Laugh To Me! Implementing Emotional Escalation on Autonomous Agents for Creating a Comic Sketch

André Carvalho, António Brisson, and Ana Paiva

INESC-ID and Instituto Superior Técnico, Universidade Técnica de Lisboa
Av. Prof. Doutor Aníbal Cavaco Silva 2744-016 Porto Salvo, Portugal
andrecarva@gmail.com
{antonio.brisson,ana.paiva}@gaips.inesc-id.pt
http://www.gaips.inesc-id.pt

Abstract. The growing interest for Interactive Storytelling has lead the research into the exploration of this new media in classical story genres. In our research, we develop autonomous agents that act in a storytelling context with a comic purpose. 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 by reasoning about emotional states in a process called Emotional Escalation. The punchline used for our test scenario, as well as the personality of the comic characters is based on the humour theory of incongruity-resolution.

Keywords: Autonomous Agents, Humour, Emotional Escalation

1 Introduction

Comedy is one of the most difficults, but yet one of the most entertaining forms of storytelling. As such, there are several examples of the pursuit to adapt the comedy genre in Interactive Storytelling [4, 3, 20, 15]. The common ground between most of the approaches is that they rely on planning formalisms that allow a character to fail. The comic situations arise from the character failing to achieve their goals. The comic effect of a failed plan is explained by the Schultz's incongruity-resolution theory [14]. Following Shultz the punchline(ending) of a joke creates an incongruity that contrasts with what was suggested by the set-up of a 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.

In our research we define the generation of incongruence as the build-up of an inconsistence or gap, which can be exploited for comic purposes. Our goal is to create a comic reality, which is a a caricature of real reality, a distorted

vision in which some features are emphasized, downplayed, or contrasted with their opposite. We do this not just by relying on action failure, but also considering other forms of incongruence such as emotional incongruence and context incongruence.

In this paper, we present the fundaments that support the implementation and evaluation of a comic sketch model for autonomous agents based on incongruence-resolution theory. In our test scenario we rely on our agents ability of creating emotional incongruence to achieve a comic effect. The agents build up the incongruent emotional context through a process called Emotional Escalation, by which they tend to set the scene passing and punchline timing.

2 Related Work

Computational Humour (the subfield of Artificial Intelligence concerned with the production of humour) has been mostly connected with works in Natural Language, since Lessard and Levison's 1992 Tom Swifties (a specific type of pun) generator[9], which used the VINCI Natural Language generator, and a subsequent work related with riddles [10].

A more ambitious approach was JAPE, developed by Binsted and Ritchie [2],[1], later used to develop STANDUP [13], a system aimed at helping children with cerebral palsy develop better language skills [22]. Another important work in computational humour was Stock and Strapparava's HAHAcronym, a simple prototype that sought to produce ironic acronyms [19]. HAHAcronym extends the lexicon with hierarchic domain information about the terms, in order to explore the incongruity between groups of concepts, for example, Sex vs. Religion.

In Interactive Storytelling the most common approach to comedy has 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 [4] was based on the sitcom Friends. The prototype resulted 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 [3]. Another implementation was that of Thawonmas et al. [20], who further noted 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 [15]. The conflicting goals here are the Coyote's aim of catching the Road Runner and Road Runner's goal of escaping. 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.

But, comedy is associated with the presence and actions of characters in a scene, and in particular, their emotions. There have been many works in Interactive Storytelling concerned about the agent's emotions. One of this works, Clark Eliott's Affective Reasoner [8] 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 [7], described via the emotions the characters felt.

Other works show how agents can choose actions deliberate to alter the emotional content of a story. One work suggested that agents using FAtiMA – the agent architecture used in this work – could support a double appraisal mechanism, in which the agent reappraises a selected action according to the emotional impact in others [12]. 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 [11]. Another extension of FAtiMA that has been proposed and is currently being implemented aims at making emotional intelligent agents [6].

A work by Pizzi et al. based on Gustave Flaubert's novel Madame Bovary[18], showed how a narrative could be described 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. The character can thus be aware of how its situation evolves: 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' "[18].

3 Background on Humor

Out of the several 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 ([14], pg. 64) the ending of a joke creates an incongruity that contrasts with what was suggested by its set-up. 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 we resolve it, by finding an explanation (is resolved) that makes the ending follow from the premise. According to this theory, failing to see an incongruence would lead to no surprise and being unable to resolve it would make the spectator puzzled, resulting in no laughter in either case.

Perret [17] and Vorhaus [21], comedians who wrote about comedy writing, both agree on the importance of incongruity. 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, whose perspective on comedy writing is strongly character-centric, introduces the concept of *strong comic perspective* in relation to comic characters ([21], 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). 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" ([17], pg. 154). In his account of what a sketch should be, Vorhaus stresses the need to create and develop a conflict between characters ([21], pp. 154-161).

4 Implementation

Fig. 1. Screenshot of the prototype, when the Client shows anger.

4.1 FAtiMA and OCC model

To implement our agents we used the FAtiMA framework [5] based on the OCC (Ortony, Clore and Collins [16]) theory as the underlying model of appraisal.

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 set of appraisal variables as well, such as *Praiseworthiness*, *Desirability*, and *Desirability for other*. The OCC model proposes a hierarchy of the different emotions according to these variables and according to the subject of the action, whether is the same one who is doing the appraisal or not.

When appraised OCC emotions have a given potential, which is represented in FAtiMA as a numeric value from 1 to 10. The intensity with which each emotion is felt equals the potential minus the threshold the agent has for that specific emotion. OCC emotion has a potential, a threshold, an intensity and a decay rate. The potential is the sheer value of 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.

In the FAtiMA architecture the personality of an agent is defined by rules derived from the OCC model. These set the values of their thresholds and decay rates for each emotion (using value between 1 and 10), and how each event is appraised in terms of appraisal variables (using value between -10 and 10). There are also a number of Action Tendencies: reactive actions activated when an emotion reaches a certain intensity.

4.2 Agents, Sketches and Incongruence

Conceptually, we divide incongruence in three types, depending on the way they relate to the agent: Context, Action and Emotional. Context refers to the environment in which the agents act, and whether or not it conflicts with their behaviour. An Action incongruence happens when the actions of the character are inconsistent, for example, with the agents goals. We consider most of the past works in Interactive Comedy described belong to this category. An Emotional incongruence arises from the personality of the agent itself, how differently an agent appraises the world considering what would normally be expected of him.

The authoring of a FAtiMA personality helps define Emotional incongruences. A character that acts or reacts in an incongruent manner is consistent with our idea of a comic character, while a character that acts in an acceptable way is a regular character.

As discussed, sketches present a specific structure. As such the actions available to the character depend on the moment of the sketch: the beginning, where the conflict is established, the middle where the action is developed, and the ending, which may be a punchline. Our implementations uses a scripted (in the sense of a predefined sequence of actions) beginning and punchline. The timing in which to activate the punchline, however, is defined by a set of preconditions.

4.3 Emotional Goals and Guidelines

To further explore the incongruence that results from the authoring of the character, we define Emotional Goals and Guidelines. This allows the agent to be active in how he will elicit certain emotions that not only explore this incongruence, but also buildup towards the punchline. As such the agents are not mere characters but active actors as well, that influence how the scene develops.

Emotional Guidelines are a function of the emotion over story time. In our implementation we define this story time as the actions of the agent itself, while the value of the emotion is given by the potential, as defined in FAtiMA's implementation of OCC. Thus each Emotional Guideline defines a desired potential for a given emotion at a point in time. The agent actively tries to select an action that evokes the emotions set by the Emotional Guidelines. The agent is capable of simulating an action and comparing its emotional output to that value.

The absolute difference between the values of the simulated and desired emotion potential is taken in account in the heuristic used by the agent to select an action. The agent tries to minimize this difference. In mixing each guideline it also prioritizes the Guidelines that present the higher desired values, given the point in time the heuristic function is being run in.

Emotional Goals group Emotional Guidelines with a set of preconditions. When these preconditions check, the Guidelines are considered activated. The agent will then actively consider these Guidelines in the action selection, according to the heuristic described above.

5 Scenario

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 bounds 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. We defined the Seller would want to make the Client angry, and as such it seemed fitting to make him trying to sell anti-depressants when the Client reached that state. After the Client refuses to buy the pills, the Seller, failing to see the inconsistence of his own actions, blames the crisis for the fact he did not sell the other product.

6 Authoring

Authoring in FAtiMA is done by defining a set of actions and goals available to the agent and its personality, meaning the emotional reaction rules and thresholds, as well as their reactive behaviour. Our scenario comprises two agents, and

the comicality of their behaviour depends on this authoring as well: how their personality relates to what we would expect of a regular character, and what their Emotional Goal for the scene is.

The actions have a precondition, `sketchMoment`, that filters the actions available in the beginning, middle and end of the sketch. The goals define several aspects: their generic goal of interacting with each other, their Emotional Goals that promote the Emotional Escalation, and the goal of activating a punchline. The preconditions of these punchline are what define when the moment is right to trigger it - the timing. Following from our scenario, this moment is defined as when the Client achieves a high value for anger, through its reactive behaviour.

6.1 The "Seller"

Since the Seller is a comic character, its personality, following from our model, should be incongruent with that of a regular character. Our Seller appraises actions such as `Insult` as desirable, even though they conflict with the goal of a normal salesman, of pleasing and making business with the client.

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.

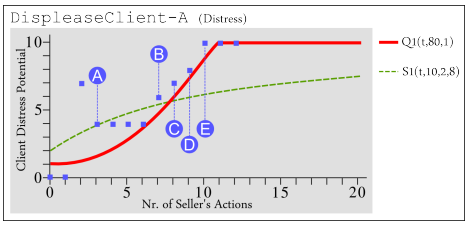

Fig. 2. Emotional Goal DispleaseClient-A. Letters represent actions: A-RaiseMoralIssues, B-Reason, C-WarnHinderAppetite, D-MakeSarcasticRemark, E-MakeFatPeopleJoke, F-FormallyComplain

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,

according to the OCC theory. Note also 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. The actions (represented by order in the x-axis) are of the Seller.

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.

6.2 The "Client"

The Client is a regular character. As such his reactions are more in line with what should be expected of someone in that situation, appraising events such as `Insult` as highly undesirable. Since its behaviour is not incongruent, the Client could perhaps be authored without emotional goals, 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.

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`.

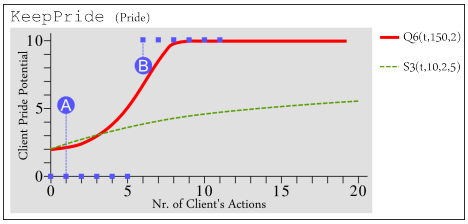

Fig. 3. Emotional Goal KeepPride. Pride guidelines only, Joy guideline not shown. Letters represent actions: A-AskCakeOrCandy, B-FormallyComplain

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 precondition that requires the subject of the action to feel Reproach towards the target.

7 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 multiple response questions discussed here are summarized in table 1.

The participants clearly identified Happiness as the initial feeling of the Seller (60% of the answers to Q1) while a significative number did not identify the feeling as any present in the list (25, 3%). The middle section of the sketch (Q2) presents less clear results, and Worry gathers only 36% of the responses and 28% are unable to find in the choice list a word that could express the Seller's feelings. As for the ending part of the sketch (Q3) opinions are divided between answers associated with Sadness (49, 3%) and Disappointment (42, 7%).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

Question no.	Question
Q1	How did the Seller feel in the beginning of the sketch?
Q2	How did the Seller feel in the middle of the sketch?
Q3	How did the Seller feel in the end of the sketch?
Q4	[Do you agree] The Seller behaved as expected, given the situation.
Q7,Q8,Q9,Q10	Same as 1-4, but in respect to the Client
Q13	[Do you agree] The sketch was too long.
Q15	[Do you agree] The sketch had a good ending.
Q16	[Do you agree] The ending should be better explained.
Q18	[Do you agree] The sketch was funny

Table 1. Questions stated on the online questionnaire.

is similar to the Seller's, with Happiness being the mode answer to Q7 (57%). 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 – Q8 – of the sketch, and $70,7\%$ – Q9 – in the ending part).

The majority of the participants (76%) totally disagree the Seller character behaved according to expectations (Q4). We can thus say the Seller was recognized as the *incongruent* character. In contrast, participants agree the Client behaved as expected.

The answers on whether the viewers thought the sketch was funny (Q18) was not conclusive with 3 as the median value selected. Some correlations with other questions may provide a better insight on *why* the participants deemed the sketch funny or unfunny.

A Spearman correlation test indicates an inverse relation between perceived length (Q13) and funniness, with a correlation factor (rho) of $-0,366$ significant at the $0,01$ level. This helps make the case that pacing is indeed an important subject in Interactive Comedy. Spearman-rho correlation tests also indicate funniness of the sketch relates directly with the quality of the ending (Q15, rho of $0,597$, significant at the $0,01$ level) as well as inversely with the need of a better explanation for the ending (Q16, rho of $-0,356$ significant at the $0,01$ level). This relation stresses the importance of the punchline of the sketch, and the way it derives from the buildup. Our model accounts for the link between buildup and punchline 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. Taking in account how the perception of the sketch's length contributes to humour, we can also consider that more jokes are needed to be triggered in the intermediate part of the sketch to enrich the buildup.

8 Conclusions

Our proposed model divides a sketch structurally into three parts, in which a conflict is introduced, developed and finished. The development of the action in

the sketch is based on the concept of Emotional Escalation. 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.

This work contributes to how Interactive Storytelling may mingle with the comedy genre, and how that can be tied to autonomous affective agents. Our model relies heavily on authoring, both for the characters personalities and on the Emotional Guidelines. With further understanding of how the evolution of emotions of characters takes place in comedy, the agent itself could use this knowledge, reducing the authoring and improving 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 suggests as well that the actions selected during the sketch should be more coherent and provided better context. Also, since humour is so connected to our social interactions, integrating the possibility of interacting with comic agents could probably also improve its comic value.

Acknowledgements Partially supported by FCT (INESC-ID funding) PIDDAC Prog. and FCT grant (SFRHBD/37476/2007).

References

1. K. Binsted. *Machine humour: An implemented model of puns.* PhD thesis, The University of Edinburgh: College of Science and Engineering: The School of Informatics, 1996.
2. K. Binsted and G. Ritchie. An implemented model of punning riddles. In *AAAI Proceedings of the 12th National Conference on Artificial intelligence*, volume 1, pages 633–638, 1994.
3. M. Cavazza, F. Charles, and S. J. Mead. Generation of humorous situations in cartoons through plan-based formalisations. In *Proceedings of CHI-2003 Workshop: Humor Modeling in the Interface*, 2003.
4. 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*, pages 609–615, 2001.
5. 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.

6. J. Dias and A. Paiva. Agents with emotional intelligence for storytelling. In *Affective Computing and Intelligent Interaction*, volume 6974 of *Lecture Notes in Computer Science*, pages 77–86. Springer Berlin / Heidelberg, 2011.
7. 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.
8. 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*, pages 181–188, 1998.
9. G. Lessard and M. Levison. Computational modelling of linguistic humour: Tom swifties. In *Proceedings of the 1992 ACH/ALLC Joint Annual Conference*, pages 175–158, 1992.
10. G. Lessard and M. Levison. Computational modelling of riddling strategies. In *Proceedings of the 1993 ACH-ALLC Joint Annual Conference*, pages 120–122, 1993.
11. S. Louchart and R. Aylett. Evaluating synthetic actors, 2007.
12. S. Louchart and R. Aylett. From synthetic characters to virtual actors. In *Proceedings of the Third Artificial Intelligence and Interactive Digital Entertainment Conference*, 2007.
13. 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*, 22(9):841–869, 2008.
14. R. A. Martin. *The Psychology of Humor: An Integrative Approach.* Elevier Academic Press, 2007.
15. 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*, pages 145–152, 2009.
16. A. Ortony, G. L. Clore, and A. Collins. *The Cognitive Structure of Emotions.* Cambridge University Press, 1988.
17. G. Perret. *Comedy Writing Step By Step.* Writer's Digest Books, 1982.
18. D. Pizzi, M. Cavazza, and J.-L. Lugrin. Extending character-based storytelling with awareness and feelings. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, pages 12:1–12:3, 2007.
19. O. Stock and C. Strapparava. Hahacronym: a computational humor system. In *Proceedings of the ACL 2005 on Interactive poster and demonstration sessions*, pages 113–116, 2005.
20. Thawonmas, Ruck, K. Tanaka, Hassaku, and Hiroki. Extended hierarchical task network planning for interactive comedy. In *Intelligent Agents and Multi-Agent Systems*, volume 2891, pages 205–213. Springer Berlin / Heidelberg, 2003.
21. J. Vorhaus. *The Comic Toolbox: How To Be Funny Even If You're Not.* Silman-James Press, 1994.
22. 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*, 1(3):16:1–16:27, 2009.