

Witty, Affective, Persuasive (and possibly Deceptive) Natural Language Processing

Carlo Strapparava

FBK-Irst - Istituto per la ricerca scientifica e tecnologica
strappa@fbk.eu

Motivations

- Exploration of some aspects (e.g. creative genres) of natural language traditionally considered outside the scope of any computational modeling
 - Emotions in texts
 - Computational Humor
 - Persuasive Language
 - Deceptive Language
 - ...
- Huge quantity of texts available
- Stress of what we can do using just texts (as we find in web)

Motivations

- Dealing with these topics is a relatively new area of research (e.g. Affective NLP)
- Important for many NLP applications
 - Opinion mining
 - Market analysis
 - Affective user interfaces
 - E-learning environments
 - ...
- Goal of the talk: overview some of our explorations for addressing these issues

Outline

1. Witty language
 - Humor generation
 - Humor recognition
2. Affective Text
 - Lexical resources
 - Annotation of emotions in text
 - Colors of emotions in texts
 - Dancing with words
3. Persuasive NLP
 - Analyzing political speeches along with audience reactions (e.g. applauses)
 - How to evaluate persuasive language ?
4. Deceptive Language recognition
 - Is it possible to recognize when people are lying, just using the produced text ?

Outline

1. Witty language
 - Humour generation
 - Humour recognition
2. Affective Text
 - Lexical resources
 - Annotation of emotions in text
 - Colors of emotions in texts
 - Dancing with words
3. Persuasive NLP
 - Analyzing political speeches along with audience reactions (e.g. applauses)
 - How to evaluate persuasive language ?
4. Deceptive Language recognition
 - Is it possible to recognize when people are lying, just using the produced text ?

Computational humour



Society needs humour

- Humor is a powerful generator of emotions
- It has an impact on people's psychological state, directs their attention, influences the processes of memorization and of decision-making (i.e. companies hire 'humour consultants')
- E.g. the persuasive effect of humor and emotions is well known and widely employed in advertising.
- Computational Humour can deliver something useful
- Deep modelling of humour in all its facets is not for the near future: humour is AI-complete
- Complete modelling of humour processes is not always required
- CH leads to falsifiable theories: can be tested on human subjects

⇒ *Humour is infectious: contagious laughter in Tanganyika, started in a group of schoolgirls and rapidly rose to epidemic proportions, infecting adjacent communities. It required the closing of the schools and it lasted for six months.*

Is computational humour realistic?

- Deep modelling of humour in all its facets is not for near future
- But not always complete modelling of humour processes is required
 - E.g. wordplays, lexicon-based semantic opposition, ambiguity, ...

A bit of Marxism in the sense of Marx Brothers :-)

Mrs Teasdale: This is a gala day for you.

*Firefly (Groucho): Well, a gal a day is enough for me.
I don't think I could handle any more.*

Computational humour for edutainment and IT

- To provide comic relief/reward
- To stimulate the attention
- To favor long-term memorization
- To enhance learning experience (positive feelings towards learning when humor is included)
- To stimulate creativity

Theories of humor

- Cognitive (incongruity, contrast)

Focus: *stimulus*

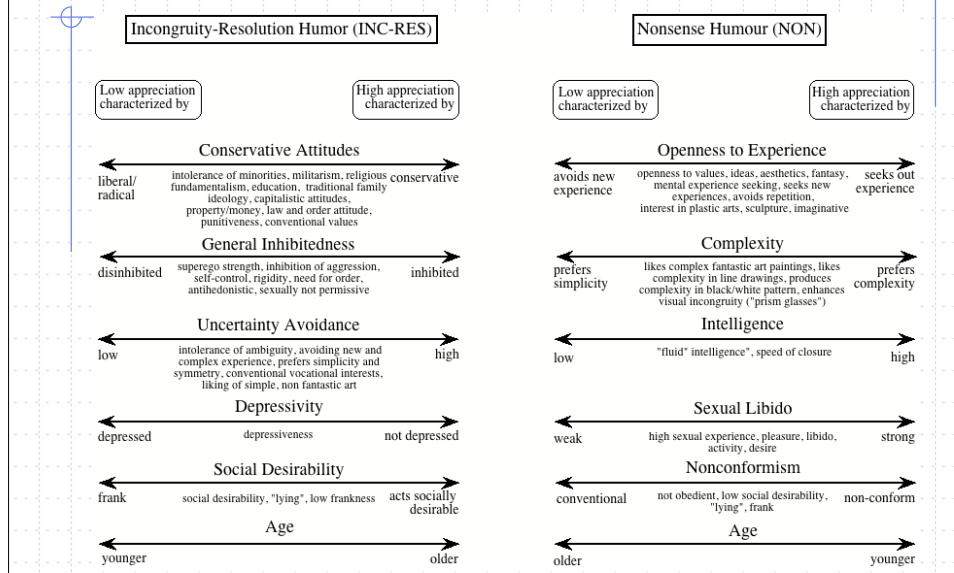
- Social (superiority, hostility, derision, disparagement)

Focus: *interpersonal effects*

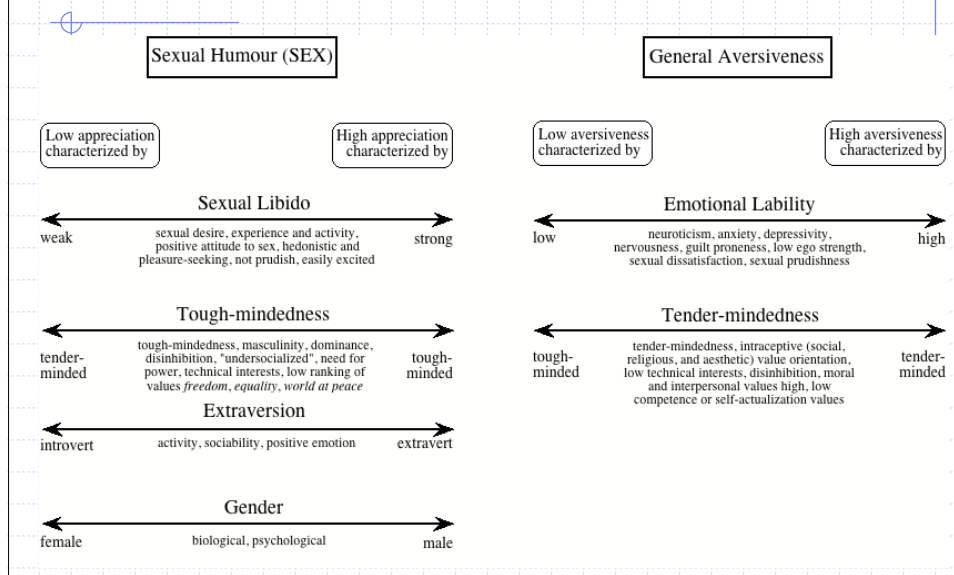
- Psychoanalytical (relief, release, liberation, sublimation)

Focus: *audience's reaction*

Individual differences Personality studies (see W. Ruch)



Individual differences (2)



Work on computational humour

- Research on linguistics and pragmatics of humor [e.g. Attardo and Raskin]
- Speculative writings in AI [e.g. Minsky, Hofstadter]
- Some efforts on building computational humor prototypes. For example:
 - Humour Production
 - ◆ JAPE [Binsted & Ritchie] It generates punning riddles, from a linguistic model of pun schemata, e.g. "*What do you call a murderer with fiber? A cereal killer*"
 - Humour Recognition
 - ◆ [Mihalcea & Strapparava 2005] investigated machine learning techniques to distinguish between humorous and non-humorous text

HAHAcronym

- HAHAcronym has been a Future Emerging Technologies (FET) European project (<http://haha.itc.it>)
- Goal: realization of an acronym re-analyzer and generator as proof of concept in a focalized but non restricted context
- various existing resources for NLP adapted for humor + some strategies for yielding humorous output

⇒ O. Stock & C. Strapparava "Getting serious about the development of computational humor" Proceedings of 18th IJCAI

HAHAcronym Resources

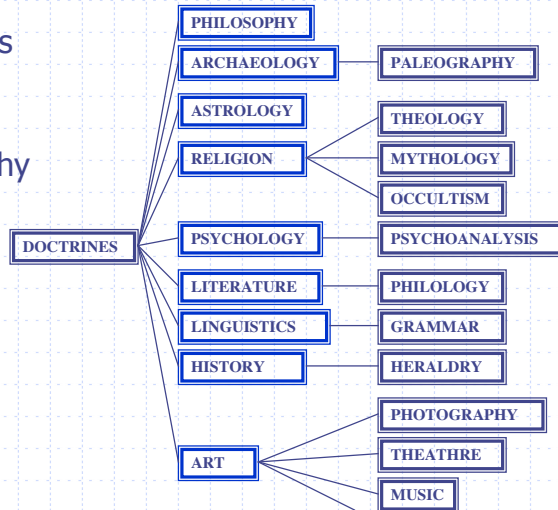
- Lexicon (full English lexicon)
- Lexical knowledge base (WordNet Domains)
- Pronunciation dictionary
- Parser and grammar
- Algorithms (for humour effects)
- Slanting dictionary
- ...

WordNet as a lexical knowledge base

- WordNet is an on-line lexical reference system whose design is inspired by *psycholinguistic* theories of human lexical memory
- Developed at Princeton University by George Miller's team. WordNet is a public domain resource
<http://www.cogsci.princeton.edu/~wn/>
- Synonym sets, representing underlying concepts (~100.000). Different relations link the synonym sets.
- IRST extensions
 - Multilinguality (synset-aligned)
 - Domain labels on synsets (e.g. **Medicine**, **Architecture**, **Sport**)

Domain label organization

- 250 Domain labels collected from dictionaries
- Four level hierarchy (Dewey Decimal Classification)




Domain labels annotation in WordNet

- Integrate taxonomic and domain oriented information
 - Cross hierarchy
 - ◆ doctor#2 [Medicine] --> person#1
 - ◆ hospital#1 [Medicine] --> location#1
 - Cross category relations: operate#3 [Medicine]
 - Cross language information
- Reduce polysemy

Use of domain label annotations

- Theories of humour suggest:
 - incongruity, semantic field opposition, apparent contradiction, absurdity
- We have defined:
 - an independent structure of domain opposition i.e. Religion vs. Technology, Sex vs. Religion, etc...
 - algorithms to detect semantic mismatches between word meaning and sentence meaning (i.e. acronym and its expansion)

Bipolar adjective structure



Rhymes

- The HAHAcronym prototype takes into account the rhyme structure of words
- CMU pronouncing dictionary, reorganized with a suitable indexing
- Over 125,000 words and their transcriptions
- Mappings from words to their pronunciations in the given phoneme set

Slanting Dictionary

- A collection of hyperbolic, epistemic, emotive adjectives, adverbs and nouns
 - ♦ Ex. abnormally, abstrusely, adorably, exceptionally, exorbitantly, exponentially, extraordinarily, voraciously, weirdly, wonderfully ...
- Useful when it is not possible to exploit other more meaningful strategies

Heuristics in HAHacronym

- Using WordNet
 - *Semantic field opposition*: e.g. Technology vs. Religion
 - *Antonymy* (for adjectives): e.g. "high" vs. "humble"
 - *Exploiting the hierarchy*:
 - ◆ e.g. detecting geographic names/adjectives
 - ◆ hyperonyms/hyponyms in the generation phase

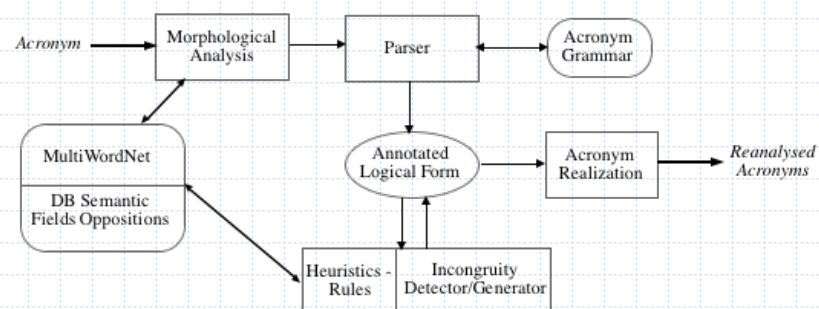
Heuristics in HAHacronym (2)

- Using general lexical resources
 - Strict rhyme and "light" rhyme
 - slanting dictionaries
- Syntactic strategies
 - e.g. keep the main head fixed


Acronym re-analysis

1. Acronym parsing and construction of logical form
2. Choice of what to keep unchanged
3. Look up for possible substitutions, e.g. exploiting semantic field oppositions
4. Granting phonological analogy and rhyme
5. Exploitation of WordNet antonymy clustering
6. Use of slanting dictionary as a last resource

Acronym re-analysis: the architecture



Examples: re-analysis



FBI - Federal Bureau of Investigation
=> Feral Bureau of Intimidation

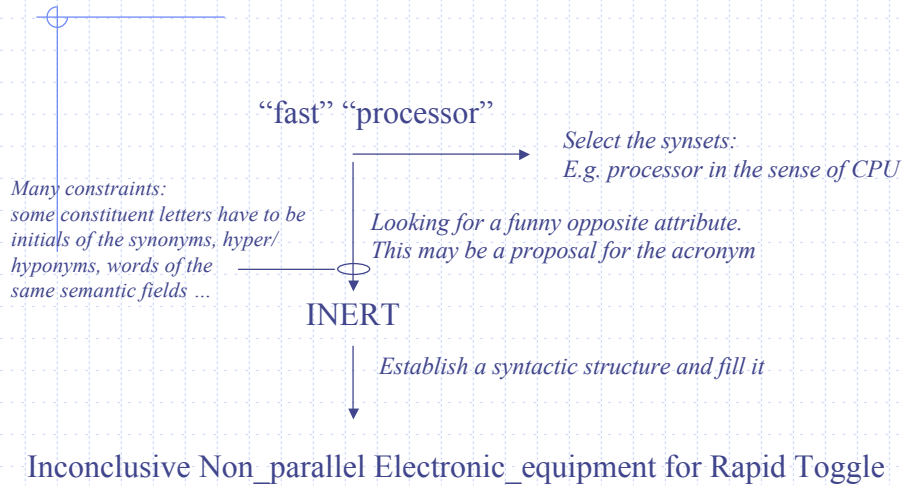
GPD - Gross Domestic Product
=> Godless Dietetic Product

PDA - Personal Digital Assistant
=> Penitential Demoniactal Assistant

Acronym generation

- Additional constraint: resulting acronyms to be words of the dictionary (APPLE is good, IBM not)
- Input: WN synsets and some minimal structural indication (e.g. the semantic head)
- Primary strategy: consider as potential acronyms words that are in ironic relation with input concepts
- Impose a syntactic structure and expand the acronym preserving coherence among semantic fields

An example: generation



Evaluation

- Success thresholds stated in the project proposal
- Evaluation carried out by Salvatore Attardo at Youngstown University
- A panel of 40 university students, all native speakers of English, homogeneous for age, and mixed for gender and race

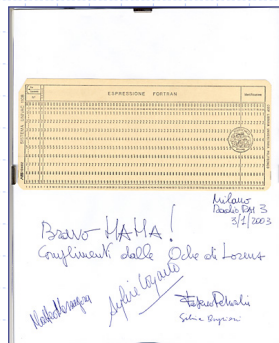
Evaluation results

- About 80 reanalyzed and 80 generated acronyms were tested.
- Also a test with randomly generated acronyms (only syntactic rules were operational)

Acronyms	Successful	Success Threshold
Generation	52.87%	45%
Re-analysis	69.81%	60%
Random re-analysis	7.69%	

HAHAcronym competes with humans

- HAHAcronym participated in a contest about (human) production of acronyms, organized by RAI, the Italian National Broadcasting Service
- The system won the jury's special prize !



Possible developments of practical impact

- Educational software for children: word-meanings exploration;
- A system that uses humor as means to promote products and to get user's attention in electronic commerce;
- An explorative environment for advertising professionals (e.g. "thirst come, thirst served" for a soft drink);
- A names generator for products and merchandise

More acronyms

- NATO - *North Atlantic Treaty Organization*
Noisy Anglophilic Torpidity Organization
- NSF - *National Science Foundation*
National Somnolence Foundation
National Science Flirtation
National Somnolence Fornication
- AAA - *American Automobile Association*
Antediluvian Automobile Association
- IBM - *International Business Machine*
Illusional Baroqueness Machine
- GSMC - *Global System for Mobile Communication*
Gastronomical System for Male Consolation

HAHAcronym goes to AI conferences

- AAAI - American Association for Artificial Intelligence
=> Antediluvian Association for Artificial Imprudence
- IJCAI - International Joint Conference on Artificial Intelligence
=> Irrational Joint Conference on Antenuptial Intemperance

Further work on computational humour

- Try to exploit further lexical semantics techniques
 - E.g. (Bucaria, 2004) "Lexical and syntactic ambiguity as a source of humor"
- ⇒ Lexical, syntactic, phonological ambiguity
- Advertisement, News Headlines

The case of news headlines

■ *Lexical ambiguity*

- Men recommended more **clubs** for wives

Club = (1) association of person;
(2) a heavy stick that is larger at one end

- Stadium air conditioning fails - **Fans** protest

Fan = (1) a device for creating
a current of air;
(2) a sport enthusiastic

- Doctor testifies in horse **suit**

Suit = (1) a set of garments;
(2) a comprehensive term for
any proceeding in a court of law

- **Queen Mary** having bottom scraped

■ *Syntactic ambiguity*

- Lawyers give poor free legal advice
- Babies are what the mother eats
- Man eating piranha mistakenly sold as pet fish

■ *Referential ambiguity*

- Autos killing 110 a day: let's resolve to do better

Initial steps for producing humorous expressions

■ Overall goal:

Realization of an environment for the production of
creative and humorous expressions: e.g. newspaper
titles, advertisements, ...

■ Current achievements:

- some basic and general techniques for *automatic*
creation of *emotional language*;
- indications for humorous expressions as variation of
existing texts

Outline

1. Witty language
 - Humour generation
 - Humour recognition
2. Affective Text
 - Lexical resources
 - Annotation of emotions in text
 - Colors of emotions in texts
 - Dancing with words
3. Persuasive NLP
 - Analyzing political speeches along with audience reactions (e.g. applauses)
 - How to evaluate persuasive language ?
4. Deceptive Language recognition
 - Is it possible to recognize when people are lying, just using the produced text ?

Questions:

1. Can we build very large data sets of humorous texts?
2. Are humorous and serious texts separable?
 - Can we automatically distinguish between humorous and non-humorous texts?
 - Does this hold for different data sets?
3. What are the distinctive features of humour?
 - Can we identify salient features of verbal humour?
 - Do they hold across data sets?
4. Can humour improve human-computer interaction?

R. Mihalcea and C. Strapparava. *Learning to laugh (automatically): Computational models for humor recognition*.
Journal of Computational Intelligence, 2006.

Data for humour recognition

- Required to *learn* and *test* models of humour
- Positive examples = humorous text
- Negative examples = non-humorous text
- Desiderata:
 - Large data sets
 - ◆ To test variation of performance with data
 - Humorous text should differ only in comic effect – force classifiers to identify humour-specific features
 - ◆ Chose non-humorous data similar in content and style with humorous data
 - Different data sets
 - ◆ To test consistency

Humorous data (1/2)

- Focus on a specific type of humour
- One-liners
 - “He who smiles in a crisis has found someone to blame”
 - ◆ Short sentence, simple syntax
 - ◆ Deliberate use of rhetoric devices (alliteration, rhyme)
 - ◆ Frequent use of creative language
 - ◆ Comic effect
- How to get 10,000+ one-liners
 - Websites or mailing lists typically include no more than 10 – 100 one-liners
- Web-based bootstrapping
 - ◆ Start with a few manually selected seeds
 - ◆ Identify a list of Web pages including at least one seed
 - ◆ Parse Web pages and find new one-liners
 - ◆ Repeat

Serious data

- EVERYWHERE! (almost)
- Data similar in structure and composition to the humorous text
 - Make the humour-recognition task more difficult (& real)
 - Allow the classifiers to identify humour-specific features
- For the one-liners:
 - Sentences of 10 – 15 words
 - Similar to one-liners with respect to creativity and intent
 - Mix of *Reuters* titles, *proverbs*, *British National Corpus*, sentences from *Open Mind Common Sense*

Serious data

- Reuters titles
 - Phrased to catch the readers attention
 - Reuters newswire 1996 – 1997
 - "Silver fixes at two-month high, but gold lags".
- Proverbs
 - From online proverb collection
 - Memorable sayings, considered true by many people
 - "Beauty is in the eye of the beholder".
- Text
 - British National Corpus
 - Most similar sentences, using vectorial similarity with tf.idf weighting
 - "The train arrives three minutes early".

Learning to recognize humour

- Hypothesis: “We can apply machine learning techniques to distinguish between humorous and non-humorous text”
 - Data
 - ◆ Positive / negative examples
 - Features
 - ◆ Content / Style
 - Learning algorithms
 - ◆ Naïve Bayes / SVM / ...

Features

- Style:
 - Rhetorical devices
 - Attention-catching sounds
- Content:
 - Specific vocabulary

Stylistic features

- Inspired from linguistic theories of humour
 - (Attardo 1994)
- Focus on features that can be implemented with current resources
 - Alliteration
 - Antonymy
 - Slang

Alliteration

- Phonetic properties: alliteration, word repetition, rhyme, producing a comic effect
 - Similar devices are used in wordplay, newspaper headlines, advertisement
- Examples
 - "Veni, Vidi, Visa: I came, I saw, I did a little shopping".
 - "Infants don't enjoy infancy like adults do adultery".
- Identify and count alliteration/rhyme chains using the CMU pronunciation dictionary

Antonymy

- Humor often relies on incongruity and contradiction
 - Antonymy is a form of incongruity that can be identified
- Examples:
 - "A **clean** desk is a sign of a **cluttered** desk drawer".
 - "Always try to be **modest** and be **proud** of it".
- Identify antonyms using WordNet:
 - Nouns, verbs, adjectives, adverbs

Adult slang

- A popular form of humour
- Can be identified through the detection of sexual-oriented vocabulary
- Examples:
 - "The **sex** was so good that even the neighbors had a cigarette"
 - "Artificial **insemination**: **procreation** without recreation"
- Use WordNet – Domains to build a lexicon with all synsets marked with the domain "sexuality"
 - Remove words with high polysemy (> 3)

Experiments

- Data set: 16,000 one-liners + 16,000 "serious" sentences
- Apply the stylistic features to humour-recognition
- Features act as heuristics
 - Require a threshold
- Learn a decision tree using 1000 (x 2) positive and negative examples
- Evaluate on remaining 15,000 (x 2) examples
- 10 trials

Results

Heuristic	Oneliners Reuters	Oneliners BNC	Oneliners Proverbs
Alliteration	74.31%	59.34%	53.30%
Antonymy	55.65%	51.40%	50.51%
Adult slang	52.74%	52.39%	50.74%
ALL	76.73%	60.63%	53.71%

- Notes:
 - A combination of features provides the best results
 - Alliteration is the most useful feature
 - Reuters titles are the most different with respect to one-liners
 - Proverbs are the most similar

Context-based features

- Formulate humour recognition as a text classification problem
 - Data
 - ◆ Positive (humorous) / negative (serious) examples
 - Features
 - ◆ N-grams
 - Learning algorithms
 - ◆ Naïve Bayes / SVM
 - ◆ 10-fold cross-validation

Humorous data (2/2)

- Daily news stories from: "The Onion"
 - "the best source of humour out there" (Jeff Grienfield, CNN)
 - ◆ Canadian Prime Minister Jean Chrétien and Indian President Abdul Kalam held a subdued press conference in the Canadian Capitol building Monday to announce that the two nations have peacefully and sheepishly resolved a dispute over their common border. "We are - well, I guess proud isn't the word - relieved, I suppose, to restore friendly relations with India after the regrettable dispute over the exact coordinates of our shared border," said Chrétien, who refused to meet reporters' eyes as he nervously crumpled his prepared statement. "The border that, er... Well, I guess it turns out that we don't share a border after all." Chrétien then officially withdrew his country's demand that India hand over a 20-mile-wide stretch of land that was to have served as a demilitarized buffer zone between the two nations."
 - 1,125 news articles from August 2005 – March 2006
 - ◆ 1,000-10,000 characters


Serious data

- For the news articles:
 - Documents with a length of 1,000-10,000 characters
 - Mix of *Los Angeles Times*, *Foreign Broadcast Information Service*, *British National Corpus*

Classification results

Classifier	One-liners	News articles
Naïve Bayes	79.69%	88.00%
SVM	79.23%	96.80%

- Significant improvement over the 50% baseline
- Better discrimination for news stories – longer size



Characteristics of humour

- What are the distinctive features of humour?
 - Identify the most salient features for humorous text
 - Classify these features into categories
- Feature list
 - Start with the score generated by the Naïve Bayes classifier
 - *Humorous score* = score in humorous text / total score
 - ◆ Score close to 1 => features specific to the humorous text
 - ◆ Score close to 0 => features specific to the non-humorous text
 - Extract the 1,500 most discriminatory features
 - ◆ Occurring at least 100 times in the entire corpus

Characteristics of verbal humour

- Observed by analyzing the features extracted from the one-liners
- Human-centric vocabulary
 - *you, I, man, woman, guy*
 - ◆ *you* occurs in more than 25% of the one-liners
 - ◆ "You can always find what you are not looking for."
- Negation
 - *doesn't, isn't, don't*
 - ◆ "If at first you don't succeed, skydiving is not for you."
- Negative orientation
 - words with negative orientation: *bad, illegal, wrong*
 - ◆ "When everything comes your way, you are in the wrong lane."



Characteristics of verbal humour

- Professional communities
 - *lawyers, programmers, policemen*
 - ◆ “It was so cold last winter, that I saw a lawyer with his hands in his own pockets.”
- Human “weakness”
 - *ignorance, stupidity, trouble, beer, drink, lie*
 - ◆ “Only adults have trouble with child-proof bottles.”



Two main features


- Human centeredness
 - Human-centric vocabulary
 - Professional communities
 - Human weakness
- Polarity orientation
 - Negation
 - Negative orientation
 - Human weakness




Human centeredness

- Measure the weight of the most salient features with respect to a semantic class
 - Score of semantic class = sum of the corresponding features normalized with the size of the class
 - E.g. I (0.88), me (0.65), myself (0.55) => 0.69
- Top 1,500 most discriminatory features
- Four semantic classes
 - **Persons:** WordNet hierarchy subsumed by *person#n#1*
 - **Social Groups:** hierarchy subsumed by *social_group#n#1*
 - **Social Relations:** hierarchies of *relative#n#1* and *relationship#n#1*
 - **Personal Pronouns**

Human centeredness: One-liners



Human centeredness: News articles



Humour for computer applications

- Find the most appropriate joke for a given context
 - Text semantic similarity
 - LSA, WordNet-based
- Determine the affective orientation of text
 - Avoid the use of humorous text for negative/sad situations
 - Automatic classification of affect

Fun Email

- Add humorous one-liners to email
 - Modification of Squirrel Mail email client
 - find the text's semantic orientation and ignore the email if adding humor would be inappropriate
 - ◆ Automatic classification of text as happy/sad
 - extract the last 30 percent of text from the email body
 - ◆ similarity is computed with respect to the topic of the last part of the email
 - compare the email's LSA vector with those of the one-liners, and identify the the most similar one-liner

From: Priscilla Rasmussen
Date: 28 November 2006
To: Carlo Strapparava
Subject: Call for Papers: Computational Approaches to Figurative Language

HLT-NAACL 2007 Computational Approaches to Figurative Language: Call for Papers

Figurative language, such as metaphor, metonymy, idioms, personification, simile among others, is in abundance in natural discourse. It is an effective apparatus to heighten effect and convey various meanings, such as humor, irony, sarcasm, affection, etc. Figurative language can be found not only in fiction, but also in everyday speech, newspaper articles, research papers, and even technical reports.

[....]

Important Dates:

Paper submission deadline: January 18, 2007

Notification of acceptance for papers: February 22, 2007

Camera ready papers due: March 1, 2007


Workshop Date: April 26, 2007

You will be six months behind schedule on your first day

Fun Email

- 10 emails covering different topics
- Add motto
 - Version 1: basic (none)
 - Version 2: random one-liner addition
 - Version 3: contextualized one-liner addition
- 13 users ranked the emails on a 10-point scale on four dimensions:
 - entertainment (the email was entertaining)
 - appropriateness (the motto was appropriate)
 - intelligence (the email program behaved intelligently)
 - adoption (I would use the email program myself)

Fun Email





Outline

1. Witty language
 - Humor generation
 - Humor recognition
2. Affective Text
 - Lexical resources
 - Annotation of emotions in text
 - Colors of emotions in texts
 - Dancing with words
3. Persuasive NLP
 - Analyzing political speeches along with audience reactions (e.g. applauses)
 - How to evaluate persuasive language ?
4. Deceptive Language recognition
 - Is it possible to recognize when people are lying, just using the produced text ?