

# **Affective Story-morphing: Manipulating Shelley's *Frankenstein* under program control using Emotionally Intelligent Agents**

Clark Elliott  
College of Computing and Digital Media  
DePaul University, 1 DePaul Center, Chicago, IL 60604  
elliott@depaul.edu

**Abstract:** We present a theoretical model for the automated generation of plot-consistent, novel, engaging narratives based on a broad, computable model of emotion. We review the background theory, relevant to the morphing of narratives, composed of 28 emotion categories, 24 emotion intensity variables, and ~400 channels for emotion expression, and which has been implemented in an AI program called the Affective Reasoner. We argue that what is primarily of interest in narratives is the emotion fabric in the interaction between characters—much of which can be manipulated by the Affective Reasoner—and that while keeping the plot the same or similar we can, under computational control, create novel, interesting, consistent new stories that make sense to human observers. We present and explain preliminary examples and then apply the story-morphing techniques to a passage from Mary Shelley's *Frankenstein*.

**Keywords:** Affective computing, intelligent agents, emotion, stories, narrative, gaming AI.

## **Introduction and motivation**

This is a theoretical position paper arguing that we can create novel stories under computational control suitable for many contexts based on sound AI emotion reasoning principles. We first argue that using a highly-computable model of emotion allows us to extract an essential structure in stories which is independent of the narrative context. We then show that we can automatically manipulate these structures with emotionally-intelligent *story-morphing* AI agents to produce novel emotion tapestries that are, themselves, the basis of good, new, stories in the original setting. We introduce the basics of the emotion theory relative to this work, and a few illustrative examples, and then analyze a short section from Mary Shelley's *Frankenstein*—or the *Modern Prometheus* to illustrate the techniques.

Narrative structure and stories are an essential element of human cognition, including metaphors such as those of *time sequencing* (Lakoff, 1993), and are part of the what Fodor calls *The Language of Thought* (Fodor, 1975) which lies at the

symbol-processing core of what makes us human. In the realm of neuroscience we see that we take in much of our permanent life experience in the form of meaningful *episodes*, which then later, and slowly, migrate from the hippocampus and elsewhere into collections of semantic meaning in the neocortex (McClelland, McNaughton, & O'Reilly, 1995). Stories from our lives are retold over and over again, and can even dominate the rest of our lives when reliving, for example, PTSD episodes. And even these sad, catastrophic conditions can sometimes be addressed through the *re-telling* of such stories (Gray & Bourke, 2015). Decades ago, Bruner (Bruner J. S., 1957) showed us that our very (phenomenological) *perception* is based, in part, on previous episodes in our lives that color how we see the world, and that as a result we know that even the (later stages of the) vision system are cognitively penetrable by these stored episodic memories<sup>i</sup>. In addition, Bruner later emphasized the importance of narrative in absorbing a culture's folk psychology (Bruner J. , 1990). Stories are powerful in manipulating our belief systems as well: for example, when we repeatedly hear stories that are based on lies, our natural processing finds ways to integrate them into our worldview as being true (De Keersmaecker, et al., 2020).

At an analytical level, various claims have been made about the essential nature of stories. For example, it is not uncommon to categorize stories based on themes such as *good versus evil*, *love*, and *redemption* (MasterClass, 2021). Shank has argued that even the simplest story must have a *point*, and that a point is generated by a failure of expectation (Shank, 1990). There are many theories of plot development (e.g., (Kim, Padó, & Klinger, 2017). But *computational* models of world knowledge are still hugely lacking because of the overwhelming symbolic complexity of the real world. While it is a noble goal and some progress has been made (Reagan, 2017), *sentiment analysis* of pure text runs into the problem that without *natural language understanding* the clues in the ambiguous statements and utterances that humans traffic in remain opaque.

By contrast, the orthogonal position we take in this paper is that the simplest story is generated when something happens, *and someone cares about it*, and that the underlying ways in which people care creates narrative fabrics that are rich, complex and—most importantly—narratively consistent. Many of the important elements of plots, and all-important themes, are all based on such caring, which yields a vast number of stories that are, irrespective of the events themselves, based on specifically unique, but generally identifiable emotion patterns. We have previously shown that these techniques are effective, and that users said that the new stories were plausible and made sense (Elliott, Brzezinski, Sheth, & Salvatoriello, 1998).

In the general case, building a content theory of plot manipulation is tantamount to building a content theory of the world. In our work, we go to great lengths to avoid this currently insurmountable task.

## EMOTION CATEGORIES

Clark Elliott, 2021  
after Ortony, et al., 1988

GROUP	SPECIFICATION	CATEGORY LABEL AND EMOTION TYPE
Well-Being	appraisal of a situation as an <i>event</i> relative to one's <i>goals</i>	<b>joy</b> : pleased about an <i>event</i> ; <i>achieved or retained goal</i> <b>distress</b> : displeased about an <i>event</i> ; <i>blocked goal</i>
Fortunes-of-Others (Well-Being)	appraisal based on presumption of how a situation is appraised by another as an <i>event</i> relative to <i>their goals</i>	<b>happy-for</b> : pleased about an <i>event</i> desirable for another <b>gloating</b> : pleased about an <i>event</i> undesirable for another <b>resentment</b> : displeased about an <i>event</i> desirable for another <b>jealousy*</b> : <i>resentment</i> over a desired mutually exclusive goal <b>envy*</b> : <i>resentment</i> over a desired non-exclusive goal <b>sorry-for</b> : displeased about an <i>event</i> undesirable for another
Prospect-based (Well-being)	appraisal of a situation as a prospective <i>event</i> relative to one's <i>goals</i> .	<b>hope</b> : pleased about a prospective desirable <i>event</i> <b>fear</b> : displeased about a prospective undesirable <i>event</i>
Confirmation (Well-being)	appraisal of a situation as an <i>event</i> confirming or disconfirming an expectation relative to one's goals	<b>satisfaction</b> : pleased about a confirmed desirable <i>event</i> <b>relief</b> : pleased about a disconfirmed undesirable <i>event</i> <b>fears-confirmed</b> : displeased about a confirmed undesirable <i>event</i> <b>disappointment</b> : displeased about a disconfirmed desirable <i>event</i>
Attribution	appraisal of a situation as an accountable <i>act</i> of some agent, relative to one's <i>principles</i>	<b>pride</b> : approving of one's own <i>act</i> <b>admiration</b> : approving of another's <i>act</i> <b>shame</b> : disapproving of one's own <i>act</i> <b>reproach</b> : disapproving of another's <i>act</i>
Attraction	appraisal of a situation as containing an [un]attractive <i>object</i> re. one's <i>tastes</i>	<b>liking</b> : finding an <i>object</i> appealing; relative to <i>desires/tastes</i> <b>disliking</b> : finding an <i>object</i> unappealing; ~ <i>desires/tastes</i>
Well-being/ Attribution	compound emotions	<b>gratitude</b> : admiration+joy <b>anger</b> : reproach+distress <b>gratification</b> : pride+joy <b>remorse</b> : shame+distress
Attraction/ Attribution	compound emotion extensions	<b>love</b> : admiration+liking <b>hate</b> : reproach+disliking <b>[frustration*]</b> : hope+disappointment] <b>[unfairness*]</b> : resentment+reproach--stdrd unfairly applied, self/other]

\*Non-symmetric additions necessary for some stories.

**Fig. 1.** Ortony, et al., (Ortony, Clore, & Collins, 1988), modified Elliott 2015 & 2021: The structure of appraisal within the content theory of emotions used as the basis for the *dispositional* component of the Affective Reasoner's emotionally intelligent agents.

Drawing on *Love of Chair* (Wikipedia, 2021)—the Eletric Company spoof of daytime soap operas—we might say, “The boy was sitting in the chair,” and it is very hard to argue that this is a story. If we say, “The boy was sitting in the chair, and felt guilty about it,” our analysis will change. The boy cares, and we might want to know why. Or we might sympathize with him because we have ourselves sat in a chair and known that we shouldn't have. Or, we are reminded of being outside the principal's office, feeling that we shouldn't be there. And, too, there are dozens of flavors of that guilt the boy is feeling that might be of interest to us.

We also have lots of stories where we know the outcome, and it is expected. “The boy sat in the chair and felt guilty about it. He just couldn’t stop himself from taking the candy.” There is not too much that is out of the ordinary here, but it is still a story. Yet we could say, “The boy sat in the chair, which had fourteen legs and was painted pink,” which does not meet our *expectations*, but is also not a story, because no one cares.

If we want to build a robust computational models of stories, we have to start with a computational model of emotion.

## Story Morphing in The Affective Reasoner

In the Affective Reasoning system we view particular—and typically essential—narrative components based on emotion as being eminently computable. Our model includes twenty-eight *emotion categories* with multiple intensities and qualities within each category (Elliott, The Affective Reasoner: A process model of emotions in a multi-agent system, 1991). (For example, the category *anger* includes annoyance, indignation, exasperation, outrage, and so on.) We typically compute at least three different intensities for many of the emotions. The model also includes over four-hundred *channels* for expressing emotions—roughly twenty channels tweaked for each emotion. For example, a *somatic* channel for the expression of *love* would include turning red, and pulse increasing, while a *verbal other-directed emotion modulation* expression of *love* might include saying something sweet to encourage the object of your attentions to respond in kind.

The essential structure of the appraisal mechanisms composing what we call the *disposition* of agents is contained in Figure 1 giving the description of twenty-eight emotions, based, originally, on the seminal work of Orton, Clore and Collins (Ortony, Clore, & Collins, 1988).

In our work we have used the construct of Emotionally Intelligent (computer) Agents to take the place of characters in stories (Elliott, Brzezinski, Sheth, & Salvatoriello, 1998). For each such agent we can ask, *how does this agent feel about the events which are unfolding?* and *how might this agent express those feelings?*

Such agents are designed and implemented with two components: a *disposition* which controls the way they interpret situations that unfold in a story, and a *temperament* which controls how they express any emotions that may arise.

Using these simple mechanisms, we can alter the disposition of the agents, such that they differently appraise situations that arise, and also alter their temperaments such that they express their emotions differently. In this way, within the constraints of not altering *what happens* in the plot of the story, we can still greatly change the emotion structure of the story and thus change the story itself. Because the agents are internally consistent within the content theory of emotion, their manifested emotions are as well, and the new stories generated make *sense* within the context of the new characters being portrayed. This is the

crux of the matter: computational control of adding and manipulation complexity, while retaining the elegance of natural human interaction.

For example, let us consider a simple example partially using what Ortony has called the *fortunes-of-others* emotions (Elliott & Ortony, Point of view: Modeling the emotions of others, 1992), (Ortony, Clore, & Collins, 1988):

Plot: Lisa has a brother Jake who has a dog Scout. Scout gets out of the house and unbeknown to Jake eats a dead squirrel she has found at the beach. Lisa visits Jake who comments that Scout seems subdued even though it is dinnertime and she shows no interest in her dinner.

*Story one:* Lisa feels sorry for her brother Jake, with whom she is close. She knows that Jake is very protective of Scout and will be worrying about her, and is simultaneously mad at Scout for getting out. Jake is indeed worried about Scout. He feels guilty that he let her get out. Scout does not feel well because of the rotten squirrel.

*Story two:* Lisa is gloating over her brother Jake, with whom she is very competitive. She feels reproach for Jake who does not know how to take care of his dog. Jake is indeed worried about Scout, but he also admires his intrepid escape-artist dog, and is proud to be her owner. Scout does not feel well because of the rotten squirrel.

*Story three:* Lisa is jealous of her brother Jake because Scout loves him. She is reproachful of Jake whom she feels should want to take better care of his dog. She makes a plan to scold him later. Jake is gloating over his stupid dog, because she has obviously done something wrong and is now sick because of it. Scout does not feel well because of the rotten squirrel.

*Story four:* Lisa does not feel much of anything—it is not her problem. Jake is furious at his dog Scout for getting out of the house and getting into trouble. He is speaking in a really loud voice. Scout is unhappy that Jake is mad at her. She is afraid that Jake will punish her. She is happy about having eaten the dead squirrel which was the high point of her day. Her stomach hurts but she doesn't care much about that.

*Story five:* Lisa does not feel much of anything—it is not her problem. Jake is furious at his dog Scout for getting out of the house and getting into trouble. He is trembling and red in the face, but not saying anything. Scout has mixed emotions. She feels guilty that she ate the dead squirrel and she is afraid that Jake will start yelling at her, but she is also happy about having eaten the dead squirrel which was the high point of her day.

In this way using only the crudest of the manipulations that can be computationally controlled by the Affective Reasoner mechanisms in real time, we are nonetheless able to generate scores of stories.

Let us now examine the mechanisms used in this simple example in a little more detail.

In table 1, we see that there are four main divisions of the twenty-eight emotion categories. First is the large set of emotions that arise because of the *goals* of agents—what agents want and don’t want, and in a related way, what they believe their friends and adversaries want and don’t want. Second is the set of emotions based on the *principles* of the agents—relevant to *actions* agents believe *should* and *should not* be performed. Third is the small set of emotions based on what agents like and don’t like—their preferences. Fourth is the set of emotions based on combinations of other categories that subsume their constituent parts. Note also that in a pseudo category we have the mixed emotions—since while there is a conflict set in the manifestation of emotion (e.g., one can’t shout and be silent at the same time) there is no such conflict in the *feeling* of multiple and even contradictory emotions at the same time (e.g., sorrow over the death of a favorite uncle; joy over the fact that he has bequeathed much-needed money).

On the one hand, to control how agents respond to the circumstances—the *emotion eliciting conditions*—that unfold in a story, we have to change the way they *appraise* those circumstances. These appraisals are part of what we build into an agent’s *stable disposition*.

In the first case Lisa appraises her brother Jake’s distress as blocking one of her own fortunes-of-others goals (Lisa’s desire for Jake’s ongoing well-being): when Jake is sad, she feels sorry for him. Lisa is angry (a compound emotion) at Scout because Scout has both violated Lisa’s *principle* that dogs should not run away and get in trouble, and also blocked her goal for Jake’s well-being. Jake appraises the situation as indicative of a possible future goal of his own—keeping his dog healthy—being blocked, although this is currently uncertain (worry). He also, independently appraises the situation as an instance of him violating his own principle of keeping Scout safe (guilt). Scout’s own health well-being goal is being blocked by an uncomfortable stomach (distress).

In the second case these appraisals change, and so do the resulting emotion states. For example, in the second story Lisa’s *friendship* relationship with Jake has changed to become (in this situation) *adversarial*. So now when Jake’s ongoing well-being has taken a downturn, Lisa feels good about it.

All of the many changes in the emotions that arise in the various simple stories about Lisa, Jake, and Scout, stem only from such changes in the ways that agents appraise the otherwise identical events that that occur in the plot.

Among the twenty-odd action channels that differentially control the expression of any particular emotion we have a spectrum of paths such as *somatic, behavioral directed toward an inanimate [or animate] object, communicative [non-verbal / verbal] responses, evaluative self-directed attributions, ..., repression, suppression, reappraisal of the situation, reappraisal of one’s self, other-directed emotion modulation, ..., full plan initiation*, and so on, from the simplest body responses, to the most complex intellectual responses.

So, on the other hand, to control how agents *manifest (express)* their emotions we have to change what we build into their *temperaments*. We achieve this by changing the action-expression channels that are activated at any given moment for an agent, which in turn, taken together, yield the agent’s temperament.

For example, in story four Jake expresses his anger by shouting, a *communicative verbal* response, or possibly (because he is talking to a dog) a *behavioral response directed toward an animate object* indicating that for this temperament, those channels are activated. By contrast, in story five, Jake is trembling and red in the face which are *somatic responses*.

Having introduced the basics, we can now look at some additional ways in which we can control the emotion structure of stories.

*Relationships:* First, as hinted above, we have four relationships that we model, between agents: *friendship*, *adversarial*, *cognitive-unit*, and *no-relationship*. A *Friendship* relationship is intended to collect together all relationships wherein, e.g., agent Lisa will feel good when good things happen to agent Jake, and bad when bad things happen to Jake. An *adversarial* relationship is the opposite: Lisa will feel bad when good things happen to Jake, and good when bad things happen to Jake. A *cognitive-unit* relationship is when Lisa feels exactly what Jake feels, as though she were in his shoes. (For example, a mother may feel frightened *with* her son when he is forgetting his lines during the school play.) These relationships are unilateral (and even when bi-lateral might not be symmetric).

It is possible to change the relationships that agents have with one-another, and in this way affect the fortunes-of-others emotions that will arise in the system (sorry-for, happy-for, gloating, resentment).

*Emotion intensity variables:* Next we can change the *emotion-intensity-variables* which contribute to the particular (intensity of) emotion that is generated within any one of the emotion categories, and also, thus, subsequently any change in action-responses that are dependent on emotion intensity.

In the Affective Reasoner we can manipulate up to twenty-four different variables that affect how strongly situations elicit emotional responses in the agents, divided into three categories (Elliott & Siegle, Variables influencing the intensity of simulated affective states, 1993).

The first category of such variables, the *simulation-event variables*, are those that are contained in the (simulation of the) *situation* itself. For story-morphing there are constraints on the usefulness of this set of variables because they are bound to the external plot and description of the story itself, which changes we always want to keep to a minimum. That is, these are always the actual plot changes, albeit possibly representing purely local changes that do not affect the remainder of the plot. For example, if we change the amount of money a patron leaves as a tip, the waiter might also appraise the event of getting the money differently—which is what we want with story-morphing—but we have to be very careful that such a change does not alter the plot in ways requiring real-world knowledge to control, which is explicitly beyond the scope of the AR's capabilities. Nonetheless, within the story-morphing context, some changes are possible within these constraints. A special non-theoretical subset of simulation-event-variables is the manifestations of emotions that agents create, and which are themselves events within the system to which other agents may respond. To the extent that they do not change the plot, they are allowed, but with constraints.

The values in the simulation-event variables change independently of an agent's interpretation of them, and one change in a single place can conceivably affect all of the agents in the system at once. Among these variables are how much a goal is realized or blocked, the extent of the blameworthiness or praiseworthiness an action is *as it is performed* (e.g., how drunk the driver was), how *certain* the situation is, how *real* it seems, how *surprising* it is, how *deserving* the agent is of the situation, and so on. For example, if an adversary is particularly deserving of her bad fortune, an agent observing that bad fortune may gloat in a particularly strong way.

The second category, the *stable disposition variables* are those variable values that are internal to each agent. Changes in one of these will not affect any other agent's interpretation of situations. Unlike the simulation-event variables these values can generally be changed at will, and thus are easy to use with story-morphing. That is, how an agent is disposed to see situations that arise is, in essence, up to them. Among these variables are how important it is to achieve a goal or keep it from getting blocked; or uphold a principle, or not have it violated. For example, after losing a game three times to a rival, the *importance* of the goal of winning might become increasingly important. Included in this category would be the emotional interrelatedness of two agents. The more (unilaterally) interrelated they are seen to be by one of the agents, the stronger the emotions (of that agent) generated in the context of this relationship.

There is room for some finesse here, as well. For example, we have independent variables for how an agent sees the importance of upholding of a principle, and for the importance of not violating it. In this way one version of an agent's disposition might have strong emotions over the admirability of hearing romantic passion in music, but not care much at all if someone does not. An alternate version of the agent's disposition might find the agent greatly disdainful of those that cannot hear romantic passion in music, but not feel much admiration those that do—taking it to be expected of them. Or, the agent might feel strongly in both instances, or only mildly experience emotions in both cases.

Next, we have *mood variables*, which are intended to vary over time and affect both agents' dispositions, and also their temperaments. *Non-relationship mood variables* which include changes in values like *arousal* and *physical well-being*, can make emotions stronger or weaker. Also included (among others) are a bias toward negative or positive emotions across a strength spectrum (a good mood, a bad mood), and anxiety-invincibility, which affects the strength of prospect-based emotions. *Relationship mood variables* affect how an agent is disposed toward judging another agent, or, differently, other agents as a whole. This affects how an agent is biased toward praising the actions of another, or toward condemning them.

*Concerns of others:* Lastly, because emotions are sometimes generated on behalf of agents according to how they believe others to interpret situations—stored in what we call “concerns of others” structures (COOs)—we can alter these beliefs and thus change the appearance and strength of the fortunes-of-others emotions.



*External constraints on story morphs:* When a character's emotion changes using the story-morphing techniques, it may alter, reduce or—in the degenerate case, be incompatible with—the *motivation* for further steps in the plot: if a story character morphs to feels sorry for another, instead of being angry at them, it doesn't make sense for them attack the other agent. In practice, we allow such constraints to be externally recorded as part of the morph instructions. Typically, these are not the burden they appear to be: (a) A simple control on the *valence* of the emotion an agent feels at a particular plot point is generally enough to reasonably avoid the problem. (b) Humans are very forgiving of plots that involve the emotional inconsistency we perceive in others. We are often mystified by the emotions of others and have libraries full of explanations of unexpected behavior. But—and this is important—the behavior must be consistently explainable and the personalities must remain reasonably consistent throughout the narrative (though we can manipulate the *moods* of agents in theoretically consistent ways as part of the story-morph). And, (c) as a last resort we can simply mark critical plot points as locked, and not allow story morphing of that particular situation within the larger plot.

To generate new stories using the story-morphing system, in any context, we must give some thought to the nature of each part of the plot. In particular, we have to decide which parts of the plot—"what happened"—fall into one of three general categories: (1) those parts that are independent of how agents feel about them (e.g., the invaders overrun the castle walls), (2) those parts that allow for emotional freedom, but within constraints (e.g., an agent has to feel *positively* (for *some* reason) about the prospect of meeting a stranger because they are next going to pursue a meeting in the plot, and (3) those parts of the plot that require specific emotions such that those emotions cannot be changed in any significant way. To meet these constraints we have to give some thought to possible appraisals that agents might make for each of the situations that arise during the unfolding of the plot. In general, we currently keep these local, but it is possible to imagine a global collection of emotion constraints that allows for, say, a change in the valence of emotions at *plot point 7* as long as some matching change in the emotions under specific constraints occurs at *plot point 12*.

Within this context we typically still find a wide range of concerns that we can build for our agents, and a wide range of ways in which they express their resulting emotions, yielding dozens, and sometimes hundreds, of valid variations of the narrative.

One serendipitous reason that story-morphing works is that we do not require a perfect product. Humans consider the reading of the emotions of others (and even themselves) an inexact science: we are often willing to jump to conclusions and provide our own (possibly incorrect) abductive explanations about what someone is feeling, and why they are feeling it. This is a normal part of the human condition.

## How development proceeds

The story-morphing system is a structure-based system that can apply in many contexts, and development for each will be different. The handling of text is an important consideration. The Affective Reasoner has no built-in text-generation capabilities, or language understanding capabilities. However, it is a system designed to *drive* such processes by providing computer-controlled manipulation and relatively sophisticated *understanding* of the underlying emotion fabric of stories, which we claim is one of the most critical elements of most stories. Specifically detailed themes arise from the emotion fabric, and can thus be transferred from one context to another, and from one story to another. Memorable scenes in stories are often dependent on the underlying emotion structure. Identification with characters often comes from identifying with the emotions they are experiencing, independent of the context of the character's lives.

To illustrate how development might proceed, let us suppose that we wish to manipulate characters within a computer-controlled presentation of our stories in some mode. For the purposes of this position paper we need not be more specific. Such underlying work would apply to a number of applications: We might, for example, wish to generate teaching stories as part of an automated tutoring system (e.g., (Elliott, Rickel, & Lester, Lifelike pedagogical agents and affective computing: An exploratory synthesis, 1999)). We might wish to play out our stories through virtual actors, using emotion-appropriate background music. We might wish to address the *content bottleneck* problem by generating real-time characters as part of a computer game that act in novel (surprising, internally consistent...) ways according to the current configuration of their personalities.

To build a platform that supports such systems we encode the plot as a series of events unfolding in a simulation. These *sim-events* (simulation events) contain ground instances of "what happens" in our modeled world, played out as a discrete series of states within the system. Included in this series of sim-events is what the characters themselves—our agents—do. Actions of agents generated by the emotion system are inserted into the plot as additional sim-events. Our agents contain specialized internal, matching, versions of all those sim-events that are important to them for one reason or another (based on their appraisal of the events being relevant to their goals, standards and preferences). These internal structures, which we call *appraisal-frames* form the basis of each agent's disposition. The appraisal-frames support complex unification matching algorithms containing variables and functions, and themselves form the left hand side of rules such that when a match succeeds between an appraisal-frame and the current sim-event the rule fires and an emotion is generated. And, notably, the emotion is generated along with all the variable bindings that were created during the match process. In this way, for example, an agent doesn't just get angry, she gets angry *at* the other agent that was bound to the "agent-that-violated-my-principle" variable during the match of the offending sim-event.

Once an emotion is generated, it is expressed through the particular temperament channels that are activated as part of the agent's current

personality. The resulting expression of the emotion is formed as a new sim-event, and placed in the simulator's event queue.

In this way all of the events in the plot are simulated, along with all of the emotion events that the plot generates. To play the original story we configure our agent personalities with (our interpretation of) the concerns and temperaments of the original characters. To create a story-morph we variously alter the agent-frames (comprising an agent's disposition) to embody different concerns of the characters, different moods, different interpretations of the concerns of others, and different relationships, and also alter the activation of the emotion-expression channels (comprising an agent's temperament) to embody different ways of the agents expressing themselves. Then, we re-run the simulation to generate a new story.

An important feature of this system is that novel, new, internally-consistent stories can be generated *automatically* by the system without intervention by human authors.

#### **Some additional comments:**

*Humor:* Certain types of stories are humorous because of their particular emotion structure (Elliott, Why boys like motorcycles: using emotion theory to find structure in humorous stories, 1999). When this structure can be fully captured by our emotion theory, or minor extensions of it, we can also generate humorous stories.

*Case-based reasoning:* Agents have case-based reasoning intelligence that allows them to dynamically choose different internal representations (COOs) for how they *believe* others see the world. These are based on the features of eliciting situations that arise, and the responses of agents to those situations, used as reminders for how others they have known, or they themselves, see the world.

*Applications:* A computational emotion-based theory of stories is widely applicable. One obvious application is in addressing what is known as the *content bottleneck* in computer games. In the gaming industry there is a delicate balance between (a) making a computer game too easy to "figure out," so that interest is not maintained for long, and (b) making it too complex so that it is too opaque to understand, and interest is never piqued. Appropriately complex content is expensive to develop, and varied plots are both difficult to generate and burdensome to make cohesive and interesting. Using story-morphing techniques we claim that highly complex and novel game-stories can be automatically generated based on how the characters *feel* about what is unfolding—possibly controlling their dispositional behavior—rather than on the external complexities of the plot.

Other applications include role-playing therapy, story-telling and story-understanding systems, applications for children, and as a component part of emotionally-intelligent agents.

Lastly, and importantly, users are treated the same as agents within the real-time simulation. Input from users creates sim-events, agents reason about users' emotions and motivations in the same way they reason about other agents, and they may have relationships with users in the same way they have relationships with other agents. This allows for the possibility of rich, novel interaction with users in general, and specifically with players of computer games, students in tutoring systems, role-playing therapy games and etc.

## **Morphing the Monster**

Let us now look at a more extended example which will help to illustrate the richness of the story-morphing pallet. We borrow a short passage from Chapter Five of Mary Shelley's *Frankenstein—or the Modern Prometheus*.

*A paraphrase of the original narrative:*

- Background: Victor Frankenstein has created an artificial human being which is just now coming to life.
- Victor is in a high state of arousal because the monster is finally coming to life.
- Physical beauty is important to Victor. Victor likes correct design proportions in the human body. Victor likes clear beautiful skin, and dislikes yellow pasty skin and yellow eyes.
- Victor has worked very hard to make a beautiful artificial human.
- Victor was hoping to celebrate creating new life, but is now disappointed that it is ugly.
- Victor leaves to get some sleep, and dreams of his wife.
- The Monster gets up, goes to see Victor, and loves him. He wants approval. He smiles at Victor and reaches out his hand.
- Victor wakes, is overcome with disgust, and rushes downstairs.
- Victor is bitterly disappointed that the new life he has created is not beautiful.

Using only those story-morphing techniques that manipulate computable aspects of characters' personalities within the context of the narrative we can create different meaning, using the same action steps that take place in the original plot. For the purposes of illustration we present only snippets of what would be significantly more extensive passage in a full story-morph treatment.

*Story morph snippet one:*

Victor feels mixed emotions. He loves his creation (admires the struggle for life; feels the newly created life itself is beautiful). Victor feels that parents *should* love their children and see them as attractive no matter what (a principle), and he is the parent of the monster. However, Victor is repulsed by the ugliness of the monster.

- Victor is ashamed that he does not see his creation as beautiful.
- The monster is afraid that Victor will not like him.
- The monster's fears are confirmed when Victor runs away from him.
- The monster is angry at Victor for hurting his feelings, and for not taking care of him.

*Story morph snippet two:*

- Victor makes an assessment of the monster. He feels nothing but scientific curiosity. He is tired and leaves to get some sleep.
- The monster comes to life and is desperate for the affection of his creator.
- When Victor does not respond the monster is sad and thrown into depression.
- Victor feels reproach for the monster for being so emotional and leaves.
- The monster feels abandoned.

*Story morph three:*

- Victor loves his monster. He fears that others will harm his monster because they will see the monster as ugly. Victor feels guilty he did not make a beautiful creation and it is his fault that others will harm his creation.
- Victor admires the monster's strength.
- Victor leaves to sleep.
- The monster is curious and goes to see Victor.
- Victor is hoping to see signs of love in the monster's eyes, but sees none.
- The monster feels nothing for Victor.
- Victor feels rejected by the monster and this leads to bitter disappointment because Victor has been hoping for two years to build someone that will love him. He has invested a great deal of effort in this project.

- Victor feels shame that he has not provided his monster with a family where members love one another. He can't bear his shame and leaves.

*Story morph four:*

- Victor is in a strong adversarial relationship with (toward) the monster he has created. His goal is to create life that he can mistreat with impunity. The monster has a strong friendship relationship with (toward) Victor.
- Victor is gloating because his monster has ugly skin and he will be able to use this against the monster. He is sad that the monster has good proportions. He is afraid that because the monster is so strong he will not be able to mistreat him very extensively.
- Victor looks forward to the moment when the monster realizes that Victor despises him.
- Victor leaves to get some rest.
- The monster comes to life and has a strong desire for human contact. He is lonely, but is hopeful of being loved by others.
- He finds Victor and is satisfied to find human company. He feels love for Victor and reaches out to him.
- Victor expresses his disgust at the monster. He is very excited about the impending feelings of rejection the monster will feel.
- The monster is now terribly sad to be rejected by Victor.
- Victor gloats over the monster's distress.
- The monster gets angry at Victor for behaving so badly.
- Victor fears that the monster will hurt him and runs away.

*Story-morph snippet five:*

- Victor loves his monster very much. He believes that parents should love their children and also that they should always find their children beautiful despite their faults. He is proud of loving his creation despite his ugliness, but he is remorseful that he finds the monster repulsive. He does not express his love strongly because such emotions are *repressed* in his temperament, but his temperament is also such that disgust is shown in a *communicative-verbal* way. He calls out at the monster saying "Disgusting!"
- The monster, who desperately wants love, does not see that Victor loves him, but only that Victor is repulsed by him.

*Story-morph snippet six:*

- As above in story-morph five, but...
- The monster is very happy to find himself alive. This puts him in a very good mood, and he is predisposed to appraising the world in a positive light.
- The monster ignores Victor's comment that he, the monster, is disgusting, but notices Victor's *obsessive attentional focus* on him as an object of love. He is satisfied to feel Victor's love.

*Story-morph snippet seven:*

[Using the case-based reasoning capabilities of agents]

The monster knows that he has pasty skin and yellow eyes. He sees that Victor is disgusted. Taking these features together he changes his his internal representation of how he believes Victor sees the world. He feels sorry for Victor because he, the monster, is so ugly, and now believes this makes Victor unhappy.

*Story-morph snippet eight:*

[Using extensions for humor]

- The monster sees Victor as an authority figure. He believes Victor holds everyone to high standards of behavior. He believes Victor will hold him, the monster, to high standards. Victor has created a monster that is ugly, thus violating one his own standards. The monster observes that Victor knows the monster knows that Victor has violated his own standards. The monster finds Victor's chagrin funny, and laughs.
- Victor is embarrassed. He resents the monster laughing at him.

*Story-morph snippet nine:*

[Using extensions for altering Concerns-of-Others structures]

- Victor looks at his sleeping monster and feels pity for him because he assumes the monster will feel very badly about being so ugly.
- Victor leaves to sleep.
- The monster wakes and is happy to be alive. This happiness trumps all other feelings.

- Later when the monster comes, Victor realizes that the monster is happy, and changes how he believes the monster perceives himself. He stops feeling pity and now feels happy for the monster that he enjoys being alive.

*Some finer-grained variations:*

- The monster is now terribly sad to be rejected by Victor:
- The intensity of this rejection is increased because this is unexpected by the monster, and the *surprisingness* contributes to intensity.
- The *importance* of not being rejected is very high for the monster, and this increases the intensity.
- Feeling rejected causes the monster to express this by *reappraising [himself]* as being ugly and unlovable.
- The disgust on Victor's face is taken by the monster to be a very strong indication of intense dislike, and this *[sim-event variable]* contributes to the intensity.
- Victor's temperament is manipulative and he expresses his disgust through *Other-directed emotion modulation* channels to make the monster feel as badly as possible.

Each of these snippets from different story-morphs is based exclusively on components of emotion that the Affective Reasoner system can manipulate. In addition, because they are based entirely on the logical structure of how emotions arise, and are expressed, they come with sophisticated explanations, embodied as explicit values in what might be hundreds of details for each emotion generated.

#### **Implementation:**

Current technical development is focused on putting the agents on the web using AWS Linux, perl, python3, PHP, the AI engine in ABCL/SBCL Common LISP, Java networking, the Google speech engine, Google speech recognition, Chrome browser, websockets, Javascript, Midi-to-MP3 for automatic music generation and browser-based SVG for morphing faces. The current design focus is on building a corpus of modifiable, common, emotion-story schemas as a basis for constructing compassionate software agents.

#### **Conclusions and summary:**

Complex, but precise, emotion structure can be teased out of all stories, based on how the characters appraise the events that arise in the narrative. This emotion structure is essential to what makes a story a story. The emotion structure is both portable (it can be repeated in an entirely different context) and subject to manipulation. The Affective Reasoner, which does manipulate such emotion



structures, can be used to automatically generate new, and novel, stories which are nonetheless internally consistent, because of the consistent nature of the artificial personalities that are dynamically constructed by the computer.

We have discussed the nature of what makes a story a story, and not simply a collection of plot events, claiming that how a character *cares* about unfolding events is critical, and also, necessarily, gives rise to emotions. We were introduced to the underlying emotion theory that drives the Affective Reasoner. We worked through several snippets of stories to see how story-morphing works, and concluded with a more extended example from chapter five of Mary Shelley's *Frankenstein*.

## Works Cited

- Bruner, J. (1990). Culture and human development: A new look. *Human development*, 33(6), 344-355.
- Bruner, J. S. (1957). On perceptual readiness. *Psychological review*, 64, 123.
- De Keersmaecker, J., Dunning, D., Pennycook, G., Rand, D. G., Sanchez, C., Unkelbach, C., & Roets, A. (2020). Investigating the robustness of the illusory truth effect across individual differences in cognitive ability, need for cognitive closure, and cognitive style. *Personality and Social Psychology Bulletin*, 46, 204-215.
- Elliott, C. (1991). *The Affective Reasoner: A process model of emotions in a multi-agent system*. Ph.D. dissertation, PhD thesis, Northwestern University, May 1992. The Institute for the \_\_\_\_.
- Elliott, C. (1997). I picked up catapia and other stories: A multimodal approach to expressivity for \_\_emotionally intelligent\_\_ agents. *Proceedings of the first international conference on Autonomous agents*, (pp. 451-457).
- Elliott, C. (1999). Why boys like motorcycles: using emotion theory to find structure in humorous stories. *Unpublished paper, School of Computer Science, DePaul University, Chicago*.
- Elliott, C., & Ortony, A. (1992). Point of view: Modeling the emotions of others. *Proceedings 14th Annual Conference of the Cognitive Science Society*, (pp. 809-814).
- Elliott, C., & Siegle, G. (1993). Variables influencing the intensity of simulated affective states. *AAAI Spring Symposium on Reasoning about Mental States: Formal Theories and Applications*, (pp. 58-67).
- Elliott, C., Brzezinski, J., Sheth, S., & Salvatoriello, R. (1998). Story-morphing in the affective reasoning paradigm: generating stories semi-automatically for use with \_\_emotionally intelligent\_\_ multimedia agents. *Proceedings of the Second International Conference on Autonomous Agents*, (pp. 181-188).

- Elliott, C., Rickel, J., & Lester, J. (1999). Lifelike pedagogical agents and affective computing: An exploratory synthesis. In *Artificial intelligence today* (pp. 195-211). Springer.
- Fodor, J. A. (1975). *The language of thought* (Vol. 5). Harvard university press.
- Gray, R. M., & Bourke, F. (2015). Remediation of intrusive symptoms of PTSD in fewer than five sessions: a 30-person pre-pilot study of the RTM Protocol. *Journal of Military, Veteran and Family Health*, 1, 13-20.
- Kim, E., Padó, S., & Klinger, R. (2017). Investigating the relationship between literary genres and emotional plot development. *Proceedings of the Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, (pp. 17-26).
- Lakoff, G. (1993). The contemporary theory of metaphor.
- MasterClass. (2021, 02 15). *Complete Guide to Literary Themes: Definition, Examples, and How to Create Literary Themes in Your Writing*. Retrieved 02 15, 2021, from MasterClass: <https://www.masterclass.com/articles/the-complete-guide-to-narrative-theme-in-literature-definition-examples-and-writing-how-to-what-is-a-literary-theme>
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102, 419.
- Ortony, A., Clore, G. L., & Collins, A. (1988). *The cognitive structure of emotions*. Cambridge university press.
- Pylyshyn, Z. (1999). Is vision continuous with cognition?: The case for cognitive impenetrability of visual perception. *Behavioral and brain sciences*, 22, 341-365.
- Reagan, A. J. (2017). Towards a science of human stories: using sentiment analysis and emotional arcs to understand the building blocks of complex social systems.
- Shank, R. C. (1990). Tell me a story: Narrative and intelligence. *Evanston, IL: Northwestern University Process*.
- Wikipedia. (2021, February 15). *Love of Chair*. Retrieved 02 15, 2021, from Wikipedia: [https://en.wikipedia.org/wiki/Love\\_of\\_Chair](https://en.wikipedia.org/wiki/Love_of_Chair)

---

<sup>i</sup> With a tip of the hat to Pylyshyn's cognitively impenetrable early vision module (Pylyshyn, 1999).