

OMG UR Funny!

Computer-Aided Humor with an Application to Chat

Miaomiao Wen
Carnegie Mellon
University

Nancy Baym, Omer Tamuz, Jaime Teevan, Susan Dumais, Adam Kalai
Microsoft Research

Abstract

In this paper we explore Computer-Aided Humor (CAH), where a computer and a human collaborate to be humorous. CAH systems support people's natural desire to be funny by helping them express their own idiosyncratic sense of humor. Artificial intelligence research has tried for years to create systems that are funny, but found the problem to be extremely hard. We show that by combining the strengths of a computer and a human, CAH can foster humor better than either alone. We present CAHOOTS, an online chat system that suggests humorous images to its users to include in the conversation. We compare CAHOOTS to a regular chat system, and to a system that automatically inserts funny images using an artificial humor-bot. Users report that CAHOOTS made their conversations more enjoyable and funny, and helped them to express their personal senses of humor. Computer-Aided Humor offers an example of how systems can algorithmically augment human intelligence to create rich, creative experiences.

Introduction

Can a computer be funny? This question has intrigued the pioneers of computer science, including Turing (1950) and Minsky (1984). Thus far the answer seems to be, "No." While some computer *errors* are notoriously funny, the problem of creating Computer-Generated Humor (CGH) systems that *intentionally* make people laugh continues to challenge the limits of artificial intelligence.

State-of-the-art CGH systems are generally textual. CHG systems have tried to do everything from generating word-play puns (Valitutti 2009) (e.g., "What do you get when you cross a fragrance with an actor? A smell Gibson") and identifying contexts in which it would be funny to say, "That's what she said," (Kiddon and Yuriy 2011) to generating I-like-my-*this*-like-my-*that* jokes (Petrovic and David 2013) (e.g., "I like my coffee like I like my war, cold") and combining pairs of headlines into tweets such as, "NFL: Green Bay Packers vs. Bitcoin – live!"¹ However, none of these systems has demonstrated significant success.

Despite the challenge that computers face to automatically generate humor, humor is pervasive when people use computers. People use computers to share jokes, create funny videos, and generate amusing memes. Humor and



Figure 1. Images suggested by CAHOOTS in response to chat line, "why u late?" (a), (b), and (e) are from image search query "funny late", (f) is from query "funny why", (c) is a canned reaction to questions, and (d) is a meme generated on-the-fly.

laughter have many benefits. Online, it fosters interpersonal rapport and attraction (Morkes et al. 1999), and supports solidarity, individualization and popularity (Baym 1995). Spontaneous humor production is strongly related to creativity, as both involve making non-obvious connections between seemingly unrelated things (Kudrowitz 2010).

Computers and humans have different strengths, and therefore their opportunity to contribute to humor differs. Computers, for example, are good at searching large data sets for potentially relevant items, making statistical associations, and combining and modifying text and images. Humans, on the other hand, excel at the complex social and linguistic (or visual) processing on which humor relies. Rather than pursuing humor solely through a CGH strategy, we propose providing computational support for humorous interactions between people using what we call *Computer-Aided Humor* (CAH). We show that by allowing the computer and human to work together, CAH systems can help people be funny and express their own sense of humor.

We explore the properties of this form of interaction and prove its feasibility and value through CAHOOTS (Computer-Aided Hoots), an online chat system that helps people be funny (Figure 1). CAHOOTS supports ordinary text chat, but also offers users suggestions of possibly funny

¹ <http://www.twitter.com/TwoHeadlines>

images to include based on the previous text and images in the conversation. Users can select choices they find on-topic or humorous and can add funny comments about their choices, or choose not to include any of the suggestions. The system was designed iteratively using paid crowd workers from Amazon Mechanical Turk and interviews with people who regularly use images in messaging.

We compare CAHOOTS to CGH using a chat-bot that automatically inserts funny images, and to ordinary chat with no computer humor. The bot uses the same images that CAHOOTS would have offered as suggestions, but forcibly inserts suggestions into the conversation. Compared to these baselines, CAHOOTS chats were rated more fun, and participants felt more involved, closer to one another, and better able to express their sense of humor. CAHOOTS chats were also rated as more fun than ordinary chat. Our findings provide insights into how computers can facilitate humor.

Related Work

In human-human interaction, humor serves several social functions. It helps in regulating conversations, building trust between partners and facilitating self-disclosure (Wanzer et al. 1996). Non-offensive humor fosters rapport and attraction between people in computer-mediated communication (Morkes et al. 1999). It has been found that five percent of chats during work are intended to be primarily humorous (Handel and James 2002), and wall posts in Facebook are often used for sharing humorous content (Schwanda et al. 2012). Despite the popularity and benefits of humorous interaction, there is little research on how to support humor during computer-mediated communication. Instead, most related work focuses on computationally generating humor.

Computational Humor

Computational humor deals with automatic generation and recognition of humor. Prior work has mostly focused on recognizing or generating one specific kind of humor, e.g. one-liners (Strapparava et al. 2011). While humorous images are among the most prominent types of Internet-based humor (Shifman 2007), little work addresses computational visual humor.

Prior work on CGH systems focus on amusing individuals (Dybala 2008; Valitutti et al. 2009). They find humor can make user interfaces friendlier (Binsted 1995; Nijholt et al. 2003). Morkes et al. (1998) study how humor enhances task-oriented dialogues in computer-mediated communication. HumoristBot (Augello et al. 2008) can both generate humorous sentences and recognize humoristic expressions. Sjobergh and Araki (2009) designed a humorous Japanese chat-bot. However, to the best of our knowledge, no prior research has studied collaboratively being funny using humans and computers.

Creativity Support Tools

CAH is a type of creativity support tool aimed specifically at humor generation within online interaction. Shneiderman (2007) distinguishes creativity support tools from productivity support tools through three criteria: clarity of task domain and requirements, clarity of success measures, and nature of the user base.

Creativity support tools take many forms. Nakakoji (2006) organizes the range of creativity support tools with three metaphors: running shoes, dumbbells, and skis. Running shoes improve the abilities of users to execute a creative task they are already capable of. Dumbbells support users learning about a domain to become capable without the tool itself. Skis provide users with new experiences of creative tasks that were previously impossible. For users who already utilize image-based humor in their chats, CAHOOTS functions as running shoes. For the remaining users, CAHOOTS serves as skis.

System Design

Our system, CAHOOTS, was developed over the course of many iterations. At the core of the system lie a number of different algorithmic *strategies* for suggesting images. Some of these are based on previous work, some are the product of ideas brainstormed in discussions with comedians and students who utilize images in messaging, and others emerged from observations of actual system use. Our system combines these suggestions using a simple reinforcement learning algorithm for ranking, based on R-Max (Brafman and Tennenholtz 2003), that learns weights on strategies and individual images from the images chosen in earlier conversations. This enabled us to combine a number of strategies.

User Interface

CAHOOTS is embedded in a web-based chat platform where two users can log in and chat with each other. Users can type a message as they would in a traditional online chat application, or choose one of our suggested humorous images. Suggested images are displayed below the text input box, and clicking on a suggestion inserts it into the conversation. Both text and chosen images are displayed in chat bubbles. See Figure 2 for an example. After one user types text or selects an image, the other user is provided with suggested image responses.

The Iterative Design Process

We initially focused on text-based humor suggestions based on canned jokes and prior work (Valitutti et al. 2009). These suffered from lack of context, as most human jokes are produced within humorous frames and rarely communicate meanings outside it (Dyner 2009). User feedback was negative, e.g., “The jokes might be funny for a three year old” and “The suggestions are very silly.”




Figure 2. The CAHOOTS user interface in a chat, with user's messages (right in white) and partner's (left in blue). All text is user-entered while images are suggested by the computer. The system usually offers six suggestions.

Based on the success of adding a meme image into suggestions, we shifted our focus to suggesting funny images. In hindsight, image suggestions offer advantages over text suggestions in CAHOOTS for multiple reasons: images are often more open to interpretation than text; images are slower for users to provide on their own than entering text by keyboard; and images provide much more context on their own, i.e., an image can encapsulate an entire joke in a small space.

Image Suggestion Strategies

In this section, we describe our most successful strategies for generating funny image suggestions based on context.

Emotional Reaction Images and gifs

Many chat clients provide emoticon libraries. Several theories of computer-mediated communication suggest that emoticons have capabilities in supporting nonverbal communications (Walther and Kyle 2001). Emoticons are often used to display or support humor (Tossell et al 2012). In popular image sharing sites such as Tumblr², users respond to other people's posts with emotional reaction images or gifs. In CAHOOTS, we suggest reaction images/gifs based on the emotion extracted from the last sentence.

Previous work on sentiment analysis estimates the emotion of an addresser from her/his utterance (Forbes-Riley and Litman 2004). Recent work tries to predict the emotion of the addressee (Hasegawa et al. 2013). Following this work, we first use a lexicon-based sentiment analysis to predict the emotion of the addresser. We adopt the widely used NRC Emotion Lexicon³. We collect reaction images and

² <http://www.tumblr.com>

³ <http://saifmohammad.com/WebPages/lexicons.html>



Figure 3. In response to text with positive sentiment, we suggest a positive emotional reaction image.




Figure 4. In response to the utterance, the user chooses a suggestion generated by Bing image search with the query "funny desert".

their corresponding emotion categories from reacticons.com. We collect reaction gifs and their corresponding emotion categories from reactinggifs.com. Then we suggest reaction images and gifs based on one of five detected sentiments: anger, disgust, joy, sadness, or surprise. An example of an emotional reaction is shown in Figure 3.

Image Retrieval

We utilize image retrieval from Bing image⁴ search (Bing image) and I Can Has Cheezburger⁵ (Cheezburger) to find funny images on topic. Since Bing search provides a keyword-based search API, we performed searches of the form "funny *keyword(s)*," where we chose *keyword(s)* based on the last three utterances as we found many of the most relevant keywords were not present in the last utterance alone. We considered both individual keywords and combinations of words. For individual words, we used the term frequency-inverse document frequency (tf-idf) weighting, a numerical statistic reflecting how important a word is to a document in a corpus, to select which keywords to use in the query. To define tf-idf, let $f(t, s_{-i})$ be 1 if term t occurred in the i^{th} previous utterance. Let U be the set of all prior utterances and write $t \in u$ if term t was used in utterance $u \in U$. Then weighted tf and tf-idf are defined as follows:

$$\text{wtf} = .7f(t, s_{-1}) + .2f(t, s_{-2}) + .1f(t, s_{-3})$$

⁴ <http://www.bing.com/images>

⁵ <http://icanhas.cheezburger.com>




Figure 5. An example of an utterance that generated a keyword combination *cat gerbil*, and a resulting image retrieved for the search *funny cat gerbil*.

$$\text{idf} = \log \frac{N}{|u \in U: t \in u|}$$

$$\text{tf-idf} = \text{wtf} * \text{idf}.$$

Here $N = 43,370$ is the total number of utterances collected during prototyping. The weights are designed to prioritize words in more recent utterances. An example of a single keyword for Bing is shown in Figure 4.

Combinations of keywords were also valuable. Humor theorists argue humor is fundamentally based on unexpected juxtaposition. The images retrieved with a keyword combination may be funnier or more related to the current conversation than images retrieved with a single keyword. However, many word pairs were found to produce poor image retrieval results. Consequently, we compiled a list of common keywords, such as cat and dog, which had sufficient online humorous content that they often produced funny results in combination with other words. If a user mentioned a common funny keyword, we randomly pick an adjective or a noun to form a keyword combination from the last three utterances. An example of a query for a combination of keywords is shown in Figure 5.

Memes

Meme images are popular forms of Internet humor. Coleman (2012) defines online memes as, “viral images, videos, and catchphrases under constant modification by users”. A “successful” meme is generally perceived as humorous or entertaining to audiences.

Inspired by internet users who generate their own memes pictures through meme generation website and then use them in conversations in social media sites like Reddit or Imgur, our meme generation strategy writes the last utterance on the top and bottom of a popular meme template. A meme template is an image of a meme character without the captions. The template is chosen using a machine-learning trained classifier to pick the most suitable meme template image based on the last utterance, as in Figure 1(d), with half of the text on the top and half on the bottom. To train our classifier to that match text messages to meme template, we collected training instances



Figure 6. A “Doge” meme example.

from the Meme Generator website⁶. This website has tremendous numbers of user-generated memes consisting of text on templates. In order to construct a dataset for training machine learning models, we collected the most popular one hundred meme templates and user generated meme instances from that site. To filter out the memes that the users find personally humorous, we only keep those memes with fifty or more “upvotes” ($N = 7,419$). We use LibLinear (Fan et al. 2008), a machine learning toolkit, to build a one-vs-the-rest SVM multi-class classifier (Keerthi et al. 2008) based upon Bag-of-words features. Even though this is multi-class classification with one hundred classes, the classifier trained in this simple way achieved 53% accuracy (compared with a majority-class baseline of 9%).

The fact that the meme's text often matched exactly what the user had just typed often surprised a user and led them to ask, “are you a bot?” Also note that we have other strategies for generating different types of image memes, which modify the text, such as the Doge meme illustrated in Figure 6.

Canned Responses

For certain common situations, we offer pre-selected types of funny images. For example, many users are suspicious that they are actually matched with a computer instead of a real person (which is partly accurate). As mentioned, we see users asking their partner if he/she is a bot. As a canned response, we suggest the results of keyword-search for “funny dog computer,” “funny animal computer,” or “funny CAPTCHA”.

Responding to Images with Images

We observed users often responding to images with similar images. For example, a picture of a dog would more likely be chosen as a response to a picture of a dog. Hence, the respond-in-kind strategy responds to an image chosen from a search for “funny *keyword*” with a second image from the same search, for any keyword.

Another strategy, called the rule-of-three, will be triggered after a user selects a respond-in-kind. The rule-of-three will perform an image search for “many *keyword*” or “not

⁶ <http://memegenerator.net>

keyword". An example is shown in Figure 7. The rule-of-three is motivated by the classic comic triple, a common joke structure in humor (Quijano 2012). Comedians use the first two points to establish a pattern, and exploit the way people's minds perceive expected patterns to throw the audience off track (and make them laugh) with the third element. In our system, when the last two images are both Bing image retrieved with the same keyword, e.g. funny dog images, we will suggest a Bing funny image with "many"+ keyword (e.g. "many dog") or "no" + keyword (e.g. "no dog") image as the third element.

In response to images, "LOL", "amused" or "not-amused" reaction images and gifs were suggested to help users express their appreciation of humor.

Ranking Suggestions using Reinforcement Learning

The problem of choosing images to select fits neatly into the paradigm of Reinforcement Learning (RL). Our RL algorithm, inspired by R-Max (Brafman and Tenenholz 2003), maintains counts at three levels of specificity, for number of times a suggestion was offered and number of times it was accepted. The most general level of counts is for each of our overall strategies. Second, for specific keywords, such as "dog," we count how many times, in general, users are offered and choose an image for a query such as "funny dog." Finally, for some strategies, we have a third level of specific counts, such as a pair for each of the fifty images we receive from Bing's API. We use the "optimistic" R-Max approach of initializing count pairs as if each had been suggested and chosen five out of five times. The score of a suggestion is made based on a back-off model, e.g., for a Bing query "funny desert": if we have already suggested a particular image multiple times, we will use the count data for that particular image, otherwise if we have sufficient data for that particular word we will use the

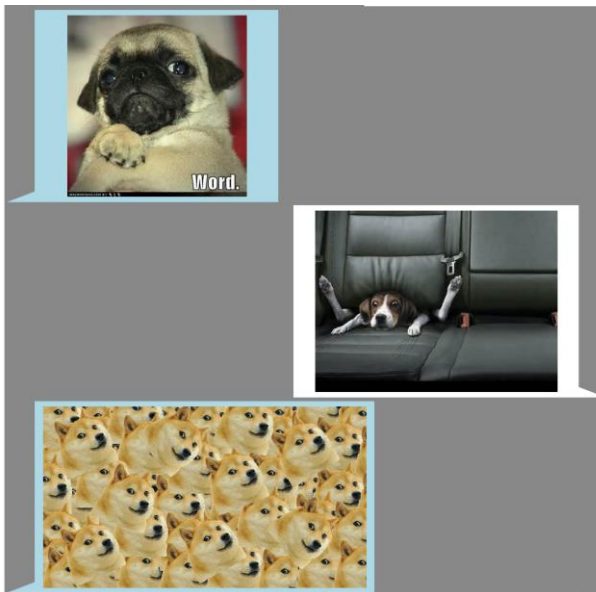


Figure 7. The "rule of three" strategy suggests putting an image of many dogs after two dog images.

count data for that word, and otherwise we will appeal to the count data we have for the general Bing query strategy.

Experiments

To test the feasibility of CAH we performed a controlled study. Before the experiment began, we froze the parameters in the system and stopped reinforcement learning and adaptation.

Methodology

Participants were paid Mechanical Turk workers in the United States. Each pair of Turkers chatted for 10 minutes using: 1) CAHOOTS, our CGH system, 2) plain chat (no image suggestions), or 3) a CGH system with computer-generated images, all using the same interface. In the CGH system, whenever one user sends out a message, our system automatically inserted the single top-ranking funny image suggestion into the chat, with "computer." inserted above the message, as is common in systems such as WhatsApp. Assignment to system was based on random assignment.

We also varied the number of suggestions in CAHOOTS. We write CAH_n to denote CAHOOTS with n suggestions. We use CAHOOTS and CAH_6 interchangeably (6 was the default number determined in pilot studies). The systems experimented with were CGH, plain chat, CAH_1 , CAH_2 , CAH_3 , CAH_6 , CAH_{10} .

A total of 738 participants (408 male) used one of systems, with at least 100 participants using each variant. Pairs of participants were instructed how to use the system and asked to converse for at least 10 minutes. After the chat, participants were asked to fill out a survey to evaluate the conversation and the system. We asked participants to what extent they agree with four statements (based on Jiang et al. 2011), on a 7 point Likert scale. The four statements were:

- *The conversation was fun.*
- *I was able to express my sense of humor in this conversation.*
- *I felt pretty close to my partner during the conversation.*
- *I was involved in the conversation.*

Experiments

Averaged over the chats where our system made suggestions ($CAH_{1,2,3,6,10}$) participants selected an image in 31% of the turns. In contrast, a field study found emoticons to be used in 4% of text messages (Tossell et al. 2012).

System Variant

Figure 8 summarizes participants' responses for the four Likert questions. Results are shown for chat, CGH, and two variants of CAHOOTS. P-values were computed using a one-sided Mann-Whitney U test.




Figure 8. Mean Likert ratings with Standard Error. 7 is strongly agree, 1 is strongly disagree, and the statements were 1. The conversation was fun. 2. I was able to express my sense of humor in this conversation. 3. I felt pretty close to my partner during the conversation. 4. I was involved in the conversation.

CAHOOTS vs. CGH

Participants rated CAHOOTS conversations better on average than CGH with p-values less than 0.05 for all four questions -- more fun, able to express sense of humor, closer to partner, and more involved in conversation

It is also interesting to compare CAH₁ to CGH as this reflects the difference between one image automatically into the conversation and one image offered as a suggestion. Here CAH₁ got higher response for fun, involvement, and closeness than CGH again with $p < .05$. Curiously, participants using CAH₁ felt somewhat less able to express their senses of humor.

CAHOOTS vs. plain chat

CAHOOTS was also rated more fun than plain chat ($p < .05$), and CAHOOTS participants also reported being able to express their own sense of humor better than plain chat participants ($p < .05$). For the other two questions CAHOOTS was not statistically significantly better than plain chat.

Note that while it may seem trivial to improve on plain chat by merely offering suggestions, our earlier prototypes (especially with text but even some with image suggestions) were not better than plain chat.

Number of Suggestions

Figure 9 shows responses to the fun question for different numbers of suggestions in CAHOOTS. In general, more suggestions makes the conversation more fun, though ten suggestions seemed to be too many. This may be because of the cognitive load required to examine ten suggestions or simply that with many suggestions scrolling is more likely to be required to see all image suggestions.

Effective Image Generation Strategies

As described earlier we used several different strategies for generating images. Table 1 shows how often each type was shown and how often it was selected. The rule-of-three




Figure 9. Mean and SE for "the conversation was fun" as we vary the number of suggestions, with 0 being plain chat.

(inspired by our meetings with comedians) was suggested less often than some other techniques, but the rate at which it was selected was higher. Reaction images/gifs were the next most frequently selected image strategy.

	# suggestions	% chosen
Bing Images	44,710	10%
Reaction Images and gifs	4,375	19%
Meme	709	13%
Rule-of-three	698	24%
Cheezburger	537	7%

Table 1 Selection rate of the top five strategies.

Limitations

Since we evaluate our system with paid workers, we have only tested the system between anonymous strangers whose only commonality is that they are US-based Mechanical Turk workers. We also asked workers to indicate with whom they would most like to use CAHOOTS: a family member, a close friend, an acquaintance, a colleague, or a stranger. Workers consistently answer that CAHOOTS would be best when chatting with a close friend who “can understand their humor.”

Also, we cannot compare CAHOOTS to every kind of CGH. For example, it is possible that users would prefer a CGH system that interjects images only once in a few turns or only when it is sufficiently confident.

Qualitative Insights

We analyzed the content of the text and image messages as well as worker feedback from both prototyping and experimentation phases. Note participants often remarked to one another, quite candidly, about what they liked or problems with our system, which helped us improve.

Anecdotally, feedback was quite positive, e.g., “It should be used for online speed dating!” and “When will this app be available for phones and whatnot? I want to use it!” Also, note that when we offered a small number of suggestions, feedback called for more suggestions. In contrast, feedback for CGH was quite negative, such as “The pictures got kind



Figure 10. Two workers start to talk about Bill Murray after using a reaction gif featuring Bill Murray.

of distracting while I was trying to talk to him/her.” We now qualitatively summarize the interactions and feedback.

Humorous Images Bring New Topics to the Conversation

Without CAHOOTS image suggestions, most of the chats focused on working in Mechanical Turk, which they seemed to find interesting to talk about. With suggestions, however, workers chose an image that suited their interests and naturally started a conversation around that image. Common topics included their own pets after seeing funny animal images, and their own children and family, after seeing funny baby images. As one worker commented: “great for chatting with a stranger, starts the conversation.” An example is shown in Figure 10, where two workers start to talk about Bill Murray after using a reaction gif featuring Bill Murray.

Image Humor is Robust

We found CAHOOTS robust in multiple ways. First, participants had different backgrounds which made them understand images differently. For example, one participant might complain that our memes were outdated, while the other participant’s feedback would indicate that they didn’t even recognize that the images were memes in the first place. Nonetheless, the latter could still find the images amusing even if they didn’t share the same background.

Second, we found CAHOOTS robust to problems that normal search engines face. For example, a normal search engine might suffer from ambiguity and therefore perform *word-sense disambiguation*, whereas humor is often heightened by ambiguity and double-entendres. While we didn’t explicitly program in *word-sense ambiguity*, it often occurs naturally.

Yes, and...

A common rule in improvisational comedy, called the *yes and* rule, is that shows tend to be funnier when actors accept one another’s suggestions and try to build them into something even funnier, rather than changing the direction even if they think they have a better idea (Moshavi 2001).

Many CAHOOTS’s strategies lead to yes-and behaviors. An example is shown in Figure 11. On the top, the computer suggestions directly addresses the human’s remark to makes the conversation funnier.




Figure 11. An example of man-machine riffing.

Users Tend to Respond with Similar Images

Humor support, or the reaction to humor, is an important aspect of interpersonal interaction (Hay 2001). With CAHOOTS, we find that users tended to respond to a funny image with a similar image to contribute more humor, show their understanding and appreciation of humor. When one user replied to her partner’s image message with an image, 35% of the time the other user chose an image generated by the same strategy. Compared with two random images in a conversation, the chance that they are generated by the same strategy is 22%.

Conclusion

In this paper we introduce the concept of Computer-Aided Humor, and describe CAHOOTS—a chat system that builds on the relative strengths of people and computers to generate humor by suggesting images. Compared to plain chat and a fully-automated CGH system, people using found it more fun, enabled them to express their sense of humor and more involvement.

The interaction between human and computer and their ability to riff off one another creates interesting synergies and fun conversations. What CAHOOTS demonstrates is that the current artificial intelligence limitations associated with computational humor may be sidestepped by an interface that naturally involves humans. A possible application of CAH would be an add-on to existing chat clients or Facebook/Twitter comment box that helps individuals incorporate funny images in computer-mediated communication.

References

Augello, A., Saccone, G., Gaglio, S., and Pilato, G. 2008. Humorist bot: Bringing computational humour in a chat-bot system. In *Complex, Intelligent and Software Intensive*

Systems, 703–708.

Baym, N. 1995. The performance of humor in computer-mediated communication. *Journal of Computer-Mediated Communication*, 1(2).

Binsted, K. 1995. Using humour to make natural language interfaces more friendly. In *Proceedings of the IJCAI Workshop AI and Entertainment*.

Brafman, R., and Tenenbholz, M. 2003. R-max: a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, 3:213–231.

Coleman, G. 2012. Phreaks, hackers, and trolls: The politics of transgression and spectacle. *The social media reader*, 99–119.

Dybala, P., Ptaszynski, M., Higuchi, S., Rzepka, R., and Araki, K. 2008. Humor prevails! - implementing a joke generator into a conversational system. In *Australasian Joint Conference on AI*, 214–225.

Dynel, M. 2009. Beyond a joke: Types of conversational humour. *Language and Linguistics Compass*, 3(5):1284–1299.

Fan, R.-E. Chang, K.-W., Hsieh, C.-J., Wang, X.-R., and Lin, C.-J. 2008. Liblinear: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874.

Forbes-Riley, K. and Litman, D. 2004. Predicting emotion in spoken dialogue from multiple knowledge sources. In *HLT-NAACL*, 201–208.

Hasegawa, T., Kaji, N., Yoshinaga, N., and Toyoda, M. 2013. Predicting and eliciting addressee's emotion in online dialogue. In *Proceedings of the ACL, Vol. 1*, 964–972.

Jiang, L., Bazarova, N., and Hancock, J. 2011. The disclosure-intimacy link in computer-mediated communication: An attributional extension of the hyperpersonal model. *Human Communication Research*, 37(1):58–77.

Keerthi, S., Sundararajan, S., Chang, K.-W., Hsieh, C.-J. and Lin, C.-J. 2008. A sequential dual method for large scale multi-class linear SVMs. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, 408–416.

Kiddon, C. and Brun, Y. 2011. That's what she said: double entendre identification. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: short papers-Volume 2*, 89–94.

Morkes, J. Kernal, H., and Nass, C. 1998. Humor in task-oriented computer-mediated communication and human-computer interaction. In *Human Factors in Computing Systems*, 215–216.

Morkes, J., Kernal, H., and Nass, C. 1999. Effects of humor in task-oriented human-computer interaction and computer-mediated communication: A direct test of SRCT theory.

Human-Computer Interaction, 14(4):395–435.

Moshavi, D. 2001. Yes and...: introducing improvisational theatre techniques to the management classroom. *Journal of Management Education*, 25(4):437–449.

Nijholt, A., Stock, O., Dix, A., and Morkes, J. 2003. Humor modeling in the interface. In *Human Factors in Computing Systems*, 1050–1051.

Petrovic, S. and Matthews, D. 2013. Unsupervised joke generation from big data. In *ACL (2)*, 228–232.

Quijano, J. 2012. Make your own comedy. In *Capstone*.

Ritchie, G. 2001. Current directions in computational humour. *Artificial Intelligence Review*, 16(2):119–135.

Sosik, V., Zhao, S., and Cosley, D. 2012. See friendship, sort of: How conversation and digital traces might support reflection on friendships. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*, 1145–1154.

Shifman, L. 2007. Humor in the age of digital reproduction: Continuity and change in internet-based comic texts. *International Journal of Communication*, 1(1):23.

Sjobergh, J., and Araki, K. 2009. A very modular humor enabled chat-bot for Japanese. In *Pacling 2009*, 135–140.

Strapparava, C., Stock, O., and Mihalcea, R. 2011. Computational humour. In *Emotion-oriented systems*, 609–634.

Tossell, C., Kortum, P., Shepard, C., Barg-Walkow, L., Rahmati, A., and Zhong, L. 2012. A longitudinal study of emoticon use in text messaging from smartphones. *Computers in Human Behavior*, 28(2):659–663.

Valitutti, A., Stock, O., and Strapparava, C. 2009. Graphlaugh: a tool for the interactive generation of humorous puns. In *Affective Computing and Intelligent Interaction and Workshops, 2009*. 1–2.

Valitutti, A., Toivonen, H., Doucet, A., and Toivanen, J. 2013. “Let everything turn well in your wife”: Generation of adult humor using lexical constraints. In *ACL (2)*, 243–248.

Walther, J. and D’Addario, K. 2001. The impacts of emoticons on message interpretation in computer-mediated communication. *Social science computer review*, 19(3):324–347.

Wanzer, M., Booth-Butterfield, M., and Booth-Butterfield, S. 1996. Are funny people popular? an examination of humor orientation, loneliness, and social attraction. *Communication Quarterly*, 44(1):42–52.