmen to shoot a prisoner. The marksmen refuse the order. The commander insists that the marksmen shoot. They shoot at the prisoner and he dies.

We modeled the commander as an agent with an explicit goal of killing the prisoner, and the marksmen as having no goals related to the prisoner but will be punished if they do not follow the commander’s order. Using our appraisal model, from the prisoner’s perspective, the marksmen hold responsibility for his/her death because they are the persons who directly perform the action. Further, the prisoner simulates decision-making process of the marksmen and finds out that the marksmen are coerced because their utilities will be hurt if they perform any action other than shooting. The commander acts right before the marksmen in the scenario and therefore is identified as a potential coercer for the marksmen. Using Algorithm 3, the prisoner can see that if the commander chose a different action, the marksmen are not coerced to shoot. Assuming the prisoner does not find a coercer for the commander, he/she will now believe that the commander holds full responsibility for his/her death. This prediction is consistent with the prediction from Mao’s model of social attribution and the data collected from human subjects to validate that model [14].

6 Discussion and Future Work

In this paper we demonstrated how appraisal dimensions can be directly derived from the utilities calculated in the agent’s decision-making processes. For decision-making and mental model update, comparison among expected utilities is the key. Comparison of expected utilities is also the central piece for deriving appraisal dimensions. As demonstrated in our computational model, none of the appraisal dimensions require additional calculation of utilities other than what has already been performed in the agent’s lookahead reasoning process.

We have also given examples of applying this appraisal model to several scenarios. We demonstrate situations in which the agent’s appraisal would change if it performed a different level of reasoning in the previous step, in which the agent assigns partial responsibility to causal agents, and in which the agent reasons about control of self and others.

We modeled appraisal dimensions within Thespian framework for interactive narrative. Thespian agents are decision-theoretic goal-based agents with theory of mind modeling. Compared to computational models of appraisal based on other types of agents, our model has two key advantages. First, the fact that Thespian agents have a theory of mind capability enables them to simulate others and reason about social emotions. This ability allows us to simulate agents’ different emotional response to the same situation and an agent’s potential misexpectations about other agents’ emotional states. For example, if Granny believes that the hunter can always kill the wolf successfully and the hunter believes that he can only kill the wolf successfully 60% of the time, Granny’s control when being eaten by the wolf will be evaluated differently from Granny’s and the hunter’s perspective. Secondly, we explicitly model the depth of reasoning in agents as the number of steps they project into the future. As shown in scenario 1, different depths of reasoning lead to different appraisals. This feature enables us to map the agent’s emotions not only to its knowledge and states, but also to its decision-making process itself.

Our future work is planned in two directions. First we want to add the emotion-cognition interaction to Thespian agents. Currently interactions among agents are generated based on “cold” reasoning, and emotion is modeled as the agents’ responses to what have happened in the environment. We want to improve our system by modeling how emotion affects the agent’s decision-making and belief update processes. Secondly, we want to enrich the current model with a more complete set of appraisal dimensions, such as urgency and emotion-focused coping potential, and consider the overlap among dimensions proposed by different cognitive appraisal theories. We plan to extend this model into a flexible framework that can be used with different theories of cognitive appraisal, and further as a platform for evaluating these theories.

7 Conclusion

In this work we provide a computational model of appraisal for POMDP based agents. We focused on five key appraisal dimensions for virtual agents: moti-

vational relevance, motivational congruence, accountability, control and novelty. We demonstrate the derivation of appraisal dimensions from the information gathered in the agent's decision-making process. Through this demonstration, we illustrate the tight coupling between emotion and cognitive decision-making. We also provide various examples of applying this model to social agents in different scenarios.

This model is built using social agents within the Thespian framework for interactive narratives. Thespian agents are decision-theoretic goal-driven agents with modeling of theory of mind. Compared to other computational models without explicit modeling of agents' subjective beliefs about each other, our model can more easily reason about utilities of actions from different agents' perspectives and therefore potentially provide a more realistic model of appraisal in social interactions.

8 Acknowledgments

This work was sponsored by the U.S. Army Research, Development, and Engineering Command (RDECOM), and the content does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

References

1. Marsella, S.C., Johnson, W.L., Labore, C.: Interactive pedagogical drama for health interventions. In: AIED. (2003)
2. Aylett, R., Dias, J., Paiva, A.: An affectively-driven planner for synthetic characters. In: ICAPS. (2006)
3. Si, M., Marsella, S.C., Pynadath, D.V.: Thespian: An architecture for interactive pedagogical drama. In: AIED. (2005)
4. Riedl, M.O., Saretto, C.J., Young, R.M.: Managing interaction between users and agents in a multi-agent storytelling environment. In: AAMAS. (2003) 741–748
5. Cavazza, M., Charles, F., Mead, S.J.: Agents' interaction in virtual storytelling. In: Proceedings of the International WorkShop on Intelligent Virtual Agents. (2001) 156–170
6. Szilas, N.: IDtension: a narrative engine for interactive drama. In: the 1st International Conference on Technologies for Interactive Digital Storytelling and Entertainment, Darmstadt Germany (2003)
7. Cassell, J., Bickmore, T., Vilhjálmsón, H., Yan, H.: More than just a pretty face: Affordances of embodiment. In: IUI, New Orleans, Louisiana (2000) 52–59
8. Roseman, I.: Cognitive determinants of emotion: A structural theory. *Review of Personality and Social Psychology* **2** (1984) 11–36
9. Smith, C.A., Ellsworth, P.C.: Patterns of appraisal and emotion related to taking an exam. *Personality and Social Psychology* **52** (1987) 475–488
10. Ortony, A., Clore, G.L., Collins, A.: *The Cognitive Structure of Emotions*. Cambridge, UK: Cambridge University Press (1998)
11. Smith, C.A., Lazarus, R.S.: Emotion and adaptation. In Pervin, L.A., ed.: *Handbook of personality: Theory and research*. Guilford, New York (1990)

12. Scherer, K.: Appraisal considered as a process of multilevel sequential checking. In Scherer, K., Schorr, A., Johnstone, T., eds.: *Appraisal Processes in Emotion: Theory, Methods*. Oxford University Press., Oxford (2001)
13. Gratch, J., Marsella, S.: A domain-independent framework for modeling emotion. *Cognitive Systems Research* **5**(4) (2004) 269–306
14. Mao, W., Gratch, J.: Social causality and responsibility: Modeling and evaluation. In: IVA, Kos, Greece (2005)
15. El Nasr, M.S., Yen, J., Ioerger, T.: Flame: Fuzzy logic adaptive model of emotions. *Autonomous Agents and Multi-Agent Systems* **3**(3) (2000) 219–257
16. Moffat, D., Frijda, N.: Where there's a will there's an agent. In: *ECAI-94 Workshop on Agent Theories, Architectures, and Languages*, Amsterdam, The Netherlands (1995)
17. Velasquez, J.: Modeling emotions and other motivations in synthetic agents. In: *AAAI*. (1997)
18. Elliott, C.: The affective reasoner: A process model of emotions in a multi-agent system. PhD thesis, Northwestern University Institute for the Learning Sciences (1992)
19. Reilly, W.S., Bates, J.: Building emotional agents. Technical Report CMU-CS-92-143, Carnegie Mellon University (1992)
20. Marsella, S.C., Pynadath, D.V., Read, S.J.: PsychSim: Agent-based modeling of social interactions and influence. In: *Proceedings of the International Conference on Cognitive Modeling*. (2004) 243–248
21. Ito, J.Y., Pynadath, D.V., Marsella, S.C.: A decision-theoretic approach to evaluating posterior probabilities of mental models. In: *AAAI-07 Workshop on Plan, Activity, and Intent Recognition*. (2007)
22. Pynadath, D.V., Marsella, S.: Fitting and compilation of multiagent models through piecewise linear functions. In: *AAMAS*. (2004) 1197–1204
23. Weiner, B.: *The Judgment of Responsibility: A Foundation for a Theory of Social Conduct*. The Guilford Press (1995)