

This is a self-archived version of an original article. This version may differ from the original in pagination and typographic details.

Author(s): Zhang, Jiayi Eurus; Hilpert, Bernhard; Broekens, Joost; Jokinen, Jussi P. P.

Title: Simulating Emotions With an Integrated Computational Model of Appraisal and Reinforcement Learning

Year: 2024

Version: Published version

Copyright: © 2024 Copyright held by the owner/author(s).

Rights: CC BY 4.0

Rights url: <https://creativecommons.org/licenses/by/4.0/>

Please cite the original version:

Zhang, J. E., Hilpert, B., Broekens, J., & Jokinen, J. P.P. (2024). Simulating Emotions With an Integrated Computational Model of Appraisal and Reinforcement Learning. In F. F. Mueller, P. Kyburz, J. R. Williamson, C. Sas, M. L. Wilson, P. T. Dugas, & I. Shklovski (Eds.), CHI '24 : Proceedings of the CHI Conference on Human Factors in Computing Systems (Article 703). ACM. <https://doi.org/10.1145/3613904.3641908>



Simulating Emotions With an Integrated Computational Model of Appraisal and Reinforcement Learning

Jiayi Eurus Zhang
jiayi.eur.zhang@jyu.fi

Cognitive Science, University of Jyväskylä
Finland

Joost Broekens
joost.broekens@gmail.com
Leiden University
Netherlands

Bernard Hilpert
b.hilpert@liacs.leidenuniv.nl
Leiden University
Netherlands

Jussi P.P. Jokinen
jussi.p.p.jokinen@jyu.fi
Cognitive Science, University of Jyväskylä
Finland




Figure 1: This paper develops and evaluates a computational model of emotion that is based on temporal difference reinforcement learning and appraisal theory. It predicts emotional responses to events by integrating reward processing and cognitive appraisal. In an illustration of a task carried out by human participants, the user tries to achieve a goal using a computer. Multiple successful attempts at the problem result in a self-evaluated feeling of happiness, which our computational model matches. We model interaction as a decision process, where an evaluation of interactive events results in a value prediction update. The appraisal process of emotion is modeled based on different computations that are carried out during this evaluation.

ABSTRACT

Predicting users' emotional states during interaction is a long-standing goal of affective computing. However, traditional methods based on sensory data alone fall short due to the interplay between users' latent cognitive states and emotional responses. To address this, we introduce a computational cognitive model that simulates emotion as a continuous process, rather than a static state, during interactive episodes. This model integrates cognitive-emotional appraisal mechanisms with computational rationality, utilizing value predictions from reinforcement learning. Experiments with human participants demonstrate the model's ability to predict and explain the emergence of emotions such as happiness, boredom, and irritation during interactions. Our approach opens the possibility of designing interactive systems that adapt to users' emotional states, thereby improving user experience and engagement. This work also deepens our understanding of the potential of modeling the

relationship between reward processing, reinforcement learning, goal-directed behavior, and appraisal.

CCS CONCEPTS

• **Human-centered computing** → HCI theory, concepts and models.

KEYWORDS

Emotion Modeling, Reinforcement Learning, Appraisal Theory, Computational Rationality

ACM Reference Format:

Jiayi Eurus Zhang, Bernard Hilpert, Joost Broekens, and Jussi P.P. Jokinen. 2024. Simulating Emotions With an Integrated Computational Model of Appraisal and Reinforcement Learning. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3613904.3641908>



This work is licensed under a Creative Commons Attribution International 4.0 License.

CHI '24, May 11–16, 2024, Honolulu, HI, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0330-0/24/05
<https://doi.org/10.1145/3613904.3641908>

1 INTRODUCTION

Emotions have a significant influence on interpersonal dynamics and outcomes in daily interactions. Similar effects are also present in human-computer interaction (HCI) [3], where users exhibit emotions akin to face-to-face interactions [46]. Consequently, emotions shape perceptions of interactive systems and impact the success of interactions [5, 11, 13]. It is therefore a long-standing goal of HCI

to understand and predict a user's emotions. This is a challenging problem because while humans have an innate ability to recognize emotions in others, and make inferences and reason about them [38], computers lack this capacity. They require an explicit emotion model in order to make sense of and adapt to users' emotions.

Many attempts to enhance computers' emotion detection focus on analyzing psychophysiological signals stemming from the user's autonomic nervous system [18, 19, 34]. However, the challenge of automated emotion detection is difficult due to the interplay between emotions and cognition [2, 41]. Cognitive processes are unobservable, limiting machines to interpreting emotions based on observable behavior and physiological changes. Yet, if humans can deduce emotions from minimal observations, why can't machines? This paper posits that discerning user emotions requires a theory bridging cognition and emotions. Models accomplishing this implement psychological theories of human cognitive-emotional processes, aiming to deduce emotions from sparse data using model-informed biases [12, 28, 43, 45]. While several such models have emerged recently, their integration into HCI remains limited.

The main contribution of our paper is the adaptation of the *temporal difference reinforcement learning model of appraisal* [54] to HCI. We assess the model's predictive capabilities in an interactive task, and expand it to capture the dynamic nature of emotions during interactions. The model's key innovation is merging a reward processing mechanism with appraisal theory, using a unified reinforcement learning (RL) framework. Yet, it hasn't so far been adapted to interactive tasks, nor assessed with any real-life tasks involving human emotions. In this study, our focus is on examining and modeling three emotions: happiness, boredom, and irritation. These were selected due to their frequent occurrence in HCI and their substantial impact on user behavior, engagement, and the overall user experience [25]. The selected emotions represent a spectrum from positive (happiness), via neutral (boredom), to negative (irritation) [20, 21, 26].

Figure 1 illustrates the model at work: 'Lucy' strives to achieve her objectives in an interactive task. Each progressive step elicits positive feedback, leading to positive value estimates. As the task advances, Lucy gains confidence in her goal attainment. When prompted about her emotions after the task, she expresses happiness, but also a hint of boredom due to the task's simplicity. She doesn't feel frustrated. The bar graphs in Figure 1 display human self-report results from a relatively straightforward, rewarding task alongside the model's predictions. The alignment between the two stems from a computational cognitive emotion model that estimates the user's likely emotions given interactive events during the task. Predictions are made by applying a computationally grounded cognitive model, not by observing physiological signals or learning a model from human responses.

Existing computational cognitive emotion models fall short in predicting the scenarios we present here, primarily because they do not incorporate a simulation of an autonomous agent capable of evaluating and selecting actions to optimize anticipated outcomes. We foresee multiple applications of this approach. First, affective computing researchers could integrate our work to existing models on physiological signals, improving the accuracy of emotion detection. Second, machines equipped with a model-based understanding of their users' emotions can simulate, *in silico*, alternative

courses of action, deciding on one that is best predicted to achieve the desired emotional outcome [14, 43].

2 BACKGROUND

Understanding and predicting the user's emotions is a long-standing objective in HCI. To that end, *affective computing* studies and develops systems designed to recognize, interpret, simulate, and respond to human emotions [51]. Since the founding of this field [33], the main research lines revolve around the detection and interpretation of affective and social signals from humans [50], modeling the different facets of human-agent interaction [23, 35, 44], as well as computational simulation of emotion processes based on psychological theories [7], especially appraisal theories [4, 10, 24]. Affective computing has produced a number of key techniques that use sensors to infer emotional states, often based on either basic emotion [15] or core affect theories [36]. In contrast, appraisal theory stands out as a promising foundation for computational cognitive emotion models due to its dedication to explaining emotions within integrated cognitive-affective processes inherent to humans [43].

Cognitive models of appraisal delineate an evaluative process (appraisal) by which specific situations evoke particular emotional responses, given the subject's goals [29, 31, 42, 47]. For instance, the component process model (CPM) proposes a set of sequential cognitive checks, which assess situational stimuli based on characteristics like novelty, intrinsic pleasantness, goal relevance, and coping [30, 37]. The CPM predicts that the collective effect of these evaluations results in an emotion-specific outcome profile. Most commonly occurring profiles are called 'modal', and are associated with an emotion word, such as happiness, joy, or anger [40]. Such appraisal models provide a detailed and empirically verifiable account of the cognitive mechanics underpinning the appraisal process and its associated emotion [41]. Moreover, these models enable their specifics to be formalized in computationally implementable terms [27, 37], making them suitable for creating machines that understand their users' emotions. While this allows for a clear step-by-step analysis of how a specific emotion may have been elicited by a given situational stimulus, especially for computational implementations of this model, it only provides a framework for the static, momentary assessment of emotion elicitation, i.e. a specific moment in time. Yet, most scenarios, especially interactive tasks, encompass an extended temporal context and repeated situational evaluations, underscoring the need for a more continuous account of computational appraisal.

Reward processing models have been used in affective computing and HCI to estimate and predict user responses, allowing systems to adapt their behaviors [6, 9, 53]. At its core, a reward processing modeling seeks to understand decision-making based on anticipated rewards, with the ultimate aim of maximizing these rewards over time[1]. The operating principle is that positive outcomes reinforce behaviors, encouraging their repetition. Yet, the approach has limitations in modeling emotions: it often oversimplifies motivations by assuming agents act purely for rewards, overlooking aspects such as cognitive processing or behavioral constraints. Furthermore, there is still a considerable gap between a reward-processing model of emotion and a realistic model of human emotions.

Computational rationality is an approach that has recently been used in modeling a variety of interactive tasks [32]. It posits that humans can be modeled as agents whose decision-making and behavior are optimal within the bounds imposed by information, computational resources, and expected outcome utility [22]. This approach has an interesting connection to the reward processing model of emotion: in computational rational modeling, RL is used to derive bounded optimal behavior policies. At the heart of computational rationality in HCI is implementing a simulated user's goals as a reward function [8, 16, 17]. This facilitates the integration of emotion into computational rationality, thereby implementing emotion as part of an emerging modeling paradigm in HCI.

However, what stands in the way of implementing a model of user's emotions within computational rationality is the aforementioned gap between reward processing models of emotion and a more realistic understanding of human emotions. It has been recently suggested that appraisal theory is a promising candidate for bridging this gap [28, 54]. In many ways, appraisal theory is well suited for modeling emotion within the computational rationality framework, if it is implemented via the reward processing carried out in bounded optimal agents. This is because both appraisal theory and computational rationality embrace the importance of goals in making predictions, and note that it is not merely the events of the environment that shape behavior, but also cognition. In a demonstration of this, a recent model integrates appraisal theory with RL [54]. While the model is promising, it is limited in being evaluated only with vignettes – textual descriptions – of everyday situations. Our goal with this paper is to review the applicability of this model in interactive tasks, and evaluate and design it further to fit this goal.

3 MODELING

In this section we review the recent model that formalizes emotional appraisal using RL [54], and develop it further. In section 3.1 we outline the foundations of RL; in sections 3.2 and 3.3 we describe the existing model that integrates appraisal into RL; and in section 3.4 we build on this model.

3.1 Sequential Decision-Making

The model's interactive episodes are formalized via a Markov decision process (MDP), a mathematical framework for modeling decision-making problems in stochastic environments [48]. It is a tuple $\langle S, A, T, R, \gamma \rangle$, where S denotes the set of states and A represents the set of actions that the agent can take. The state transition function $T(s, a, s')$ describes the probability of transitioning from state $s \in S$ to state $s' \in S$ when taking action $a \in A$. The reward function $R(s, a, s')$ defines the immediate reward an agent receives when transitioning from state s to state s' by performing action a . The discount factor γ discounts future rewards when calculating the value of actions.

In order to maximize the long-term rewards of a sequential decision-making task described with an MDP, an RL agent interacts with the environment, encoding the state transition probabilities and the reward function. The problem of RL is to derive an optimal policy π^* , which maps states to action probabilities such that behavior according to it maximizes the expected cumulative reward over

time. The *value function* of a state s under a policy π , denoted as $v_\pi(s)$, is the expected return when starting in state s and following policy π thereafter. The function $v_\pi(s)$ is the state-value function for policy π :

$$v_\pi(s) = \mathbb{E}_\pi[G_t | S_t = s], \text{ for all } s \in S, \quad (1)$$

where $G_t = \sum_{k=0}^{\infty} \gamma^k R_k$ represents the expected discounted return, and \mathbb{E}_π denotes the expected value of the policy. The value of performing an action $a \in A$ while in a state $s \in S$ is defined as:

$$q_\pi(s, a) = \mathbb{E}_\pi[G_t | S_t = s, A_t = a], \text{ for all } s \in S \text{ and } a \in A. \quad (2)$$

The agent learns the optimal policy by interacting with the environment, receiving feedback in the form of rewards, and updating its value estimates for state-action pairs. In temporal difference (TD) learning, the value estimates are based on the difference between the expected and the observed value:

$$V(s) \leftarrow V(s) + \alpha[R_{s'} + \gamma V(s') - V(s)], \quad (3)$$

where α is the learning rate. $R_{s'}$ is the reward received after moving to the new state and $V(s')$ is the estimated value for the new state. This operation updates the value $V(s)$ associated with a state s as soon as the new state s' is reached, by computing the difference between predicted and observed values. Combining equations 2 and 3 results in a form of TD learning called Q-learning [48], which can be expressed as

$$q(s, a) \leftarrow q(s, a) + \alpha[R(s, a) + \gamma \max_{a'} q(s', a') - q(s, a)]. \quad (4)$$

3.2 Appraisal Calculation

Several appraisals are discussed in the literature, including relevance, implication, coping potential, and normative significance [39]. In the RL appraisal model [54], four appraisals were considered: suddenness, goal relevance, conduciveness, and power. This choice was made because these appraisals have distinct representational capacities (regarding real-life episodes), and minimal inter-correlation, and they are suitable for integration into an RL model.

Suddenness is part of the novelty assessment of an event during appraisal. Specifically, it quantifies the frequency with which a transition to state s' occurs after action a is taken in a prior state s by the agent. Suddenness is denoted by A_s and is defined as:

$$A_s = 1 - \frac{\hat{T}(s, a, s')}{\sum_{s'' \in S} \hat{T}(s, a, s'')}, \quad (5)$$

where \hat{T} is a *world model*. It approximates the true transition function T , and is learned by the agent during interaction. The intuition of \hat{T} is that the agent learns to expect certain state-action-state transitions, and therefore encountering such a transition triggers a suddenness appraisal: how expected was this transition?

Goal relevance checks how relevant an event is, given the agent's current goal. The more goal-relevant an event is, the stronger emotional reactions there will probably be [42]. Goal relevance is operationalized as the magnitude of the TD error observed during value prediction updates:

$$A_{gr} = \min(1, |\Delta|), \quad (6)$$

where $\Delta(s, a) = \alpha[R(s, a) + \gamma \max_{a'} q(s', a') - q(s, a)]$ is TD error.

Conduciveness appraisal in the CPM evaluates if an event aids the agent’s goal attainment. Conducive events generally elicit positive emotions, while obstructive ones invoke negative ones [37]. In the RL appraisal model, conduciveness is likened to both the direction and magnitude of the discrepancy between expected and actual outcomes. This concept is quantified by standardizing its values between 0 (highly uncondusive) and 1 (very conducive), with 0.5 marking neutral events that meet expectations. The intrinsic conduciveness of an event relies on the agent’s cognitive value update, informed by prior expectations and goals. Goal conduciveness is expressed as:

$$A_{gc} = \min(\max(\Delta, -1), 1) * 0.5 + 0.5, \quad (7)$$

It is worth mentioning that goal relevance and conduciveness in emotional appraisal are not inherently correlative. Events that are goal-relevant may still be uncondusive, as observed in negative emotions like despair, irritation, and sadness. Conversely, conducive events can have low goal relevance, exemplified by scenarios eliciting boredom.

Power appraisal is part of the more general coping evaluation, asking how much an agent influences an event’s outcome. For instance, an experienced user possesses power due to their knowledge, while a novice lacks this. Power appraisal provides a means to explain why a particular event, such as an error message, might cause widely different emotions in different users (e.g. confusion or even fear in novice users, and irritation in experienced users). In the model, power reflects the agent’s ability to discern between beneficial and non-beneficial actions. When the q values for various actions differ, the agent is believed to have power. Conversely, identical q values or a singular action option denote no power. Power is quantified as:

$$A_p \propto \begin{cases} |\max_{a'} q(s') - \text{avg}(q(s'))|, & \text{if } |\min_{a'} q(s')| < \max_{a'} q(s'), \\ |\min_{a'} q(s') - \text{avg}(q(s'))|, & \text{otherwise.} \end{cases} \quad (8)$$

This formulation underscores that the essence of the agent’s power lies in its ability to identify which actions to pursue and which to avoid. In the simulations below, we standardize the power appraisal by dividing by the highest absolute q value.

3.3 Classifier

In order to predict modal emotions (emotion words such as ‘happiness’ or ‘irritation’) from the computed appraisals, the vector of the four appraisals needs to be classified. In the model, this classifier is created by connecting modal emotions to particular values or profiles of such vectors. These profiles are shown in Table 1, which summarizes textual descriptions connecting appraisal profiles and modal emotions [39]. Details of how this was done are reported in [54]. For this study, we employed simulated data to train and test our classifier. The simulated data were generated by transforming nominal appraisals from Table 1 into a range of quantitative values (Table 2). We used a linear Support Vector Machine (SVM) for classification, focusing on the penalty parameter c to balance maximizing the margin and minimizing classification errors. Our goal was to approximate human performance in the classifier. We

Table 1: Appraisal patterns for selected emotions and appraisals. Goal rel. = goal relevance, Conduc = conduciveness, obs. = obstruct.

	Suddenness	Goal rel.	Conduc.	Power
Happiness	low	medium	high	open
Boredom	very low	low	open	medium
Irritation	low	medium	obs.	medium

analyzed human precision in emotion identification from the first experiment reported below (mean: 0.626, variance: 0.048) and used it as a benchmark.

We tested 100 SVM classifiers with varying c values (0.0035 to 0.006) against the simulated data. The classifier’s precision closely matched human performance at a c value of 0.0049, with a variance of 0.0004. To account for individual differences, we trained an SVM classifier for each participant, each with a c value sampled from a normal distribution (mean = 0.0049, variance = 0.0004), reflecting the variance in human precision, thereby ensuring that our model not only matched average human performance but also captured individual variability. Importantly, the parameter c was not fitted to minimize the model’s prediction error against human emotion ratings, but to the same rating precision level as found in the human data. The goal of this procedure was to bring the variance of our modal emotion predictions more in line with human self-responses: with a c value too large, only the most intense emotion would be predicted; by lowering the value, the model predicts also other, less intense emotions. This reflects how humans are able to experience various emotions simultaneously. With the classifier, the computational appraisal model is able to predict emotion words from value computations of an RL agent, via the equations for different appraisals.

3.4 Extending The Model for Sequential Emotions

While the model presented above bridges appraisal theory and a general computational approach to modeling interactive behavior, it lacks in capturing the episodic nature of the interaction, wherein a single episode there are bound to be different emotional reactions. For instance, encountering an error multiple times during interaction should not result in multiple ‘snapshot’ instances of irritation, but rather a continuously growing feeling of irritation. In other words, emotions do not appear and disappear, but linger and interact.

To that end, we augment the model. Initially, in this paper, we implement a simple moving window average, which considers not merely the present state evaluation, but those that precede it.

$$\text{SMA}(t) = \frac{x(t) + x(t-1) + \dots + x(t-n)}{n+1}, \quad (9)$$

where, $x(t)$ is the prediction from the classification of a given emotion e at time step t , and n is the length of the window. In this paper, experiment 2, we set $n = 2$, but this number depends on the abstraction level of the simulation. In the future, we also envision a discounting factor, making emotions that occurred further in the past have less impact on the present emotional state.

Table 2: Mapping of nominal appraisal values to scales was done by setting the parameters of a normal (N) or a uniform (U) distribution to correspond to the text description.

obstruct	$N, x \geq 0, \mu = 0, \sigma = 0.05$
very low	$N, x \geq 0, \mu = 0, \sigma = 0.05$
low	$N, x \geq 0, \mu = 0, \sigma = 0.1$
medium	$N, \mu = 0, \sigma = 0.05$
high	$N, x \leq 1, \mu = 1, \sigma = 0.1$
open	$U(0, 1)$

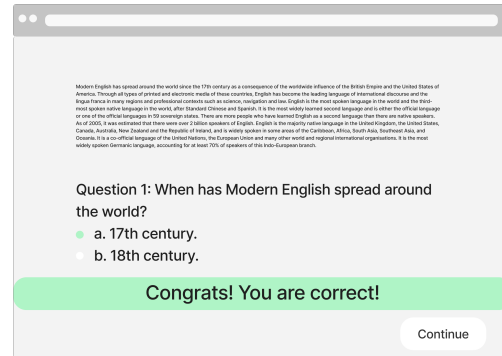
4 EVALUATION

4.1 General Method

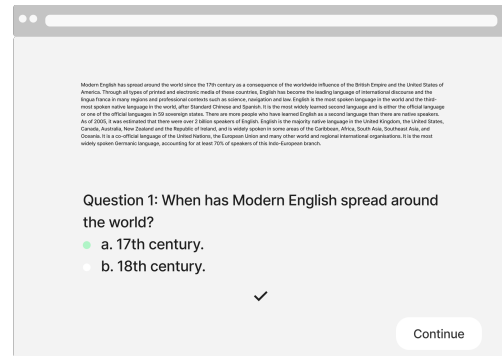
Given the goal of this paper – adapting an RL-based appraisal model to predict emotion in interactive tasks –, we focus our evaluation on users within an interactive environment. This approach differs from the vignette-based method used previously in validation. This paper introduces two original studies and assesses the model’s predictions based on their outcomes. Our focus is on three common emotions in HCI: happiness, boredom, and irritation [20, 21, 26]. Happiness reflects the fulfillment of a user’s goals or desires, and can lead to increased user engagement. Boredom signifies a lack of stimulation, possibly due to a system’s failure to maintain the user’s interest. Irritation is typically associated with frustrating events that may be due to system errors, poor design, or a failure to meet user expectations. The experimental tasks derive from appraisal theory principles, reflecting the targeted emotions’ appraisal profiles (see Table 1). For example, the happiness task featured low-suddenness, and high goal-conduciveness events, while irritation involved goal-obstructive events where participants had some power. Having formalized these appraisals computationally, we implemented identical manipulations in the computational task designs. The first experiment tests the original single-appraisal model against data collected from real emotional experiences. In the second, we test the idea of averaging emotions over longer sequences.

4.1.1 Materials: We constructed six online tasks, three for each experiment. The material was a text paragraph (about 220 words) in the English language sourced from Wikipedia, and the participants had to answer questions about the text. Multiple questions were designed from the same source text. To influence participants’ emotions, we made specific design alterations. For the happiness task, the questions were meaningful, correct answers resulted in positive feedback, and in the end the participant received a message congratulating them for good performance (Figure 2a). The boredom task featured a large number of monotonous, simple questions, and intentionally neutral feedback both for an individual task and at the end of the experiment (Figure 2b). Finally, the irritation task incorporated multiple system errors, leading to incorrect selections irrespective of user decisions, culminating in task failure and negative feedback (Figure 2c). The text and all questions are presented in full in Appendix B.

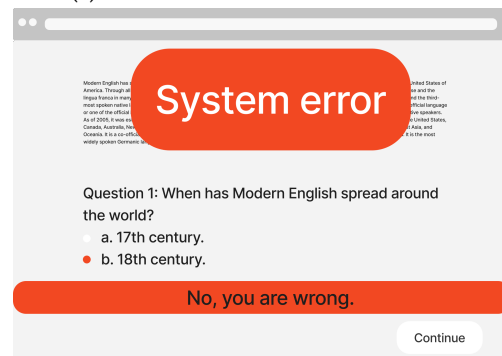
4.1.2 Participants: In the first experiment, $N = 27$ participants were recruited, 9 per task. The participants had an average age of 34 years ($sd = 10.7$), consisting of 7 men and 20 women. For the



(a) Happiness task with positive feedback



(b) Boredom task with neutral feedback



(c) Irritation task with an error

Figure 2: Screenshots from the experiments that targeted one of the three emotions. The happiness task was designed to be encouraging and rewarding, the boredom task contained a large number of repetitive tasks and neutral feedback, and the irritation task was ridden with errors and ended up failing the task.

second experiment, $N = 45$ participants were recruited, 15 for each task (average age 29 ($sd = 7.3$), 15 men and 30 women).

All participants were sourced online through Prolific, and were required to be native English speakers. With this requirement we aimed to eliminate potential biases or variations in the comprehension of the text, allowing participants to concentrate primarily on the test’s structure.

4.1.3 Procedure: Participants evaluated their emotions using a rating scale that was part of the online experiment design. This scale encompassed emotion words, asking participants to report on their current feelings (0 indicating “not experiencing this emotion at all” and 10 denoting “experiencing this emotion intensely”). Besides the primary emotions of interest (happiness, boredom, irritation), we introduced two other emotions (joy, sadness) to divert concentrated attention from the manipulated emotions. In the first experiment, the self-report was administered once post-task. In the second, evaluations occurred four times: initially, twice during the tasks, and upon completion. Correlations between the self-reported emotions are reported in the Appendix C. A between-subjects design was used, with participants engaging solely in one of the three emotional conditions. Different participants took part in the two experiments.


4.1.4 Data Analysis: In analyzing the self-reporting of the targeted three emotions, we first standardized them to reduce the impact of individuals interpreting the scale differently. For each participant, we normalized their ratings by dividing their rating for an emotion by the total sum of their ratings for all three emotions. This ensured that the rating for each emotion ranged between 0 and 1, with the combined ratings always summing up to 1. The rationale for this was that we expected the participants to hold a different internal standard for how strong a particular rating for an emotion is. However, what could be assumed to be common to all participants is how they rate the emotions in relation to each other.

For model predictions, we designed 6 simulated environments to represent each task using the MDP formalism. The formalized tasks are shown in Figure 3 (experiment 1) and Figure 5 (experiment 2). An RL agent was trained via tabular Q-learning to converge on an optimal policy separately for each task. From the converged models, we computed four appraisal measures using the equations of the previous section. The resulting appraisal vectors were classified into modal emotion probabilities using an SVM, which was calibrated as described in section 3.3. An overview of the data processing flow is shown in Appendix A.


4.2 Experiment 1

The first experiment aimed to elicit three emotions in three between-subjects tasks: happiness, boredom and irritation. The participants carried out tasks that manipulated these target emotions based on appraisal theory. Table 1 shows the appraisal profile.


The MDPs that formalize the three tasks are illustrated in Fig. 3. For the MDP of the happiness task, G serves as the Goal state with positive rewards, while E denotes the error state with negative rewards. Initiating from the state S , the agent can perform the exclusive action $frwd$, transitioning to $S1$. This action represents the task’s start, and $S1$ represents the state where the participant is shown a question with two options. These choices are represented by two actions: $a1$ for the selection that is correct, and $a2$ for incorrect. Electing for $a1$ offers an 80% likelihood for the agent to land in the G goal state, contrasted with a 20% chance of ending up in the E error state. This probabilistic outcome recognizes the real-world scenario where, despite intending to choose correctly, participants might inadvertently err due to incorrect knowledge or confusion. The appraisal analysis occurs at the onset of the goal state, when the agent transitions from $S1$ to G . The numerical patterns for



(a) The MDP model for Happiness



(b) The MDP model for Boredom



(c) The MDP model for Irritation

Figure 3: MDP models for various emotions. Circles are states and arrows are transitions caused by actions. Denoted are also transition probabilities and rewards, when relevant. In the Happiness model, selecting action ‘a1’ from state ‘S1’ results in a 20% probability of encountering an error ‘E’, mirroring the anticipated error rate among participants. Conversely, in the Irritation model, choosing ‘a1’ at ‘S1’ leads to error ‘E’ with an 80% likelihood, reflecting the high probability of system errors occurring during the task.

appraisal of the experiment, generated by these models, are shown in Table 3. Note the discrepancy in goal relevance of irritation between Tables 1 and 3. The reason for this is that in the experiment, we wanted to emphasize the irritation in human participants by making the obstructing task very goal-relevant.

The boredom MDP shares a similar configuration in its initial states S and $S1$. The distinction lies in the reward values: $r(s_1, a_1, s_2)$ is set to 1, while $r(s_1, a_2, s_2)$ is set to -1. This design choice makes the task less rewarding, both positively and negatively, implying that the outcome is less important. Appraisal analysis is conducted at $S2$. In the irritation task, we introduce a high likelihood of reaching the problematic state of the system P even when opting for the correct choice $a1$. This represents the frustrating event when a certainly correct action results in an unwanted state due to system errors. The appraisal analysis happens accordingly when the problem state P is encountered.

The average standardized self-rated emotions from the participants after the tasks and the model predictions are presented in Fig. 4. Overall, our model achieved a reasonable degree of fit to

Table 3: Appraisal profiles generated by our model.

	Suddenness	Goal rel.	Conduc.	Power
Happiness	0.2	0.6	0.8	0.87
Boredom	0	0	0.5	0.53
Irritation	0.2	1	0	0.3

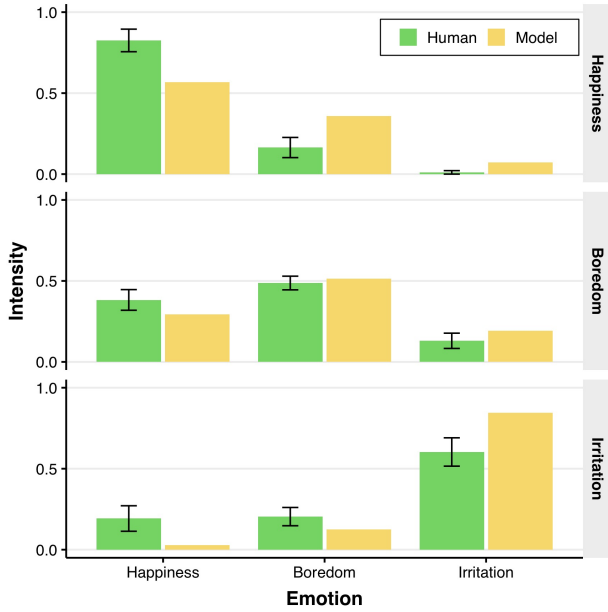


Figure 4: A comparison of human and modeling predictions for emotional ratings in each tasks. Error bars indicate 95% confidence intervals.

the data, $R^2 = 0.78$, $RMSE = 0.13$. The manipulations proved effective for both human evaluations and model outcomes. In every task, both humans and the model rated the target emotion with the highest intensity.

4.3 Experiment 2

Our second experiment aimed to expose the process nature of emotion and show that a static snapshot emotional state, either via a self-report or a model-based prediction, does not provide a full understanding of emotions during the interaction. To that end, the participants again interacted with the tasks designed to elicit one of the three emotions, but now they self-reported their emotions four times: beginning, twice during the tasks, and at the end of the experiment. While these separate self-reports are still static measurements alone, the progression of these self-reports over time can be used to evaluate the process nature of emotion and how well our model captures that.

Figure 5 illustrates the MDPs utilized in the second experiment. They bear a resemblance to those from the first experiment, but are extended to three appraisal stages. Unlike the participants, who were expected to already have some emotional experiences upon

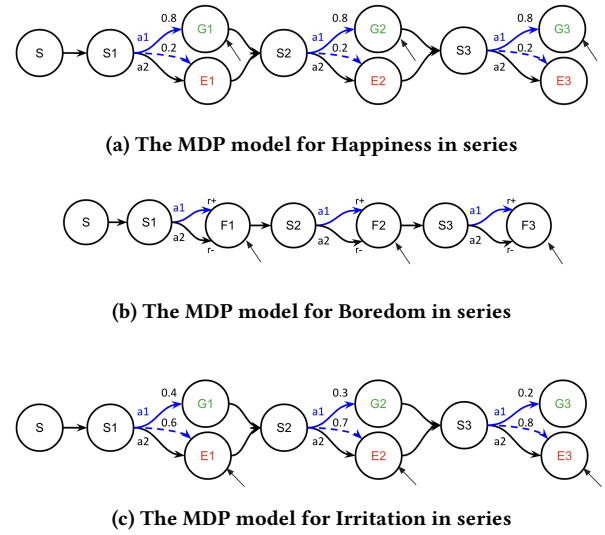


Figure 5: MDP models in series. Circles are states and arrows are transitions caused by actions. Denoted are also transition probabilities and rewards, when relevant. The error rates in the Irritation model escalate from 60% to 70%, and finally to 80%, reflecting the true the task failure rates.

starting the experiment, our model does not have the first emotion measurement. That complicates how the SMA (Eq. 9) is implemented. Thus, the initial emotions of the model were set to the average initial values obtained from human participants when they rated their emotions before the task started. At each time stage, we extracted appraisals from the model. Using the same trained SVM classifier as in the first experiment, these appraisals were then transformed into predictions of modal emotion intensities. We did not recalibrate the SVM classifier’s parameters for this new experiment. Unlike in the first experiment, we used the average of the current and previous emotion predictions to capture the process nature of emotion.

We performed a regression analysis on the human emotion values against our model’s predictions, with each experiment as a fixed term. The results yielded an $R^2 = 0.86$, $RMSE$ of 0.19, a reasonable fit between predictions and responses. From Figure 6, it is evident that the targeted emotions increase over time, while the non-targeted ones decrease. This trend is consistent in both human evaluations and model predictions, and is particularly pronounced in the happiness and irritation tasks.

However, the results from the boredom task warrant further discussion. The boredom rating sees an uptick, but there’s also a slight increase in participants’ irritation ratings. This can be attributed to the inherent challenge in designing a universally neutral interactive task. In our boredom task, the repetitiveness of the simple questions, combined with the sheer volume and limited feedback, likely led to participants becoming impatient and consequently irritated. Furthermore, the self-reported happiness levels in the boredom task were higher than boredom ratings, potentially due to

participant response bias where the participants want to provide good feedback to the experimenters and are generally favorable of their task designs [49]. This might have resulted in over-reporting happiness in a task designed primarily to elicit boredom. This is also possible for the happiness task in the first experiment, where the participants did not report as much boredom as predicted by the model. However, a more general view of all the graphs over these time stages clearly shows a decline in happiness and a rise in boredom. This points to the need for a more in-depth data analysis, focusing on the temporal shifts in emotion ratings as a contextual reference for participants' emotional responses.




Figure 6: A comparison of human and model predictions for emotional ratings in each task at each time stage. Targeted emotion magnitudes are reported with numbers.

5 DISCUSSION AND CONCLUSION

Empirical evaluation: The goal of the paper was to adapt and demonstrate a computational cognitive emotion model that simulates emotion in response to goal-driven interaction events. With some exceptions, the model closely mirrored emotional self-reports collected from human participants engaging in interactive tasks. This accuracy stemmed from incorporating appraisal theory into an RL computational framework, facilitated by a theory linking reward prediction errors with emotional responses [28, 54]. Furthermore, our findings support the validity of appraisal theory, as both the human and computational experiment task designs were rooted in appraisal-centric hypotheses.

While our validating experiments included human participants self-reporting their emotions, future research should extend its scope into more emotions and tasks. The designed interventions intentionally exerted a pronounced emotional impact, and the tasks do not represent the wide range of tasks typically performed by humans on computers. In the future, adapting the model for more involved interactive scenarios beyond the simple interactions discussed here is crucial. The MDP framework utilized in this paper

has simulated various interactions, such as multitasking while driving [17], touchscreen typing [16], and GUI-based decision-making [8]. Given that these tasks elicit emotions, our model should be able to predict them. Furthermore, while happiness, boredom, and irritation are prevalent emotions in interaction, there are other emotions relevant in HCI. We limited the amount of emotion to these key ones to focus on testing the model. However, there is a broader spectrum of emotions that the model should be able to predict. The challenge for future research lies in either controlling the experiments carefully to elicit targeted emotions, or collecting a large naturalistic dataset that considers a wide range of lived human emotions. Finally, this paper used a simple moving window average to capture the persistence of emotion throughout an episode. In the future, more complicated formulas should be considered and the complex time-dependent dynamic of emotion investigated.

Implications: In the future, we foresee our model used in interactive systems that anticipate and adapt to their users' states [14], including emotional responses. Even in its current theoretical form, the model can provide designers with insights by allowing them to examine how variations in task progression or user goals influence emotional outcomes. This implies that the model discussed herein should be adjustable individually, and tailored to a specific user's objectives and proficiency. By inferring a user's underlying cognitive states, the model's alignment with the human user can be enhanced, potentially boosting the accuracy and validity of its predictions.

Conclusion: With the increase in automated and intelligent machines that interact with their users, it becomes imperative that collaborative agents possess an understanding of their human users. A crucial aspect of this understanding is emotion. Alignment between humans and technologies embedded with artificial intelligence is risky, if the latter cannot predict their users' emotional responses to interactive events. With computational cognitive models like those developed here, it becomes possible to implement an explicit understanding of emotion into artificial agents. This understanding is not merely an ability to predict, given observed behavior or physiological signals, but to provide reasons for the causes of predicted emotions and internally simulate various 'what if' experiments to facilitate fluent interaction with the user. For the purposes of open science as recommended by [52] we present the model code and data from the experiments freely available at https://gitlab.jyu.fi/zhangjy/simulating_emotions_chi.

ACKNOWLEDGMENTS

This research has been supported by the Academy of Finland (grant 330347) and partially sponsored by the Hybrid Intelligence project (grant number 024.004.022).

REFERENCES

- [1] Hyungil Ahn and Rosalind W Picard. 2005. Affective-cognitive learning and decision making: A motivational reward framework for affective agents. In *International conference on affective computing and intelligent interaction*. Springer, 866–873.
- [2] Lisa Feldman Barrett, Ralph Adolphs, Stacy Marsella, Aleix M Martinez, and Seth D Pollak. 2019. Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. *Psychological science in the public interest* 20, 1 (2019), 1–68.
- [3] Russell Beale and Christian Peter. 2008. The role of affect and emotion in HCI. *Affect and emotion in human-computer interaction: From theory to applications*

- (2008), 1–11.
- [4] Christian Becker-Asano. 2008. *WASABI: Affect simulation for agents with believable interactivity*. Vol. 319. IOS Press.
 - [5] Chirag Bhuvaneshwara, Manuel Anglet, Bernhard Hilpert, Lara Chehayeb, Ann-Kristin Meyer, Daksitha Withanage Don, Dimitra Tsovaltzi, Patrick Gebhard, Antje Biermann, Sinah Auchtor, et al. 2023. MITHOS-Mixed Reality Interactive Teacher Training System for Conflict Situations at School. *Proceedings of the International Society of the Learning Sciences* (2023).
 - [6] Djallel Bouneffouf, Irina Rish, and Guillermo A Cecchi. 2017. Bandit models of human behavior: Reward processing in mental disorders. In *Artificial General Intelligence: 10th International Conference, AGI 2017, Melbourne, VIC, Australia, August 15–18, 2017, Proceedings 10*. Springer, 237–248.
 - [7] Joost Broekens. 2021. Emotion. *The Handbook on Socially Interactive Agents: 20 Years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 1: Methods, Behavior, Cognition* (2021), 349–384.
 - [8] Xiuli Chen, Sandra Dorothee Starke, Chris Baber, and Andrew Howes. 2017. A cognitive model of how people make decisions through interaction with visual displays. In *Proceedings of the 2017 CHI conference on human factors in computing systems*. 1205–1216.
 - [9] Daniel Gaspar-Figueiredo, Marta Fernández-Diego, Silvia Abrahao, and Emilio Insfran. 2023. A Comparative Study on Reward Models for UI Adaptation with Reinforcement Learning. *methods* 13 (2023), 14.
 - [10] Patrick Gebhard. 2005. ALMA: a layered model of affect. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*. 29–36.
 - [11] Patrick Gebhard, Dimitra Tsovaltzi, Tanja Schneeberger, and Fabrizio Nunnari. 2022. Serious games with SIAs. In *The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application*. 527–546.
 - [12] Jonathan Gratch and Stacy Marsella. 2005. Evaluating a computational model of emotion. *Autonomous Agents and Multi-Agent Systems* 11 (2005), 23–43.
 - [13] Bernhard Hilpert, P Gebhard, and T Schneeberger. 2021. Employing Virtual Agents for Building Trust in Driving Automation: A Qualitative Pilot Study. In *Workshop on Robo-Identity: Artificial identity and multi embodiment at the 16th International Conference on Human-Robot Interaction (HRI'21)*.
 - [14] Andrew Howes, Jussi PP Jokinen, and Antti Oulasvirta. 2023. Towards machines that understand people. *AI Magazine* (2023).
 - [15] Carroll E Izard. 2007. Basic emotions, natural kinds, emotion schemas, and a new paradigm. *Perspectives on psychological science* 2, 3 (2007), 260–280.
 - [16] Jussi Jokinen, Aditya Acharya, Mohammad Uzair, Xinhui Jiang, and Antti Oulasvirta. 2021. Touchscreen Typing As Optimal Supervisory Control. In *Proc CHI*. 1–14.
 - [17] Jussi PP Jokinen, Tuomo Kujala, and Antti Oulasvirta. 2021. Multitasking in driving as optimal adaptation under uncertainty. *Human factors* 63, 8 (2021), 1324–1341.
 - [18] Sylvia D Kreibitz. 2010. Autonomic nervous system activity in emotion: A review. *Biological psychology* 84, 3 (2010), 394–421.
 - [19] Sylvia D Kreibitz, Gunnar Schaefer, and Tobias Brosch. 2010. Psychophysiological response patterning in emotion: Implications for affective computing. *A blueprint for an affectively competent agent: Cross-fertilization between emotion psychology, affective neuroscience, and affective computing* 23 (2010), 105–30.
 - [20] Jonathan Lazar, Adam Jones, and Ben Shneiderman. 2006. Workplace user frustration with computers: An exploratory investigation of the causes and severity. *Behaviour & Information Technology* 25, 03 (2006), 239–251.
 - [21] Louis Leung. 2020. Exploring the relationship between smartphone activities, flow experience, and boredom in free time. *Computers in Human Behavior* 103 (2020), 130–139.
 - [22] Richard L Lewis, Andrew Howes, and Satinder Singh. 2014. Computational rationality: Linking mechanism and behavior through bounded utility maximization. *Topics in cognitive science* 6, 2 (2014), 279–311.
 - [23] Birgit Lugrin, Catherine Pelachaud, and David Traum. 2022. *The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application*. ACM.
 - [24] Stacy C Marsella and Jonathan Gratch. 2009. EMA: A process model of appraisal dynamics. *Cognitive Systems Research* 10, 1 (2009), 70–90.
 - [25] Marion Martin, Gaynor Sadlo, and Graham Stew. 2006. The phenomenon of boredom. *Qualitative Research in Psychology* 3, 3 (2006), 193–211.
 - [26] Elisa D Mekler and Kasper Hornbæk. 2016. Momentary pleasure or lasting meaning? Distinguishing eudaimonic and hedonic user experiences. In *Proceedings of the 2016 chi conference on human factors in computing systems*. 4509–4520.
 - [27] Ben Meuleman and Klaus R Scherer. 2013. Nonlinear appraisal modeling: An application of machine learning to the study of emotion production. *IEEE Transactions on Affective Computing* 4, 4 (2013), 398–411.
 - [28] Thomas M Moerland, Joost Broekens, and Catholijn M Jonker. 2018. Emotion in reinforcement learning agents and robots: a survey. *Machine Learning* 107 (2018), 443–480.
 - [29] Agnes Moors, Phoebe C Ellsworth, Klaus R Scherer, and Nico H Frijda. 2013. Appraisal theories of emotion: State of the art and future development. *Emotion Review* 5, 2 (2013), 119–124.
 - [30] Marcello Mortillaro, Ben Meuleman, and Klaus R Scherer. 2012. Advocating a componential appraisal model to guide emotion recognition. *International Journal of Synthetic Emotions (IJSE)* 3, 1 (2012), 18–32.
 - [31] Andrew Ortony, Gerald L Clore, and Allan Collins. 2022. *The cognitive structure of emotions*. Cambridge university press.
 - [32] Antti Oulasvirta, Jussi PP Jokinen, and Andrew Howes. 2022. Computational rationality as a theory of interaction. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–14.
 - [33] Rosalind W Picard. 2000. *Affective computing*. MIT press.
 - [34] Rosalind Wright Picard et al. 1995. Affective computing. (1995).
 - [35] Anna Lea Reinwarth, Tanja Schneeberger, Fabrizio Nunnari, Patrick Gebhard, Uwe Altmann, and Janet Wessler. 2023. Look What I Made It Do-The ModelIT Method for Manually Modeling Nonverbal Behavior of Socially Interactive Agents. In *GENEA: Generation and Evaluation of Non-verbal Behaviour for Embodied Agents Workshop 2023*.
 - [36] James A Russell. 2003. Core affect and the psychological construction of emotion. *Psychological review* 110, 1 (2003), 145.
 - [37] David Sander, Didier Grandjean, and Klaus R Scherer. 2005. A systems approach to appraisal mechanisms in emotion. *Neural networks* 18, 4 (2005), 317–352.
 - [38] Rebecca Saxe and Sean Dae Houlihan. 2017. Formalizing emotion concepts within a Bayesian model of theory of mind. *Current opinion in Psychology* 17 (2017), 15–21.
 - [39] Klaus R Scherer. 2001. Appraisal considered as a process of multilevel sequential checking. In *Appraisal processes in emotion: Theory, methods, research*. Cambridge University Press, 92–120.
 - [40] Klaus R Scherer. 2005. What are emotions? And how can they be measured? *Social science information* 44, 4 (2005), 695–729.
 - [41] Klaus R Scherer. 2009. The dynamic architecture of emotion: Evidence for the component process model. *Cognition and emotion* 23, 7 (2009), 1307–1351.
 - [42] Klaus R Scherer. 2009. Emotions are emergent processes: they require a dynamic computational architecture. *Philosophical Transactions of the Royal Society B: Biological Sciences* 364, 1535 (2009), 3459–3474.
 - [43] Klaus R Scherer. 2019. Towards a prediction and data driven computational process model of emotion. *IEEE Transactions on Affective Computing* 12, 2 (2019), 279–292.
 - [44] Tanja Schneeberger. 2018. Transfer of Social Human-Human Interaction to Social Human-Agent Interaction. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. 1778–1780.
 - [45] Tanja Schneeberger, Mirella Hladky, Ann-Kristin Thurner, Jana Volkert, Alexander Heimerl, Tobias Baur, Elisabeth André, and Patrick Gebhard. 2023. The Deep Method: Towards Computational Modeling of the Social Emotion Shame driven by Theory, Introspection, and Social Signals. *IEEE Transactions on Affective Computing* (2023).
 - [46] Tanja Schneeberger, Mirella Scholtes, Bernhard Hilpert, Markus Langer, and Patrick Gebhard. 2019. Can social agents elicit shame as humans do?. In *2019 8th International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 164–170.
 - [47] Craig A Smith and Leslie D Kirby. 2009. Putting appraisal in context: Toward a relational model of appraisal and emotion. *Cognition and Emotion* 23, 7 (2009), 1352–1372.
 - [48] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. 2nd ed. MIT press.
 - [49] Aditya Vashistha, Fabian Okeke, Richard Anderson, and Nicola Dell. 2018. "You Can Always Do Better!" The Impact of Social Proof on Participant Response Bias. In *Proceedings of the 2018 chi conference on human factors in computing systems*. 1–13.
 - [50] Johannes Wagner, Florian Lingensfelder, Tobias Baur, Ionut Damian, Felix Kistler, and Elisabeth André. 2013. The social signal interpretation (SSI) framework: multimodal signal processing and recognition in real-time. In *Proceedings of the 21st ACM international conference on Multimedia*. 831–834.
 - [51] Yan Wang, Wei Song, Wei Tao, Antonio Liotta, Dawei Yang, Xinlei Li, Shuyong Gao, Yixuan Sun, Weifeng Ge, Wei Zhang, et al. 2022. A systematic review on affective computing: Emotion models, databases, and recent advances. *Information Fusion* 83 (2022), 19–52.
 - [52] Janet Wessler, Tanja Schneeberger, Bernhard Hilpert, Alexandra Alles, and Patrick Gebhard. 2021. Empirical research in affective computing: an analysis of research practices and recommendations. In *2021 9th International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 1–8.
 - [53] Zhiguo Zeng, Yi-Ping Fang, Qingqing Zhai, and Shijia Du. 2021. A Markov reward process-based framework for resilience analysis of multistate energy systems under the threat of extreme events. *Reliability Engineering & System Safety* 209 (2021), 107443.
 - [54] Jiayi Zhang, Joost Broekens, and Jussi Jokinen. 2023. Modeling Cognitive-Affective Processes with Appraisal and Reinforcement Learning. arXiv:2309.06367 [cs.HC]

A DATA PROCESSING FLOW

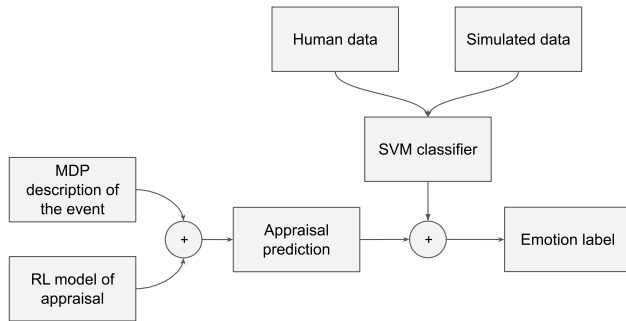


Figure 7: A reinforcement learning agent is trained in a Markov Decision Process-based task environment. Learning signals (TD errors from the converged model) are converted into appraisal predictions, which are classified into modal emotion predictions using a pre-trained SVM. The classifier is fine-tuned to have human-like spread between emotion ratings.

B EXPERIMENTAL SETUP (EXPERIMENTS 1 AND 2)

The text used in the tasks (from https://en.wikipedia.org/wiki/English_language)

Modern English has spread around the world since the 17th century as a consequence of the worldwide influence of the British Empire and the United States of America. Through all types of printed and electronic media of these countries, English has become the leading language of international discourse and the lingua franca in many regions and professional contexts such as science, navigation and law. English is the most spoken language in the world and the third-most spoken native language in the world, after Standard Chinese and Spanish. It is the most widely learned second language and is either the official language or one of the official languages in 59 sovereign states. There are more people who have learned English as a second language than there are native speakers. As of 2005, it was estimated that there were over 2 billion speakers of English. English is the majority native language in the United Kingdom, the United States, Canada, Australia, New Zealand and the Republic of Ireland, and is widely spoken in some areas of the Caribbean, Africa, South Asia, Southeast Asia, and Oceania. It is a co-official language of the United Nations, the European Union and many other world and regional international organisations. It is the most widely spoken Germanic language, accounting for at least 70% of speakers of this Indo-European branch.

Questions for the Happiness and Irritation task:

- Question 1: When has Modern English spread around the world?
 a. 17th century.
 b. 18th century.
- Question 2: Which of the following countries influenced the spread of modern English?
 a. The Great Britain.

- b. France.
- Question 3: How did English become the world's leading language?
 a. Through word of mouth of the British Empire and the US.
 b. Through printed and electronic media of the British Empire and the US.
- Question 4: Is English the most spoken language in the world?
 a. Yes, it is.
 b. No, it isn't.
- Question 5: Is English the most spoken native language in the world?
 a. Yes, it is.
 b. No, it isn't.
- Question 6: What is the most spoken native language in the world?
 a. English.
 b. Chinese.
- Question 7: What is the most learned second language in the world?
 a. English.
 b. Spanish.
- Question 8: Which of the following categories has a larger number?
 a. People who speak English as a native language.
 b. People who learn English as a second language.
- Question 9: How many people speak English as of 2005?
 a. 2 billion.
 b. 3 billion.
- Question 10: Which of the following countries has English as a majority native language?
 a. India.
 b. New Zealand.
- Question 11: Which of the following languages belongs to the official languages of the United Nations?
 a. English.
 b. Hindi.
- Question 12: What type of language is English?
 a. Germanic language.
 b. Latin language.
- Question 13: What is the most spoken Germanic language?
 a. German.
 b. English.
- Question 14: What percentage of Germanic speakers are English speakers?
 a. 70%.
 b. 80%.
- Question 15: Which of the following countries influenced the spread of modern English?
 a. Australia.
 b. The United States of America.
- Question 16: Are there more people who have learned English as a second language than native speakers?
 a. Yes.
 b. No.
- Question 17: Which of the following organizations recognizes English as a co-official language?
 a. African Union.
 b. European Union.
- Question 18: In which professional contexts is English the leading language?
 a. Art, music, and sports.

- b. Science, navigation, and law.
- Question 19: Which of the following countries has English as a majority native language?
- a. The Republic of Ireland.
 - b. The Republic of Congo.
- Question 20: How many sovereign states have English as an official language?
- a. 39.
 - b. 59.

Questions for the Boredom task:

- Question 1: When has Modern English spread around the world?
- a. 17th century.
 - b. 18th century.
- Question 2: Which of the following countries influenced the spread of modern English?
- a. The Great Britain.
 - b. France.
- Question 3: How did English become the world's leading language?
- a. Through word of mouth of the British Empire and the US.
 - b. Through printed and electronic media of the British Empire and the US.
- Question 4: Is English the most spoken language in the world?
- a. Yes, it is.
 - b. No, it isn't.
- Question 5: Is English the most spoken native language in the world?
- a. Yes, it is.
 - b. No, it isn't.
- Question 6: What is the most spoken native language in the world?
- a. English.
 - b. Chinese.
- Question 7: What is the most learned second language in the world?
- a. English.
 - b. Spanish.
- Question 8: Which of the following categories has a larger number?
- a. People who speak English as a native language.
 - b. People who learn English as a second language.
- Question 9: How many people speak English as of 2005?
- a. 2 billion.
 - b. 3 billion.
- Question 10: Which of the following countries has English as a majority native language?
- a. India.
 - b. New Zealand.
- Question 11: Which of the following languages belongs to the official languages of the United Nations?
- a. English.
 - b. Hindi.
- Question 12: What type of language is English?
- a. Germanic language.
 - b. Latin language.
- Question 13: What is the most spoken Germanic language?
- a. German.
 - b. English.
- Question 14: What percentage of Germanic speakers are English speakers?

- a. 70%.
 - b. 80%.
- Question 15: Which of the following countries influenced the spread of modern English?
- a. Australia.
 - b. The United States of America.
- Question 16: Which of the following statements about English is true?
- a. It is the most spoken native language.
 - b. It is the most learned second language.
- Question 17: Which of the following statements about English is true?
- a. It is the official language in many professions.
 - b. It is a leading language of international discourse.
- Question 18: When has Modern English spread around the world?
- a. 16th century.
 - b. 17th century.
- Question 19: Which of the following countries has English as a majority native language?
- a. The Republic of Ireland.
 - b. The Republic of Congo.
- Question 20: How many sovereign states have English as an official language?
- a. 39.
 - b. 59.
- Question 21: How did English become the world's leading language?
- a. Through word of mouth of the British Empire and the US.
 - b. Through printed and electronic media of the British Empire and the US.
- Question 22: Is English a co-official language of the United Nations?
- a. Yes, it is.
 - b. No, it isn't.
- Question 23: Is English the most spoken Germanic language?
- a. Yes, it is.
 - b. No, it isn't.
- Question 24: Does English have the most speakers?
- a. Yes, it does.
 - b. No, it doesn't.
- Question 25: Does the US influence the spread of modern English?
- a. Yes, it does.
 - b. No, it doesn't.
- Question 26: Is English widely used in many professions?
- a. Yes, it is.
 - b. No, it isn't.
- Question 27: Is English the native language in Australia?
- a. Yes, it is.
 - b. No, it isn't.
- Question 28: Is English widely spoken in South America?
- a. Yes, it is.
 - b. No, it isn't.
- Question 29: Is English a Latin language?
- a. Yes, it is.
 - b. No, it isn't.
- Question 30: Which type of language does English belong to?
- a. Germanic language.

- b. Latin language.
- Question 31: Which of the following languages belongs to the official languages of the United Nations?
 - a. Swedish.
 - b. English.
- Question 32: What is the most spoken native language in the world?
 - a. Chinese.
 - b. English.
- Question 33: What is the most spoken Germanic language?
 - a. German.
 - b. English.
- Question 34: Which of the following area is English widely spoken?
 - a. South America.
 - b. South Asia.

The questionnaire at the end of experiment 1, and at the four stages of experiment 2

Please reflect the emotions you experienced during the task:
 0 = not experiencing this emotion at all, 5= experiencing this emotion moderately, 10 = experiencing this emotion intensely

	0	1	2	3	4	5	6	7	8	9	10
Joy	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Happiness	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Boredom	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Sadness	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Irritation	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

C CORRELATION BETWEEN DIFFERENT EMOTIONS

Pearson correlation coefficients between the self-reported emotions are displayed.

Experiment 1:

Table 4: Emotion response correlation in Experiment 1

	Happiness	Joy	Boredom	Sadness	Irritation
Happiness	1	0.825	-0.145	-0.358	-0.684
Joy	0.825	1	-0.27	-0.517	-0.56
Boredom	-0.145	-0.279	1	0.11	0.05
Sadness	-0.358	-0.517	0.112	1	0.43
Irritation	-0.685	-0.560	0.052	0.430	1

Experiment 2:

Table 5: Emotion response correlation in Experiment 2 Section 1

	Happiness	Joy	Boredom	Sadness	Irritation
Happiness	1	0.933	-0.201	-0.363	-0.514
Joy	0.933	1	-0.096	-0.287	-0.469
Boredom	-0.201	-0.096	1	0.229	0.311
Sadness	-0.362	-0.287	0.229	1	0.527
Irritation	-0.514	-0.469	0.312	0.527	1

Table 6: Emotion response correlation in Experiment 2 Section 2

	Happiness	Joy	Boredom	Sadness	Irritation
Happiness	1	0.958	-0.391	-0.389	-0.5
Joy	0.958	1	-0.352	-0.331	-0.489
Boredom	-0.391	-0.352	1	0.138	0.232
Sadness	-0.389	-0.331	0.138	1	0.664
Irritation	-0.5	-0.489	0.232	0.664	1

Table 7: Emotion response correlation in Experiment 2 Section 3

	Happiness	Joy	Boredom	Sadness	Irritation
Happiness	1	0.929	-0.249	-0.532	-0.572
Joy	0.929	1	-0.224	-0.562	-0.544
Boredom	-0.249	-0.224	1	0.153	0.499
Sadness	-0.532	-0.562	0.153	1	0.503
Irritation	-0.572	-0.544	0.499	0.503	1