Getting Serious about the Development of Computational Humor

Oliviero Stock and Carlo Strapparava
ITC-irst, 1-38050 Povo Trento Italy
{stock, strappa}@itc.it

Abstract

Society needs humor, not just for entertainment. In the Web age, presentations become more and more flexible and personalized and they will require humor contributions for electronic commerce developments (e.g. product promotion, getting selective attention, help in memorizing names, etc...) more or less as it happened in the world of broadcasted advertisement. Even if deep modeling of humor in all of its facets is not something for the near future, there is something concrete that has been achieved and that can help in providing attention to the field. The paper refers to the results of HAHACRONYM, a project devoted to humorous acronym production, a circumscribed task that nonetheless requires various generic components. The project opens the way to developments for creative language, with applications in the world of advertisement.

1 Introduction

Future human-machine interaction, according to a widespread view, will emphasize naturalness and effectiveness and hence the incorporation of models of possibly all human cognitive capabilities, including the handling of humor. There are many practical settings where computational humor will add value. Among them: business world applications (e-commerce to name one), general computer-mediated communication and human-computer interaction [Morkes *et al,* 1999], educational and edutainment systems. There are important prospects for humor also in automatic information presentation. In the Web age presentations become more and more flexible and personalized and they will require humor contributions for electronic commerce developments (e.g. product promotion, getting selective attention, help in memorizing names, etc...) more or less as it happened in the world of broadcasted advertisement.

Yet deep modeling of humor in all of its facets is not something for the near future; the phenomena are too complex, humor is one of the most sophisticated forms of human intelligence. It is *AI-complete:* the problem of modeling it is as difficult to solve as the most difficult Artificial Intelligence problems. But some steps can be followed to achieve results. In the general case, in order to be successfully humorous, a computational system should be able to: recognize situations appropriate for humor; choose a suitable kind of humor for the situation; generate an appropriately humorous output; and, if there is some form of interaction or control, evaluate the feedback.

We are concerned with systems that automatically produce humorous output (rather than systems that appreciate humor). Some of the fundamental competencies are within the range of the state of the art of natural language processing. In one form or in another humor is most often based on some form of incongruity. For verbal humor this means that different interpretations of utterances must be possible (and must not be detected before the culmination of the humorous process) or must cause perception of specific forms of opposition. Natural language processing research has often dealt with ambiguity in language. A common view is that ambiguity is an obstacle for deep comprehension. Most current text processing systems attempt to reduce the number of possible interpretations of the sentences, and a failure to do so is seen as a weakness of the system. From a different point of view, the potential for ambiguity can be seen as a positive feature of natural language. Metaphors, idioms, poetic language and humor use all the multiple senses of texts to suggest connections between concepts that cannot, or should not, be stated explicitly.

The work presented here is based on various resources for natural language processing, adapted for humor. It is a small step, but aiming at an appreciable concrete result. It has been developed within HAHACRONYM[1], a project devoted to computational humor. A visible and evaluable result was at the basis of the deal. We proposed a situation that is of practical interest, where there is no domain restriction and many components are present, but simpler than in more extended scenarios. The goal is a system that makes fun of existing acronyms, or, starting from concepts provided by the user, produces a new acronym, constrained to be a word of the given language. And, of course, it has to be funny.

The project was meant to convince about the potential of computational humor, through the demonstration of a working prototype and an assessment of the state of the art and

[1]European Project IST-2000-30039. HAHACRONYM has been the first EU project devoted to computational humor. The consortium included ITC-irst, as coordinator, and the University of Twente. See http://haha.itc.it.

of scenarios where humor can add something to existing information technologies. The results of the project put us in a better position to move forwards in introducing computational humor in more complex scenarios.

2 AI and Computational Humor

So far only very limited effort has been put on building computational humor prototypes. Indeed very few working prototypes that process humorous text and/or simulate humor mechanisms exist, and mostly they worked in very restricted domains.

There has been a considerable amount of research on linguistics of humor and on theories of semantics and pragmatics of humor [Attardo, 1994; Attardo and Raskin, 1991; Giora and Fein, 1999; Attardo, 2002]; however, most of the work has not been formal enough to be used directly for computational humor modeling. An effort toward formalization of forced reinterpretation jokes has been presented by Ritchie [2002].

Within the artificial intelligence community, most writing on humor has been speculative [Minsky, 1980; Hofstadter *el al.,* 1989]. Minsky made some preliminary remarks about how humor could be viewed from the artificial intelligence/cognitive science perspective, refining Freud's notion that humor is a way of bypassing our mental "censors" which control inappropriate thoughts and feelings. Utsumi [1996] outlines a logical analysis of irony, but this work has not been implemented. Among other works: De Palma and Weiner [1992] have worked on knowledge representation of riddles, Ephratt [1990] has constructed a program that parses a limited range of ambiguous sentences and detects alternative humorous readings. A formalization, based on a cognitive approach (the belief-desire-intention model), distinguishing between real and fictional humor has been provided by Mele [2002].

Probably the most important attempt to create a computational humor prototype is the work of Binsted and Ritchie [1997]. They have devised a formal model of the semantic and syntactic regularities underlying some of the simplest types of punning riddles. A punning riddle is a question-answer riddle that uses phonological ambiguity. The three main strategies used to create phonological ambiguity are syllable substitution, word substitution and metathesis.

Almost all the computational approaches try to deal with the incongruity theory at various level of refinement [Raskin, 1985; Attardo, 1994]. The incongruity theory focuses on the element of surprise. It states that humor is created out of a conflict between what is expected and what actually occurs in the joke. This accounts for the most obvious features of a large part of humor phenomena: ambiguity or double meaning.

Specific workshops concerned with Computational Humor have taken place in recent years and have drawn together most of the community active in the field. The proceedings of the most comprehensive events are [Hulstijn and Nijholt, 1996] and iStock *et al,* 2002]. Ritchie [2001] has published a survey of the state of the art in the field.

3 HAHACRONYM

3.1 Resources

In order to realize the HAHACRONYM prototype, we have refined existing resources and we have developed general tools useful for humorous systems. As we will see, a fundamental tool is an incongruity detector/generator, that makes the system able to detect semantic mismatches between word meaning and sentence meaning (i.e. in our case the acronym and its context). For all tools, particular attention was put on reusability.

The starting point for us consisted in making use of some "off-the-shelf resources, such as WORDNET DOMAINS [Magnini *et al.,* 2002] (an extension of the well-known English WORDNET) and standard parsing techniques. The tools resulting from the adaptation will be reusable for other applications, and are portable straightforwardly to other languages (e.g. WORDNET DOMAINS is multilingual).

Wordnet and Wordnet Domains

WORDNET is a thesaurus for the English language inspired by psycholinguistics principles and developed at the Princeton University by George Miller [Fellbaum, 1998]. It has been conceived as a computational resource, therefore improving some of the drawbacks of traditional dictionaries, such as circularity of definitions and ambiguity of sense references. Lemmata (about 130,000 for version 1.6) are organized in synonym classes (about 100,000 *synsets).* WORDNET can be described as a "lexical matrix" with two dimensions: a dimension for *lexical relations,* that is relations holding among words, and therefore language specific, and a dimension for *conceptual relations,* which hold among senses (the synsets) and that, at least in part, we consider independent from a particular language. A synset contains all the words by means of which it is possible to express the synset meaning: for example the synset {horse, Equus-caballus} describes the sense of "horse" as an animal, while the synset {knight, horse} describes the sense of "horse" as a chessman. The main relations present in WORDNET are *synonymy, antonymy, hyperonymy-hypony my, meronymy-holonymy, entailment, troponymy.*

Augmenting WordNet with Domain information

Domains have been used both in linguistics (i.e. Semantic Fields) and in lexicography (i.e. Subject Field Codes) to mark technical usages of words. Although this is useful information for sense discrimination, in dictionaries it is typically used for a small portion of the lexicon. WORDNET DOMAINS is an attempt to extend the coverage of domain labels within an already existing lexical database, WORDNET (version 1.6). The synsets have been annotated with at least one domain label, selected from a set of about two hundred labels hierarchically organized. (Figure 1 shows a portion of the domain hierarchy.)

We have organized about 250 domain labels in a hierarchy (exploiting Dewey Decimal Classification), where each level is made up of codes of the same degree of specificity: for example, the second level includes domain labels such as BOTANY, LINGUISTICS, HISTORY, SPORT and RELIGION, while at the third level we can find specialization such as

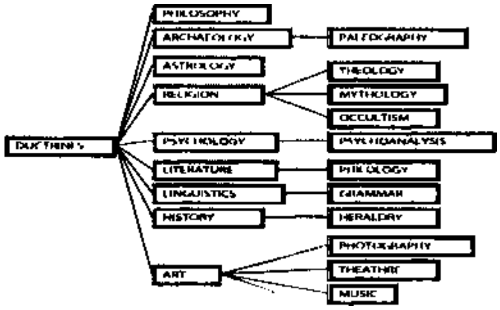

Figure 1: A portion of the domain hierarchy

AMERICAN_HISTORY, GRAMMAR, PHONETICS and TENNIS.

Information brought by domains is complementary to what is already present in WORDNET. First of all a domain may include synsets of different syntactic categories: for instance MEDICINE groups together senses from Nouns, such as doctor#l and hospital#l, and from Verbs such as operate#7. Second, a domain may include senses from different WORDNET sub-hierarchies For example, SPORT contains senses such as athlete#l, deriving from life_form#l, game equipment # l, from physical_object#l, sport#l from act#2, and playing_field#l, from location#l.

Opposition of semantic fields

On the basis of well recognized properties of humor accounted for in many theories (e.g. incongruity, semantic field opposition, apparent contradiction, absurdity) we have modelled an independent structure of domain opposition, such as RELIGION VS. TECHNOLOGY, SEX VS. RELIGION, etc... We exploit these kind of opposition as a basic resource for the incongruity generator.

Adjectives and Antonymy Relations

Adjectives play an important role in modifying and generating funny acronyms. WORDNET divides adjectives into two categories. *Descriptive adjectives* (e.g. big, beautiful, interesting, possible, married) constitute by far the largest category. The second category is called simply *relational adjectives* because they are related by derivation to nouns (i.e. electrical in electrical engineering is related to noun electricity). To relational adjectives, strictly dependent on noun meanings, it is often possible to apply similar strategies as those exploited for nouns. Their semantic organization, though, is entirely different from the one of the other major categories. In fact it is not clear what it would mean to say that one adjective "is a kind of' (ISA) some other adjective. The basic semantic relation among descriptive adjectives is antonymy. WORDNET proposes also that this kind of adjectives is organized in clusters of synsets associated by semantic similarity to a focal adjective. Figure 2 shows clusters of adjectives around the direct *antonyms fast I slow.*

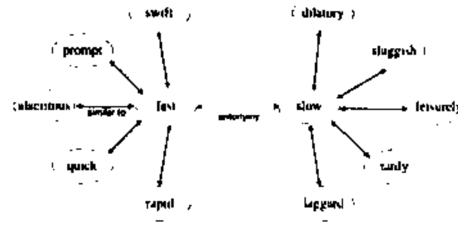

Figure 2: An example of adjective clusters linked by antonymy relation

Exploiting the hierarchy

It is possible to exploit the network of lexical and semantic relations built in WORDNET to make simple ontological reasoning. For example, if a noun or an adjective has a geographic location meaning, the pertaining country and continent can be inferred.

Rhymes

The HAHACRONYM prototype takes into account word rhymes and the rhythm of the acronym expansion. To cope with this aspect we got and reorganized the CMU pronouncing dictionary (http: //www. speech.es.cmu.edu/cgi-bin/cmudict) with a suitable indexing. The CMU Pronouncing Dictionary is a machine-readable pronunciation dictionary for North American English that contains over 125,000 words and their transcriptions.

Its format is particularly useful for speech recognition and synthesis, as it has mappings from words to their pronunciations in the given phoneme set. The current phoneme set contains 39 phonemes; vowels may carry lexical stress.

Parser, grammar and morphological analyzer

Word sequences that are at the basis of acronyms are subject to a well-defined grammar, simpler than a complete noun phrase grammar, but complex enough to require a nontrivial analyzer. We have decided to use a well established nondeterministic parsing technique. As far as the dictionary is concerned, we use the full WORDNET lexicon, integrated with an ad-hoc morphological analyzer.

Also for the generation part we exploit the grammar as the source for syntactic constraints.

All the components are implemented in Common Lisp augmented with nondeterministic constructs.

Other resources

An "a-semantic" dictionary is a collection of hyperbolic/epistemic/deontic adjective/adverbs. This is a last resource, that some time can be useful in the generation of new acronyms. Some examples are: *abnormally, abstrusely, adorably, exceptionally, exorbitantly, exponentially, extraordinarily, voraciously, weirdly, wonderfully.* This resource is hand-made, using various dictionaries as information sources.

Other lexical resources are: a euphemism dictionary, a proper noun dictionary, lists of typical foreign words commonly used in the language with some strong connotation.

3.2 Architecture and Implementation

To get an ironic or "profaning" re-analysis of a given acronym, the system follows various steps and strategies. The main elements of the algorithm can be schematized as follows:

- acronym parsing and construction of a logical form
- choice of what to keep unchanged (typically the head of the highest ranking NP) and what to modify (e.g. the adjectives)
- look up for possible substitutions
- exploitation of semantic field oppositions
- granting phonological analogy: while keeping the constraint on the initial letters of the words, the overall rhyme and rhythm should be preserved (the modified acronym should sound similar to the original as much as possible)
- exploitation of WORDNET antonymy clustering for adjectives
- use of the a-semantic dictionary as a last resource

Figures 3 and 4 show a sketch of the system architecture.

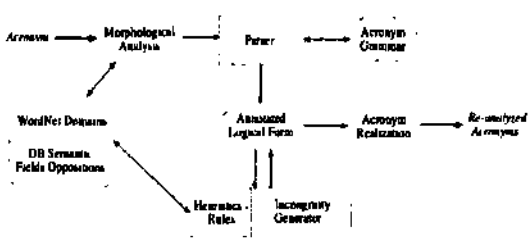

Figure 3: Acronyms Reanalysis

In our system, making fun of existing acronyms amounts to an ironical rewriting, desecrating them with some unexpectedly contrasting, but otherwise consistently sounding expansion.

As far as acronym generation is concerned, the problem is more complex. To make the task more attractive - and difficult - we constrain resulting acronyms to be words of the dictionary (APPLE would be good, IBM would not). The system takes in input concepts (actually synsets, possibly the result of some other process, for instance sentence interpretation) and some minimal structural indication, such as the semantic head. The primary strategy of the system is to consider as potential acronyms words that are in ironic relation with the input concepts. By definition acronyms have to satisfy constraints - to include the initial letters of some lexical realization of the inputs words synsets, granting that the sequence of initials satisfy the overall acronym syntax. In this primary strategy, the ironic reasoning comes mainly at the level of acronym choice in the lexicon and in the selection of the fillers of the *open slots* in the acronym.

For example, giving as input "fast" and "CPU", we get static, torpid, dormant. (Note that the complete synset for "CPU" is {processor#3, CPU#1, central..processing.unit#l, mainframe#2}. So

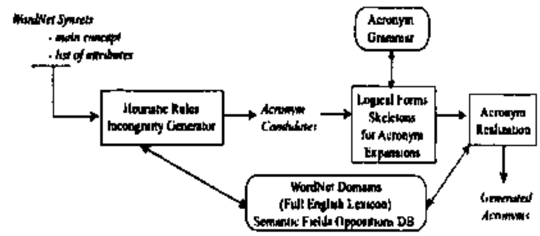

Figure 4: Acronyms Generation

we can use in the acronym expansion a synonym of "CPU". The same happens for "fast"). Once we have an acronym proposal, a syntactic skeleton has to be filled to get a correct noun phrase. For example given in input "fast" and "CPU", the system selects TORPID with the possible syntactic skeleton:

$< adv >_T < adj >_O$ **Rapid Processor** $< prep > < adj >_I < noun >_D$

or

$< adj >_T < adj >_O$ **Rapid Processor** $< prep > < noun >_I < noun >_D$

where "rapid" and "processor" are synonyms respectively of "fast" and "CPU" and the notation <Part_of_Speech>Le^tr means a word of that particular part_of.speech with *Letter* as initial. Then the system fills this syntactic skeleton with strategies similar to those described for re-analysis.

The system is fully implemented in Common Lisp, exploiting CLOS and the Meta-Object protocol to import WORDNET DOMAINS and to implement the navigation/reasoning algorithms.

3.3 Examples

Here below some examples of acronym re-analysis are reported. As far as semantic field opposition is concerned we have slightly tuned the system towards the domains FOOD, RELIGION and SEX. We report the original acronym, the re-analysis and some comments about the strategies followed by the system.

ACM - Association for Computing Machinery
-->Association for Confusing Machinery
FBI - Federal Bureau of Investigation
—► Fantastic Bureau of Intimidation

The system keeps all the main heads and works on the adjectives and the PP head, preserving the rhyme and/or using the a-semantic dictionary.

CRT - Cathodic Ray Tube
—► Catholic Ray Tube
ESA - European Space Agency
—> Epicurean Space Agency
PDA - Personal Digital Assistant
—►Penitential Demoniacal Assistant
—> Prenuptial Devotional Assistant
MIT - Massachusetts Institute of Technology
—► Mythical Institute of Theology

Some re-analyses are RELIGION oriented. Note the rhymes.

As far as generation from scratch is concerned, a main concept and some attributes (in terms of synsets) are given as input to the system. Here below we report some examples of acronym generation.

Main concept: *processor* (in the sense of CPU);
Attribute: *fast*

OPEN - On-line Processor for Effervescent Net

PIQUE - Processor for Immobile Quick Uncertain Experimentation

TORPID - Traitorously Outstandingly Rusty Processor for Inadvertent Data_proeessing

UTMOST - Unsettled Transcendental Mainframe for Off-line Secured Tcp/ip

We note that the system tries to keep all the expansions of the acronym coherent in the same semantic field of the main concept (COMPUTER _SCIENCH). At the same time, whenever possible, it exploits some incongruity in the lexical choices.

3.4 Evaluation

Testing the humorous quality of texts is not an easy task. There have been relevant studies though, such as (Ruch, 19961. For HAHAc RONYM, a simpler case, an evaluation was conducted under the supervision of Salvatore Attardo at Youngstown University, Ohio. Both reanalysis and generation have been tested according to criteria of success stated in advance and in agreement with the European Commission, at the beginning of the project.

The participants in the evaluation were forty students. They were all native speakers of English. The students were not told that the acronyms had been computer generated. No record was kept of which student had given which set of answers (the answers were strictly anonymous). No demographic data were collected. However, generally speaking, the group was homogeneous for age (traditional students, between the ages of 19 and 24) and mixed for gender and race.

The students were divided in two groups. The first group of twenty was presented the reanalysis and generation data. We tested about 80 reanalyzed and 80 generated acronyms (over twice as many as required by the agreement with the European Commission). Both the reanalysis module and the generation module were found to be successful according to the criteria spelled out in the assessment protocol. The acronyms reanalysis module showed roughly 70% of acronyms having a score of 55 or higher (out of a possible 100 points), while the acronym generation module showed roughly 53% of acronyms having a score of 55 or higher. The thresholds for success established in the protocol were 60% and 45%, respectively.

A set of randomly generated acronyms were presented to a different group of twenty students. A special run of the system was performed in which the semantic filters and heuristics had been disabled, while only the syntactical constraints were operational. (If the syntactical rules had been disabled as well the output would have been gibberish and it would be difficult to compare with the regular system production). In this test, less than 8% of the acronyms passed the 55 points score test, we conclude that the output of HAHACRONYM is significantly better than random production of reanalysis.

Acronyms	*scored > 55*	*Success Thresholds*
Generation	52.87%	45%
Reanalysis	69.81%	60%
Random	7.69%	

Table 1: Evaluation Results

A curiosity that may be worth mentioning: HAHACRONYM participated to a contest about (human) production of best acronyms, organized in December 2002 by RAI, the Italian National Broadcasting Service. The system won a jury's special prize.

4 Prospects for advertisement

Humor is the healthy way of creating 'distance* between one's self and the problem, a way of standing back and looking at the problem with perspective. Humor reveals new aspects, disarms and relaxes. It is also infectious and it is an important way to communicate ideas. On the cognitive side humor has two very important properties:

- it helps getting and keeping people's attention. Type and rhythm of humor may vary and the time involved in building the effect may be different in different cases: some times there is a context - like joke telling - that from the beginning let you expect for the humorous climax, which may occur after a long while: other times the effect is obtained in almost no time, with one perceptive act - for instance in static visual humor, funny posters or in cases when some well established convention is reversed with an utterance;
- it helps remembering. For instance it is a common experience to connect in our memory some knowledge we have acquired to a humorous remark or event. In a foreign language acquisition it may happen that an involuntary funny situation is created because of so called "false friends" - words that sound similar in two languages and may have the same origin but have a very different meaning. The humorous experience is conducive to remembering the correct use of the word.

No wonder that humor has become one of the favorite strategies used by creative people involved in the advertising business. In fact among the various fields in which creative language is used, advertising is probably the area of activity in which creativity (and hence humor) is practiced with most precise objectives.

From an applied AI point of view, we believe that an environment for proposing solutions to advertising professionals can be a realistic practical development of computational humor and one of the first attempts in dealing creative language.

Some examples of the huge array of opportunities that language offers and that existent NLP techniques can cope with: rhymes, wordplays, popular sayings or proverbs, quotations, alliterations, triplets, chiasmus, neologism, non sequitur, adaptation of existing expressions, adaptation of proverbs, personification, synaesthesia (two or more senses combined).

An intelligent system devoted to advertisement, based on an apparatus that extends the one we have described, can play with these expressions, alter them, change the context, use them as a platform for completely new ideas (e.g. create the slogan: *Thirst come, thirst served* for a soft drink).

In the longer run the aim is to go for individual-oriented advertisement. Let us leave aside issues of privacy for the moment, though we know this is a critical theme. The availability of individual profiles, automatic inferencing of interest and behavior models, information about the location of the individual, combined with reasoning on the offer side (what is advisable to promote in a given context of business) will all provide the basic input for advertising things worth getting the interest of a specific individual in a given context. Commercial advertisement will be just one case along this line - one can think of influencing the individual behaviour on the basis of the availability of any possible resources such as, for instance, cultural goods. A prediction is that when such situation-oriented advertisements will become widespread, the same features that prevail now with broadcasted material will be sought after for the new case. So, forms of humorous advertisement for the individual (even playing on personal aspects) will become a plus.

An important aspect to be taken into account is how humor is appreciated by different individuals. Personality studies regarding this specific theme give important indications [Ruch, 20021. One option will also be to develop humor for conversational systems, based on embodied agents. The work of Nijholt [2002] Andre and Rist [2000] and Cassell [2001] could provide the starting point for introducing dynamic humor.

For obtaining good results (there are not many things that are worse than bad humor!) for these longer term objectives, much work will be needed, especially further integration with ontologies, common sense reasoning and pragmatics.

References

[Andre and Rist, 2000] E. Andre and T. Rist. Presenting through performing: On the use of multiple lifelike characters in knowledge-based presentation systems. In *Proc. of the Second International Conference on Intelligent User Interfaces (IUI 2000),* pages 1-8,2000.

[Attardo and Raskin, 1991] S. Attardo and V. Raskin. Script theory revis(it)ed: joke similarity and joke representation model. *Humour,* 4(3):293-347,1991.

[Attardo, 1994] S. Attardo. *Linguistic Theory of Humor.* Mouton de Gruyter, Berlin, 1994.

[Attardo, 2002] S. Attardo. Formalizing humor theory. In *[Stock et, 2002],* 2002.

[Binsted and Ritchie, 1997] K. Binsted and G. Ritchie. Computational rules for punning riddles. *Humor,* 10(1), 1997.

[Cassell, 2001] J. Cassell. Embodied conversational agents: Representation and intelligence in user interface. *AI Magazine,* 22(3):67'-83,2001.

[Ephratt, 1990] M. Ephratt. What's in a joke. In M. Golumbic, editor, *Advances in AI: Natural Language and Knowledge Based Systems,* pages 406-450. Springer Verlag, 1990.

[Fellbaum, 1998] C. Fellbaum. *WordNet. An Electronic Lexical Database.* The MIT Press, 1998.

[Giora and Fein, 1999] R. Giora and O. Fein. Irony: Context and salience. *Metaphor and Symbol,* 14:241-257,1999.

[Hofstadter et al., 1989] D. Hofstadter, L. Gabora, V. Raskin, and S. Attardo. Synopsis of the workshop on humor and cognition. *Humor,* 2(4):293-347, 1989.

[Hulstijn and Nijholt, 1996] Joris Hulstijn and Anton Nijholt, editors. *Proc. of International Workshop on Computational Humour (TWLT 12),* University of Twente, Enschede, 1996.

[Magnini et al., 2002] B. Magnini, C. Strapparava, G. Pezzulo, and A. Gliozzo. The role of domain information in word sense disambiguation. *Journal of Natural Language Engineering,* 8(4):359-373,2002.

[Mele, 2002] F. Mele. Real and fictional ridicule. In [*Stock* etal., *2002]*, 2002.

[Minsky, 1980] M. Minsky. Jokes and the logic of the cognitive unconscious. Technical report, MIT Artificial Intelligence Laboratory, 1980. AI memo 603.

[Morkes *et al,* 1999] J. Morkes, H. Kernal, and C. Nass. Effects of humor in task-oriented human-computer interaction and computer-mediated communication. *Human-Computer Interaction,* 14:395-435,1999.

[Nijholt, 2002] A. Nijholt. Embodied agents: A new impetus to humor research. In *[Stock* et al., *2002],* 2002.

[Palma and Weiner, 1992] P. De Palma and E. J. Weiner. Riddles: accessibility and knowledge representation. In *Proc. of the 15th International Conference on Computational Linguistics (COLING-92),* Nantes, 1992.

[Raskin, 1985] V. Raskin. *Semantic Mechanisms of Humor.* Dordrecht/Boston/Lancaster, 1985.

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

IRitchie, 2002] G. Ritchie. The structure of forced interpretation jokes. In *[Stock* et al., *2002],* 2002.

[Ruch, 1996] W. Ruch. Special issue: Measurement approaches to the sense of humor. *Humor,* 9(3/4), 1996.

[Ruch, 2002] W. Ruch. Computers with a personality? lessons to be learned from studies of the psychology of humor. In *[Stock* et al., *2002],* 2002.

[Stock *et al.,* 2002] O. Stock, C. Strapparava, and A. Nijholt, editors. *Proceedings of the The April Fools Day Workshop on Computational Humour (TWLT20),* Trento, 2002.

[Utsumi, 1996] A. Utsumi. Implicit display theory of verbal irony: Towards a computational theory of irony. In *[Hulstijn and Nijholt, 1996],* 1996.