

A Domain-Independent Approach of Cognitive Appraisal Augmented by Higher Cognitive Layer of Ethical Reasoning

Suman Ojha (Suman.Ojha@student.uts.edu.au)

Jonathan Vitale (Jonathan.Vitale@uts.edu.au)

Mary-Anne Williams (Mary-Anne.Williams@uts.edu.au)

University of Technology Sydney – Centre For Artificial Intelligence
Innovation and Enterprise Research Lab (The Magic Lab), 15 Broadway, Ultimo NSW 2007, Australia

Abstract

According to cognitive appraisal theory, emotion in an individual is the result of how a situation/event is evaluated by the individual. This evaluation has different outcomes among people and it is often suggested to be operationalised by a set of rules or beliefs acquired by the subject throughout development. Unfortunately, this view is particularly detrimental for computational applications of emotion appraisal. In fact, it requires providing a knowledge base that is particularly difficult to establish and manage, especially in systems designed for highly complex scenarios, such as social robots. In addition, according to appraisal theory, an individual might elicit *more than one emotion at a time* in reaction to an event. Hence, determining which emotional state should be attributed in relationship to a specific event is another critical issue not yet fully addressed by the available literature. In this work, we show that: (i) the cognitive appraisal process can be realised without a complex set of rules; instead, we propose that this process can be operationalised by knowing only the positive or negative perceived effect the event has on the subject, thus facilitating extensibility and integrability of the emotional system; (ii) the final emotional state to attribute in relation to a specific situation is better explained by ethical reasoning mechanisms. These hypotheses are supported by our experimental results. Therefore, this contribution is particularly significant to provide a more simple and generalisable explanation of cognitive appraisal theory and to promote the integration between theories of emotion and ethics studies, currently often neglected by the available literature.

Keywords: Cognitive appraisal theory; computational emotion model; emotion combination; ethics

Introduction

The attribution of an emotional state to self or others can occur when a complex state of the organism is accompanied by variable degrees of awareness, often referred to as *appraisal* (Scherer, 2001). Two levels of appraisal can be distinguished (Lambie & Marcel, 2002): a first-order phenomenological state and a conscious second-order awareness. Both states can be either self-directed (first-person perspective) or world-directed (third person perspective) (Vitale, Williams, Johnston, & Boccignone, 2014). The present work will be concerned in discussing the nature of the conscious second-order appraisal process, known as *cognitive appraisal process of emotion*.

Traditional literature in emotional processing studies suggests that this cognitive appraisal process may underlie the evaluation of a set of variables called *appraisal variables* (Ortony, Clore, & Collins, 1990; Lazarus, 1991; Roseman, Spindel, & Jose, 1990; Scherer, 2001). Appraisal variables can be understood as the criteria used to assess a situation in relation to emotion elicitation process. For example, in

appraisal theory of Ortony et al. (1990), core appraisal variables¹ considered are *desirability* - which assesses how desirable an event is, *praiseworthiness* - which measures how praiseworthy the action of an agent is and *appealingness* - which measures how appealing is the agent to the appraising individual. Appraisal theories suggest that individuals converge to an emotional state depending on the evaluation of these variables. This position is further supported by the majority of existing computational explanations of cognitive appraisal (Dias & Paiva, 2005; El-Nasr, Yen, & Ioerger, 2000; Velasquez, 1997). However, the proposed accounts offer limited perspectives addressing only domain specific situations and making use of knowledge shaped as a set of pre-defined rules (Dias & Paiva, 2005; El-Nasr et al., 2000). Thus, (i) *the available literature in cognitive appraisal theory currently does not provide a clear computational explanation for domain-independent cognitive appraisal mechanisms*. This is a significant research problem for both cognitive science and computer science research communities; in fact, on one hand, having a computational theory of domain-independent cognitive appraisal mechanisms can assist cognitive science researchers in addressing open research gaps in emotional processing studies, and, on the other hand, this computational account can be more easily integrated in disparate intelligent systems without the need of defining a complex set of domain-dependent rules.

However, this is not the only limitation presented by currently available explanations of cognitive appraisal theory. According to cognitive appraisal theory of emotion, an event can elicit more than one emotions simultaneously with varying intensities (Ortony et al., 1990). Nevertheless, (ii) *it is not clear yet what is the best strategy to select an emotional state for attribution following this appraisal process*. This is again a significant research problem. In particular, having a mechanism able to determine the final optimal emotional state is a highly desirable feature for intelligent systems interacting with humans, such as social robots (Williams, 2012), since this is a necessary skill for being proficient in emotional intelligence (Mayer & Salovey, 1993). For example, it has been widely documented that the appraised emotional state of an individual has direct impact on decision making and action selection (Isen & Means, 1983; Loewenstein & Lerner,

¹Note that there are other appraisal variables proposed by the theory. Describing all the appraisal variables and their meanings is out of the scope of this paper.

2003). Thus, without an appropriate mechanism able to determine the final optimal emotional state, the intelligent system cannot take socially acceptable and ethical actions (Vitale, Williams, & Johnston, 2014).

This paper aims to present a computational model of emotion processing that adds a higher layer of cognition to appraisal mechanism. The significance of this paper is further increased by this novel approach going beyond the domain of emotion theories and embracing the strengths of ethical theories. Although, the literature includes previous studies suggesting interactions between theories of emotions and ethics (Callahan, 1988; Gaudine & Thorne, 2001), to our knowledge, there are no computational explanations addressing the interactions between ethics and emotion processing mechanisms (Ojha & Williams, 2016). Therefore, in this paper we aim to:

- (i) Provide a computational model of cognitive appraisal of emotion able to elicit appropriate emotional states without the need of defining pre-determined rules, but rather by using a general domain-independent approach facilitating easy extensibility of the emotionally intelligent systems;
- (ii) Provide a novel computational process inspired by ethical theories for the selection of the optimal emotional state among the elicited ones.

The offered outcomes will provide valuable insights to gather a better understanding on how integrating ethical theories in emotion processing mechanisms can improve existing computational models of emotions. This in turn will advance the understanding of the role of cognition in emotion.

Computational Models of Cognitive Appraisal

Theories from cognitive science and psychology have been implemented in various computational models of cognition. In this section, we will present some of the computational models of emotions implementing cognitive appraisal theory of emotion that are related to our discussion and identify their current limitations.

The models available in literature use evaluation criteria called appraisal variables (Ortony et al., 1990; Lazarus, 1991; Roseman et al., 1990; Scherer, 2001) to assess or evaluate the events for the generation of emotion. The choice of appraisal variables depends on the appraisal theory used and also on the application of the model. One common limitation of the existing accounts is their heavy specificity to the considered application domain and the determination of the elicited emotional states by means of pre-defined rules (Dias & Paiva, 2005; El-Nasr et al., 2000). This approach likely leads to low extensibility of the system.

One available account is Fuzzy Logic Adaptive Model of Emotions (FLAME), a fuzzy logic based computational model of emotion (El-Nasr et al., 2000) inspired by appraisal theories suggested by Ortony et al. (1990) and Roseman et al. (1990). The main strategy used by FLAME is the evaluation of *if-then* rules in order to assess the considered appraisal

variables. As we already discussed, this approach leads to a particularly poor extensibility of the system, since adding a new rule would require to consequently revise and adapt the entire knowledge base.

EMotion and Adaptation (EMA) (Gratch & Marsella, 2004; S. C. Marsella & Gratch, 2009) borrows the ideas from the cognitive motivational appraisal theory of Lazarus (1991). It stands out from other existing computational models of emotion in that it is able to compute emotions irrespective of the experiment domain. However, this model is not able to achieve this only by using the perceived positivity or negativity of an event like our model, which will be discussed later.

Another related account is Fearnot AffecTive Mind Architecture (FATiMA), a computational model of emotion proposed by Dias and Paiva (2005). It is significantly inspired by appraisal theory of Ortony et al. (1990). FATiMA considerably uses domain specific scenarios built on top of pre-defined appraisal rules in order to appraise the desired situation without clearly suggesting how to easily generalise the proposed appraisal mechanisms for different domains.

Beside not providing a valid and easy strategy to integrate the suggested computational model in disparate application domains, the available accounts do not offer an effective way to determine the final emotional state in response to an event in a specific situation. This is still an open research problem since most appraisal theories do not explain how this can be achieved (see, for example, Ortony et al. (1990); Scherer (2001)). Some strategies propose to select the emotional state exhibiting (i) *highest intensity* (Gratch & Marsella, 2004) or driven by the higher motivational state (i.e. hunger, thirst, pain, and fatigue) (El-Nasr et al., 2000), whereas other strategies propose to (ii) *blend the elicited intensities* of multiple emotions in order to determine the final emotional state (see Reilly (2006) for more details on the strategy used). In the Evaluation section, we shall discuss why these approaches might not be desirable methods to reach to a final emotional state.

As previously discussed, an emotion processing model developed by using a rule-based approach is unlikely to offer easy extensibility and high integrability in disparate emotionally intelligent systems among different application domains. Thus, in this paper we provide computational mechanisms general enough to be used in different domains without the need of re-implementing or adapting the proposed model, but at the same time able to appraise the appropriate emotional states for the considered situation. In addition, we suggest to use ethical theories to determine the final emotional state among the ones elicited by the cognitive appraisal process. Determining this state is particularly important to drive socially acceptable behaviours.

Hypotheses

Consider a social interaction between two subjects. In this work we will call *sender* the subject producing a behavioural

response directed to the other subject, which we will call *receiver*. Denote with $S_{receiver}^{(B,C)}$ a value determining how negative or positive the behaviour \mathbf{B} of the sender is perceived by the receiver in a given context \mathbf{C} . $S_{receiver}^{(B,C)}$ is a plausible computational representation summarising within a single valanced value the *somato-visceral reactions of the body* to the given situation (\mathbf{B}, \mathbf{C}) following the first-order phenomenological stage of emotional processing (Bechara, Damasio, & Damasio, 2000). As previously mentioned in the introduction of this paper, this work is not concerned with discussing the implementation of first-order phenomenological processes.

Denote with $C(S_{receiver}^{(B,C)})$ a cognitive appraisal process able to appraise the intensities $\mathbf{I} = \{i_{e_1}, \dots, i_{e_n}\}$ of a set of n considered emotional states $\{e_1, \dots, e_n\}$ given the first-order phenomenological reaction of the receiver $S_{receiver}^{(B,C)}$. Thus, our first hypothesis is that:

Hypothesis 1 *The value $S_{receiver}^{(B,C)}$ is a sufficient information to perform a cognitive appraisal process C able to elicit the intensities of the considered emotional states and consequently promoting the selection of a final emotional state resembling human cognitive appraisal.*

Importantly, the value $S_{receiver}^{(B,C)}$ is completely independent from other pre-existing values S already available by the system and concerning different behaviours and contexts. In other words, adding a new value S to our model, thus extending the knowledge of the system, will not require to adapt the pre-existing knowledge and it will not necessitate to modify the parameters of the computational model.

Denote with $\mathcal{E}(\mathbf{I}, \theta^{ethics})$ and with $\mathcal{E}(\mathbf{I})$ two processes able to provide a final emotional state given the set of the elicited emotion intensities \mathbf{I} realised by the cognitive appraisal process C . $\mathcal{E}(\mathbf{I}, \theta^{ethics})$ includes parameters operationalising ethical theories, whereas $\mathcal{E}(\mathbf{I})$ uses a generic strategy without considering the ethical dimension of the given situation. Therefore, our second hypothesis is that:

Hypothesis 2 *The cognitive appraisal process augmented by ethical reasoning mechanisms $\mathcal{E}(\mathbf{I}, \theta^{ethics})$ converge to more accurate emotional states compared to cognitive appraisal processes augmented by generic reasoning mechanisms $\mathcal{E}(\mathbf{I})$.*

In the remainder of this paper we will offer the functional level description of our computational account and experimental results validating our hypotheses.

Model Implementation

Cognitive Appraisal Process. The process of emotion generation in our computational model² is shown in Figure 1. As mentioned earlier, when an event occurs, its appraisal (evaluation) is done by using a set of variables called appraisal variables. Ortony et al. (1990) state that these appraisal variables are computed based on the goals, standards and attitudes of

the individual. In the context of our computational model of emotion, if we denote these goals, standards and attitudes as an internal parameter θ^{int} and the perceived knowledge of the environment that the system receives when an event occurs as K^{env} , then a function for computing appraisal variable can be represented as:

$$v_i = \mathcal{V}_i(K^{env}, \theta^{int}) \quad (1)$$

Which means that the quantitative value of an appraisal variable is the function of the event knowledge gathered from the environment (K^{env}) and the internal goals, standards and attitudes (θ^{int}). This computation is done by an Appraisal Mechanism component, as shown in Figure 1. Each computed appraisal variable contributes in the generation of one or more emotions (Ortony et al., 1990) and helps in estimating the intensities of the considered emotions³.

The majority of available computational models of emotion compute v_i by using domain-specific rule-based functions (Dias & Paiva, 2005; El-Nasr et al., 2000; Velasquez, 1997). Because of this, when the application domain or input parameters (K^{env}) change in those models, the internal representation of goals, standards and attitudes (θ^{int}) also needs to be changed. In our model, K^{env} is modelled as a set of valanced scores S providing an interpretation of the negative or positive connotation of the experienced events. Importantly, the scores S are completely independent from θ^{int} . Thus, extending our model with new knowledge or adapting previous one will not necessitate to modify the model's parameters θ^{int} . In this paper we will not provide implementation details and we consequently limit our contribution to this functional description, since this is sufficient for the validation of the proposed hypotheses. The detailed mechanism of computation of appraisal variables in our computational model can be found in another paper (Ojha & Williams, 2017).



Figure 1: General Appraisal Mechanism.

Emotional State Selection. Next crucial step is determining the final emotional state of the system. Our proposition is that when more than one active emotions are generated, then a final emotional state is best determined by a higher cognitive layer of ethical reasoning.

Figure 2 shows more details of the Emotion Combination Mechanism included in Figure 1 and suggests the mechanism to determine the final emotional state for attribution. The emotions e_1, e_2, e_3, \dots and e_n with respective intensities $i_{e_1},$

²Our computational model is inspired by the work of S. Marsella, Gratch, and Petta (2010) but implemented with completely different computation mechanism.

³Currently, our model can generate and express eight different emotions described by (Ortony et al., 1990) namely Joy, Distress, Appreciation, Reproach, Gratitude, Anger, Liking and Disliking.




Figure 2: Ethical Emotion Combination Mechanism.

i_{e_2}, i_{e_3}, \dots and i_{e_n} output from the Affect Generation component are processed by applying the concepts of *deontological* and *consequentialist* ethics (Hooker, 1996) in order to determine the final emotional state. Deontological ethics says that one should satisfy one's duties before making a choice of action/decision and consequentialist ethics says that one should consider the consequences to all the relevant parties before making a decision (Hooker, 1996). Functionally, our ethical emotion combination mechanism is shown in 2.

$$e_{ethical} = \mathcal{E}(\mathbf{I}, \theta^{ethics}) \quad (2)$$

Where, $e_{ethical} \in \{e_1, e_2, e_3, \dots, e_n\}$ is the final emotional state. \mathbf{I} is the set of emotion intensities and θ^{ethics} represents ethical standards.

$$e_{high} = \mathcal{E}_{high}(\mathbf{I}) \quad (3)$$

$$e_{blended} = \mathcal{E}_{blended}(\mathbf{I}) \quad (4)$$

Equations 3 and 4 represent the functions computing respectively the final emotional states for highest intensity approach and blended intensity approach, which were introduced earlier. Clearly, these functions only take the intensities of various emotions for the determination of the final emotional state. However, our model reaches to a final emotional state with the help of higher cognitive mechanism of ethical reasoning (as shown in 2).

The emotional responses of our computational model based on: (1) Highest Intensity Approach, (2) Intensity Blending Approach and (3) Ethical Reasoning Approach will be compared with emotion data obtained from human participants in the Evaluation section.

Evaluation

In order to operationalise our model and to consequently validate our hypotheses, we designed two sets of web-based surveys requiring two tasks: an action scoring task and a mind-reading task. In both the experimental conditions we provided a set of stories concerning social exchanges between two individuals, a sender and a receiver, as previously denoted.

Participants covering a broad set of countries were invited on Facebook or through mailing lists to take our surveys. The surveys were completely anonymous. We received a total of

153 responses (male=82, female=71). Importantly, the subjects were randomly attributed to either the action scoring task or the mind-reading task.

Scenario Design

In order to avoid ad-hoc scenarios facilitating our model, we did not design the scenarios ourselves. Rather, we requested 4 naïve adults, without any knowledge about the objectives of the present research, to cooperate in designing six scenarios under the following conditions:

- The scenario shall include the interactions of two subjects, one of them denoted as sender and the other as receiver;
- A minimum of 5 and a maximum of 10 actions of the sender directed to the receiver describing a plausible social interaction between two persons shall be provided;
- At the beginning, each scenario shall provide the contextual information about the designed situation and the two considered subjects. Moreover, additional contextual information could be provided during the development of the described social exchanges, whenever this information is necessary to contextualise the remaining interactions;
- No contextual information suggesting the potential emotional state of the receiver shall be provided for individual interactions, with the exception of the contextual information provided at the beginning of the scenario.

The result of this process was a set of scenarios used during both the action scoring and the mind-reading tasks mentioned earlier. The scenarios included interactions between (1) two strangers (a male and a female) interacting on a bench of a park, (2) two close friends (both males) meeting at a beach, (3) a husband and a wife having an argument about forgetting the birthday, (4) an elderly woman affected by dementia and her nurse (both females) experiencing a distressful moment, (5) a guy having argument with his brother, and (6) an interaction between a customer of a café and a waiter (both males). In total, the scenarios included 48 social exchanges of the senders directed to the receivers.

Action Scoring Task

The experimental subjects participating in the action scoring task were asked to guess, for each scenario, *how positive or negative each social exchange performed by the sender would be perceived by the receiver in that specific context*. The rating was based on 7-point Likert scale: Extremely Negative, Very Negative, Negative, Neither Negative Nor Positive, Positive, Very Positive, Extremely Positive. We numerically evaluated the responses by attributing a weight to each point of the scale (*i.e.* -1, -0.66, -0.33, 0, 0.33, 0.66 and 1 respectively). We averaged the responses obtaining a value \mathcal{S} for each of the considered social exchanges \mathbf{B} in the specific context \mathbf{C} . In this way we were able to provide the




Figure 3: Results of the experiments. (a) The cumulative rank-distances of the models’ predictions from human assessment. (b) The rank-distances of the models’ predictions from human assessment.

necessary input knowledge to our system (*i.e.* a set of numeric scores $\mathcal{S} \in [-1, +1]$) and to consequently perform cognitive appraisal processes estimating the emotional state of the receivers in each considered scenario. Recall that this process did not require any changes to our computational model, which provides a valid domain-independent approach of cognitive appraisal process.

Mind-Reading Task

In order to compare the emotional response of our computational model, we asked to the subjects participating in the mind-reading task to guess, for each interaction of the sender, *what would have been the chances that the receiver would happen to be in a particular emotional state, based on the just happened interaction and the previously occurred social exchanges and contextual information.* Therefore, for each of the eight considered emotional states the rating was based on 6-point Likert scale: Not at all, Very Low, Low, Medium, High and Very High. The additional rating “Not at all” was necessary to allow the participants to express no chances to attribute such emotional state to the receiver. We numerically evaluated the responses by attributing a weight to each point of the scale (*i.e.* 0.0, 0.2, 0.4, 0.6, 0.8 and 1.0 respectively). Average score given by the participants to various emotions was calculated by performing the weighted average of the ratings.

Results

Based on the results of the mind-reading task, the emotions for each interaction of each scenario were ranked from 1 to 8, with the emotion having the highest average score ranked as 1 and the one with lowest score ranked as 8.

We considered three strategies to computationally predict the final emotional state of each interaction: choosing the emotional state with highest intensity, blending the emotional intensities to determine the final emotional state as described

by Reilly (2006), and our suggested approach based on ethical reasoning. Each of these three strategies followed a common domain-independent cognitive appraisal process, as discussed in the model implementation section. We compared these computational predictions against the gathered human assessments (*i.e.* emotions ranked 1) by computing their rank-distances, suggesting how close the computational model was compared to human assessment. The results are summarised in Table 1.

Table 1: Descriptive statistics of the gathered results.

| | Mean | Median | Std |
|--------------------------|--------|--------|--------|
| High intensity | 2.4167 | 2 | 2.3232 |
| Blended emotion | 2.3125 | 2 | 2.0228 |
| Ethical reasoning | 2.0833 | 1 | 2.3140 |

In order to demonstrate that the proposed common stage of domain-independent cognitive appraisal was able to elicit emotional intensities similarly to human cognitive appraisal process (Hypothesis 1), we analysed the human responses of the mind-reading task. We noticed that for most of the considered interactions some of the emotions resulted in very close averaged scores. Therefore, given $\epsilon = 0.1$, for each interaction we counted the number of emotions having an average score of greater than or equal to the score of highest scored emotion minus ϵ for that interaction. ϵ was chosen to be equal to half of the score attributed to each point of the Likert scale (*i.e.* 0.2), thus being able to group emotions plausibly ranked with similar likelihood by most of the human assessors. The average number of similarly rated emotional states among all the 48 interactions was 3.2, thus suggesting that on average human cognitive appraisal promoted 3 comparable emotional states to attribute to the receiver. From Figure 3a it is clear to see that for distances less than 2 ranks to the human assessment (*i.e.* predictions among the first 3 higher scored emo-

tions) our cognitive appraisal model was able to promote the selection (using all the three considered strategies) of approximately 70% of the emotional states plausibly attributed by humans participants to the receivers described in the mind-reading task scenarios.

In addition, we can also observe that the cognitive process augmented by the proposed ethical reasoning mechanism converges to more accurate emotional states compared to the other investigated strategies (Hypothesis 2). Figure 3b further suggests that the proposed ethical reasoning mechanism reduces average rank-distances from human appraisal. Therefore, the present results support both the proposed hypotheses.

Conclusion and Future Work

In this paper, we presented our computational model of emotion based on appraisal theory that is able to generate emotions using the expected degree of positivity or negativity associated with an action/event. This allowed our model to be completely independent of the application domain and efficiently appraise a situation for the elicitation of various emotions. In addition, our model adds a higher layer of cognition in the emotion mechanism by integrating an ethical reasoning capability for the determination of the final emotional state when more than one emotions are generated by the model. Experimental results support our first hypothesis proposing that cognitive appraisal is possible without prior domain knowledge and second hypothesis suggesting that ethical reasoning is a better strategy to explain human emotional state attribution process.

Yet, our computational model still has some room for improvement. For example, it is important to consider that people with different personality generate emotions in different ways. Thus, in the future, we aim to use the concept of personality and examine how the difference in personality makes difference in ethical standards and hence in emotion generation.

References

Bechara, A., Damasio, H., & Damasio, A. R. (2000). Emotion, decision making and the orbitofrontal cortex. *Cerebral cortex*, 10(3), 295–307.

Callahan, S. (1988). The role of emotion in ethical decision-making. *Hastings Center Report*, 18(3), 9-14.

Dias, J., & Paiva, A. (2005). Feeling and reasoning: A computational model for emotional characters. In *Progress in artificial intelligence* (p. 127-140). Springer.

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

Gaudine, A., & Thorne, L. (2001). Emotion and ethical decision-making in organizations. *Journal of Business Ethics*, 31(2), 175-187.

Gratch, J., & Marsella, S. (2004). A domain-independent framework for modeling emotion. *Cognitive Systems Research*, 5(4), 269-306.

Hooker, J. (1996). *Three kinds of ethics*.

Isen, A. M., & Means, B. (1983). The influence of positive affect on decision-making strategy. *Social cognition*, 2(1), 18-31.

Lambie, J. A., & Marcel, A. J. (2002). Consciousness and the varieties of emotion experience: a theoretical framework. *Psychological review*, 109(2), 219.

Lazarus, R. S. (1991). *Emotion and adaptation*. New York: Oxford University Press.

Loewenstein, G., & Lerner, J. S. (2003). The role of affect in decision making. *Handbook of affective science*, 619(642), 3.

Marsella, S., Gratch, J., & Petta, P. (2010). Computational models of emotion. In *A blueprint for an affectively competent agent: Cross-fertilization between emotion psychology, affective neuroscience, and affective computing* (p. 21-46). Oxford: Oxford University Press.

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

Mayer, J. D., & Salovey, P. (1993). The intelligence of emotional intelligence. *Intelligence*, 17(4), 433–442.

Ojha, S., & Williams, M.-A. (2016). Ethically-guided emotional responses for social robots: Should i be angry? In *International conference on social robotics* (p. 233-242). Springer.

Ojha, S., & Williams, M.-A. (2017). Emotional appraisal : A computational perspective. In *Annual conference on advances in cognitive systems*.

Ortony, A., Clore, G. L., & Collins, A. (1990). *The cognitive structure of emotions*. Cambridge University Press.

Reilly, W. S. N. (2006). Modeling what happens between emotional antecedents and emotional consequents. In *Eighteenth european meeting on cybernetics and systems research* (p. 607-612). Vienna, Austria: Austrian Society for Cybernetic Studies.

Roseman, I. J., Spindel, M. S., & Jose, P. E. (1990). Appraisals of emotion-eliciting events: Testing a theory of discrete emotions. *Journal of Personality and Social Psychology*, 59(5), 899.

Scherer, K. R. (2001). Appraisal considered as a process of multilevel sequential checking. *Appraisal processes in emotion: Theory, methods, research*, 92, 120.

Velasquez, J. D. (1997). Modeling emotions and other motivations in synthetic agents. In *Aaai/iaai* (p. 10-15).

Vitale, J., Williams, M.-A., & Johnston, B. (2014, October). Socially impaired robots: Human social disorders and robots' socio-emotional intelligence. In *6th international conference on social robotics* (p. 350-359).

Vitale, J., Williams, M.-A., Johnston, B., & Boccignone, G. (2014). Affective facial expression processing via simulation: A probabilistic model. *Biologically Inspired Cognitive Architectures*, 10, 30-41.

Williams, M.-A. (2012). Robot social intelligence. In *Social robotics* (pp. 45–55). Springer.