
Am I Really Scared? A Multi-phase Computational Model of Emotions

Jason R. Wilson

JASONWILSON2017@U.NORTHWESTERN.EDU

Kenneth D. Forbus

FORBUS@NORTHWESTERN.EDU

Matthew D. McLure

MCLURE@U.NORTHWESTERN.EDU

Qualitative Reasoning Group, 2133 Sheridan Road, Evanston, IL 60208 USA

Abstract

We present a *multi-phase model of emotions* that integrates automatic and fast affective reactions with slower cognitive appraisals to create a final emotional appraisal that is grounded in previous emotional experiences. *Affective reactions* are the product of an analogical comparison between the current situation and similar cases in memory. This is followed by a *cognitive appraisal* to do a deeper affective appraisal of the situation. Finally, a *retrospective appraisal* of experiences generates affective memories that are available for future situations. We demonstrate this model on a problem solving task and show that the inclusion of affective reactions is successful in beneficially altering the selection of strategies in a problem solving task and captures broad patterns of emotional dynamics found in the literature.

1. Introduction

You are driving along a scenic highway and stop at one of the many scenic rest stops. You get out and approach a cliff overlooking a picturesque valley with a quaint river running through a grove of old trees. Standing on the edge of a cliff, an immediate feeling of dizziness occurs, and you become afraid of falling and getting severely hurt. Then you realize that you are holding onto a guardrail, likely preventing any fall or pending doom. The feelings of fear shift to that of happiness as you realize you have security and enjoy the pleasant surroundings below you.

Most models of emotions focus on the conscious evaluation, or appraisal, of a situation or an event. These models do not account for the seemingly illogical emotions that may arise. They ignore emotions that are automatic and happen before any conscious thought. Additionally, these models cannot demonstrate how an initial emotion felt in a situation can be overridden once a conscious evaluation of the situation has been done. A separation of affective reactions and a conscious cognitive appraisal is necessary to demonstrate these phenomena.

Similar emotional events are felt by students while working on solving problems. Upon initially reading a problem, a fear might arise because the problem seems hard and unfamiliar. Then they are reminded of a similar problem that was solved. Relief comes over the student, who is now hopeful that this problem, too, can be solved. The work described in this paper captures the emotional dynamics that occur when a computer simulated student attempts to solve various physics problems. Previously, Klenk & Forbus demonstrated the ability of a Companion to solve physics problems using analogical reasoning (Klenk & Forbus, 2009). Using that system as a starting point, we extend it with a new *multi-phase model of emotions*, and show how the

system's problem solving efforts are influenced by the emotional dynamics occurring during the task. These emotions include both the reactive emotions, *affective reactions*, and the emotions arising from conscious thought, *cognitive appraisals*.

The experiments described in this paper have been conducted using the Companions architecture (Forbus et al. 2009), which was also used for the previous transfer learning research. We extend it via ideas from a computational model of cognitive appraisal, namely EMA (Gratch & Marsella, 2004; Marsella & Gratch; 2009). The Companions architecture includes the FIRE reasoning engine (Forbus et al. 2010), which provides reasoning capabilities used to re-implement the rules used in EMA. FIRE also incorporates SME (Forbus et al., 1994), which performs analogical mapping, and MAC/FAC (Forbus et al., 1995), a model of analogical retrieval. In our model, these analogical reasoning capabilities are used to rapidly retrieve previous experiences and construct an initial automatic affective reaction via analogical projection of emotions from those previous experiences. First-principles reasoning using rules from EMA and other sources provides the cognitive appraisal. These two sources of information are then reconciled by identifying conflicts, re-evaluating key appraisal variables, and completing a final cognitive appraisal.

We begin by describing prior research on emotions in Section 2. Section 3 describes our three-phase model of emotional processing in detail, including what emotions and their effects on cognition that we model, and how affective memories are generated from the system's experiences. Section 4 describes a simulation experiment that shows the multi-phase model improves problem-solving performance, and Section 5 compares its emotional dynamics with patterns from the literature. Finally, we discuss future work.

2. Background

Affective reactions and cognitive appraisals have been explored extensively. There are also many sources in the literature that describe a dual-process system. The most relevant work is summarized below.

2.1 Affective Reaction

The inclusion of affective reactions in describing emotions is far from a new concept in psychology. Zajonc (1980) presents evidence that affective reactions are primary and often independent from cognition. Subjects in many experiments were able to judge an object positively or negatively even without being able to recognize it. He suggests that there are preferences that are too vague and ambiguous for immediate cognitive processing (e.g. recognition) but are sufficient for producing affective reactions. The input used for affective reactions in our model is not the result of any inferences and is purely the result of environmental states. The same base information is used for cognitive appraisals but additional inferences are done. Additionally, further reasoning is necessary to complete the cognitive task at hand.

There are varying views on the level of complexity in an affective reaction. One perspective is that it is simply a valenced reaction to basic stimuli (Ortony et al., 1990; Ortony et al., 2005). Others argue that "automatic affective responses ... may contain information that is useful enough to alter subsequent cognitions and behavior." (Baumeister et al., 2007). It is this latter view that is the basis of our model.

Neuroscience researchers have also found evidence for faster, affective responses. LeDoux (1986) describes two possible pathways to creating an emotional response from stimuli. The high-road passes through the thalamus, to the cortex, and finally the amygdala. The low-road, however, skips the cortex and is able to create a significantly faster response. Schaefer et al. investigate the neural correlates of affective responses, or “hot” emotions. They describe a schema - an abstract representation of sensory, perceptual, and semantic information - that gets activated to create the emotional response. Our model similarly uses a pattern of contextual information that gets recalled and directs the emotional response.

Thagard and Shelley describe three classifications of interactions between emotion and analogies, one of which is categorized as “Analogies that Transfer Emotions” (2001). This type of analogy is the focus of the study done by Bliznashki and Kokinov (2009). In these analogies, an object in the analogy has some emotional connotation associated to it, though typically this affective expression is never explicitly stated. Thagard and Shelley show an example using the negative connotation of divorce. Emotions associated to divorce do not need to be made explicit, as the recipient of the analogy would already have this association. Similarly, Bliznashki and Kokinov rely on (and demonstrate) the negative emotions associated with spiders and the positive emotions associated with rabbits. Both works assume that these emotional associations can be treated like any other facts in an analogy, including being transferred.

Despite the similarities of these works, there is a stark contrast in the intended effect of the emotion transfer. Bliznashki and Kokinov show that the emotions associated to one object can alter the emotions associated with another. Specifically, spiders can be judged less negatively after being seen in an analogy with a rabbit, and vice versa. In this scenario, the target object already has an associated emotion. In contrast, Thagard and Shelley use examples for emotional persuasive analogies to show the potential of emotions getting transferred to an object or event that may not already have an emotion clearly associated to it. Instead, the analogy is used to introduce or suggest an emotion appropriate to the target object or event in the analogy. In their example using divorce, the persuasive argument compares it to the separation of Quebec from Canada. The intent is for the negative connotation of divorce to be transferred to the prospect of a Quebec separation from Canada.

The use of affective reactions in our model is more similar to work of Thagard and Shelley. An analogy is formed between a previous situation that has a known emotional association and a new situation that does not yet have an associated emotion. Given a suitable mapping between the situations, emotions may be transferred to the new situation. However, further cognitive processing may contradict or not completely agree with the analogically inferred emotion. As a result, the emotion may not get fully transferred and only altered in the direction of the emotion in the base of the analogy. Thus the reconciliation process proposed in our model provides a possible explanation for the results shown by Bliznashki and Kokinov.

2.2 Cognitive Appraisal

There are numerous existing computational models of emotion appraisal. These models focus on higher-order cognitive processing of stimuli to appraise a situation and determine the applicable emotions. See recent reviews (Marsella & Gratch, 2010; Lin et al. 2012) for a discussion of a wide range of approaches to computational models of emotion.

None of the models reviewed utilize emotions stored in long term memory. Some use an encoding of recent events, stored in short term memory, as a source of information for the

appraisal. For example, EMA uses a causal representation of recent events to encode memory, which is used for evaluating appraisal variables like how to attribute a cause to the current situation (Marsella & Gratch, 2006). SoarEmote uses mood to affect appraisals, where mood is simply a summary of recent situations (Marinier & Laird, 2007).

Some of the rules for cognitive appraisal used in EMA are the basis for the rules used in the present work. EMA models seven emotions (joy, hope, sadness, fear, guilt, anger, and surprise) and also includes a model of coping mechanisms to alter behavior based on the appraised emotions (Gratch & Marsella, 2004; Marsella & Gratch, 2009). To compute the emotion appraisal it creates a series of appraisal frames for each proposition to be appraised. A set of appraisal variables are evaluated with respect to each frame. These variables are then used by the rules for each of the emotions to infer which emotions are true for the given proposition and frame.

2.3 Multiple Processes

A similar model to the one proposed here comes from the literature on attitude change. The associative-propositional evaluation (APE) model combines two types of reasoning to do an evaluation (Gawronski & Bodenhausen, 2006). The associative processes are the basis of implicit attitudes, and evaluations using associative processes are characterized as automatic affective reactions. The propositional processes use truth statements about a given situation to construct judgments. The propositional processes are the basis for explicit attitudes. Additionally, the propositional processes are responsible for maintaining cognitive consistency. Their two process model is similar to ours, except we use a model of analogy instead of association to create the affective reaction. Our models of cognitive appraisal provide much of the propositional reasoning. Furthermore, the appraisal rules used in our model are involved in reconciling differences between the affective reaction and the conscious cognitive appraisal.

Our model is also similar to the dual emotional processes described by Baumeister and colleagues. They make a strong argument for the need for processes for automatic affective reactions that are simple and rapid, and conscious evaluation of emotions using more cognitive processes (Baumeister et al., 2007). They are primarily interested in the role of emotions in shaping behavior and provide evidence for conscious emotion as a feedback mechanism that promotes learning, such as anticipatory emotions causing behavior that avoids potential negative consequences. Our model also uses emotional memories of past problem solving attempts and applies them to new situations. Moreover, one of the behaviors modeled in the simulations is avoidance of steps that have previously been associated with negative emotions.

Another similar model describes the multiple processes as reactive, routine, and reflective levels of information processing (Ortony et al., 2005). Though affect is present at all levels, only “proto-affect”, i.e., a simple positive/negative signal, is available at the lowest level. The middle level, routine, uses perceptual input to reason about emotions and is able to represent past, present, and future emotions. The reactive level of Ortony, et al. is perhaps most similar to the primary processes of Panksepp, and it also is similar to the vague negative or positive reactions to stimuli as described by Zajonc. As mentioned above, this basic valence-only content of proto-affect differs from our model of reactive emotions; we utilize a broader range of emotional content in affective reactions that allow for separate emotions of similar valence to be produced (e.g. fear and sadness).

The dual-process model of Kahneman (2011) divides information processing into fast, unconscious reasoning (System 1) and slower, conscious reasoning (System 2). System 1 is characterized by judgments based on intuition, effortless and automatic operation, and influenced by memories and emotions. System 2 is characterized by logical reasoning and fact-based rule following. Our model of affective reactions is a System 1 operation. Emotions are usually reserved for System 1, but our model of cognitive appraisal uses logical reasoning incorporating facts and rules. The emotions arising from System 2 are consciously developed, no different than any information inferred by other System 2 processes.

The idea of multiple emotional processes can be found in the neuroscience literature as well. A dual memory model of schema and propositional reasoning has been investigated to find their neural correlates (Schaefer et al., 2003). “Hot” emotions rely on schema activation and are associated with the ventromedial prefrontal cortex. Propositional processing handles “cold” emotions and is correlated with the lateral prefrontal cortex. LeDoux (1996) has the notion of a high-road and a low-road, where the high-road passes through the cortex and is able to do higher-order reasoning but the low-road bypasses the cortex and produces faster responses. Panksepp (1998) describes three sets of processes: primary (innate processes involved sensory processing, body homeostasis, and primary emotions), secondary (learning and memory processes), and tertiary (higher-order cognitions, including social emotions, that occur in the cortex). He focuses on the primary processes as he investigates the neurological basis for innate emotions in animals. These innate emotions form the basis for more complex emotions that are involved in the secondary and tertiary processes. The affective reactions described in our model map to the secondary processes, and the cognitive appraisals are best associated with the tertiary processes.

The computational model of emotions that comes closest to demonstrating affective reactions and a multi-process appraisal approach is WASABI (Becker-Asano & Wachsmuth, 2009). It creates appraisals of primary and secondary emotions. The primary emotions, which they call “core affect,” are similar to the affective reactions in our model. It is an automatic and non-conscious appraisal creating emotional impulses. Secondary emotions, on the other hand, involve a conscious appraisal. In contrast to our model, primary emotions in WASABI are modeled using the dimensional PAD model, which associates emotions to a space along the dimensions of Pleasure-Arousal-Dominance (Russell & Mehrabian, 1977). However, the secondary emotions in WASABI are the prospect-based emotions of hope, fears-confirmed, and relief. This use of a taxonomy of emotions is more similar to our model, and we believe this to be an important similarity due to the contrasting effects of emotions that are similar along one dimension (Lerner & Keltner, 2001). Also like our model, the primary and secondary emotions in WASABI go through a final reappraisal and end with different coping behaviors.

3. Three-Phase Emotional Processing

Our model consists of three phases: affective reaction, cognitive appraisal, and reconciliation. We describe each in turn. The input is a current event. For example, in the simulations below, the event is the current state of work in solving a physics problem.

At each step in the solving of the physics problem, the entire process of determining emotions is conducted. Emotions are calculated for both the overall problem and for the current problem-solving step. Changes in the emotions relating to the overall problem tend to change slowly and often do not reveal any interesting behavior. The exceptions are when solving the problem

becomes too boring and increases in challenge. Otherwise, the agent typically remains hopeful of solving the problem until it either succeeds (joy) or fails (sad).

Each problem solving step is also assessed. Before the step is started, the agent assesses emotions related to this step. Here we get drastically varying emotions from step to step. Some steps are simply more costly or have a lower likelihood of succeeding. These steps result in fear and/or guilt, which need to be coped with.

3.1 Affective Reaction

The initial phase models the emotions that arise in an affective reaction. Our model of affective reactions uses analogical reasoning. Analogy is an ideal mechanism for affective reactions because it has been shown to be effective in transfer of emotions (Bliznashki & Kokinov, 2009) and other information (Gentner & Markman, 1997; Hummel & Holyoak, 1997).

Given the current event, MAC/FAC is used to retrieve previous events that are similar. Previous events are generated by prior runs of the system itself, and are stored in a case library for later retrieval. Each case is a step from a previous problem, and it includes both the domain content and emotional content. The emotional content includes emotions that were felt and the justifications for them, in terms of the values of appraisal variables. (Importantly, we do not assume that these justifications are inspectable by processes other than the ones described here, i.e. a Companion is not able to “drill down” into the causes of its emotions as a problem-solving operation.) MAC/FAC retrieves from the case library a set of potential matches, which are then analyzed by SME. The analogical processing of SME assesses the structural similarities of the

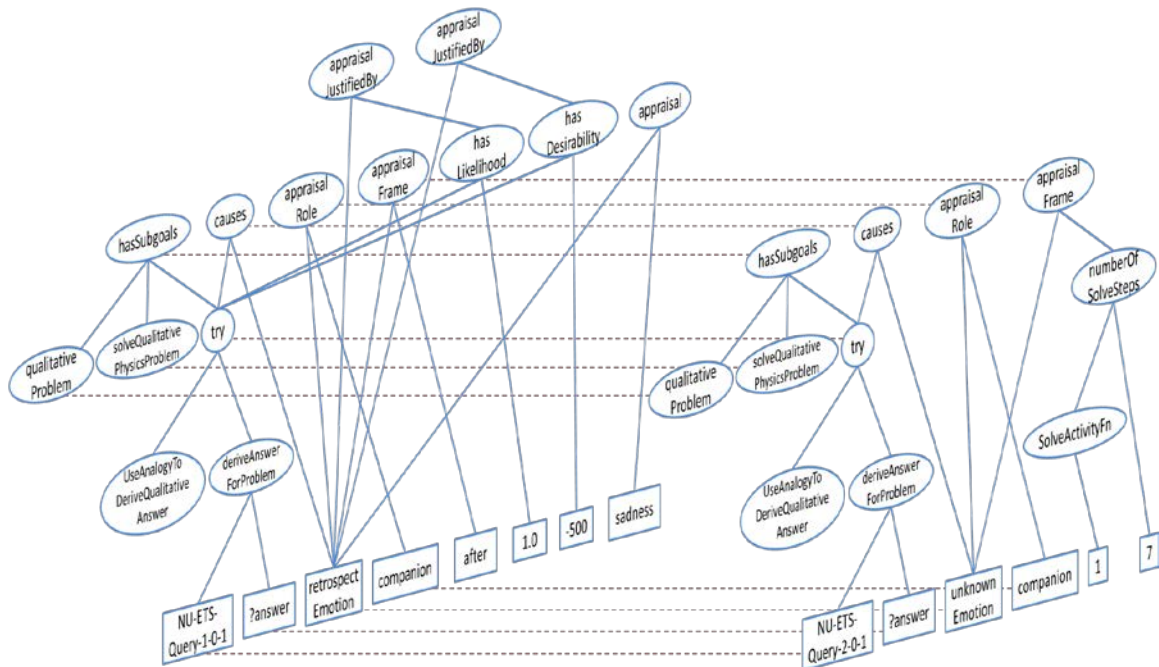


Figure 1: Retrieved events are mapped to the current situation. The graph on the left is an affective memory of a step. The graph on the right represents the next proposed step in the problem solving. Red dashed lines indicate correspondences.

events and produces one or more *mappings*. A mapping consists of a set of correspondences between the base and target (here, the retrieved and current event, respectively), a numerical similarity score, and a (possibly empty) set of *candidate inferences*. Based on the mapping between the base (retrieved event, left side in Figure 1) and target (current event, right side in Figure 1), the candidate inferences constitute surmises about the target based on projecting information from the base onto it. Finally, the similarity scores determine which of the retrieved events are closest, and up to three of these reminders are returned by MAC/FAC. Our model then iterates over the set of candidate inferences to find any that include emotional content. These inferences represent emotions from past events projected onto the current event.

The probe event - the current situation - has incomplete information because it contains only the immediately observable data. It does not include any information that would be the result of higher-level reasoning. In the context of solving physics problems, the immediate observables include the cost of the next step, the form of the next step, and the number of steps taken thus far. It does not include appraisal variables such as likelihood or desirability. Furthermore, it does not include any inferred emotion. If the reminding has content for the appraisal variables and the resulting emotion, then a mapping between the structures of the situations leads to a candidate inference for the probe being generated. These candidates suggest that the information present in the reminding may be true in the current situation. Assuming that they are true generates an initial affective reaction before any cognitive appraisal is conducted.

3.2 Cognitive Appraisal

The second phase of the simulation constructs a cognitive appraisal. Like EMA, we view this process as being involuntary, and perhaps even proceeding in parallel with the first phase, but we assume that it is slower than that memory-based reactive process. This is not in contradiction to Zajonc's affective reaction primacy because the reaction process is faster and is able to utilize less, and more ambiguous, information. Higher-order cognitive processes are likely occurring in parallel to acquire more detailed information that can be used to make these cognitive appraisals.

At each step in the problem solving task, the set of emotions for that step are determined by using rules that define the emotions (see Table 1 below). The reasoning is done in a context associated with that step, providing a localization equivalent to an appraisal frame. These rules are implemented via Horn clauses that ground out in appraisal variables. Appraisal variables, in turn, are computed based on further reasoning. For example, each step in the solving of a physics problem has associated with it a rough estimate of its cost. The cost estimate is also part of the function to determine desirability (higher cost tends to be less desirable). Thus this information is not available to the earlier reactive processing, since it needs reasoning to compute.

We currently model a small set of emotions that are particularly relevant to problem-solving. The cognitive appraisal rules for the emotions joy, hope, sadness, and fear have been mostly adapted from EMA (Gratch & Marsella, 2004; Marsella & Gratch, 2009). We extended the definition of sadness to include desirable events with no likelihood of occurring. In the taxonomy defined by Ortony et al. (1988), this is classified as disappointment. For simplicity, we do not distinguish here between sadness and disappointment.

The emotions of guilt and anger that are modeled in EMA are more social emotions and are not as directly applicable to a student solving a physics problem on its own. We introduce models for bored and challenged based on the dimensions defined by Smith & Ellsworth (1985). They defined the key characteristics of boredom to be low effort, high certainty, situational control, and

low attention. The attention dimension seems mostly an effect of boredom and will be addressed later in this paper. We map high certainty to the likelihood appraisal variables used in the rules adapted from EMA. Similarly, we map situational control to low controllability. Low effort required the introduction of a new appraisal variable that estimates the amount of anticipated effort remaining in a task.

Table 1. Rules used for the cognitive appraisal of joy, hope, sadness, fear, challenged, and bored.

Emotion	Rules
Joy	(<== (joy ?agent ?p ?frame) (desirability ?agent ?p ?des ?frame) (greaterThan ?des 0) (likelihood ?agent ?p ?lik ?frame) (equals ?lik 1))
Hope	(<== (hope ?agent ?p ?frame) (desirability ?agent ?p ?des ?frame) (greaterThan ?des 0) (likelihood ?agent ?p ?lik ?frame) (lessThan ?lik 1) (greaterThan ?lik 0))
Sadness	(<== (sadness ?agent ?p ?frame) (desirability ?agent ?p ?des ?frame) (lessThan ?des 0) (likelihood ?agent ?p ?lik ?frame) (equals ?lik 1)) (<== (sadness ?agent ?p ?frame) (desirability ?agent ?p ?des ?frame) (greaterThan ?des 0) (likelihood ?agent ?p ?lik ?frame) (equals ?lik 0))
Fear	(<== (fear ?agent ?p ?frame) (desirability ?agent ?p ?des ?frame) (lessThan ?des 0) (likelihood ?agent ?p ?lik ?frame) (lessThan ?lik 1) (greaterThan ?lik 0))
Challenged	(<== (challenged ?agent ?p ?frame) (desirability ?agent ?p ?des ?frame) (greaterThan ?des 0) (likelihood ?agent ?p ?lik ?frame) (greaterThan ?lik challengeLikelihood) (effort ?agent ?p ?cont ?frame) (greaterThan ?cont challengeEffort))
Bored	(<== (bored ?agent ?p ?frame) (effort ?agent ?p ?des ?frame) (lessThan ?des boredEffort) (likelihood ?agent ?p ?lik ?frame) (greaterThan ?lik boredLikelihood) (controllability ?agent ?p ?cont ?frame) (different ?cont low))

Smith & Ellsworth find that challenged is defined by many of the same dimensions as boredom – effort, certainty, and attention – but on the opposite end of each of these dimensions. Challenged also is associated with desirable goals, though Smith & Ellsworth do not have this as a dimension. They do, however, note that subjects feeling challenged by an entity “strongly desired to attend to it.” Additionally, they note that subjects described desirable goals in relation to challenge. For these reasons, our model of challenged includes the appraisal variable desirability.

3.3 Reconciliation and Estimation

The third phase of processing occurs after both the initial affective reaction and the cognitive appraisal are available. Its purpose is to reconcile them, if needed, and generate one or more final emotions for the current state that will be used to influence domain processing. Suppose the emotions generated by affective reaction and cognitive appraisal match. In that case no further reasoning is necessary, that emotion is considered the final emotion of the situation. However, if the emotions do not match, then they must be reconciled.

Reconciliation begins by analyzing the appraisal variables for each of the emotions generated. Recall that the reminders that generated the affective reaction can include candidate inferences which provide values for some of the appraisal variables, stored in the justifications for the emotion in the memory of the past event. These projected appraisal variable values are compared to the values for appraisal variables found via the cognitive appraisal process. When a mismatch is found, the appraisal variables are re-estimated.

Re-estimation occurs via taking an average of the conflicting values. These include the values projected by all of the reminders (recall that MAC/FAC can produce up to three reminders for each probe) and the value computed by the cognitive appraisal process. For example, if there were three reminders that have likelihood values of 1.0, 0.0, and 0.0, and the cognitive appraisal used an estimated likelihood of 0.8, then the re-estimated likelihood is 0.45.

Once all of the appraisal variables have been re-estimated, a final appraisal is done. This appraisal relies mostly on the values that were just generated, though it may still infer other values where necessary. For example, some appraisal variable values are not available from the affective memories. These variables, such as controllability, are not associated with the retrospective emotions that are stored in memory and thus are not available when it is retrieved. The retrospective emotions currently included in the model are joy and sadness, which are justified by likelihood and desirability. Only the appraisal variables that justify the appraised emotion are stored with the associated memory. Values for the other variables not available in the affective memories are inferred in the identical manner as the original appraisal. The final appraisal generates the emotion(s) that then influences the cognitive behavior in the task at hand. Multiple emotions may occur simultaneously, as is often the case with challenged and hope (defined below).

3.4 Affective memories

A retrospective analysis of each step in the solving of the problem involves conducting another emotional appraisal, based on the success of the step. Thus, a successful step will be retrospectively viewed positively and with joy, whereas a failed step will be remembered with

sadness. This appraisal, along with its justifications and context, are stored as a case in a case library. These cases are then available as a reminding by MAC/FAC on future problem solving tasks. An example of an affective memory and its associated facts are shown below.

```
(appraisalFrame thisEmotion after)
(aoNodeControlStatus (AO-NodeFn 7 1) Closed)
(aoNodeCost (AO-NodeFn 7 1) 210.0)
(appraisal thisEmotion sadness)
(appraisalRole thisEmotion companion)
(causes
  (try UseAnalogyToDeriveQualitativeAnswer
    (deriveAnswerForProblem NU-ETS-Query-2-0-1 ?answer))
  thisEmotion)
(appraisalJustifiedBy thisEmotion
  (hasDesirability
    (try UseAnalogyToDeriveQualitativeAnswer
      (deriveAnswerForProblem NU-ETS-Query-2-0-1 ?answer)) -500))
(appraisalJustifiedBy thisEmotion
  (hasLikelihood
    (try UseAnalogyToDeriveQualitativeAnswer
      (deriveAnswerForProblem NU-ETS-Query-2-0-1 ?answer)) 1.0))
(hasDesirability
  (try UseAnalogyToDeriveQualitativeAnswer
    (deriveAnswerForProblem NU-ETS-Query-2-0-1 ?answer)) -500)
(hasLikelihood
  (try UseAnalogyToDeriveQualitativeAnswer
    (deriveAnswerForProblem NU-ETS-Query-2-0-1 ?answer)) 1.0)
(hasSubgoals
  (try UseAnalogyToDeriveQualitativeAnswer
    (deriveAnswerForProblem NU-ETS-Query-2-0-1 ?answer))
  (TheList (qualitativeProblem NU-ETS-Query-2-0-1)
    (solveQualitativePhysicsProblem NU-ETS-Query-2-0-1 ?answer)))
```

3.5 Cognitive Effects of Emotions

Emotions affect cognition in a variety of ways. Here we discuss the interactions explored in our current simulation. Currently, there are three possible actions for the agent to choose before attempting each step. The default behavior is to proceed in executing the next task. Alternatively, the agent may choose to avoid the next task. Lastly, it may opt to not only execute the next task but increase attention devoted to the task.

If a problem-solving step evokes a feeling of sadness or fear, then it will attempt to avoid that step. Avoidance is a coping mechanism commonly used when an undesirable situation is encountered (Lazarus, 1991). The agent attempts to avoid a step by adding it back onto the queue of steps to try. If there are no other steps or all steps in the queue have been tried, avoidance is no longer possible. Eventually, when no other options are available, the feared step will be attempted. A more sophisticated model of avoidance would determine which of the avoided steps to try, how long to avoid the steps, and whether to completely avoid them by giving up on the

task. Further emotion modeling using intensity of the emotions associated to each of the steps would enhance the capability in choosing which steps to stop avoiding, if any.

An increase in attention is associated with feeling challenged, and a decrease in attention is associated with boredom (Smith & Ellsworth, 1985). While Smith & Ellsworth do not define the attention dimension to be an effect of these emotions, it seems intuitively plausible. The increase of attention has been modeled by extending the resources allocated for the task. Arguably, if resources had been exhausted, the agent would cease performing the task and find a different task. Instead, when feeling challenged, the increase in available resources allows the agent to persevere in a challenging but rewarding task.

4. Simulations

The models developed in this work were tested using eight of the physics problems used by Klenk & Forbus (2009). Attempts to solve each problem were done three times. The first attempt did not use any emotions or effects of emotions. The next attempt in solving used cognitive appraisals and effects of the inferred emotions. The final attempt to solve the problem encompassed the entire multi-phase appraisal by using affective reactions, cognitive appraisals, and reconciliation & estimation. The emotion resulting from the final reconciliation was then used to influence how the Companion proceeded with the next step of the task.

The emotional memories used in generating affective reactions were created at the completion of an attempt at problem solving. The Companion does a retrospective analysis of all of the steps considered and whether each succeeded or not. Each step and its associated context and emotional content are recorded. Experiments were done with memories from one and two previously attempted problems. Due to the similarities in the problems used in the simulation, increasing the number of previously attempted problems (and their associated affective memories) had no interesting effect. Future work that uses a wider variety of problems will need to use more contextual information and more problems to learn from. The additional contextual information will be required to aid in distinguishing between similar steps from different classes of problems. For example, steps associated with looking for a quantitative answer may have negative memories when the class of problem requires a qualitative answer.

Table 2 summarizes the results of the simulations. Some problems originally could not be solved because resource limits were reached. When effects of emotions are allowed (all but the baseline), avoidance of some steps is attempted. The multi-phase emotion processing reduces the average number of avoidances by more than half. Lastly, an increase in resources allocated to the task is made when feeling challenged, which results in some problems being solved that previously could not be solved.

Table 2. Multi-phase emotion alters selected problem solving strategies.

Condition	Solved	Average Avoidances	Increases Resources
Baseline	6/8	N/A	N/A
Cognitive appraisals only	6/8	170	1/8
Multi-phase appraisal	8/8	70	2/8

In solving a problem while using only cognitive appraisal, the Companion finds fear in many steps deep in the problem solving process. This is the result of these steps being less desirable because they are deeper in the problem. An example is assessing whether the problem is looking for a quantitative answer or not. Fearing this step, it attempts to avoid it, though eventually it gives in and tries the step anyway, once other options have been exhausted. Since the problem does require a quantitative answer, this step ends up being successful. When attempting a similar problem, it recalls the joy of successfully executing this step. However, the cognitive appraisal still finds fear in the step. This conflict is resolved by reconciling the appraisal variables that justify each of the emotions. The recalled experience of joy has associated to it the high likelihood of success and the highly desirable positive result. These outweigh the low desirability of a step deep in the problem solving and the moderate estimate of likelihood of the step succeeding. The revised appraisal variables allow the Companion to now see the step with hope and optimism. It no longer attempts to avoid the step but continues to avoid other less desirable steps. As a result, the recalled emotional experiences are guiding the Companion to steps in the problem solving process that it would have otherwise feared and avoided.

Some physics problems take many more steps to solve. Attempting to solve a long problem can be a daunting task, and using only the cognitive appraisal rules, most of the latter steps induce fear. Like the example above, with the influence of emotional experiences, some of these same steps create an affective reaction of joy because they have previously succeeded. In addition to the final appraisal shifting to hope like the above example, the Companion may also get a sense of challenge. The likelihood of success of a step can cross a minimum threshold for challenge once sufficient examples of success have been experienced. Furthermore, the depth in the problem solving task leads the companion to generate a high estimate of remaining anticipated effort. The high likelihood and high anticipated effort lead to a feeling of being challenged, which causes the agent to remain focused on the problem. Without the challenge emotion being present, the limit of allocated resources will be reached and the problem solving task will not be completed. Instead, the effect of the challenge emotion is to remain attending to the task at hand by allocating further resources to the problem.

5. Discussion

The simulation's behavior is consistent with much of the literature. For example, Baumeister et al. argue for emotions associated with an experienced event leading to future automatic affective responses. The simulations show that steps in the problem solving process that have failed are retrospectively appraised with negative emotions, which are then triggered as an automatic affective response the next time a similar step occurs. They also review evidence for anticipated emotions influencing an avoidance behavior when the consequences are potentially negative. Again, the simulations demonstrate that when there is fear of the potential failure of a step, an avoidance mechanism is triggered.

Emotional dynamics in complex problems solving has not consistently shown positive or negative emotions being an indicator of task performance. However, emotions are consistently a good indicator of strategies used during problem solving. Negative emotions often lead to an increased information gathering behavior (Spering et al., 2005; Zhou, 2013). Unfortunately, this information gathering behavior is not necessarily a good thing. In a study on complex

information problem solving, participants reporting negative emotions exhibited a pattern of continuous information gathering. Zhou (2013) rejects the notion that the participants are skilled information gatherers because their task performance was low, and one alternative proposed is that the negative emotions are rooted in an avoidance motivation. This is consistent with the avoidance behavior seen in our simulations. When only cognitive appraisals are used in the simulation, it can be seen that a very high number of steps are avoided, indicating that many steps are feared or evoke sadness. The resulting pattern of avoiding steps until one that is not feared (or disliked) is similar to the pattern of continuous searching reported by Zhou. This problem is mitigated with positive affective reactions based on learned successes with previous problems, altering the final appraisal to be one of hope or joy, and thus choosing to not avoid useful steps.

In addition to feelings of fear, hope also has results consistent with other findings. Subjects remembering feelings of hope anticipated exerting effort and attended to the situation at hand (Smith & Ellsworth, 1985). When the next problem-solving step in the simulations is appraised with hope, the agent is willing to exert the effort necessary to attempt the step. It does not attempt to avoid the step or the task altogether but instead maintains attention on the current step.

Maintaining attention should not be confused with the *heightened* attention seen with feelings of challenge. In this case, subjects reported an *increase* in attention (Smith & Ellsworth, 1985). Clément and Duvalllet (2010) have found increases in positive facial expressions and skin conductance rate in subgoal achievement and suggests that this leads to an increase in motivation and likelihood that the goal can be accomplished. Our simulations of challenge require a high likelihood of success and a desirable (positive affect) goal. When the final appraisal is challenge, this increase in motivation and attention is shown in the increase in resources allocated to the task. While other models of resource allocation may be functionally equivalent, this model provides a possible underlying cognitive mechanism. Furthermore, since the resulting challenged emotion is influenced by affective reactions, which are learned behaviors, this model provides a plausible explanation for how previous experiences can be used to determine the appropriate situations for the increased resource allocation. The simulations do show an increase in the appraisal of the challenged emotion when affective reactions are included, and this increase directly leads to an additional problem being solved. While this is a very limited sample size, it still demonstrates the potential of this model and the inclusion of affective reactions to get more appropriate behavior.

6. Conclusions

Using a simulation of solving physics problems, the model described here demonstrates the role of affective reactions in a multi-phase appraisal of emotions. This model shows how problem solving efforts are influenced by the emotional dynamics derived from affective reactions and cognitive appraisals. Affective reactions are based on analogical comparisons between the current step and past experiences. Emotions from these past experiences are projected onto the current situation to generate an initial emotional response. Cognitive appraisals rely on deeper reasoning about a situation to produce its own emotional response. When these responses do not match, a reconciliation process estimates new appraisal variable values based on both affective memories and cognitive processes. A reappraisal produces a final emotional response that influences behavior choices. Finally, a retrospective appraisal of the

problem solving experience facilitates learning and provides the basis for future affective reactions.

Simulations of the multi-phase model of emotions demonstrate advantages over a model that relies only on cognitive appraisal. A reduction in fear of steps in a problem solving task occurs due to the affective reactions included in the model. Affective reactions are grounded in previous experiences and provide more information for making better behavioral choices.

6.1 Future Work

A significant amount of modeling work remains. Other emotions such as surprise, interest, pride, satisfaction, relief, and regret are relevant to this domain because they would influence the selection of behaviors, learning, and cognitive processing. Also, the currently modeled emotions need further development. For example, models for adapting the thresholds for becoming bored or challenged are needed. Also, our model of avoidance is fairly simplistic and does not account for which previously avoided step is best to pursue next, if any. Complete avoidance of the problem (due to boredom, frustration, or hopelessness) is also not yet modeled, except via failure due to resource exhaustion.

An important component missing from the current model is incorporating recent emotions and mood in the retrieval of affective memories. Since emotions are facilitating a role in learning how to react to future situations, emotions should also facilitate memory (Baumeister et al., 2007). Mood-congruent recall (Bower, 1981) would retrieve affective memories that are more similar to the current mood. Emotion-related surface similarity has been shown to influence analogical memory retrieval (Hesse et al., 1997).

An addition of intensity to the model would demonstrate how the emotional reaction can completely override a cognitive appraisal. This overpowering effect of the emotional reaction, despite conscious reasoning to the contrary, can be seen in emotional disorders like depression, OCD, and PTSD (Taylor & Liberzon, 2000).

References

- Baumeister, R. F., Vohs, K. D., DeWall, C. N., & Zhang, L. (2007). How emotion shapes behavior: feedback, anticipation, and reflection, rather than direct causation. *Personality and Social Psychology Review*, 11(2), 167–203.
- Becker-Asano, C., & Wachsmuth, I. (2009). Affective computing with primary and secondary emotions in a virtual human. *Autonomous Agents and Multi-Agent Systems*, 20(1).
- Bliznashki, S., & Kokinov, B. (2009). Analogical transfer of emotions. *New frontiers in analogy research*, 45-53.
- Bower, G. H. (1981). Mood and memory. *The American psychologist*, 36(2), 129–48.
- Clément, E., & Duvallet, D. (2010). Emotional-related responses to critical events in problem solving. *Current psychology letters*, 26(2), 2011.
- Forbus, K., Ferguson, R. & Gentner, D. (1994). Incremental structure-mapping. *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, August.

- Forbus, K., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, 19, 141-205.
- Forbus, K., Hinrichs, T., de Kleer, J., & Usher, J. (2010). FIRE: Infrastructure for Experience-based Systems with Common Sense. *AAAI Fall Symposium on Commonsense Knowledge*, Arlington, VA.
- Forbus, K., Klenk, M. & Hinrichs, T. (2009, July/August). Companion Cognitive Systems: Design Goals and Lessons Learned So Far. *IEEE Intelligent Systems*, 24(4), 36-46.
- Gawronski, B., & Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: an integrative review of implicit and explicit attitude change. *Psychological bulletin*, 132(5), 692–731.
- Gentner, D., & Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist*, 52, 45-56.
- Gratch, J., & Marsella, S. (2004). A Domain-independent Framework for Modeling Emotion, *Journal of Cognitive Systems Research*, 5(4), 269–306.
- Hesse, F. W., Kauer, G., & Spies, K. (1996). Effects of emotion-related surface similarity in analogical problem solving. *The American journal of psychology*, 110(3), 357-383.
- Hummel, J., & Holyoak, K. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, 104, 427–466.
- Kahneman, D. (2011). *Thinking, fast and slow*. New York: Farrar, Straus and Giroux.
- Klenk, M. & Forbus, K. (2009). Analogical Model Formulation for AP Physics Problems. *Artificial Intelligence*, 173(18), 1615-1638.
- Lazarus, R. (1991). *Emotion and Adaptation*. NY: Oxford University Press.
- LeDoux, J. E. (1996). *The emotional brain: The mysterious underpinnings of emotional life*. New York: Simon & Schuster.
- Lerner, J. S., & Keltner, D. (2001). Fear, anger, and risk. *Journal of personality and social psychology*, 81(1), 146.
- Lin, J., Spraragen, M., & Zyda, M. (2012). Computational Models of Emotion and Cognition. *Advances in Cognitive Systems*, 1, 59–76.
- Marsella, S. C., & Gratch, J. (2009). EMA: A process model of appraisal dynamics 1 Introduction. *Journal of Cognitive Systems Research*, 10(2000), 70–90.
- Marsella, S., Gratch, J., & Petta, P. (2010). Computational models of emotion. In K. R. Scherer, T. Bänziger, & E. Roesch (Eds.), *A blueprint for an affectively competent agent: Cross-fertilization between Emotion Psychology, Affective Neuroscience, and Affective Computing* (pp. 177–210).
- Ortony, A., Clore, G. L., & Collins, A. (1988). *The cognitive structure of emotions*. New York: Cambridge University Press.
- Ortony, A., Norman, D., & Revelle, W. (2005). Affect and proto-affect in effective functioning. In J. Fellous & M. Arbib (Eds.), *Who needs emotions: The brain meets the machine* (pp. 173–202). Oxford University Press.
- Panksepp, J. (1998). *Affective neuroscience: The foundations of human and animal emotions*. Oxford University Press.

- Russell, J. A., & Mehrabian, A. (1977). Evidence for a three-factor theory of emotions. *Journal of research in Personality*, 11(3), 273-294.
- Schaefer, A., Collette, F., Philippot, P., Linden, M. V. D., Laureys, S., Delfiore, G., ... & Salmon, E. (2003). Neural correlates of “hot” and “cold” emotional processing: a multilevel approach to the functional anatomy of emotion. *Neuroimage*, 18(4), 938-949.
- Smith, C. A., & Ellsworth, P. C. (1985). Patterns of Cognitive Appraisal in Emotion. *Journal of Personality and Social Psychology*, 48(4), 813–838.
- Spering, M., Wagener, D., & Funke, J. (2005). The role of emotions in complex problem-solving. *Cognition and Emotion*, 19(8), 1252-1261.
- Taylor, S. F., & Liberzon, I. (2007). Neural correlates of emotion regulation in psychopathology. *Trends in cognitive sciences*, 11(10), 413-418.
- Thagard, P., & Shelley, C. (2001). Emotional analogies and analogical inference. *The analogical mind: Perspectives from cognitive science*, 335-362.
- Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist*, 35, 151-175.
- Zhou, M. (2013). “I am Really Good at It” or “I am Just Feeling Lucky”: the effects of emotions on information problem-solving. *Educational Technology Research and Development*, 1-16.