

Laugh To Me

(extended abstract of the MSc dissertation)

André Filipe Rodrigues de Carvalho

Departamento de Engenharia Informática

Instituto Superior Técnico

Advisor: Doutora Ana Maria Severino de Almeida e Paiva

Abstract—The growing interest for Interactive Storytelling, has taken research to explore the potential of this new media to the classical story genres. In our research, we develop autonomous agents that act in a storytelling context with a comic purpose. In this paper we present a comic sketch model for autonomous agents with affective reasoning. The agents that use this model prepare the timing of the comic punchline timing by reasoning about emotional states in a process called Emotional Escalation. The punchline used, as well as the personality of the comic characters is based on a theory of incongruity-resolution.

I. INTRODUCTION

Humour is the quality of what causes laughter or is meant to do so [1]. We can thus think of humour as the process of evoking laughter, but it is more than that. Rod A. Martin proposes a definition of humour from the psychological point of view comprising four elements: a “social context”, a “cognitive-perceptual process”, an “emotional response” and the “vocal-behavioral expression of laughter” [2]. We focus especially on the cognitive-perceptual process – why and how we laugh.

Interactive Storytelling has been a topic of growing interest from the Artificial Intelligence community, and there have been a number of works in the area such as Façade [3] or FearNot [4]. The majority of these, however, are interactive dramas. Only very few are comedies, and the topic of how humour can be modeled in the context of Interactive Storytelling has only been marginally approached. The history of Computational Humour (the subfield of Artificial Intelligence that concerns the production of humour) in its whole is somewhat more rich, but deals primarily with Natural Language systems. Interactive Storytelling is often used to build games, including serious games (games in which the main purpose is *not* entertainment) such as FearNot! [4] or IN-TALE [5]. Serious games could benefit from several of the qualities attributed to humour such as building trust [6], enhancing human-computer interaction [7], [8], [9] or enhancing agent believability [10].

In this research we consider sketches, a comedy form that consists of small scenes confined in themselves, that is, with beginning, middle and end. Our research question is thus:

“How can we endow autonomous agents with the capacity of creating a comic situation or sketch, specially concerning its pace and timing?”

To do this we propose establishing a connection between a sketch and its emotional content. We think of the agents as both characters and active actors who play the roles of those characters and contribute actively to the emotional content of the scene. Our hypothesis follows from this and states that:

“Agents can control the pacing and timing of a sketch by using emotions to define the evolution of the story and contributing actively with their actions to produce an emotional escalation.”

We propose a model in which the agents try to contribute in a controlled manner to an Emotional Escalation of the sketch. We implement a prototype system of this model as a proof-of-concept. In this document we discuss the related work in computational Humour and Affective Storytelling (section II). We then provide a very brief background on the thoughts on humour that inspired our model (section III). We discuss our proposed sketch model and comic agents (section IV and its subsequent implementation (section V). Finally we present an evaluation of the prototype (section VI) and conclusions and future work (section VII).

II. RELATED WORK

A. Computational Humour and Natural Language

Computational Humour (the subfield of Artificial Intelligence concerned with the production of humour) has been mostly connected with works in Natural Language. The first working prototype of a Computational Humour system was probably Lessard and Levison’s 1992 Tom Swifties generator [11], which used the VINCI Natural Language generator. Tom Swifties are puns in which a character’s (Tom) quoted sentence is described by an adverb that bears relation to that sentence. For example: “‘this matter is opaque’, said Tom obscurely.” Following this work, Lessard and Levison presented another prototype for making riddles [12], which also used VINCI and worked in a similar manner.

A more ambitious approach was JAPE, developed by Binsted and Ritchie at the University of Edinburgh [13],[14] was a more ambitious attempt, featuring WordNet, a much bigger lexicon than those used by previous pun generators. JAPE was able to produce certain types of simple punning riddles, of the form “What’s the difference between A and B?” and “How is A like B?” for example. JAPE was later

used to develop STANDUP [15], a system aimed at helping children with cerebral palsy develop better language skills [16].

Another important work in computational humour was Stock and Strapparava's HAHACronym, a simple prototype that sought to produce ironic acronyms [17]. HAHACronym also used WordNet, extended with hierarchic domain information about the terms. These domains are used by HAHACronym to explore the incongruity between groups of concepts, for example, Sex vs. Religion. HAHACronym can already be considered a cross between referential and verbal humour [18], that is, it is not simply a play of words (verbal), it deals also with their meaning (referential).

B. Interactive Comedy

Comedy as a genre has been pursued in various works of Interactive Storytelling. This has often been based on planning formalisms that allow the character to fail. The comic situations arise from the character failing to meet their goal.

One prototype by Cavazza et al. at the University of Teeside [19], [20] was based on the sitcom Friends. However, the sitcom genre was chosen more because the storyline is relatively simple and fitted the character-centered paradigm than as an explicit attempt at making humour. Yet, the prototype did result in some funny situations, which emerged from the failure of a character's plans. Cavazza reckons that the situation of two characters that have different conflicting goals is "likely to result in a series of comic situations and *quiproquos*." A similar work with a different planning mechanism was based on the Pink Panther cartoon [21]. In this Pink Panther prototype, the potential comic situation results of a plan failure (that is allowed by the planning formalism) and the long-range dependencies between actions, in which the character does not foresee all the consequences. Thawonmas et al. proposed a similar model [22], noting that there should be some control how the plan failure occurs.

Another important work by David Olsen and Michael Matheas further developed the plan failure method in ACME, a prototype system set in the world of the Coyote and Road-Runner cartoons [23]. The Coyote and the Road-Runner are two autonomous agents. The behaviour of the Road-Runner is fairly straightforward: he is either standing still or running away from the Coyote. The Coyote also has two states: chasing the Road-Runner and planning the use of an item in order to try catch the Road Runner. The system itself has the story goal of frustrating Coyote's plan through the occurrence of some gag, such as an anvil falling over Coyote. The occurrence of these gags take in account a level of anticipation, that grows with the number of steps the Coyote has gone through in his plan to catch the Road-Runner. This level of anticipation sets the probability of a gag occurring that causes the Coyote to fail his goal.

C. Agents with Emotions and Affective Storytelling

There have been many works in Interactive Storytelling concerned about the agent's emotions. It is a cross between

Affective Computing and Interactive Storytelling, sometimes called Affective Storytelling. According to Bates "emotion is one of the primary means to achieve believability" [24] in agents. One of this works, Clark Elliott's Affective Reasoner [25] is an appraisal system that relates emotions with story variability. Appraisal is the process by which the agent attributes emotions to his perceptions of the world. The premise of the Affective Reasoner is that two stories that are, for the most, identical, in terms of the events, are perceived as different because the appraisal of actions done by the characters is different.

Elliott suggested the use of his Affective Reasoner paradigm in the context of Computational Humour, exemplifying with a very specific type of humorous situation [26]. The situation is that of an authority figure, that defends a certain principle, and a victim or a group of victims whom it accuses of violating that very principle. Elliott gives the example of a school dean, who, in faculty meetings, criticized those who were late, saying that, as the meetings were scheduled in advance, they had no reason not to show up on time. Everyone in the meeting thought it was funny when the dean himself was late in the following meeting. Elliott proceeds to analyse this situation in terms of the Affective Reasoner. He notices that the way characters feel about the situations is the same in many different, but similar scenarios:

- 1) the authority figure is angry at the victims over one principle;
- 2) the victims feel ashamed;
- 3) the authority figure breaks that very principle;
- 4) the authority figure feels ashamed;
- 5) the victims gloat.

This is – as far as we know it – the only work (though never implemented) that makes a connection between the emotional content of a scene and the possibility of humour. Other works in agents with emotions are interesting because they also involve agents that have active roles in changing the emotional content of the scene. Some of these are also based on FATiMA, an architecture for agents with emotions that we also extended for our prototype.

One work suggested that FATiMA agents could support a double appraisal mechanism, in which the agent reappraises a selected action according to the emotional impact in others [27]. This could make agents behave more like actors and less like characters, who evaluate the dramatic interest of an action. Indeed double appraisal mechanisms have been shown to create more interesting narratives [28].

Another extension of FATiMA that has been proposed and is currently being implemented aims at making emotional intelligent agents [29]. An agent with emotional intelligence should be able to:

- use emotions in decision making;
- understand and reason about emotions (for example, Anger indicates a possible intention of harming other);
- manage emotions, this is, try to elicit certain emotions in oneself and in others.

An Interactive Storytelling in which emotions had a lot of focus was a work by Pizzi et al. based on Gustave Flaubert's novel *Madame Bovary* [30],[31],[32]. They remarked that this was a "psychological novel" and, as such, character's feelings and relationships with other characters had a greater role than the actual action outcome. They proposed describing narrative in terms of the agents feelings.

The contribution of each action to achieve the character desired emotional state is given by an heuristic function, in which a low value means that the character is closer to a desired emotional state and a high value means it is more distant. Pizzi et al. suggested the use of values of this function to measure narrative progression from the point of view of the character.

The character can thus be aware of how its situation evolves, and have specific feelings about it. For example, an initial decrease of the heuristic, which gives a character hope, followed by a prolonged increase. This can, according to Pizzi et al. "correspond (...) to the narrative notion of 'shattered hopes' "[31].

Façade, a pioneer project in Interactive Storytelling by Mateas and Stern [3], is different from other systems mentioned until now because story progression is not purely character based, but also plot based. This means that the action does not evolve simply because of the actions each character does, but also as the result of a plot that is controlled at a high level by the drama manager. This drama manager tries to steer the story according to an author defined tension arc.

III. BACKGROUND IN HUMOUR

Out of the big humour theories that were discussed across centuries, the one that is more closely related to our model is the incongruity-resolution theory. Following Thomas Shultz ([2], pg. 64) the punchline (ending) of a joke creates an incongruity that contrasts with what was suggested by the set-up of the joke. One must go back and search for an ambiguity, in the set-up and interpret it in a different way in order to get the joke. This is called an incongruity-resolution theory because first we are surprised by the incongruence and then that incongruence finds an explanation (is resolved) in the final part of the joke.

Incongruity is a concept that seems to be necessary to humour. Perret, a comedian who writes about comedy writing, states that normally "a joke comprises two different ideas that come together to form one" [33]. We consider two main uses of incongruities in comedy writing: as *comic premises* and as *punchlines*. Comic premises are the initial idea behind a comic scene or a joke, such as a man interviewing a dictator about his love of botanic. The punchline is the ending part of a joke or a scene, that resolves the incongruity and provides the humour.

John Vorhaus, another comedy writer that wrote books on the subject [34], and whose perspective on comedy writing is strongly character-centric, introduces the concept of *strong comic perspective* in relation to comic characters ([34], pg. 42). A strong comic perspective is a point of view by a

character that is very unique and related to specific traits of his or her personality. This point of view is unlike that of a normal person.

An important factor in comedy is building up the *tension* (set-up) before delivering the joke (pay-off). Perret, whose book is mainly directed at stand-up comics, also makes the point that separate one-liners can work better if sequenced and grouped in a routine ([33] pg. 85). According to Perret, it also helps to cover for weaker jokes, which are laughed at, and seem funnier, because the spectator gains momentum from previous lines. This is consistent with Provine's observations that laughter is contagious and people deem the jokes they laugh at funnier than those they do not [35].

Sketches are short, isolated scenes that develop a certain comic premise (a comic premise is the incongruity that composes the initial idea of a comic story). Perret refers to an analogy of jokes as the building blocks of comedy. He remarks a sketch is not just a collection of jokes, much like a house is not just a collection of bricks. Perret considers a good sketch must have "a premise; some complications; an ending, or in other words a beginning, a middle and an end" ([33], pg. 154). In his account of what a sketch should be, Vorhaus stresses the need to create and develop a conflict between characters ([34], pp. 154-161).

IV. MODEL

There is no consensual theory of humour nor a theory with a description precise enough to be almost directly translated to a computational system (outside of linguistic theories). Our model results mainly of the general ideas presented by the incongruity-resolution theory and the more practical knowledge of comedy writing.

We define incongruence as a gap or inconsistency, which can be exploited for comic purposes. In this sense, the comic premise, which Vorhaus defines as "the gap between comic reality and real reality" can be described as the incongruence that establishes the high concept of the sketch. A comic reality is like a caricature of real reality, a distorted vision in which some features are emphasized, downplayed, or contrasted with their opposite. For the purposes of our model we categorize incongruences into three types: action, emotional and context.

An *action incongruence* is the one in which the actions of the character conflict with the character's ultimate goal or even the simple common sense of how he should act. This has been the type of incongruence more often explored in Interactive Comedy. For example, in the Pink Panther prototype [21], an example given was that the Pink Panther could not sleep because of a leaking tap. He then went to cut off the water supply, getting dirty in the process. Before coming back to bed the Pink Panther tries to take a shower but fails because he has just cut off the water supply. The incongruence here is the gap between the intention behind the action and its actual effect.

An *emotional incongruence* can be defined as a contradiction between the characters expected and actual appraisal of events. A good example is a sketch of the British comedy

show *A Bit Of Fry And Laurie*, in which a client enters a store to buy a Jason Donovan single [36]. Jason Donovan is an Australian pop singer, considered cheesy at times. The client felt overwhelming shame when asking for Jason Donovan’s music, so he tried to cover it up, by asking for condoms as well. Buying condoms is usually not something people boast about or feel very comfortable buying. However the client speaks loud and clear when asking for the condom packets, and in a low, hindered voice when requesting Donovan’s single. The actions asking for condom packets and asking for a Jason Donovan single have unexpected and incongruent appraisals. The sketch evolves by the seller repeatedly asking the client to clarify his request of the Jason Donovan single, until the client’s growing shame and distress lead him to leave the store without accomplishing his initial goal. In this example of emotional incongruity, the disparity results from the overreaction of one of the character’s but the incongruity could also result of a gross devaluation of the event.

A *context incongruity* happens when the character acts as if it was in a different context than that it actually founds him or herself in. A common gag in sitcoms consists of showing some intervention as to stop another character from abusing drugs or other addiction, but it actually refers to something much more harmless, say a compulsive need to own Jason Donovan’s albums.

As mentioned before we use this definition of incongruities in defining aspects of the sketch, such as the premise, the punchline or the personalities of the characters involved. We define a comic character precisely as a character that acts or reacts in an incongruent manner relative to what is expected of him, given the situation. The main source of incongruence in our model, being basically a character centered approach, is a character personality. Thus the client that goes into the store to buy the Jason Donovan single is a comic character. The seller, who does not show any unusual behaviour, is a regular character.

We define a sketch as a scene structured into three parts, beginning, middle and end. The beginning sets the conflict between the characters. This conflict is related to the premise of the sketch which, as we have seen, is an incongruity. A conflict needs a context (for example, a shop), a relationship between the characters (such as that between a seller and a client who wants something) and a subject of conflict (such as something the client wants and the seller will not give to him). The end part is what we call resolution, though only in the sense it concludes the sketch, not so much as in incongruity-resolution theories in which it also means that it solves an incongruity, making the viewer rethink a joke or, in this case, the sketch. When a resolution also does this process, however, we call it a *punchline*.

Sketches only reach their full potential when the sequence of actions is defined is properly paced and timed. In between the beginning and the end, characters explore the comic premise by adding to the scene tension. The characters build up on their original conflict, until an emotional peak

is reached. A buildup, however, needs to be gradual. The conflict is explored until it reaches an extreme emotional peak that triggers the resolution. We call this progressive rising of emotions towards a peak an *Emotional Escalation*.

We propose that each agent, instead of representing not only a passive character that acts according to its traits, but that can also act to arouse, in itself and others, emotions that can lead to an Emotional Escalation and, ultimately, to a comic resolution. Thus we propose the agents have an action selection mechanism that takes in account the emotional impact of the actions in the scene. We present our conceptual model for a comic agent in figure 1.

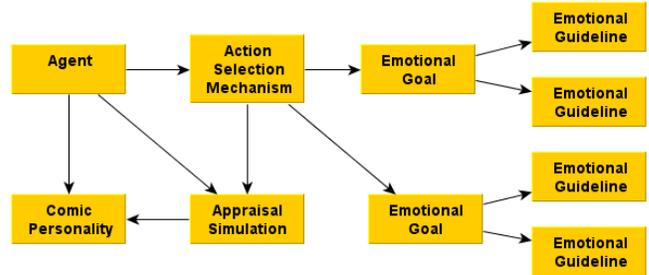


Figure 1. Comic agent conceptual model.

In figure 1 we can see each agent has an action selection mechanism that takes in account the Emotional Goals and Guidelines of the agent. This mechanism must comply with our progressive emotional escalation model. This means it should not only depend on what emotions the agent wants to evoke, but also on how he wants to evoke them over time. We propose the agents selection mechanism takes in account Emotional Guidelines. An Emotional Guideline is a monotonically increasing function that maps the scene time (which may be represented by the number of character’s actions) into the desired emotional output for a given emotion (we assume in our model that the emotion model used provides some sort of measurement of the intensity of emotions at a given time; in FATiMA it could be the emotion potential). This defines the *pacing* of the sketch, because we can make the Emotional Escalation grow faster or slower by simply changing the growth parameters of the Guidelines.

As can be seen in figure 1, we group Emotional Guidelines into Emotional Goals. Each Emotional Goal should be activated by a number of preconditions, so that there is better control to when and which Emotional Guidelines are activated at a time. This could allow the agent to shift Emotional Guidelines during the sketch, by taking into consideration changes in the world – for example, reactions from other agents.

To influence action selection in order to comply to a given Emotional Guideline, i.e., to try and select actions for which the expected emotional output is as similar to what is defined by the guidelines as possible, we propose a specific heuristic. We assume the agent is capable of appraising actions and

events to generate emotions which can be measured – we call this value the emotional output. The emotion output a given guideline defines at a certain point is the *desired emotional output*. We call our heuristic cost the *emotion weight*. The agent tries to select a plan that minimizes the value of the emotion weight.

The emotion weight defines how much an action contributes to make the scene comply to the Emotional Guidelines. In considering this there are some aspects to have in mind:

- the worst case scenario is when an action does not generate any of the emotions defined by the various guidelines;
- the same action can contribute to one guideline defined emotion and not to other that is also active;
- an action may contribute to a guideline but not to increase the current potential of the emotion;
- we need to consider how far the expected potential is from the desired emotional potential at a given time;
- the best case scenario is an action that increases or maintains the current emotion potential and has the minimum difference to the desired emotional potential.

To know whether an action contributes to a given Emotional Guideline or not, we *simulate the emotional state*. This process is similar to that of the double appraisal work described in II-C. The agent appraises the event that would be generated should he execute that action, and gets a *simulated emotional state*.

This simulated emotional state must be sufficient to know whether the action generates the emotion defined by each guideline and if so what is the *expected output* of that emotion. When calculating the heuristic, we need access to all of the currently active Emotional Guidelines (this is, guidelines that pertain to active Emotional Goals). We calculate the emotion weight for each guideline separately, taking into account special cases, such as when an action does not contribute for any active guideline.

When evaluating the contribution of the action for each active Emotional Guideline, there are several cases to consider. If an action generates the emotion defined by the guideline, we consider the absolute value of the difference between the expected emotional output and the desired emotional output. It does not matter whether the expected emotional output is above or below the desired emotional output. We want the growth to be as similar as possible to that of the emotional guideline and that is why we consider the absolute difference.

What is important to take into account, though, is whether the last action selected by the character was stronger than this one or not (since Emotional Escalation should be monotonically increasing). When calculating the emotion weight of an action, we consider whether its expected emotional potential is higher or lower than that of the last action selected. If it is smaller we add a *penalty* to the emotion weight value. This penalty needs to have some importance but not so much that it rejects the action no matter if it fits

all other guidelines better or not.

If an action does not generate the emotion defined by the guideline, we calculate the emotion weight taking into account the expected emotional output as 0 and add the expected emotional output of the *last selected action*. This rule is as penalizing as the last action was strong. This allows the agent to, at first, consider actions that only contribute to certain emotional guideline and, as the action develops, to be stricter with which actions he chooses.

Also, in our model, the contribution to each guideline is relative to its importance to the scene. Since we want our emotions to escalate, guidelines that currently have the biggest desired emotional outputs are more preponderant. As such, the contribution of each guideline is multiplied by a factor that reflects how an Emotional Guideline currently compares to others. This factor is given by calculating the fraction of the desired emotional output of the guideline currently being tested relative to the total sum of the desired emotional outputs of all the guidelines.

Finally, when an action does not contribute to any of the guidelines, it should be more penalized than it is by this last rule. In this case the emotion weight is a constant value that is equal to the maximum emotion weight that could be defined by the previous two rules. This emotion weight is multiplied it by the number of active guidelines plus one, to ensure the action (or the plan containing it) is always rejected if there are alternatives.

If the character is comic, he should have a personality that is *emotional incongruent* with what should be expected of him. The rules that define his appraisal of the world should be incompatible with what one would normally expect. That allows for him to choose comic actions and create a conflict with other character.

V. IMPLEMENTATION

A. FAtiMA and emotions

To implement our agents we used the FAtiMA framework [37] with the OCC model as the underlying model of appraisal. OCC stands for Ortony, Clore and Collins [38], who proposed an appraisal model that has been widely used to make agents with emotions in Interactive Storytelling. OCC encompasses in total 22 valenced emotion types. Valenced means that these emotions always have a negative or a positive charge: for example Joy has a positive valence while Distress has a negative valence. OCC proposes a hierarchical organization of these emotion types according to how we appraise them.

The OCC model considers several appraisal variables. The ones we make use of in our implementation are *Praiseworthiness*, *Desirability*, and *Desirability for other*. Praiseworthiness refers to the moral aspect of an action, whether it is right or wrong. Desirability defines how pleasing an event is, in general. A desirable event is considered generally desirable irrelevant of the perspective of the agent (being the subject or the target). Desirability for other concerns only how pleasing an event is for other than the agent itself.

These appraisal variables give information about the agents assessment of the world are then converted to emotions. The emotions shown in our prototype are Joy, Distress, Gloating, Pride and Anger. Distress and Joy are well-being emotions that depend solely on the Desirability appraisal variable – if the appraisal variable has a positive value it causes Joy, while a negative value causes Distress. Gloating is a *fortune of others* emotion. It depends on a positive value of Desirability and a negative value of Desirability for other.

Pride is an attribution emotion. Attribution emotions depend on the value of the Praiseworthiness and of the subject of (i.e. who performed) the action. If the agent itself performed the action it can feel either Pride (if the value of Praiseworthiness was positive) or Shame (if it was negative). If the action was performed by other, the agent feels either Admiration (if the Praiseworthiness was positive) or Reproach (if it was negative). Anger is a composed emotion that results of mixing Distress and Reproach. Thus Anger happens when someone else’s action is appraised as both undesirable and blameworthy.

An OCC emotion has a potential, a threshold, an intensity and a decay rate. The potential is the initial value attributed to the emotion after appraising. The threshold is the minimum limit beyond which we do not feel a certain emotion. The intensity is the value with which the emotion is actually felt, and is given by the difference between the potential and the threshold. Finally the decay rate defines how fast emotions fade with time.

B. Emotional Goals and Guidelines

We defined our emotional output as the potential of the emotion, and our scene time as the actions of the character after a given Emotional Goal is activated. Therefore an Emotional Guideline maps the number of a character’s actions into the desired emotional potential at that moment.

We provided three different implementations of Emotional Guidelines, represented by three mathematical functions. These were a Sigmoid function (which grows fast in the beginning and tends to a certain specific value – an horizontal asymptote); a Quadratic function (starts slow then grows quadratically to infinity); a Linear function (which grows at a constant rate). All functions have a cap of 10, which is the maximum emotional potential of an emotion in FATiMA.

Also, we did not take into account, when calculating the heuristic, the threshold value, thus taking in account actions that are not necessarily felt by the character.

C. Scenario Outline

We required for our scenario: two characters, at least one of whom should be a comic character; an object of conflict between these two characters; and a reason to keep the characters together throughout the sketch. We set our sketch in a pastry shop, involving one Client and one Seller. The Client is a regular character, while the Seller is the comic character who refuses to sell the cake (thus the cake is the object of conflict and their client/seller relationship what

binds them together). The Client is obese, and the attitude of the Seller ranges from being plain insulting to stress the fact he is overweight as a reason not to sell the cake.

The punchline would be the Seller trying to sell something else. When the Client gets angry, the Seller tries to sell anti-depressants. After the Client refuses to buy the pills, the Seller, failing to see the inconsistency of his own actions, blames the crisis for the fact he did not sell the other product. Since this is the punchline, the Seller has as a goal to anger the Client.

D. Scenario Authoring

Authoring is the process by which we define the scenario in FATiMA. It allows one to specify the personalities of the agents, the actions that are available to them, the objects that are part of the world and predicates about the agents and the objects. It is through this parametrization that different FATiMA agents are defined.

Since our model is based on a structured plot, we defined a state object, which holds the plot information, such as who has the *turn to talk*, and what the *moment of the sketch* is. The sketch moment property refers to the current moment of the sketch, as defined by the model: Beginning, Middle and Punchline (throughout this section, punchline and resolution are used synonymously). Each agent also has the specific goal of interacting, that is to make pass their turn to talk to the other agent.

The *sketch moment* property is used to filter which actions are available to the characters as the plot develops. The beginning and punchlines are scripted, meaning that there is only one course of action defined for the characters (though, theoretically, an agent could choose to activate different resolutions, once a punchline is activated he has a strict sequence of actions to follow).

As a comic character the Seller has a *comic perspective*; in our scenario the Seller is a very unpleasant person, who takes pleasure in actions that are undesirable for others. The `Insult`, `MakeFatPeopleJoke` and `MakeSarcasticRemark` actions are appraised taking this into account. When the subject is the Seller he appraises these events as desirable for him but not for others. When the action subject is someone else, however, the Seller appraises these events as undesirable. Thus, this character likes to insult other characters but does not like to be insulted. Note that the stereotypical relation between a seller and a client – “the client is always right” – is incongruent with this appraisal. Though the aforementioned actions all follow the same pattern (high desirability for self, low desirability for others), they must have different gradations of those variables.

As a regular character the Client has appraisal rules more in line with what is expected of him. The Client, appraises actions such as `Insult`, `MakeSarcasticRemark` and `MakeFatPeopleJoke` as blameworthy and undesirable both to others and to himself. While he also finds undesirable for the Seller to `RaiseMoralIssues` in order not to sell him the cake, he finds it justified and, as such,

praiseworthy. Since his patience would be put to test by the Seller, we focused on the actions the character could use to recover pride after the Seller’s negative actions. Actions such as TryCalmDown, DemandRespect and FormallyComplain as praiseworthy. While these last two actions have the same high praiseworthiness value, their desirability is different, as the Client finds more desirable the DemandRespect action than the FormallyComplain.

E. Seller Emotional Goal authoring and output

The goal of the Seller as a comic character is to annoy the Client. As such two emotions that the Seller will likely want to arouse are arousing Distress in the Client and Gloating in himself. Note Distress is the result of a negative desirability, while Gloating is caused by appraising an event as desirable, but not desirable for others. In our prototype we tested two Emotional Goals. However, here, for simplicity reasons, we consider only the Emotional Goal without Gloating. This Emotional Goal, DispleaseClient-A, consists of an initial sigmoid curve of the Client’s Distress that is followed by an exponential growth (defined by another guideline). This makes for a change of pace in the sketch, that starts slow, but escalates fast afterwards.

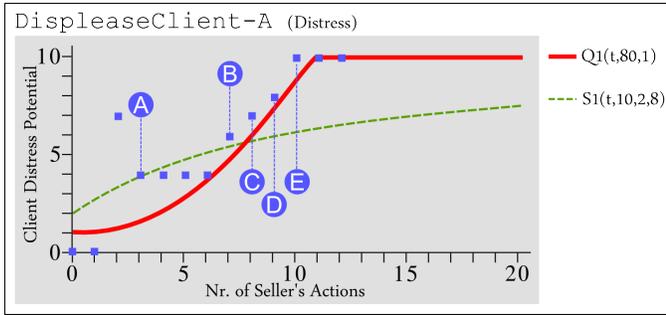


Figure 2. Emotional Goal DispleaseClient-A. Letters represent actions: A-RaiseMoralIssues, B-Reason, C-WarnHinderAppetite, D-MakeSarcasticRemark, E-MakeFatPeopleJoke

The output resulting of this Seller’s Emotional Goal can be seen in figure 2. We can see both Distress Emotional Guidelines that are part of this Emotional Goal, a Sigmoid, we dub S_1 and a Quadratic we call Q_1 . Before explaining how each action was selected, recall Distress is the result of negative Desirability. Note as well that, since the desired effect is to distress the Client, the emotional output shown in figure 2 refers to the model of other the Seller has of the Client.

At the beginning, the most influent guideline (remember the higher the emotional potential defined by the guideline, the more importance it is given by the heuristic) is the Sigmoid, S_1 . This Sigmoid has a slow growth rate, and this is why the initial action chosen – RaiseMoralIssues repeats several times, as the Quadratic guideline grows to catch it.

The sketch only proceeds when the Quadratic guideline’s value is bigger than that of the Distress

potential of action RaiseMoralIssues, choosing a more undesirable action, which is Reason. From here on, the exponential guideline gains preponderance, making the sketch evolve at a faster pace. The Seller selects the action WarnHinderAppetite, followed by MakeSarcasticRemark and finally MakeFatPeopleJoke. These actions are more and more undesirable, leading to a growing Distress of the Client. The MakeFatPeopleJoke is appraised by the Client as especially undesirable (and also as undesirable for the target of the action), which angers the Client in such a way that it triggers the punchline.

Note how the shape of the guidelines helps set the pacing of the sketch. In a second prototype that defined a more complex Emotional Goal for the Seller that included a Gloating Emotional Guideline, we also reduced the growth rates of the several guidelines. This resulted in a longer sketch, in which the Seller repeated some actions, such as MakeSarcasticRemark. By adding the Gloating guideline some actions also were not chosen for the sketch.

F. Client Emotional Goal authoring and output

The Client, being a regular character, could perhaps be purely reactive, as its actions are mainly reactions to the Seller’s inappropriate behaviour. However we defined his behaviour through an Emotional Goal. The initial goal of the Client is to accomplish his goal of getting the cake; we define this as a Joy emotional guideline. However, as the Client is provoked, he will need to react by making actions that recover his hindered pride. As such the Client also has two Pride emotional guidelines.

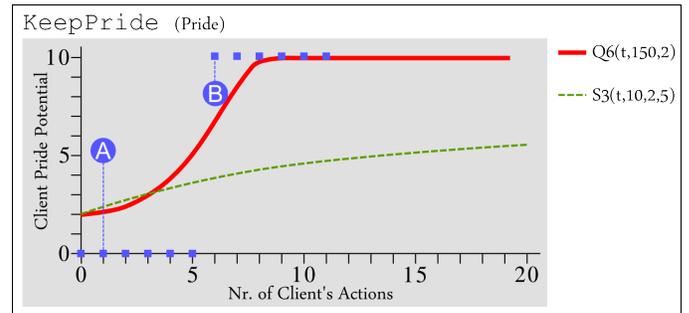


Figure 3. Emotional Goal KeepPride. Pride guidelines only, Joy guideline not shown. Letters represent actions: A-AskCakeOrCandy, B-FormallyComplain

We can see the representation of the Client’s emotional goal in image 3. The simplicity of the Emotional Goal attributed to the Client reflects in his simple behaviour, selecting only two different actions, AskCakeOrCandy and FormallyComplain. In the beginning the most relevant guideline of the Client is the sigmoid Joy. Though this guideline is not shown in figure 3, the AskCakeOrCandy action generates a level of Joy that is just below the guideline value. It also produces Pride, though the potential of the emotion that results of this action is below the threshold

level of the Client, which means the Client himself does not feel that emotion as a result of `AskCakeOrCandy`.

The second action, `FormallyComplain`, is chosen because of the rapid growth of the `Pride` guideline. There are several other actions the Client could do that are praiseworthy, for example `Reason` and `DemandRespect`. However the Client appraises reasoning as undesirable, while `DemandRespect` has an emotional pre-condition that requires the subject of the action to feel `Reproach` towards the target.

G. Animation System

Our prototype includes an animation system that showcases the actions of our agents. The animation system should:

- Provide a comic, non-serious context;
- Allow the characters to communicate through textual utterings.
- Be able to showcase the characters emotions;

To comply with the first requirement we chose an environment with a cartoonish 2D look-and-feel, inspired by comic animated series. Actions should be able to pop up speech bubbles with text, as used in graphic novels, which solved the second problem. The third property was a bigger challenge. We animated the facial expressions of the character separately – the eyebrows, mouth, eyes and face color. The purpose was to combine this into more expressions than each part is able to convey on its own. In animating these expressions we considered Paul Ekman’s observations on how human faces express emotions [39].

VI. EVALUATION

To test if the prototype complied with our model and analyze its potential we devised an online questionnaire presenting a video of the sketch. Two versions were evaluated, but here we only discuss the one that used the Seller’s emotional goal without `Gloating` (`DispleaseClient-A`). This questionnaire registered 75 responses out of which 37 were males and 38 were females.

The participants clearly identify `Happiness` as the initial feeling of the Seller (60% of the answers) while a significant number do not identify the feeling as any present in the list (25.3%). The middle section of the sketch is slightly less clear, with responses scattered across answers `Anger` (6.7%), `Disappointment`(12%), `Happiness` (6.7%) and `Sadness`(9.3%). However the prevalent feeling is the Seller gets worried as the sketch goes on (36%). 28% did not find in the choice list a word that could express the Seller’s feelings. As for the ending part of the sketch opinions are divided between answers `Disappointment` (42.7%) and `Sadness` (49.3%).

The perceived emotions are thus consistent with both the actions and expressions of the Seller character. Initially the Seller feels glad for seeing the Client, thus `Happiness` seems the most appropriate answer. As the sketch proceeds, the Seller’s smile fades to a neutral smile. Participants had

some doubts on how to interpret this, but decided the Seller was worried. In the ending part the Seller fails to sell the antidepressant pills and, as a result, he sports an extremely sad smile. Participants recognized his sadness, and inferred, from the actions and subsequent reactions, the Seller got disappointed for not selling the antidepressants.

The Client’s emotional escalation was even more straightforward than the Seller’s. Being the regular character, most of the emotional escalation of the sketch was perceivable through him. The initial perception of the Client’s feelings is similar to the Seller’s, with `Happiness` gathering 57% of the answers. The evolution of the Client’s feelings is then perceived as a growth of `Anger`. 61.3% thought the Client was angry throughout the middle section of the sketch, growing to 70.7% in the ending part.

The overwhelming majority of the participants (76%) think the Seller character did not behave according to expectations. The cumulative percentage of those who answered between 1 and 3 (in which 1 means they totally disagree and 3 they neither agree nor disagree that the Seller behaved according to expectations) is 93.3%. We can thus say the Seller was recognized as the *incongruent* character. In contrast, participants agree the Client behaved as expected.

To the question asking whether they agreed (a Likert scale from 1 to 5) the sketch was funny, the mean value of the participants’ answers was 3,17, with a 1,319 standard deviation. However these results provide more interesting information if we relate the answers to this question with answers to the other questions, trying to understand *why* participants thought the sketch was funny or unfunny.

A Spearman correlation test indicates a correlation factor (ρ) of -0.366 significant at the 0.01 level, between those who agreed the sketch was too long and those who agreed the sketch was funny. This supports the hypothesis that the feeling the sketch is too long grows in inverse proportion with how funny it is. This may make the case that pacing is indeed an important subject in Interactive Comedy.

A Spearman- ρ correlation test for the question in which participants were asked whether they agree the ending was good, and the question about funniness shows a very strong correlation (ρ of 0.597, significant at the 0,01 level). This stresses the importance of the punchline of the sketch. If we see it in relation to participants who thought the sketch was too long, we can also consider that more jokes are needed to be triggered in the intermediate part of the sketch and provide a better buildup.

The relationship between those who think the ending needs explaining and the funniness does not seem so evident. The Spearman- ρ test shows a correlation, though not as strong, with a ρ of -0.356 significant at the 0.01 level. Once again this emphasizes the importance of the punchline, as well as the need for a good relation between the buildup and the punchline. Our model accounts for this relation through the preconditions that are needed to trigger a certain punchline. However, the buildup could probably be bettered by adding a bit more context to the actions each character selects.

VII. CONCLUSIONS

From our overview of the works in Computational Humour we observed that it has mainly revolved around Natural Language systems that generate puns and other basic types of jokes. Humour in Interactive Storytelling (Interactive Comedy) has known limited evolution since its recent inception. We intended to provide a more generic model that could be adapted for a larger variety of situations.

We focused on sketch comedy writing. We based our model on the incongruity-resolution theory of humour as well as on the basics of comedy writing. Our proposed model divides a sketch structurally into three parts, in which a conflict is introduced, developed and finished. The first part is scripted. The middle part develops this conflict towards an emotional peak in which a resolution or punchline is activated. The resolution is also scripted. However, our model makes possible the existence of several different resolutions which are activated depending on the context that results from the interaction between characters. Our main contribution is in the intermediate part of the sketch. We propose agents behave not only as characters but also as actors that play characters. As such they guide their actions in relation to an Emotional Guideline, that maps the scene time into the emotional output. The pacing of the sketch can be controlled by the shape of these guidelines, and how fast or slow they contribute to the Emotional Escalation. An Emotional Escalation is the evolution of emotions towards an emotional peak in which the sketch is resolved. Also, the preliminary results of our study indicate that this is a promising start, since the viewers identified this process and the evolution of emotions in the agents.

We have implemented this model as a prototype built upon the FAtiMA agent architecture, and tied it to an animation system that is capable of expressing the agents emotions and thus portraying the emotional escalation. The assessment of the comedic value of the resulting sketch is encouraging albeit non-conclusive. The relation between the perceived length of the sketch and its funniness suggests pacing should be a topic of interest in Interactive Comedy.

Our model contributes for the area of Interactive Comedy and Interactive Storytelling in general. Being an initial effort, it also presents a margin of improvement that is left as future work. It would be useful to have a more comprehensive evaluation and validation of the model that assesses how different Emotional Escalations affect the comic value of the sketch. Also, our model relies heavily on authoring, both for the characters personalities and on the Emotional Guidelines. Ideally, some of this knowledge should be part of the agent itself, to reduce the authoring and improve the ability of the agent to change its behaviour according to his appraisal of the world and of his interaction with other agents. Our evaluation also suggests that the actions selected during the sketch should be more coherent and provided better context. Finally, since humour is so connected to our social interactions, integrating the possibility of interacting with comic agents could probably also improve its comic

value.

REFERENCES

- [1] O. Dictionaries, "<http://oxforddictionaries.com/>."
- [2] R. A. Martin, *The Psychology of Humor: An Integrative Approach*. Elsevier Academic Press, 2007.
- [3] M. Mateas and A. Stern, "Procedural authorship: A case-study of the interactive drama façade," in *In Digital Arts and Culture (DAC)*, 2005.
- [4] R. S. Aylett, S. Louchart, J. Dias, A. Paiva, and M. Vala, "Fearnot!: an experiment in emergent narrative," in *Lecture Notes in Computer Science*. London, UK, UK: Springer-Verlag, 2005, pp. 305–316.
- [5] M. O. Riedl and A. Stern, "Believable agents and intelligent story adaptation for interactive storytelling," in *In Proceedings of the 3rd International Conference on Technologies for Interactive Digital Storytelling and Entertainment (TIDSE06)*. Springer, 2006, pp. 1–12.
- [6] A. Nijholt, "Embodied agents: A new impetus to humor research," in *Proceedings Twente Workshop on Language Technology 20*, 2002, pp. 101–111.
- [7] O. Stock, "Password swordfish: Verbal humor in the interface," *Humor International Journal of Humor Research*, vol. 16, no. 3, pp. 281–295, 2003.
- [8] J. Morkes, H. K. Kernal, and C. Nass, "Humor in task-oriented computer-mediated communication and human-computer interaction," in *CHI 98 conference summary on Human factors in computing systems*, 1998, pp. 215–216.
- [9] K. Binsted, "Using humour to make natural language interfaces more friendly," in *Proceedings of the AI, A-Life, and Entertainment Workshop, Intern. Joint Conf. on Artificial Intelligence*, 1995, p. 4.
- [10] A. Dufresne and M. Hudon, "Modeling the learner preferences for embodied agents: Experimenting with the control of humor," in *Proceedings of the Workshop on Individual and Group Modelling Methods that Help Learners Understand Themselves (International Conference on Intelligent Tutoring Systems)*, 2002, pp. 43–51.
- [11] G. Lessard and M. Levison, "Computational modelling of linguistic humour: Tom swifties," in *Proceedings of the 1992 ACH/ALLC Joint Annual Conference*, 1992, pp. 175–158.
- [12] —, "Computational modelling of riddling strategies," in *Proceedings of the 1993 ACH-ALLC Joint Annual Conference*, 1993, pp. 120–122.
- [13] K. Binsted and G. Ritchie, "An implemented model of punning riddles," in *AAAI Proceedings of the 12th National Conference on Artificial intelligence*, vol. 1, 1994, pp. 633–638.
- [14] K. Binsted, "Machine humour: An implemented model of puns," Ph.D. dissertation, The University of Edinburgh: College of Science and Engineering: The School of Informatics, 1996.
- [15] R. Manurung, G. Ritchie, H. Pain, A. Waller, D. O'Mara, and R. Black, "The construction of a pun generator for language skills development," *Applied Artificial Intelligence*, vol. 22, no. 9, pp. 841–869, 2008.

- [16] A. Waller, R. Black, D. A. O'Mara, H. Pain, G. Ritchie, and R. Manurung, "Evaluating the standup pun generating software with children with cerebral palsy," *ACM Transactions on Accessible Computing*, vol. 1, no. 3, pp. 16:1–16:27, 2009.
- [17] O. Stock and C. Strapparava, "Hahacronym: a computational humor system," in *Proceedings of the ACL 2005 on Interactive poster and demonstration sessions*, 2005, pp. 113–116.
- [18] G. Ritchie, "Can computers create humor?" *AI Magazine*, vol. 30, pp. 71–81, 2009.
- [19] F. Charles, S. J. Mead, and M. Cavazza, "Character-driven story generation in interactive storytelling," in *Proceedings 7th International Conference In Virtual Systems And Multimedia*, 2001, pp. 609–615.
- [20] M. Cavazza, F. Charles, and S. J. Mead, "Planning characters' behaviour in interactive storytelling," *Journal of Visualization and Computer Animation*, vol. 13, pp. 121–131, 2002.
- [21] —, "Generation of humorous situations in cartoons through plan-based formalisations," in *Proceedings of CHI-2003 Workshop: Humor Modeling in the Interface*, 2003.
- [22] Thawonmas, Ruck, K. Tanaka, Hassaku, and Hiroki, "Extended hierarchical task network planning for interactive comedy," in *Intelligent Agents and Multi-Agent Systems*. Springer Berlin / Heidelberg, 2003, vol. 2891, pp. 205–213.
- [23] D. Olsen and M. Mateas, "Beep! beep! boom!: towards a planning model of coyote and road runner cartoons," in *Proceedings of the 4th International Conference on Foundations of Digital Games*, 2009, pp. 145–152.
- [24] J. Bates, "The role of emotion in believable agents," *Commun. ACM*, vol. 37, no. 7, pp. 122–125, 1994.
- [25] C. Elliott, J. Brzezinski, S. Sheth, and R. Salvatoriello, "Story-morphing in the affective reasoning paradigm: generating stories semi-automatically for use with "emotionally intelligent" multimedia agents," in *Proceedings of the second international conference on Autonomous agents*, 1998, pp. 181–188.
- [26] C. Elliott, "Why boys like motorcycles: using emotion theory to find structure in humorous stories," in *EBAA '99 Workshop on Emotion-Based Agent Architectures, at the Autonomous Agents '99 Conference*, May 1999.
- [27] S. Louchart and R. Aylett, "From synthetic characters to virtual actors," in *Proceedings of the Third Artificial Intelligence and Interactive Digital Entertainment Conference*, 2007.
- [28] —, "Evaluating synthetic actors," 2007.
- [29] J. Dias and A. Paiva, "Agents with emotional intelligence for storytelling," in *Affective Computing and Intelligent Interaction*, ser. Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 2011, vol. 6974, pp. 77–86.
- [30] D. Pizzi, M. Cavazza, B. S. Magerko, and M. O. Riedl, "Affective storytelling based on characters' feelings," in *Intelligent Narrative Technologies, Papers from the 2007 AAAI Fall Symposium*, 2007, pp. 111–118.
- [31] D. Pizzi, M. Cavazza, and J.-L. Lugin, "Extending character-based storytelling with awareness and feelings," in *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, 2007, pp. 12:1–12:3.
- [32] M. Cavazza, J.-L. Lugin, D. Pizzi, and F. Charles, "Madame bovary on the holodeck: immersive interactive storytelling," in *Proceedings of the 15th international conference on Multimedia*, 2007, pp. 651–660.
- [33] G. Perret, *Comedy Writing Step By Step*. Writer's Digest Books, 1982.
- [34] J. Vorhaus, *The Comic Toolbox: How To Be Funny Even If You're Not*. Silman-James Press, 1994.
- [35] R. R. Provine, *Laughter: A Scientific Investigation*. Penguin Books, 2001.
- [36] S. Fry and H. Laurie, "A bit of fry and laurie - the complete collection," 2 Entertain Video DVD, 2006.
- [37] J. Dias, S. Mascarenhas, and A. Paiva, "Fatima modular: Towards an agent architecture with a generic appraisal framework," in *Workshop in Standards in Emotion Modeling, Leiden, Netherlands*, 2011.
- [38] A. Ortony, G. L. Clore, and A. Collins, *The Cognitive Structure of Emotions*. Cambridge University Press, 1988.
- [39] P. Ekman, *Emotions Revealed: Recognizing Faces and Feelings to Improve Communication and Emotional Life*. Holt Paperbacks, 2004.