Affect Detection from Spoken and Written Text

Computational Models for Affect and Sentiment Analysis

Alexandra Balahur
European Commission Joint Research Centre
Italy

Maite Taboada
Simon Fraser University
Canada

Björn Schuller
University of Passau & Imperial College
Germany / UK

Tutorial held at IJCAI 2016, July 9th, 2016, New York

http://www.sfu.ca/~mtaboada/sentiment-tutorial/

1. Introduction and term definitions

Affect, emotion, related concepts

- *Affect*
  - "a superordinate concept that subsumes particular valenced conditions such as emotions, moods, feelings and preferences" (Ortony et al., 2005)
  - one of the four components whose interaction make the human organism "function effectively in the world" (Ortony et al., 2005), along with **motivation**, **cognition** and **behaviour**.

Affect, emotion, related concepts

- *Emotion*
  - complex phenomenon
  - no definition that is generally accepted
  - "An episode of interrelated, synchronized changes in the states of all or most of the five organismic subsystems (Information processing, Support, Executive, Action, Monitor) in response to the evaluation of an external or internal stimulus event as relevant to major concerns of the organism". (Scherer, 1987; Scherer, 2001).

Affect, emotion, related concepts

- *Feeling*
  - “The conscious subjective experience of emotion.“ (Van den Bos, 2006)
  - “(...) points to a single component of emotion, denoting the subjective experience process, and is therefore only a small part of an emotion” (Scherer, 2005)

Affect, emotion, related concepts

- *Sentiment*
  - “suggests a **settled** **opinion** reflective of one’s **feelings**.”

Affect, emotion, related concepts

- *Opinion*
  - implies a conclusion thought out yet open to dispute; it is:
  - A): a view, judgment, or appraisal formed in the mind about a particular matter; B): approval, esteem;
  - A): a belief stronger than impression and less strong than positive knowledge; B): a generally held view;
  - A): a formal expression of judgment or advice by an expert; B): the formal expression (as by a judge, court, or referee) of the legal reasons and principles upon which a legal decision is based.

Affect, emotion, related concepts

- *View*
  - suggests a subjective opinion.
- *Belief*
  - implies often deliberate acceptance and intellectual assent.
- *Conviction*
  - applies to a firmly and seriously held belief.
- *Persuasion*
  - suggests a belief grounded on assurance (as by evidence) of its truth.

Affect, emotion, related concepts

- *Attitude*
  - “hypothetical construct that represents an individual's degree of like or dislike for something.” (Breckler and Wiggins, 1992)
  - generally positive or negative views of a person, place, thing, or event — the attitude object.
  - Attitudes are judgments.

Subjectivity and attitude

- Subjectivity
  - Private states (Wiebe 1995)
    - Feelings, emotions, goals, evaluations, judgments
  - Point of view (Langacker, Lyons), in Linguistics
- Subjectivity analysis
  - Recognize subjective language, distinguish it from descriptions of facts
- Attitude – AAAI 2004 Spring Symposium on Attitude
  - “hypothetical construct that represents an individual's degree of like or dislike for something.” (Breckler and Wiggins, 1992)
  - Attitude = {affect, judgment, appreciation}
    - Appraisal
  - Used for speaker/author “intentionality”

2. Opinion mining and sentiment analysis

Knowledge-rich approaches

Creation of resources, existing resources

Affect and sentiment detection from text

- The 3-component model

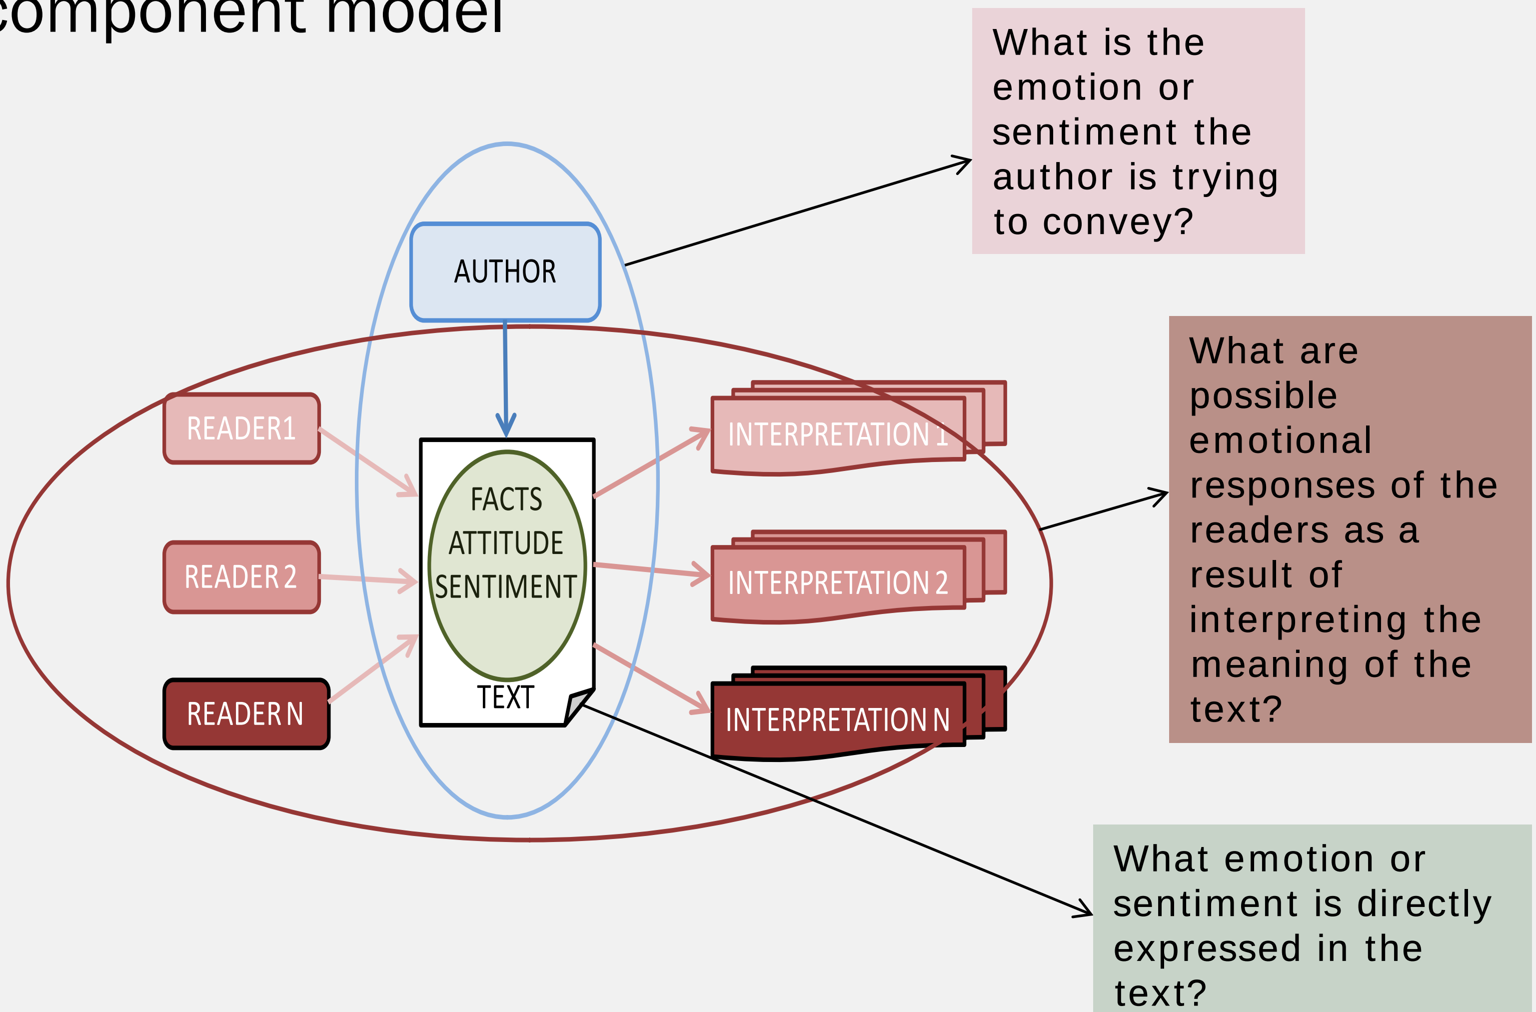

Types of tasks and applications

- The 3-component model

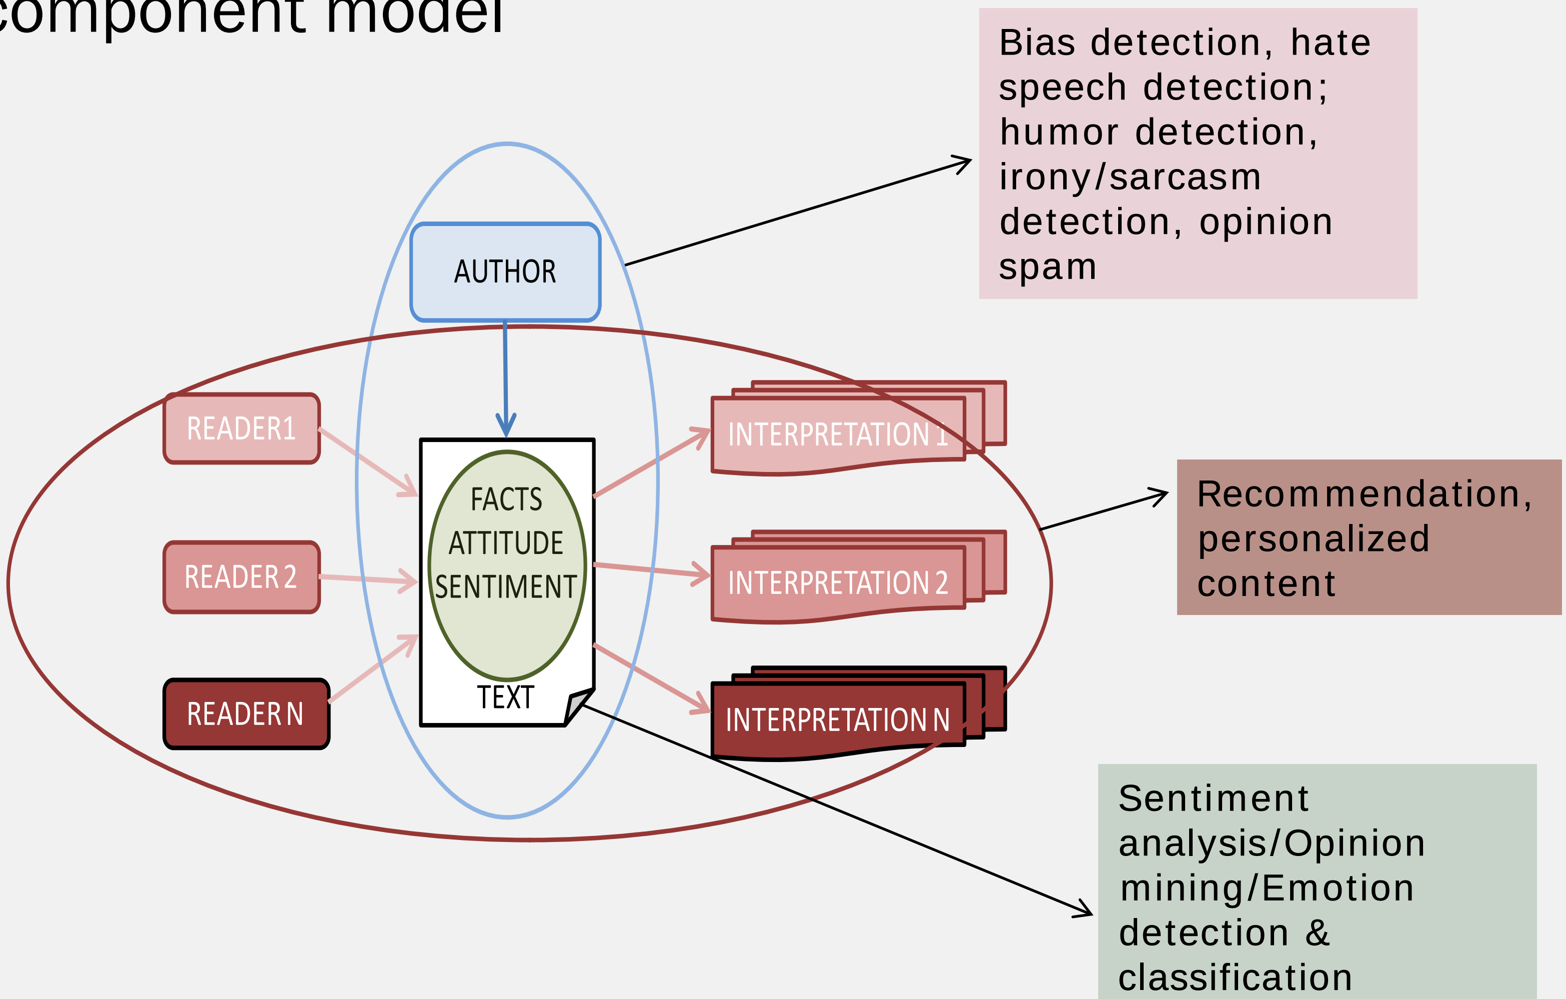

From needs to opinions

- Expressions of desires and needs > Expressions of emotions and evaluation
- The *Affective Turn*
  - In philosophy, sociology and political science (Clough & Halley 2007)
  - Affective computing (Picard 1997)
  - In social media
- We seem to express our feelings more frequently and openly
  - Increase of subjectivity in language over time (Biber 2004; Vis, Sanders & Spooren 2012)
    - More frequent use of stance markers
    - Stylistic change to a more conversational style even in writing

And then came social media

Applications

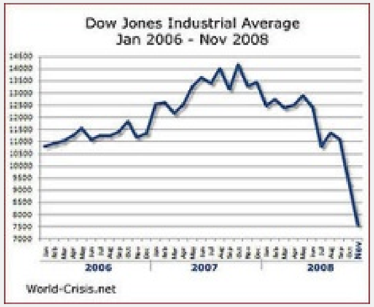

- Helps companies, customers, politicians
  - Marketing, financial studies
  - Choice of products
  - Social media analysis
  - Political view tracking & eRulemaking
  - Election results prediction
  - Policy making
  - Trend analysis

- Improves other NLP tasks
  - IE, QA, MPQA, summarization, authorship, WSD

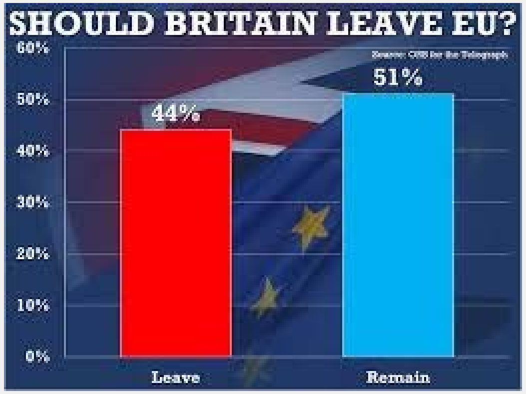

Sentiment analysis or opinion mining?

- Seems to be a matter of preference

- SA maybe more from computational linguistics and NLP (Natural Language Preference)
- Opinion mining maybe inspired by data mining

Sentiment ≠ Opinion ≠ Subjectivity

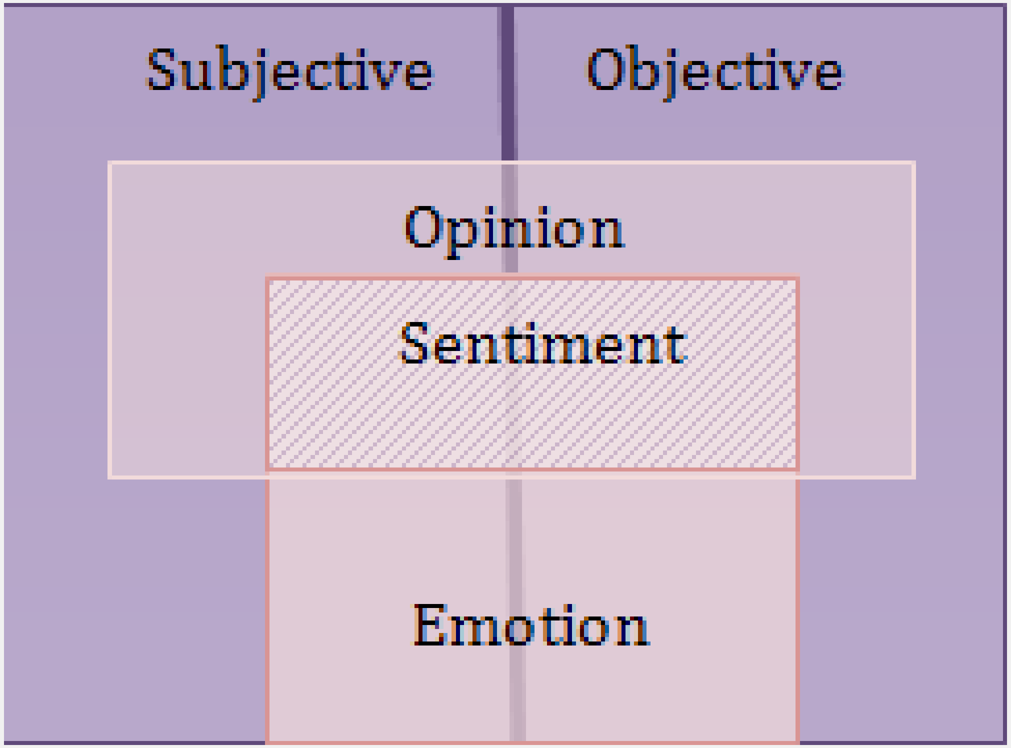

Opinion mining = Sentiment analysis ≠ Subjectivity analysis

Definition of sentiment analysis/opinion mining

- Classification of texts (documents, blog posts, tweets, sentences, headlines…) based on subjective content (=sentiment)
  - Positive
  - Negative
- Input a text, and produce a numeric value that expresses its subjective content ⟶ the text's sentiment

Interpretations of *positive* and *negative*

- Good or bad news (Ku et al. 2005)
- Likes or dislikes (Pang et al. 2002)
- Candidate likely or unlikely to win (Kim and Hovy 2005)
- Support or opposition (Bansal et al. 2008; Terveen et al. 1997)
- Pros and cons (Kim and Hovy 2006)
- Improvement or death in medical texts (Niu et al. 2005)
- Agreement or disagreement with a topic (Malaouf et al. 2005)
- Arguments in favor of or against a topic
  - Stance, argumentation mining (Somasundaran and Wiebe 2009, Hasan and Ng 2013, Stab and Gurevych 2016)

Computational treatment of sentiment

- Bing Liu, definition

$$opinion = (e, a, s, h, t)$$

entity | aspect | sentiment | holder | time

$$s = (y, o, i)$$

semantic category | orientation = polarity | intensity =strength

Example

This restaurant serves incredibly delicious food!

$$opinion = (e, a, s, h, t)$$

- e = restaurant
- a = food
- s = *(y, o, i)*
- h = @chris_smith
- t = 08/07/2016

$$s = (y, o, i)$$

- y = appreciation
- o = positive
- i = high

Methods

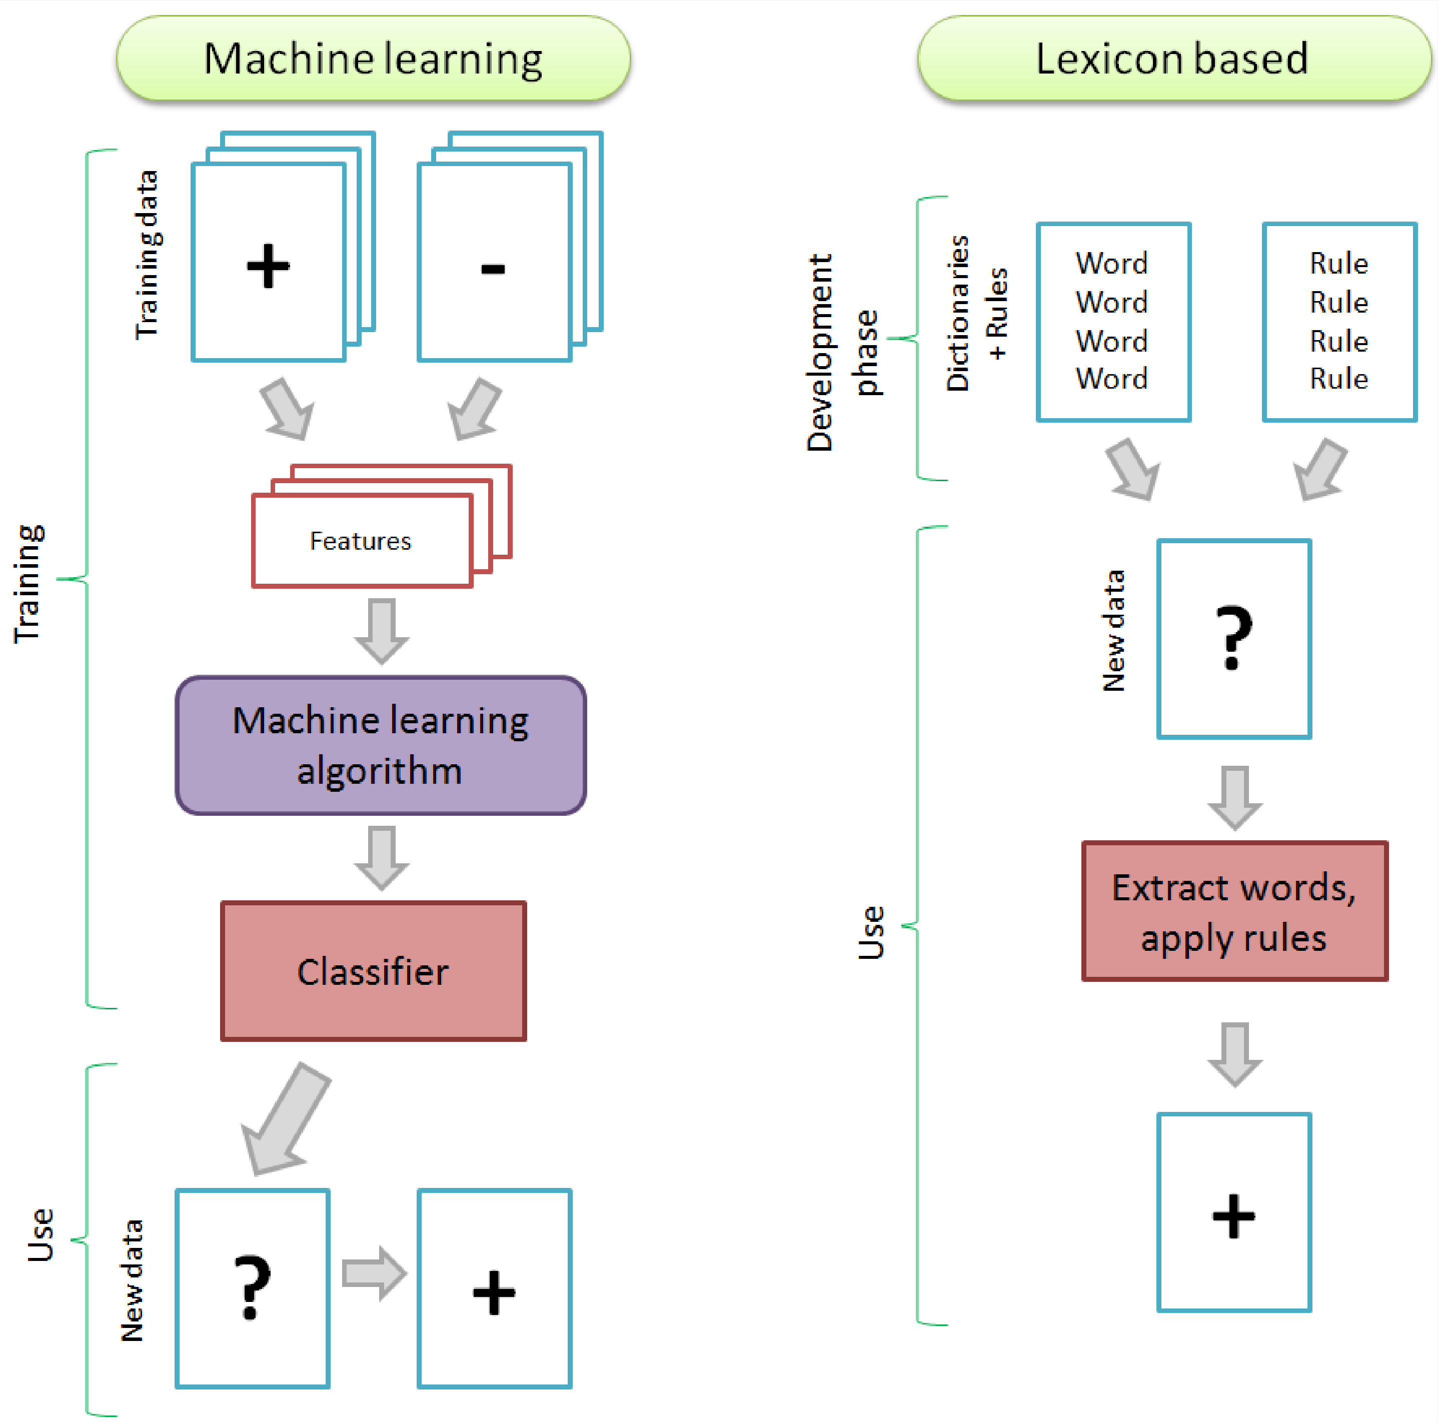

Taboada, M. (2016) Sentiment analysis: An overview from Linguistics. *Annual Review of Linguistics* 2: 325-347.

http://www.sfu.ca/~mtaboada/research/pubs.html

http://get.sfu.ca/TQiDOU

Lexicon-based approaches

- Building dictionaries
  - Manually
    - From intuition, corpora or culling existing resources (e.g., WordNet)
  - Automatically
    - By association, using seed words
    - SVM features
    - Crawl through thesauri
- Basic idea for automatic semantic orientation (SO) of words: you know the SO of a word by the company that it keeps (Hatzivassiloglou and McKeown 1997)
  - Excellent and X
  - Depressing but X

General techniques

- Pre-processing
  - Lemmatizing/stemming and stop word removal
    - Some might prove important (*for, no, and, but*)
  - Text normalization
    - Especially for microblogs, SMS
  - POS-tagging, syntactic parsing, Semantic role labelling
- Features (bags of words)
  - Term presence or term frequency (tf-idf…)
  - Parts of speech
  - Presence of opinion words
  - Presence of negators, intensifiers, downtoners
  - N-grams of different sizes
- Feature selection
  - PMI, Chi-square, LSI/LSA

Aggregation in lexicon-based approaches

- Dictionaries
  - Containing different types of words (adjectives only, adjectives, nouns, adverbs, verbs, etc.)
  - Created automatically or manually
- Process a new text
  - Extract the opinion words from it
  - What to do with the words?
  1. Average the values
  2. Find out more about the context in which the values appear
     - Beginning, middle or end of the text
     - Relating to a particular topic or aspect
       - It is loud
       - Car vs. a telephone

Problems with lexicon-based approaches

- Dictionaries are static
- New domain or new language involves creating new dictionaries
- Words may have different meanings in different contexts
  - loud
    - phone, restaurant
  - unpredictable
    - movie, computer

Machine learning approaches

- Supervised learning
  - Based on annotated corpora
  - Sentiment analysis as a classification problem (2-3-5 classes)

  - Using a plethora of algorithms:
    - Naïve Bayes, Bayesian Network, Maximum Entropy
    - Support Vector Machines
      - Traditionally considered the best
    - Neural Networks, Deep Learning
    - Decision Trees
- Unsupervised learning

Problems with statistical methods

1. **Feature learning is too specific**

- Experiment
- Polarity Dataset (Pang and Lee 2004), a set of 2,000 movie reviews
- We built an SVM classifier with the most positive and negative ngrams
  - Some of the results are predictable
    - *worst, waste, unfortunately* and *mess* – negative
    - *memorable, wonderful, laughs* and *enjoyed* – positive
  - Some are not
    - Positive: mention of performances, ending, flaws
    - Negative: mention of writer, director, plot or script
  - Names are learnt (the *Angelina Jolie* effect)
  - The unigrams *2, video* and *tv* are negative

2. **Poor cross-domain performance**

Current state of the art in sentiment analysis

- Lots of initial progress
  - 'Bags of words' approaches
  - Machine learning from labelled data
- Currently
  - The '80-20 rule' (cf. the Pareto principle)
  - We need pragmatic, discourse and other information
    - Whole-text approach
      - Argumentation structure
    - Threads (in Twitter)
    - Information about users, social networks
    - Textual entailment
    - Intent detection

Why is sentiment analysis difficult?

- Individual words used in context
- Negation and other irrealis phenomena (nonveridicality)
- Out-of-topic sentences
- Irony and sarcasm
- World knowledge
- Discourse structure and argumentation

Negation

- Basic negators: not, never
  - Not funny
- Negation raising (negative is on the main verb)
  - I don’t think it’s funny
- Other negative pronouns and modifiers
  - Nobody liked it
- Negative words (lexical negation)
  - It fails to entertain – It doesn’t entertain
- Partial negation
  - I didn’t like it until the end
  - He didn’t come until 7

Dealing with negation in SA

- Add artificial words to the typical BoW representation (NOT_x) – Pang et al. (2002)
  - I do not NOT_like NOT_this NOT_new NOT_Nokia NOT_model.

- Rules (Polanyi & Zaenen 2004; Kennedy & Inkpen 2005; Taboada et al. 2011)
  - Words have polarity values associated: excellent (+4)
  - Negation means value word *(-1) : not excellent (-4)
  - Shift negation: shift the value up or down the scale, but not reverse it: not excellent (-2)

- Semantic composition using use syntactic phrase structure trees (Moilanen and Pulman 2007)

- Heuristic rules to model *scope of negation* (Choi and Cardie 2008; Jia et al. 2009)
  - Window size after negation word,
  - First occurrence of polar expression, whole sentence

Wiegand, Balahur et al., 2010 "A survey on the role of negation in sentiment analysis"

Irrealis

- Also known as nonveridicality
  - The propositions expressed are non-factual (i.e., they are not ‘real’ or ‘veridical’)
  - Events in the future
  - Events and states modified by a modal verb
  - Conditional and hypothetical situations
- Examples
  - She will become a really good actress
  - I may like the movie, I thought
  - It could have been a really good movie, if only the director had known how to direct

Out-of-topic sentences

- Individual sentences, but also entire paragraphs that do not refer to the main topic/thing being evaluated

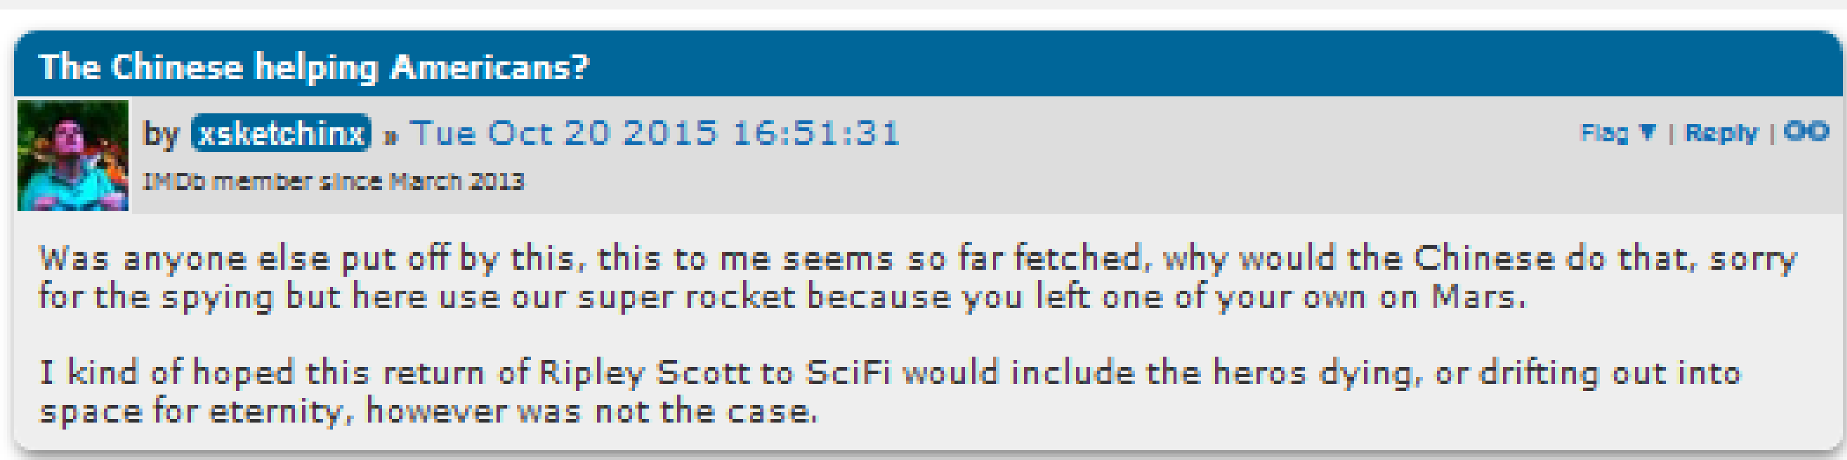

The Chinese helping Americans?

by xsketchinx » Tue Oct 20 2015 16:51:31 Flag ▼ | Reply |

IMDb member since March 2013

Was anyone else put off by this, this to me seems so far fetched, why would the Chinese do that, sorry for the spying but here use our super rocket because you left one of your own on Mars.

I kind of hoped this return of Ripley Scott to SciFi would include the heros dying, or drifting out into space for eternity, however was not the case.

A new, dynamic definition of sentiment

- Current work with Farah Benamara
  (Université Paul Sabatier, Toulouse)
- Old definition, static

$$opinion = (e, a, s, h, t)$$

- New definition, dynamic

$$\langle \Omega, \Im \rangle$$

$$\Omega = (e, a, s, h)$$

intrinsic properties

contextual updates

$$\Im = \{F_1, \ldots, F_n\}$$

$$\forall F_i \in \Im, F_i : \Omega \mapsto Update_l(\Omega)$$

What a great animated movie! I was so scared the whole time that I didn't even move from my seat.

$$\boxed{\Omega = (e, a, s, h)}$$

- great

$$\Omega_1 = Update(\Omega_1) = (movie, \_, (great, +1), author)$$

- scared

$$\Omega_2 = Update\ (\Omega_2) = (movie, \_, (scared, -1), author)$$

- so scared

$$Update_{sentence}(\Omega_2) = (movie, \_, (so\ scared, -2), author)$$

$$Update_{discourse}(\Omega_2) = (movie, \_, (so\ scared, +2), author)$$

- didn't even move from my seat

$$Update_{pragmatic}(\Omega_2) = (movie, \_,$$
$$\{(so\ scared, +2), (didn't\ move...seat, +3\}, author)$$

Benamara, F., M. Taboada and Y. Mathieu (to appear) Evaluative language beyond bags of words: Linguistic insights and computational applications. *Computational Linguistics*.

http://www.sfu.ca/~mtaboada/research/pubs.html

Main research areas

- Creation of resources
  - Lexical resources for subjectivity/polarity (subjectivity, orientation, strength)
  - Annotation schemes appropriate to each textual genre (news/blogs/product reviews)
  - Corpora labelling for training and evaluation
    - Some approaches use as gold standard already punctuated reviews (stars)

Methods to create lexicons for SA

- Seed adjectives – apply synonymy and antonymy in WN (Hu & Liu 2004)
- Seed adjectives – use conjunctions/disjunctions to deduce orientation of new words & min-cut graphs (Pang & Lee 2002; Hatzivassiloglou & McKeown, 1997)
- Terms with similar orientation tend to co-occur in documents (seed words + PMI using number of AltaVista returned results with NEAR) (Turney 2002)
- Terms with similar glosses in WordNet tend to have similar polarity (Esuli & Sebastiani 2005)

Existing resources: Lexicons

- Opinion & affect lexicons
  - WordNet Affect (Strapparava & Valitutti 2004)
    - http://wndomains.fbk.eu/wnaffect.html
  - SentiWordNet (Esuli & Sebastiani 2006, 2010)
    - http://sentiwordnet.isti.cnr.it/
  - Subjectivity indicators (MPQA) (Cardie et al. 2003)
    - http://mpqa.cs.pitt.edu/lexicons/
  - Appraisal terms (Whitelaw 2006)
  - NRC Twitter lexicons (Mohammad et al. 2014)
    - http://saifmohammad.com/WebPages/lexicons.html

Existing resources: Lexicons

Manually created lexical resources

- Dictionary of Affect (Whissell)
  - http://sail.usc.edu/dal_app.php
- Affective Norms for English Words (Bradley & Lang)
  - http://csea.phhp.ufl.edu/media.html
- Harvard General Inquirer categories (Stone etc.)
  - http://www.wjh.harvard.edu/~inquirer/
- NRC Emotion Lexicon (Mohammad & Turney) http://saifmohammad.com/WebPages/lexicons.html
- MaxDiff Sentiment Lexicon (Kiritchenko, Zhu & Mohammad) http://saifmohammad.com/WebPages/lexicons.html

Existing resources: Datasets

- Affective Text Dataset – news, headlines (Strapparava & Mihalcea)
  - http://web.eecs.umich.edu/~mihalcea/downloads.html#affective
- Affect Dataset (Alm) – classic literary tales; sentences
  - http://people.rc.rit.edu/~coagla/
- 2012 US Presidential Elections – tweets (Mohammad et al.)
  - http://saifmohammad.com/WebDocs/ElectoralTweetsData.zip
- EmotionML (Schröder et al.)
  - http://www.w3.org/TR/emotionml/
- ISEAR (Scherer 1997)
- MPQA (Wiebe et al. 2002)
- TAC/TREC data (2006-2008)
- NTCIR MOAT data (2007-2010)
- SemEval data
  - Sentiment analysis in Twitter (2013-2016)
  - Aspect-based sentiment analysis (2016)
  - Detecting stance in Tweets (2016)
  - Detecting sentiment intensity (2016)

Existing resources: Other languages

- Spanish:
  - TASS (Taller de Analisis de Sentimientos y Sujetividad) http://www.sngularmeaning.team/TASS2013/corpus.php
  - Perez-Rosas Lexicon https://web.eecs.umich.edu/~mihalcea/downloads.html#SPANISH_SENT_LEXICONS
  - iSOL (Molina-Gonzales et al. 2013)
- Dutch:
  - Framework for interpersonal communication (Vaassen & Daelemans, 2011)
  - OpeNER http://www.opener-project.eu/documentation/
- German:
  - German polarity clues: http://www.ulliwaltinger.de/sentiment/
- Chinese:
  - 2013 Chinese Microblog Sentiment Analysis Evaluation (CMSAE) Dataset of posts from Sina Weibo annotated with seven emotions:
  - http://tcci.ccf.org.cn/conference/2013/pages/page04 eva.html
- Japanese:
  - Japanese customer reviews corpus with the same eight emotions used in the Chinese Ren-CECps Corpus (Sun et al., 2014)

Competitions

- TAC 2008 Opinion Pilot
- SemEval 2007 Affect in Text
- SemEval 2013-2016 Sentiment Analysis in Twitter
- NTCIR-MOAT series
- SemEval 2016 Detecting Stance in Tweets
- SemEval 2016 Aspect-based Sentiment Analysis
- SemEval 2016 Determining Sentiment Intensity in of English and Arabic Tweets
- SemEval 2017 Tasks: Just announced!

Applications of SA

Interactive

http://text-processing.com/demo/sentiment/

http://nlp.stanford.edu:8080/sentiment/rntnDemo.html

http://demo2-opener.rhcloud.com/welcome.action

https://www.lexalytics.com/demo

http://www.citizenandscience.eu

http://www.alchemyapi.com/products/demo

https://www.csc.ncsu.edu/faculty/healey/tweet_viz/tweet_app/

Very VERY important caveat

- Know your data!
- Not all lexicons, not all methods will work well on all types of data
- Test of 24 (!) SA systems, on 18 (!) labelled datasets
    - Wild variation in performance depending on lexicon and/or data

SentiBench - a benchmark comparison of state-of-the-practice sentiment analysis methods

Filipe N Ribeiro, Matheus Araújo, Pollyanna Gonçalves, Marcos André Gonçalves and Fabrício Benevenuto

EPJ Data Science 2016 5:23 | DOI: 10.1140/epjds/s13688-016-0085-1 | © Ribeiro et al. 2016
Received: 3 February 2016 | Accepted: 19 June 2016 | Published: 7 July 2016

http://epjdatascience.springeropen.com/articles/10.1140/epjds/s13688-016-0085-1

Finally, Table 9 presents the Friedman's test results showing that there are significant differences in the mean rankings observed for the methods across all datasets. It statistically indicates that in terms of accuracy and Macro-*F*1 there is no single method that always achieves a consistent rank position for different datasets, which is something similar to the well-known 'no-free lunch theorem' [16]. So, overall, before using a sentiment analysis method in a novel dataset, it is crucial to test different methods in a sample of data before simply choose one that is acceptable by the research community.

3. Applications

News bias detection

- Complex process
  - Social, political and economic dimensions (Hamilton 2004)
- Persuasion works with
  - Specific choice of words
  - Subtle structure of sentences can persuade the reader towards one point of view or another and are sufficient to influence whether people interpret violent acts as patriotism or terrorism (Dunn et al., 2012)
- Methods
  - Usage of various parts of speech, like adjectives, adverbs and nouns and how these properties differ (Pollak et al. 2011)
  - Length of texts, headlines

Hate speech detection

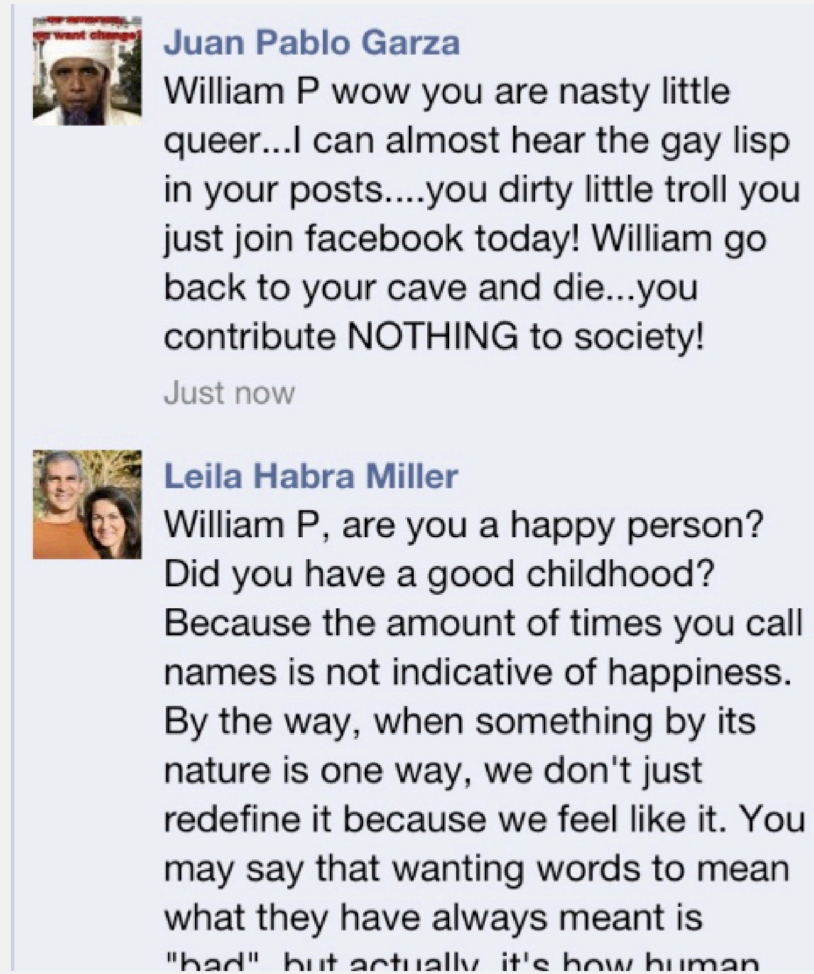

Hate speech detection

- Very little work done in this field
- Difficult to formally define the task
- Borders freedom of speech
- Methods
  - Keyword frequency (Warner & Hirschberg 2012)
  - Word embeddings and neural networks (Djuric et al. 2015)

The Guardian

The web we want

theguardian

The dark side of Guardian comments

As part of a series on the rising global phenomenon of online harassment, the Guardian commissioned research into the 70m comments left on its site since 2006 and discovered that of the 10 most abused writers eight are women, and the two men are black. Hear from three of those writers, explore the data and help us host better conversations online

by Becky Gardiner, Mahana Mansfield, Ian Anderson, Josh Holder, Daan Louter and Monica Ulmanu

Tuesday 12 April 2016 06.23 EDT | 1877 Shares

https://www.theguardian.com/technology/2016/apr/12/the-dark-side-of-guardian-comments

Humour, irony, sarcasm

- The new frontier
- Humour
  - Very much dependent on interpretation
  - Current methods: surrounding information (pragmatics, emoticons)
  - Mihalcea and Strapparava (2005); Stock and Strapparava (2003, 2005, 2006)

Irony, sarcasm

- #sarcasm
- Contrast between very positive and very negative terms in the same message (mostly Twitter)
- Most current work
  - Social network analysis
  - Tweets, retweets
  - Who's the author?
- Filatova (LREC 2012), (Riloff 2013)
- Paolo Rosso
  - http://www.slideshare.net/AINL/1-paolo-rosso-ainl-irony

Interim conclusions

- Sentiment analysis is ‘hot’
- Extracting evaluation automatically is hard
- We’ve tried the easy approaches
- Now we’ve got to tackle
  - Other aspects of the text
    - Pragmatics, discourse and context
  - Other aspects of the user
    - Social network, information culled from profile

4. Machine learning methods

Sentiment “in the Wild”

- **Sentiment “In the Wild”**
  - A/V multilingual data
  - Advertisement
  - Social Games
  - Multimedia Recommender

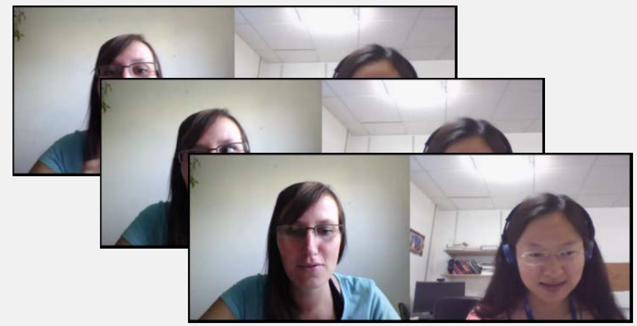

- **Big Data Analytics**
  - Multilingual text, A/V, social data
  - Social TV
  - Brand Reputation Management
  - Call Centre Operations

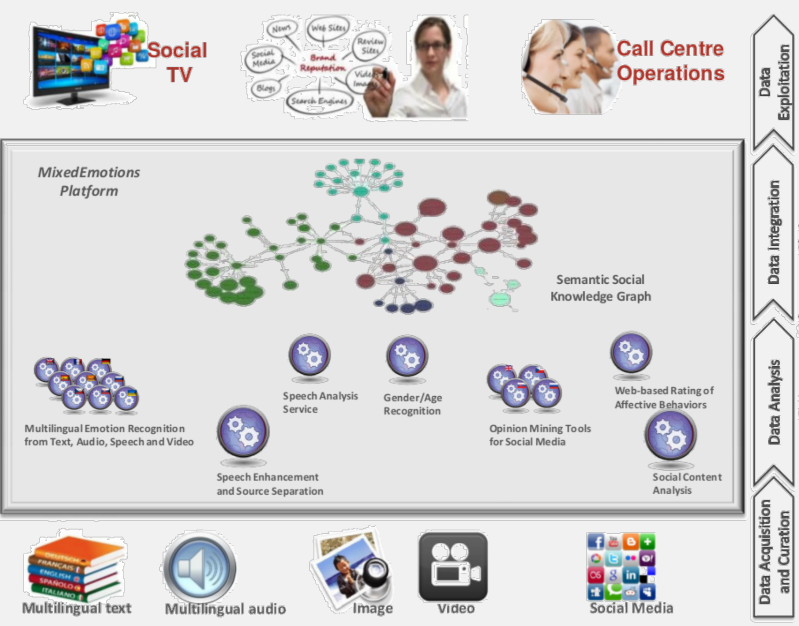

Big Data

- **Characteristics**
  - Volume – e.g., 300 hours videos / min (YouTube, Dec 2014)
  - Velocity – e.g., 500 mio Tweets / day (Twitter, Aug 2013)
  - Variety – e.g., text, audio, video, sensors, diverse formats
  - Basically all 10(+) Vs…

- **Challenges**
  - Unstructured (emails, social media, transactions, …)
  - HW limits (data: ~ x2/1.5 years, disk speed: linear…)
  - Scaling, Visualisation, Privacy, Ethics…

- **Chances**
  - Parallelisation (GPGPUs, multicore, etc.)
  - Distribution
    (Cloud MapReduce, Disco, Hadoop, Skynet, etc.)

Big Data

MixedEmotions

HORIZON 2020

DB	Volume	Velocity	Variety	Type
DW	13 TB / 8300 h	GB/h	MP3, MP4, diff. resol. / format	video, audio
PX	7 TB	$n*10^2$ h/d	raw audio, screens, web pages, WebEx videos, …	audio, text, images, video
ES	350 mio tweets/day users: 900 mio fb 400 mio google+ +1 mio LinkedIn / d	real-time	social data feed from 150 mio social media / online feeds	social media, text
PT	130 mio web pages, 12,000 company websites	crawled every 48 h	various forms/sizes	text, social media

Multimodality

- **Large-Scale Sentiment Analysis**
  - Amazon, IMDB, Twitter

% Acc.	Amazon	IMDB	Twitter
Polarity	91.9	91.6	80.8

  - YouTube & ExpoTV

% Acc.	Audio	Video	AV/Text
Polarity	64.4	61.2	73.0

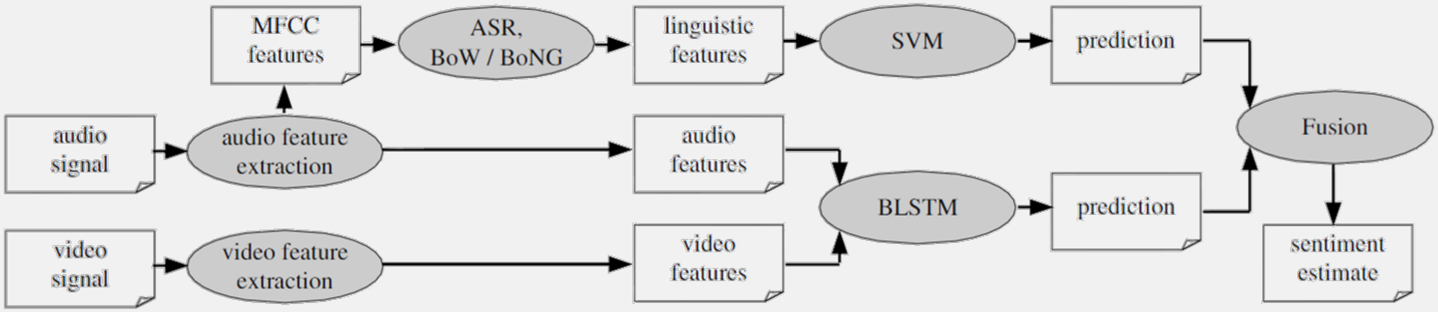

"YouTube Movie Reviews: In, Cross, and Open-domain Sentiment Analysis in an Audiovisual Context", IEEE Intelligent Systems Magazine, 2013.

Features

- **Hetergeneous Features**
  - Fast computation
  - Cross-signal
  - Speed-aware selection

RTF (#feat)	Intel i7	HTC OneM9	Galaxy S3
½ k	.01	.06	.43
6 k	.04	.23	.63

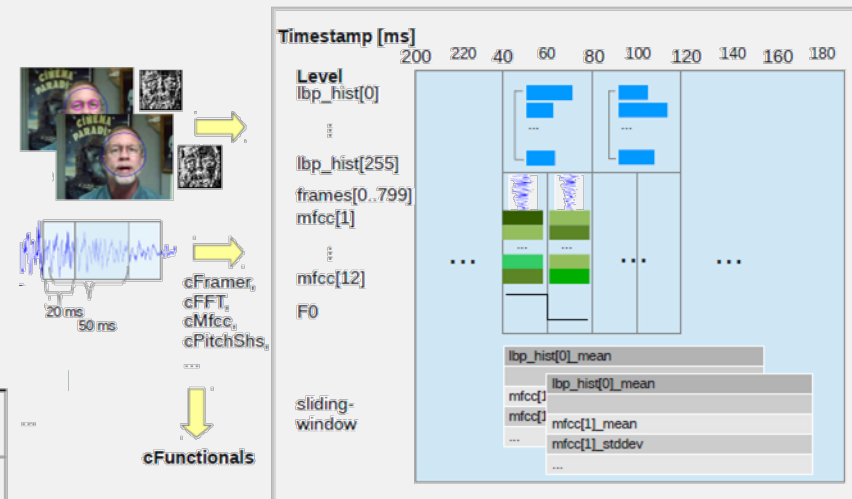

openSMILE:)

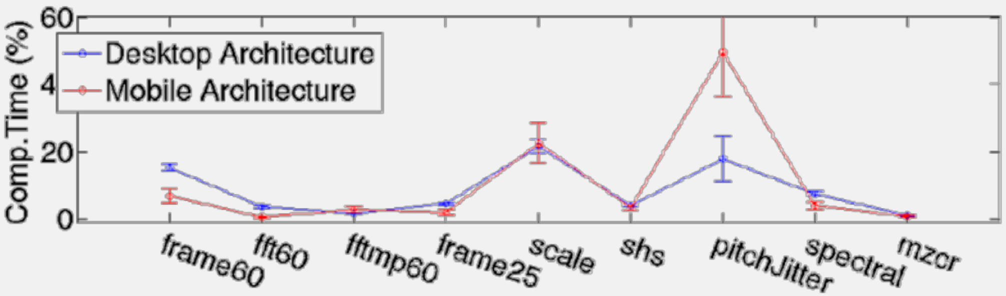

"Recent Developments in openSMILE, the Open-Source Multimedia Feature Extractor", ACM Multimedia, 2013.

Features

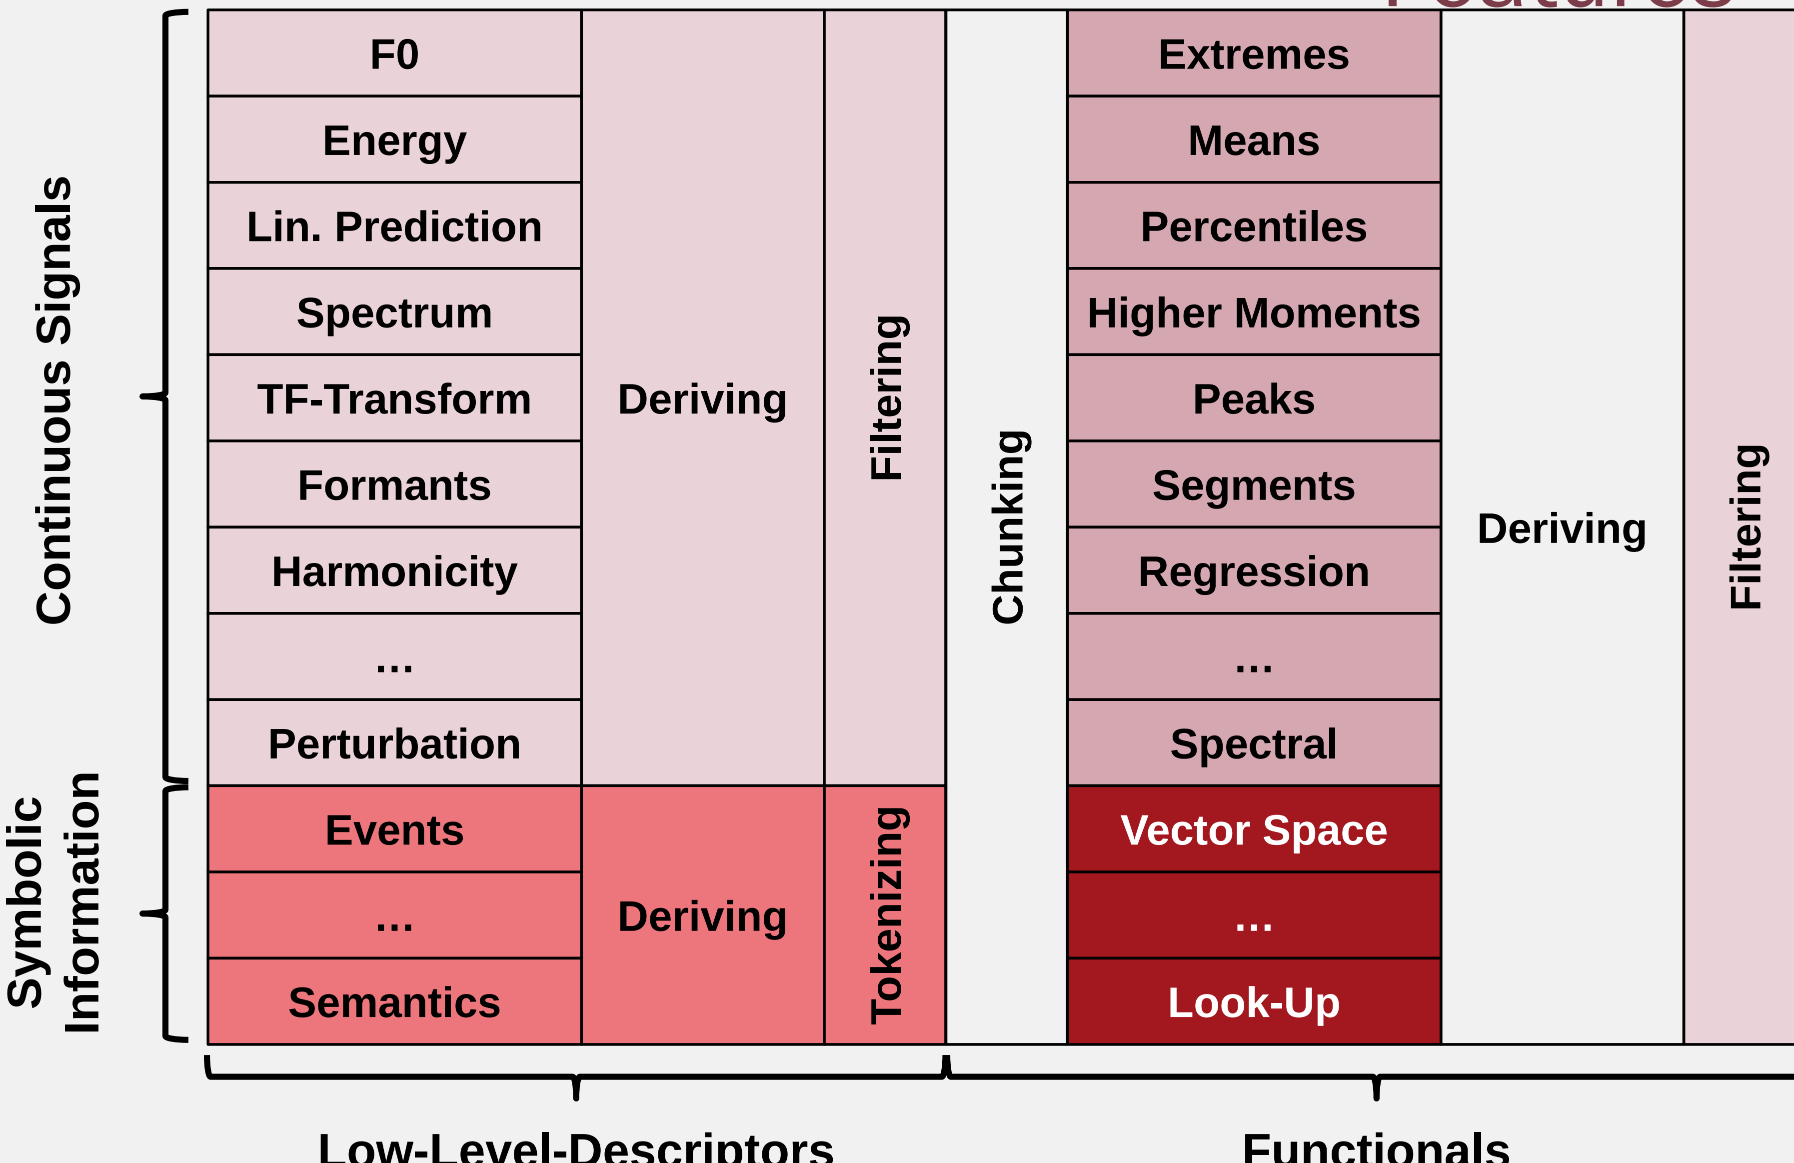

Continuous Signals
Symbolic Information
F0
Energy
Lin. Prediction
Spectrum
TF-Transform
Formants
Harmonicity
…
Perturbation
Events
…
Semantics
Deriving
Filtering
Deriving
Tokenizing
Chunking
Extremes
Means
Percentiles
Higher Moments
Peaks
Segments
Regression
…
Spectral
Vector Space
…
Look-Up
Deriving
Filtering
Low-Level-Descriptors
Functionals

Classification and Regression

Requirement	Example
Adequate Modeling	Static-/(async.)dynamic modeling
	Data-/Knowledge-driven
	Handling of missing features
	Handling of uncertainty
	Learning stability
	Model-/Instance-based
	Transparency
Optimal Accuracy	Non-linear problem handling
	Discriminative learning
	Auto-weighting of features
	Tolerance wrt. dimension
	Adaptability
	Allowance for diverse spaces
Efficiency	Real-time recognition
	Short learning/adaptation time
Economic Factors	Low computational cost
	Low memory requirement
	Low HW realization costs
	Space optimization w/o training
Optimal Integration	N-best provision
	Confidence provision
	Class-wise confidences
	Distributable

Sentiment

- **“Classical” Sentiment Analysis**
  - Text corpora

F1 [%]	AMAZ	PD2	N&BD	SPD1	AMP3	LMRD	TWIT
Vocabulary Size	10k	5k	5k	5k	10k	10k	10k
Binarized Naïve Bayes	89.98	94.72	79.70	81.24	89.44	89.09	79.26
Bernoulli Naïve Bayes	90.32	94.64	76.93	81.37	**89.62**	89.33	79.28
Multinomial Naïve Bayes	89.99	94.64	79.83	81.07	89.47	89.10	79.25
Maximum Entropy	**91.91**	**96.99**	**80.81**	**81.72**	89.22	**91.55**	**80.83**
Ordinal Regression	79.32	89.30	71.00	78.75	86.17	85.65	78.24
SoftMax Regression	75.18	84.61	66.52	75.87	75.64	83.90	77.15

Twitter Dataset[21]	400,000	400,000	Tweets	TWIT

“Sentiment Analysis and Opinion Mining: On Optimal Parameters and Performances,” WIREs Data Mining and Knowledge Discovery (5):255–263, 2015.

Deep Learning.

$$\mathbf{y} = f_{DNN}(\mathbf{x}) = o\left(\mathbf{W}^K \sigma\left(\mathbf{W}^{K-1} \sigma(\dots \mathbf{W}^1 \mathbf{x})\right)\right)$$

- Deep Neural Networks
  - DNN = MLP ?
  - Build them layer-wise
  - Less uninitialised parameters
  - Use "raw" features as input
  - Net learns own higher-level features
  - Use more training data

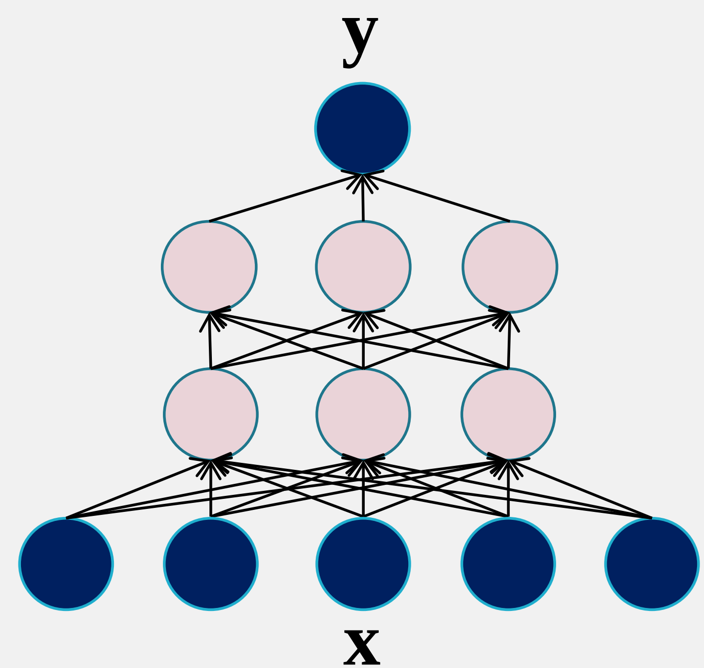

$$E_{Tr}(\mathbf{y}, \mathbf{y}^*) = \sum_{\mathbf{x} \in Tr} D(f_{DNN}(\mathbf{x}), \mathbf{y}^*)$$

$$w^{(i+1)} = w^{(i)} - \eta \frac{\partial E_B}{\partial w}(w^{(i)}),\ B \subset Tr$$

Deep Recurrent Nets

- “Very Deep” NN

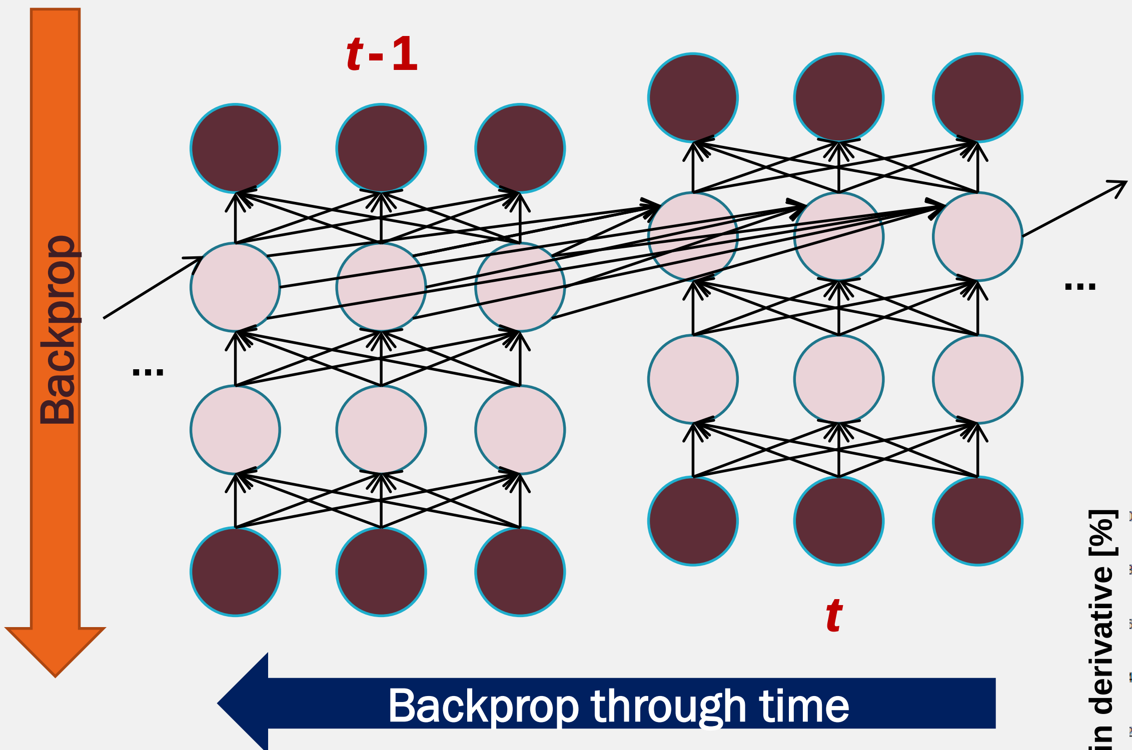

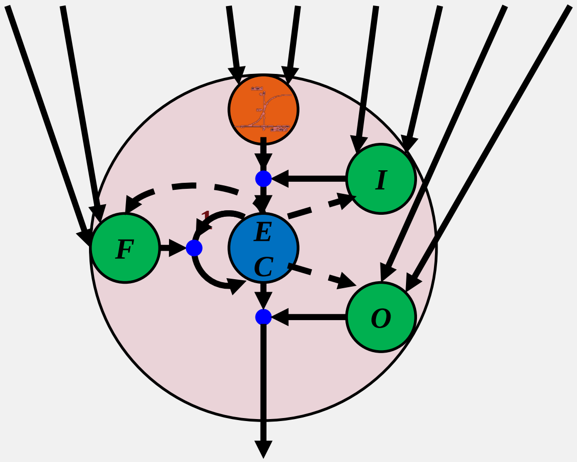

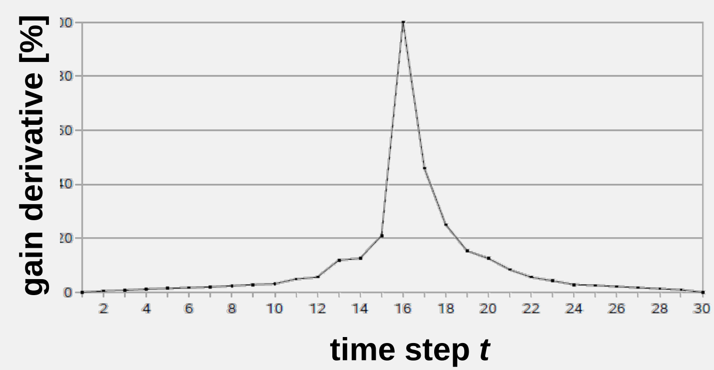

$$\mathbf{y}_t = f_{RNN}(\mathbf{x}_0, \ldots, \mathbf{x}_t)$$
$$= o\left(\mathbf{W}^o \sigma\left(\mathbf{W}\mathbf{x}_t + \mathbf{R}\,\sigma(\ldots \mathbf{W}\mathbf{x}_0 + \mathbf{R}\mathbf{h}_0)\right)\right)$$

Deep Recurrent Nets

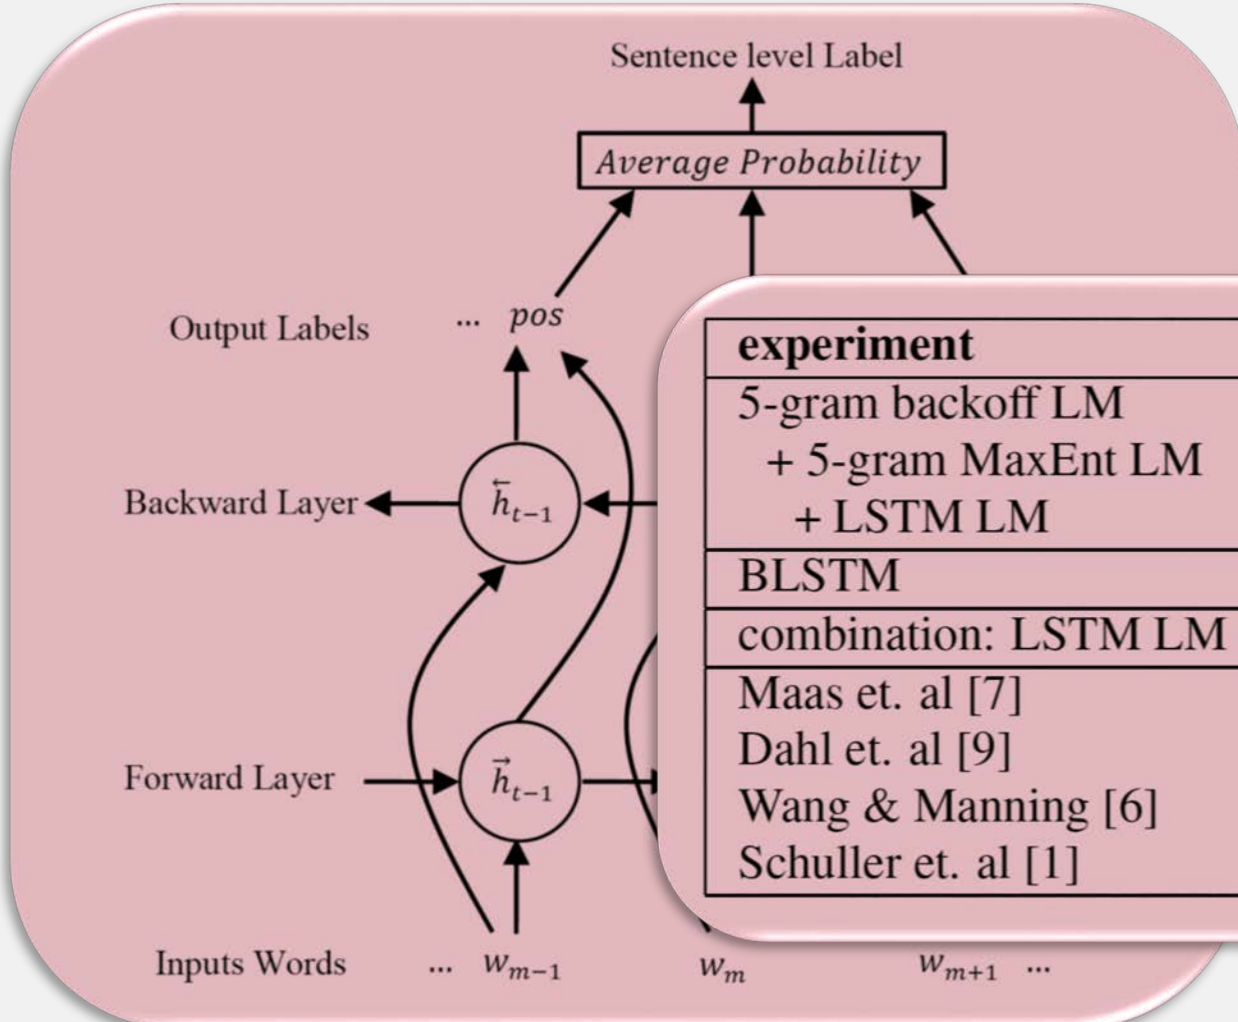

experiment	Accuracy [%]
5-gram backoff LM	90.00
+ 5-gram MaxEnt LM	90.50
+ LSTM LM	**90.80**
BLSTM	90.15
combination: LSTM LM + BLSTM	**92.25**
Maas et. al [7]	88.98
Dahl et. al [9]	89.23
Wang & Manning [6]	91.22
Schuller et. al [1]	91.55

"Long Short-Term Memory Recurrent Neural Network Language Models for Sentiment Analysis", to appear.

Word Embeddings

- Word2Vec (Mikolov et al):
  2-layer NNs
  reconstruct linguistic contexts of words

- Text corpus → high-dim space
  per word in vocab: corresponding vector in the space
  words w/ common contexts: close-by

- Distributional hypothesis:
  “words in similar contexts have similar meanings”

Holism

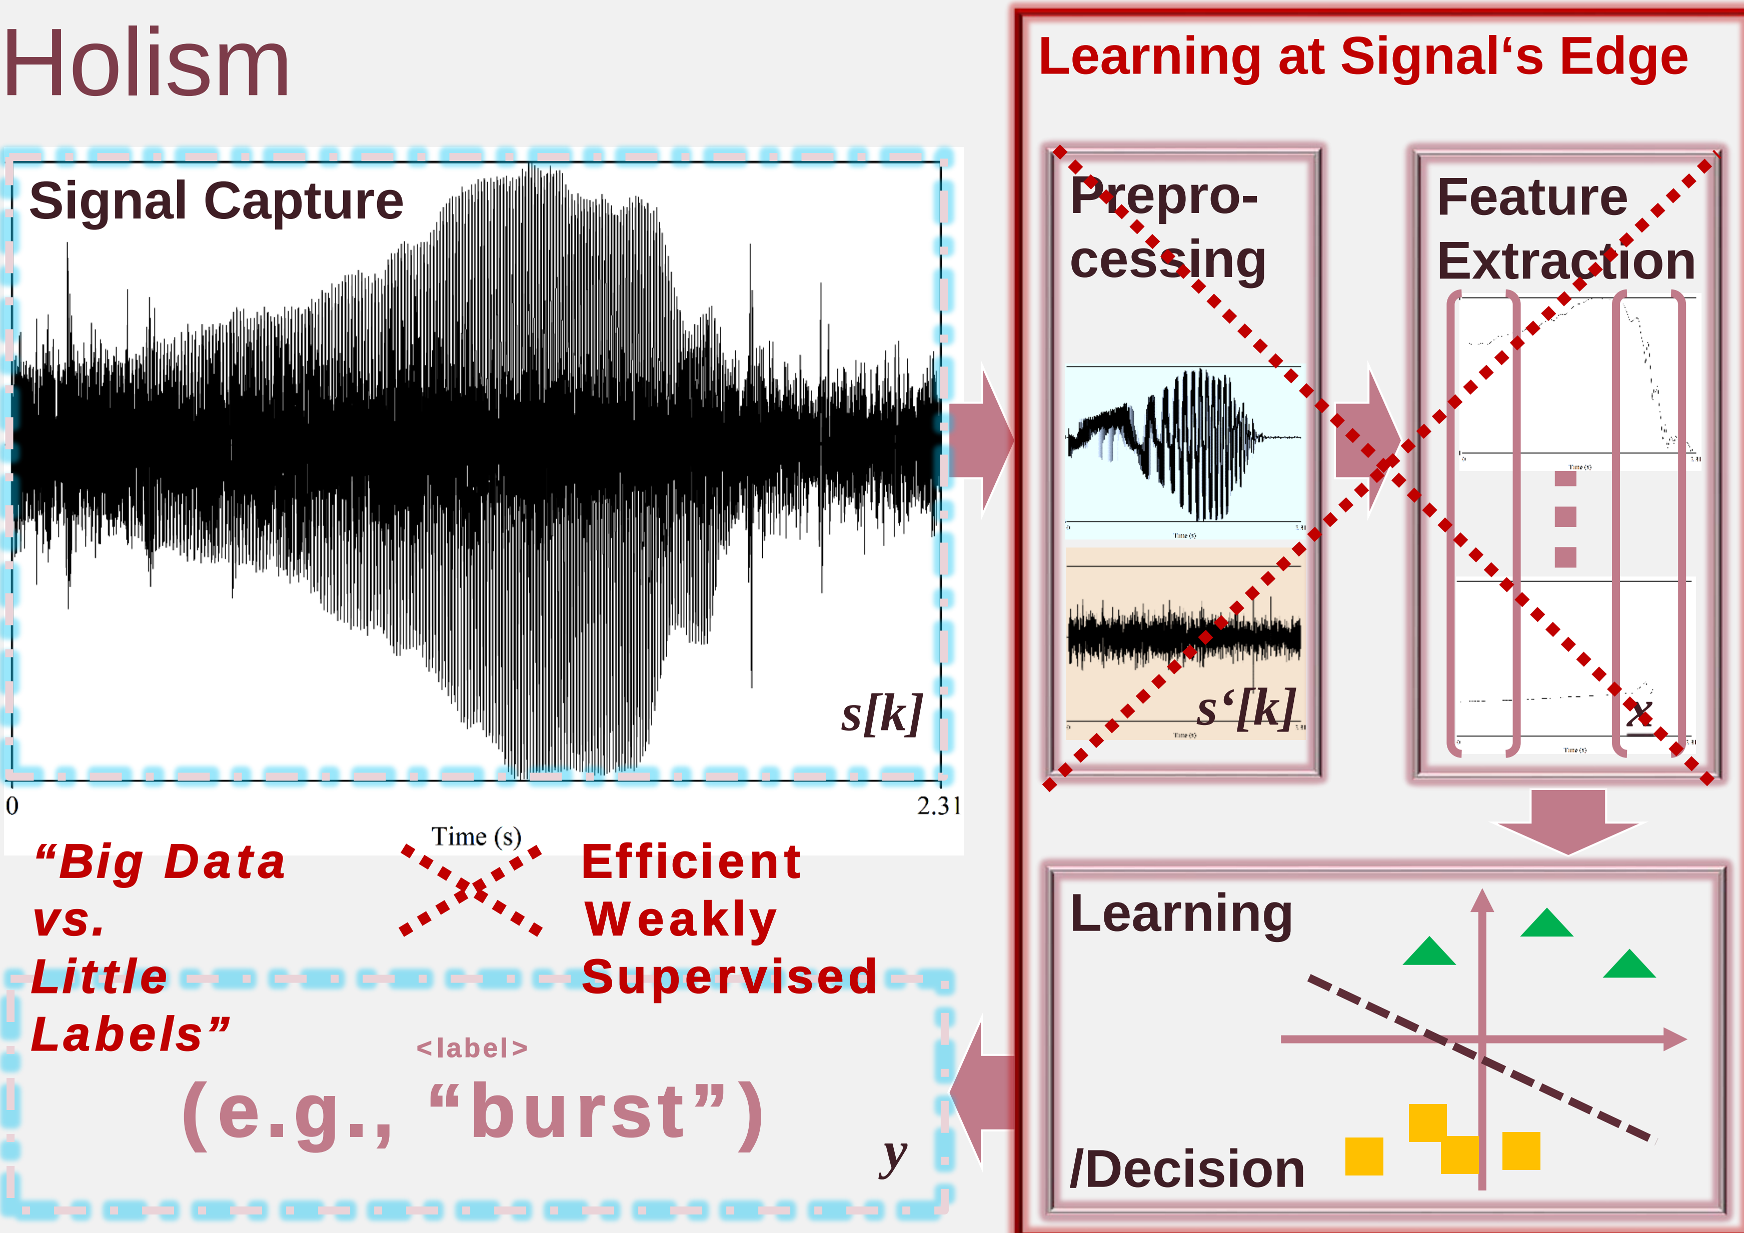
Signal Capture
s[k]
0
2.31
Time (s)
"Big Data vs. Little Labels"
Efficient Weakly Supervised
<label>
(e.g., "burst")
y
Learning at Signal's Edge
Prepro-cessing
s'[k]
Feature Extraction
x
Learning
/Decision

End-2-End Learning

- **Convolutional RNNs**

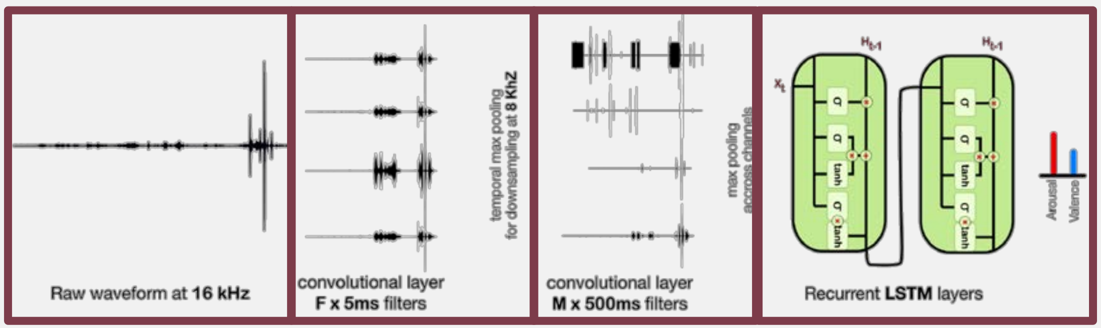

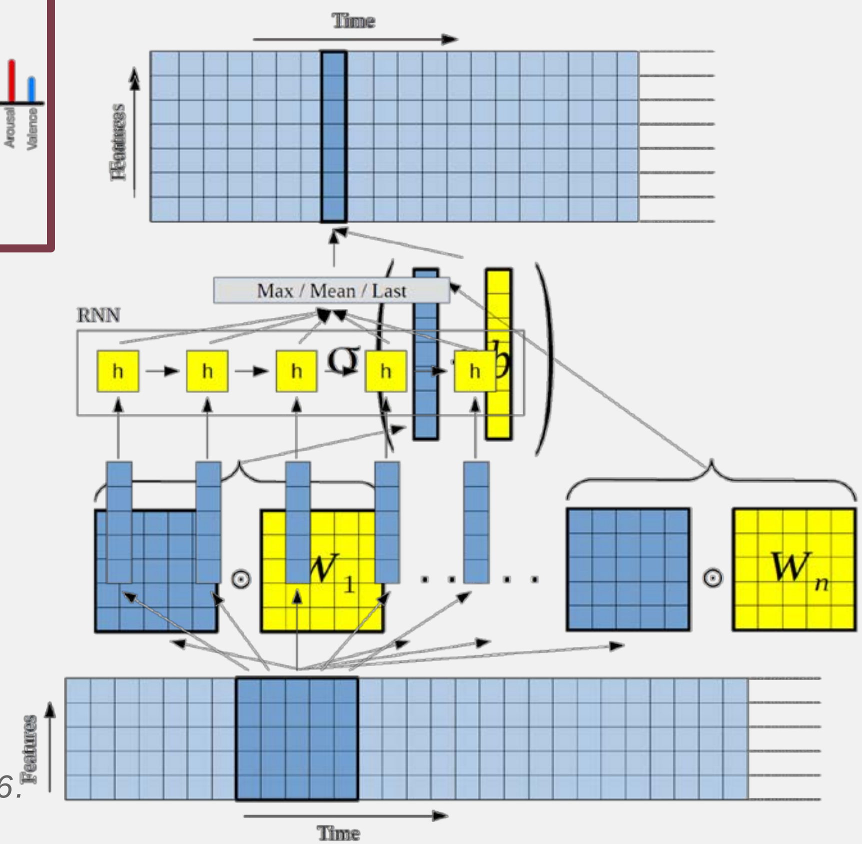

Arousal	CC
Baseline	.366
Deep CRNN	.686

"Adieu Features? End-to-End Speech Emotion Recognition using a Deep Convolutional Recurrent Network", ICASSP, 2016.

End-2-End Learning

- Example: AVEC 2016

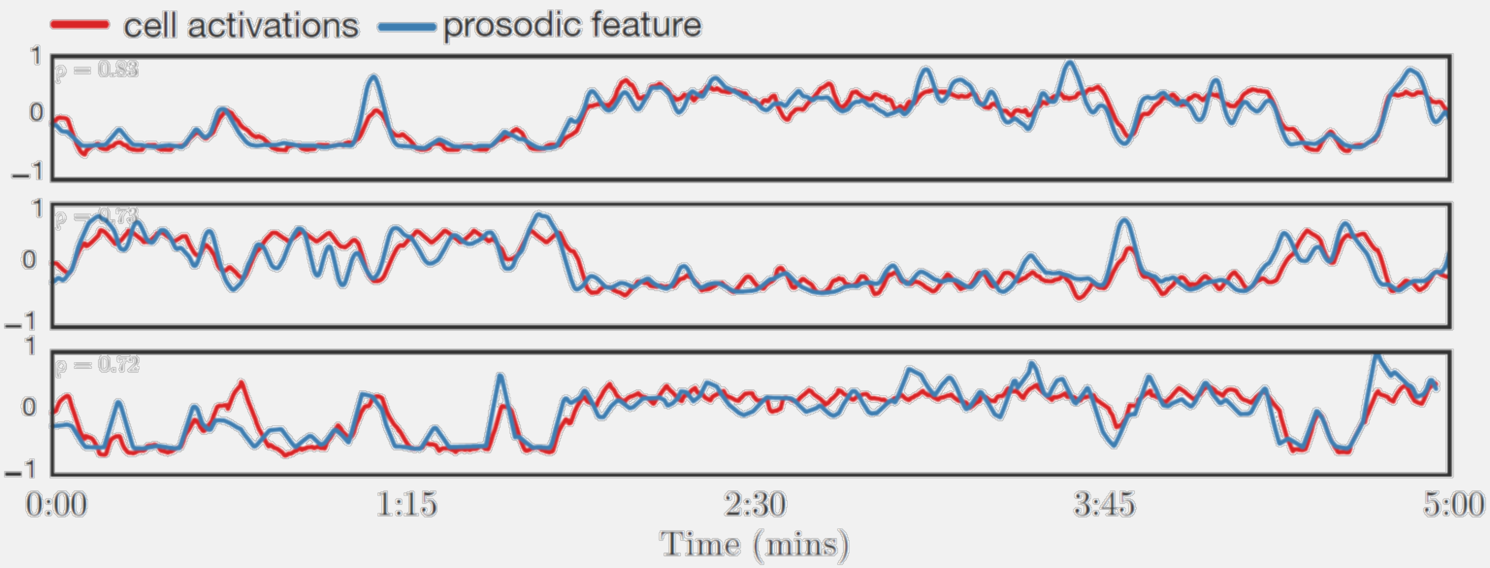

- energy range (.77), loudness (.73), F0 mean (.71)

Fast Learning

- **CURRENNT**
  - 10 – 1k LSTM cells,
  - 2k – 4Mio parameters
  - GPGPU

CHiME 2013 (#1)	RNNLIB		CURRENNT	
#Parallel seq.	1	1	10	200
Error (10 ep.)	0.138	0.138	0.135	0.144
Error (50 ep.)	0.120	0.119	0.116	0.119
Train t / epoch [s]	7 420	3 805	580	334
Speedup	(1.0)	2.0	12.8	22.2

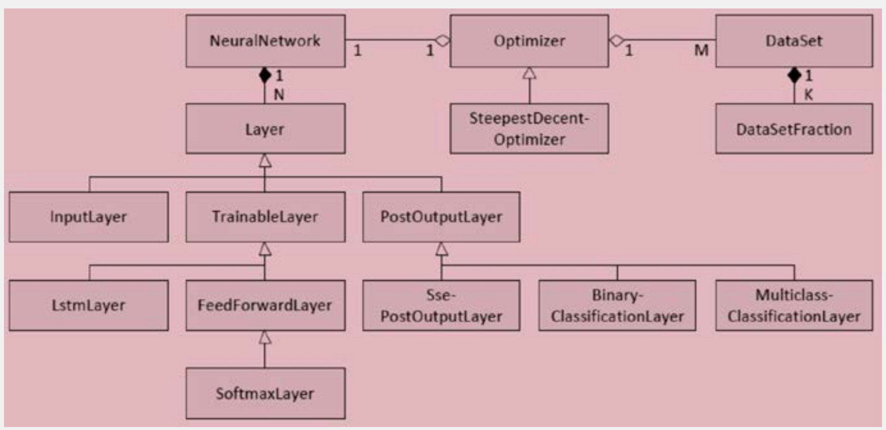

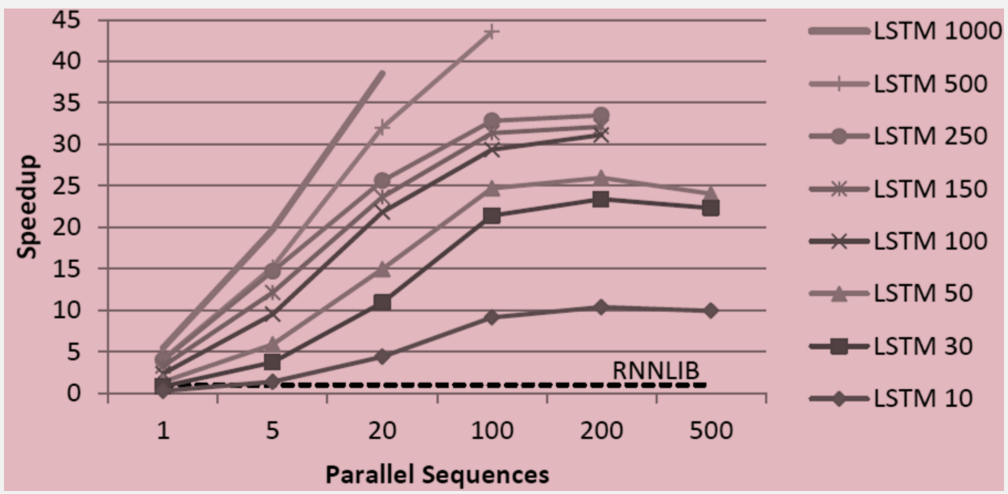

"Introducing CURRENNT - the Munich Open-Source CUDA RecurREnt Neural Network Toolkit", Journal of Machine Learning Research, 2014.

Deep Clustering

- **Deep Clustering**
  - Learn latent attribute hierarchy
  - Better representation of attribute
  - w/ lowest variability

Emotion	%UA
Baseline	60.5
DSNMF	83.2

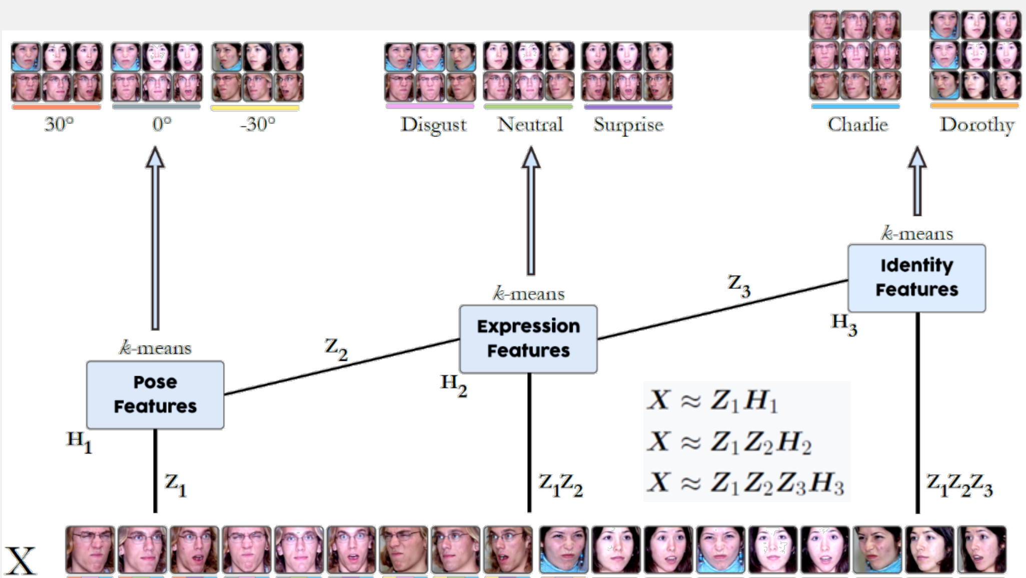

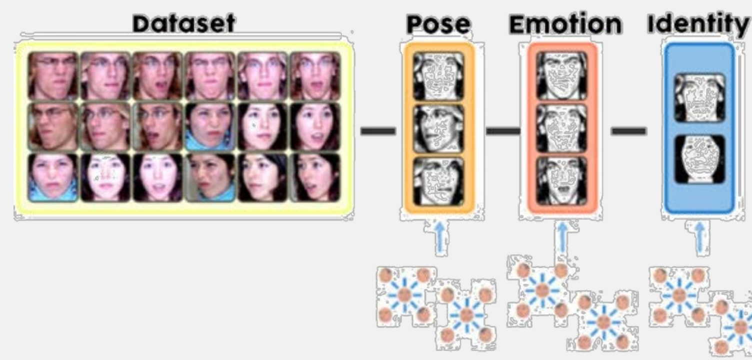

$$C_{\text{deep}} = \frac{1}{2}\|X - Z_1 Z_2 \cdots Z_m H_m\|_F^2$$

"A deep matrix factorization method for learning attribute representations", IEEE Transactions on Pattern Analysis and Machine Intelligence, 2016.

Deep Cross-Modal Warping

- **Deep Cross-Modal Warping**
  - Maximise correlation (CCA), introduce hierarchy (deep)

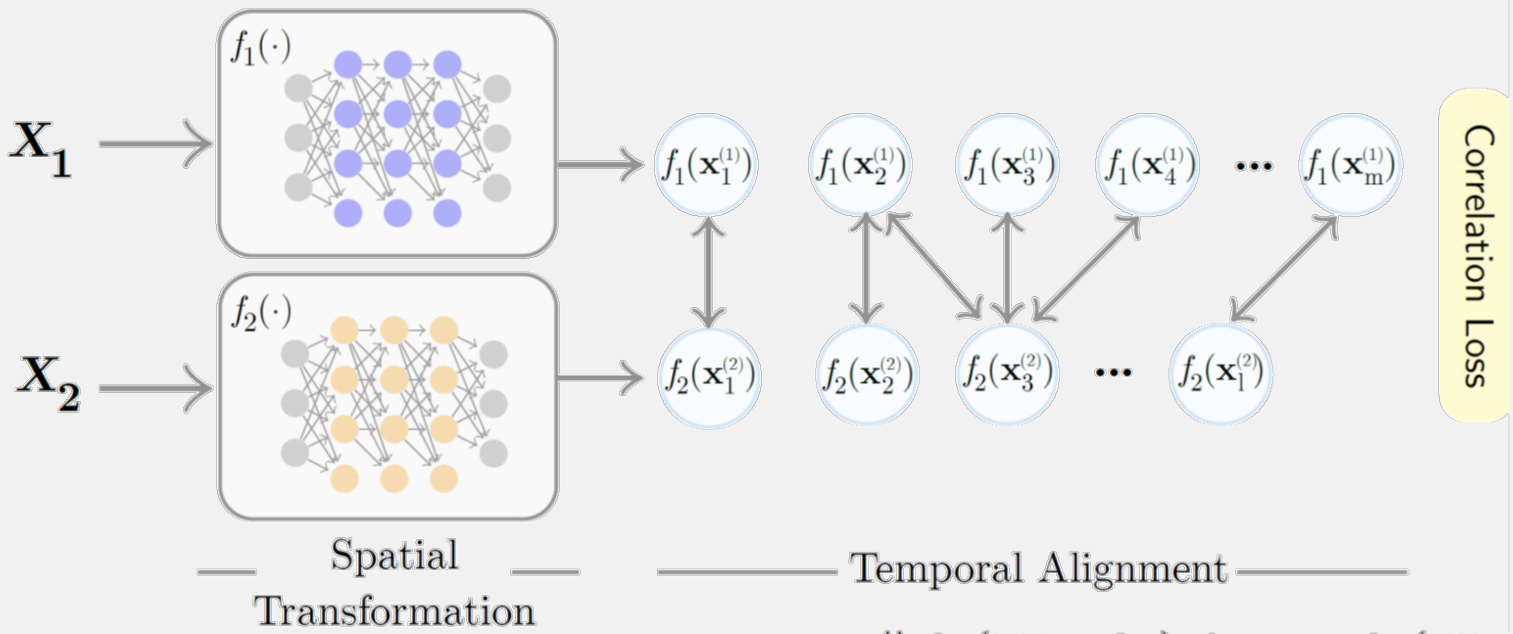

$$\underset{\theta_1, \theta_2}{\arg\min} \, \| f_1(\mathbf{X}_1; \theta_1)\mathbf{\Delta}_1 - f_2(\mathbf{X}_2; \theta_2)\mathbf{\Delta}_2 \|_F^2$$

"Deep Canonical Time Warping", CVPR, 2016.

Bag-of-X-Words

- **Example: Audio Words**
  - Base:
  - UA/WA = 54.3/61.2%
  - VQ vs SVQ
  - VAM: Valence

LLDs	Functionals
ZCR, ΔZCR	mean, standard deviation
RMS energy, ΔRMS energy	kurtosis, skewness, linear
F0, ΔF0	regression (offset, slope,

Configuration			Recognition	
# SV	SVCS [bit]	CS [bit]	UA [%]	WA [%]
2	2	4	61.6	65.1
2	4	2	60.3	62.7
3	**2**	**5**	**64.2**	**65.6**
3	2	6	63.7	65.6
3	4	2	64.6	64.6
3	4	3	63.8	63.2
4	4	3	62.1	61.7
6	3	3	63.2	63.2
8	4	3	65.3	62.7
12	3	4	63.4	64.1

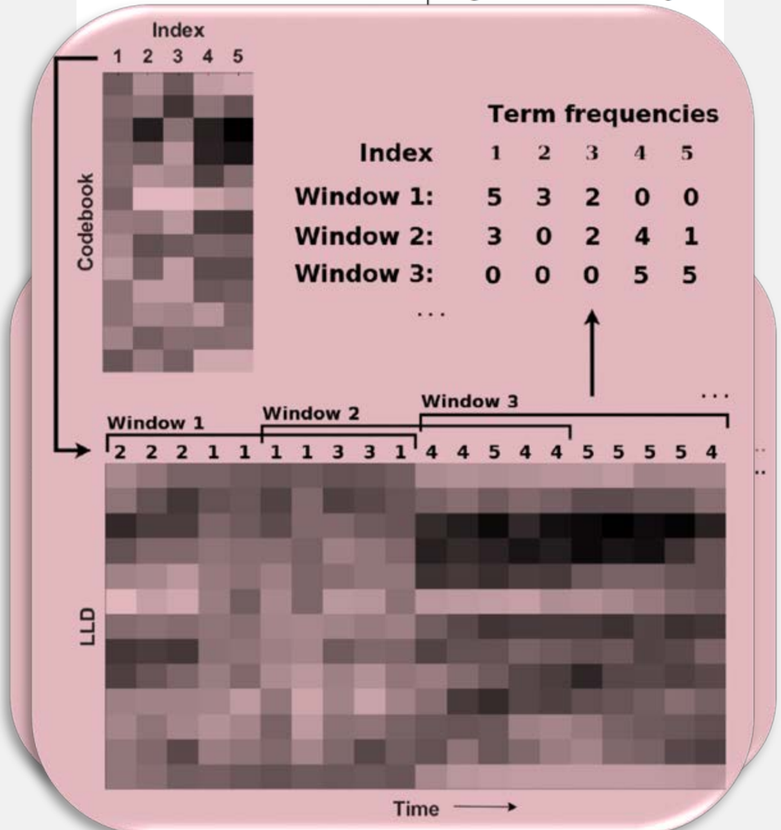

"Detection of Negative Emotions in Speech Signals Using Bags-of-Audio-Words", WASA, 2015.

Bag-of-X-Words

- Split Vector Quantisation
- + Histrogram

openXBOW – |)→

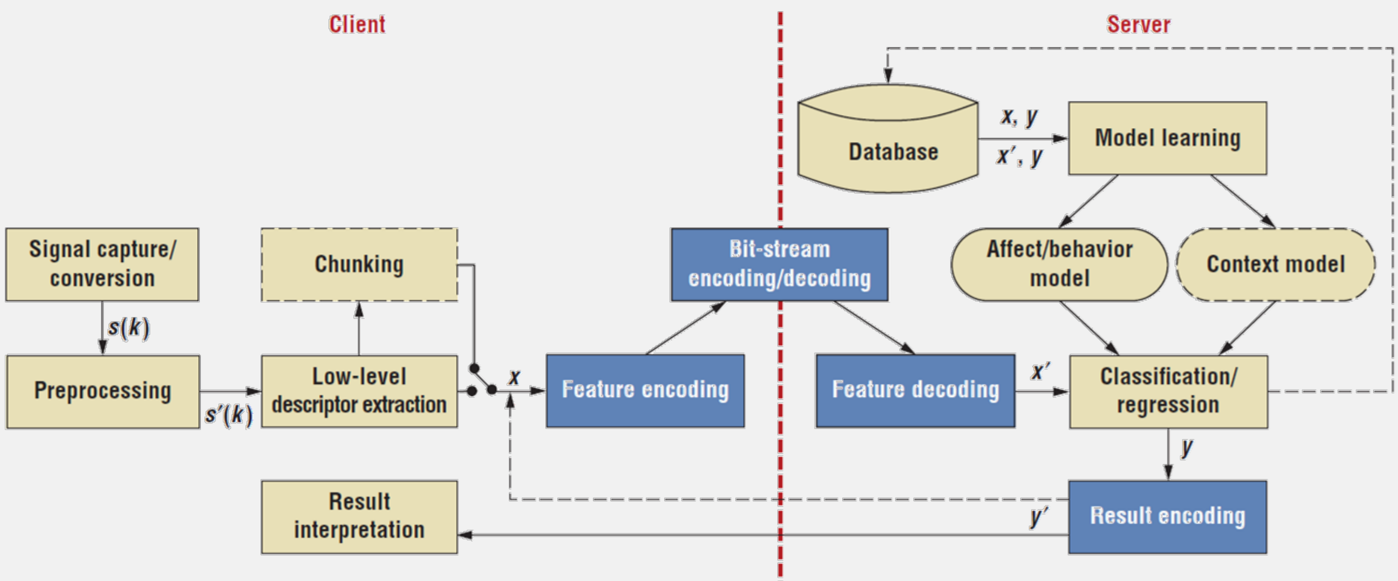

“Distributing Recognition in Computational Paralinguistics,” IEEE Transactions on Affective Computing, 2014.

Learnt Features

- Comparison on the RECOLA (AVEC 2016) task

CCC Valid/Test	Arousal	Valence
Functionals	.790/.720	.459/.402
BLSTM-RNN	.800/.???	.398/.???
CNN (e2e)	.741/.686	.325/.261
BoAW	.793/.753	.550/.430
BoAW+Fctls	.799/.738	.521/.465

openXBOW – |)→ + CURRENT?

"At the Border of Acoustics and Linguistics: Bag-of-Audio-Words for the Recognition of Emotions in Speech", Interspeech, 2016.

Transfer Learning

- Cross-Lingual Emotion Recognition

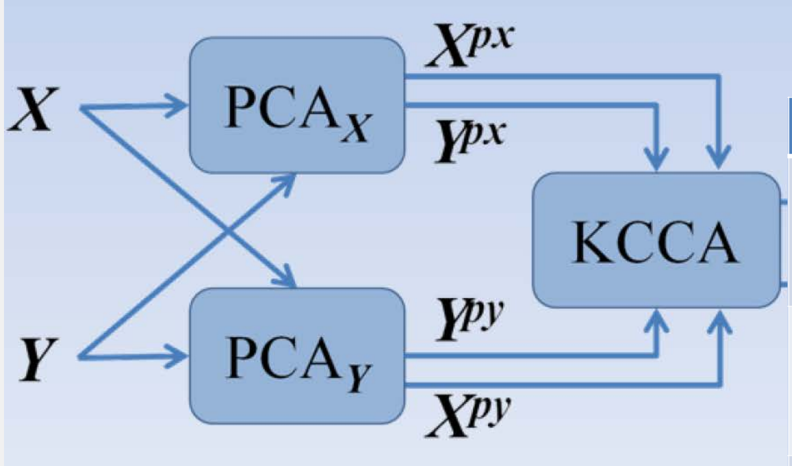

PAIR-WISE (UAR)		EMODB	SAVEE	EMOVO	Polish
EMODB (German)	Baseline		55.8	56.4	71.9
	KCCA		**65.2**	**57.9**	**72.8**
	SHLA[2]		63.7	56.8	59.7
SAVEE (English)	Baseline	62.5		54.6	65.0
	KCCA	**71.9**		**59.5**	69.4
	SHLA	67.7		59.0	**70.2**
EMOVO (Italian)	Baseline	58.0	51.2		55.3
	KCCA	62.9	**58.5**		**71.9**
	SHLA	**67.3**	58.2		64.3
Polish	Baseline	65.1	55.4	57.3	
	KCCA	70.9	**61.9**	56.1	
	SHLA	**71.1**	60.6	**58.3**	

"Cross Lingual Speech Emotion Recognition using Canonical Correlation Analysis on Principal Component Subspace", ICASSP, 2016.

Transfer Learning

- Music vs Speech

% UA / CC	Target	w/o	DAE	DAE-NN
ComParE:EC	60.4	56.3	59.2	64.2
M → S: A	.82	.05	.32	-
M → S: V	.51	.06	.16	-

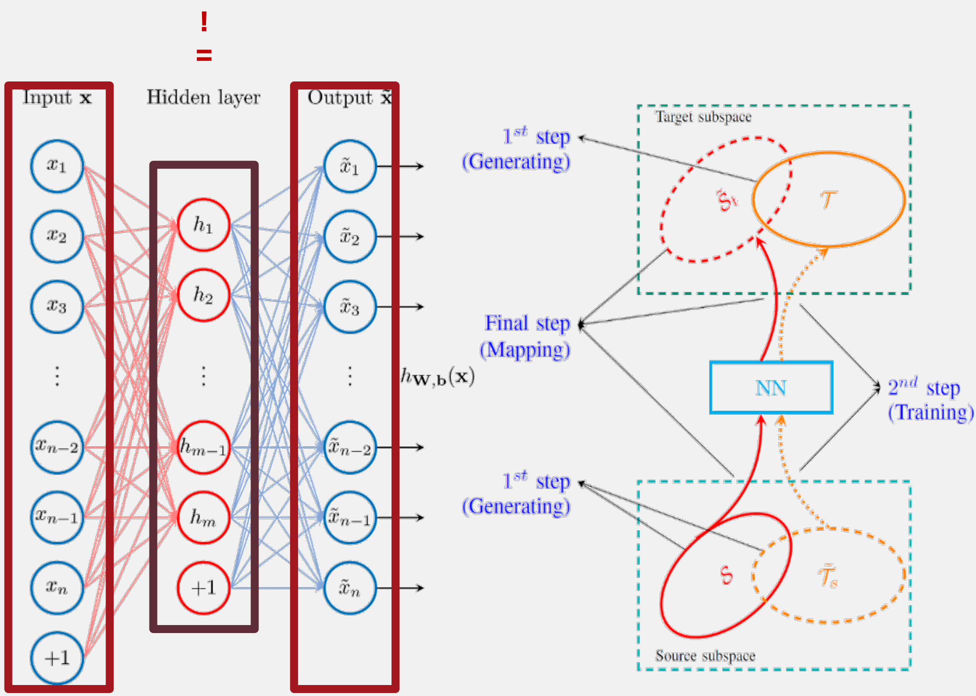

"Autoencoder-based Unsupervised Domain Adaptation for Speech Emotion Recognition", IEEE Signal Processing Letters, 2014.

Cooperative Learning

- **Cooperative Learning in aRMT**
  0) Transfer Learning
  1) Dynamic Active Learning
  2) Semi-Supervised Learning

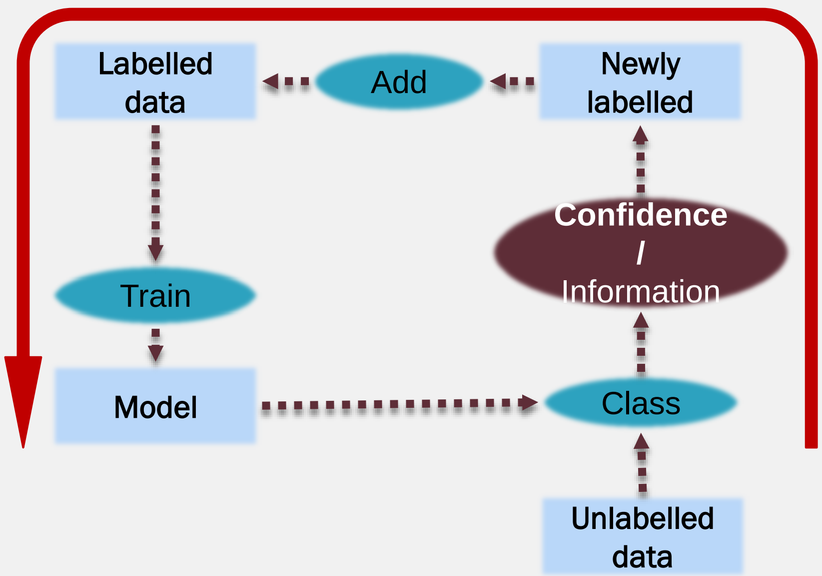

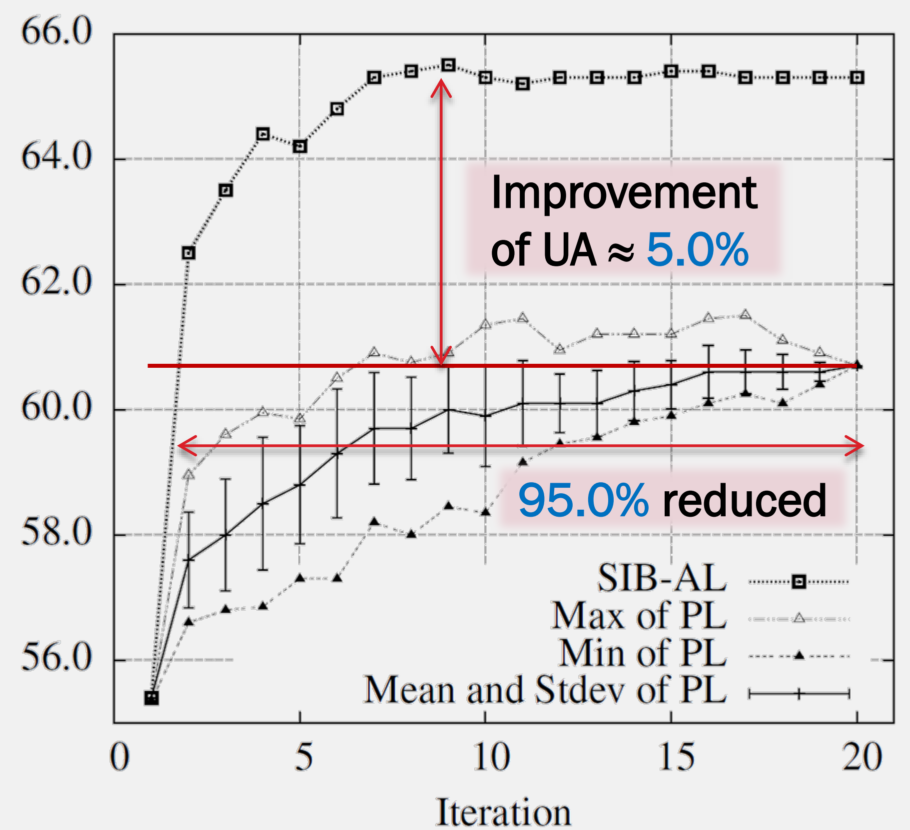

"Cooperative Learning and its Application to ESR", IEEE Transactions ASLP, 2015.

iHEARu PLAY

Crowd Sourcing

- **Playful Sourcing**
  - Gamified annotation
  - Competing w/ others
  - Gratification…

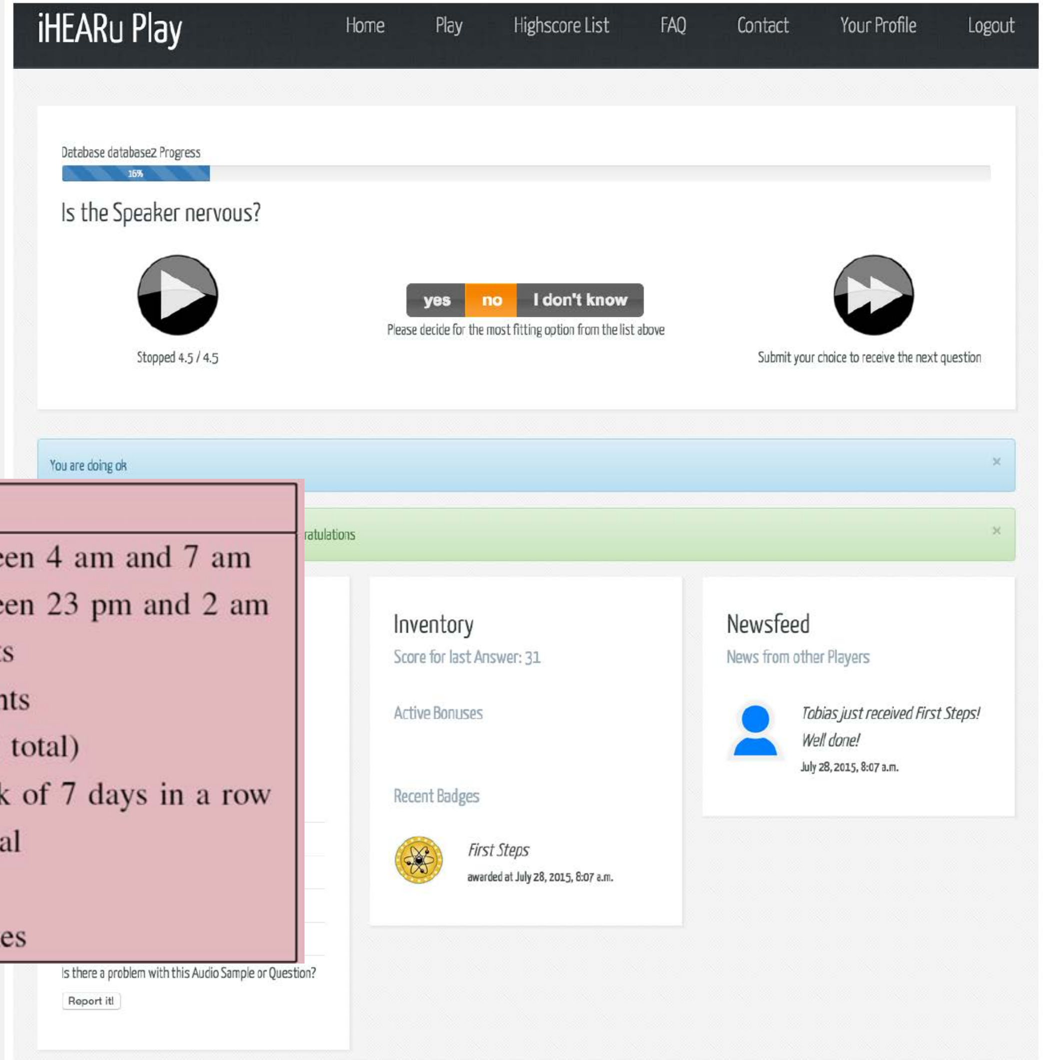

Badge Name	Conditions
Early Bird	Answer 100 questions between 4 am and 7 am
Night Owl	Answer 100 questions between 23 pm and 2 am
Expert	Reach a score of 5000 Points
Master	Reach a score of 20000 Points
Powerman	Collect 100 Bonus Items (in total)
Regular Customer	Have a constant log-in streak of 7 days in a row
Way to go	Answer 100 questions in total
Autobiographer	Fill out own bibliography
Chatterbox (hidden)	Used the contact form 5 times

“iHEARu-PLAY: Introducing a game for crowdsourced data collection for affective computing”, WASA, 2015.

Holistic Modelling

- **Multiple-Targets**

  There is just one Vocal Production Mechanism…

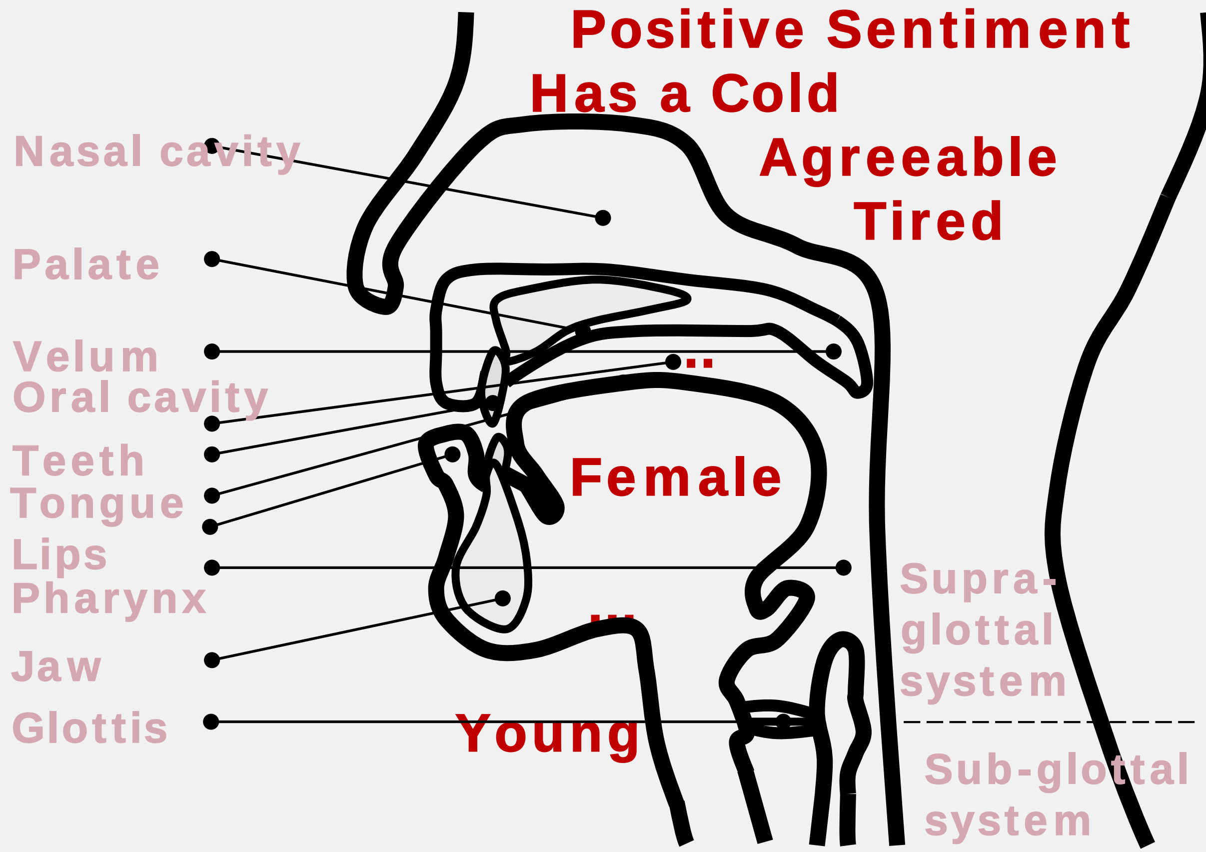

% UA	Single	Multiple
Likability	59.1	(+A,G,Cl) **62.2**
Neuroticism	62.9	(+G,OCEA, Cl) **67.5**

Cross-Task Labelling

- **Cross-Task Self-Labelling**

% UA	Likability	Emotion	Personality
Baseline	57.2	68.9	66.4
Cross-Task Labelling	60.3	69.0	66.6

Algorithm: *Cross-Task Labelling*
Repeat for each task:
Repeat until $\mathcal{U} \in \{\}$:

1. (Optional) Upsample training set $\mathcal{L}$ to even class distribution $\mathcal{L}_D$
2. Use $\mathcal{L}/\mathcal{L}_D$ to train classifier $\mathcal{H}$, then classify $\mathcal{U}$
3. Select a subset $\mathcal{N}_{st}$ that contains those instances predicted with the highest confidence values
4. Remove $\mathcal{N}_{st}$ from the unlabelled set $\mathcal{U}$, $\mathcal{U} = \mathcal{U} \setminus \mathcal{N}_{st}$
5. Add $\mathcal{N}_{st}$ to the labelled set $\mathcal{L}$, $\mathcal{L} = \mathcal{L} \cup \mathcal{N}_{st}$

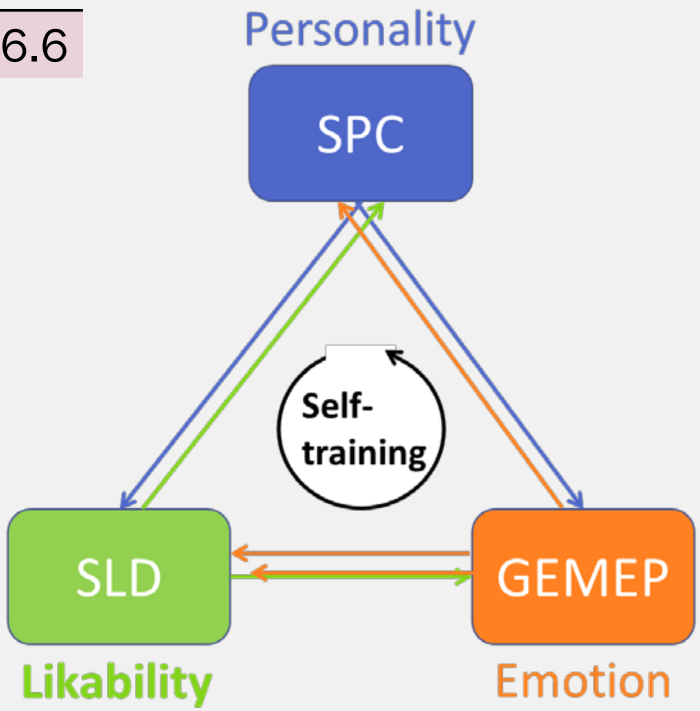

"Semi-Autonomous Data Enrichment Based on Cross-Task Labelling of Missing Targets for Holistic Speech Analysis", ICASSP, 2016.

Confidence

- **Confidence Measure**
  - Learning Recogniser Behaviour
  - Adaptation to target domain:
  - Semi-supervised learning

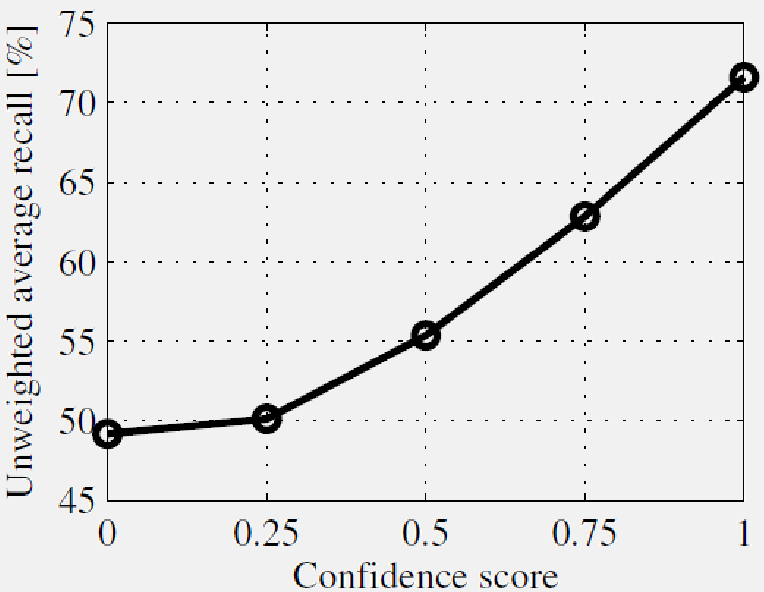

Example: ComParE:EC

	Predicted -	Predicted +
Actual -	CN	ICN
Actual +	ICP	CP

"Confidence Measures in SER Based on Semi-supervised Learning", ISCA Interspeech, 2012.

Next-Gen?

- Evolutionary Learning
- Reinforced Learning
- Analysis/Synthesis Gap

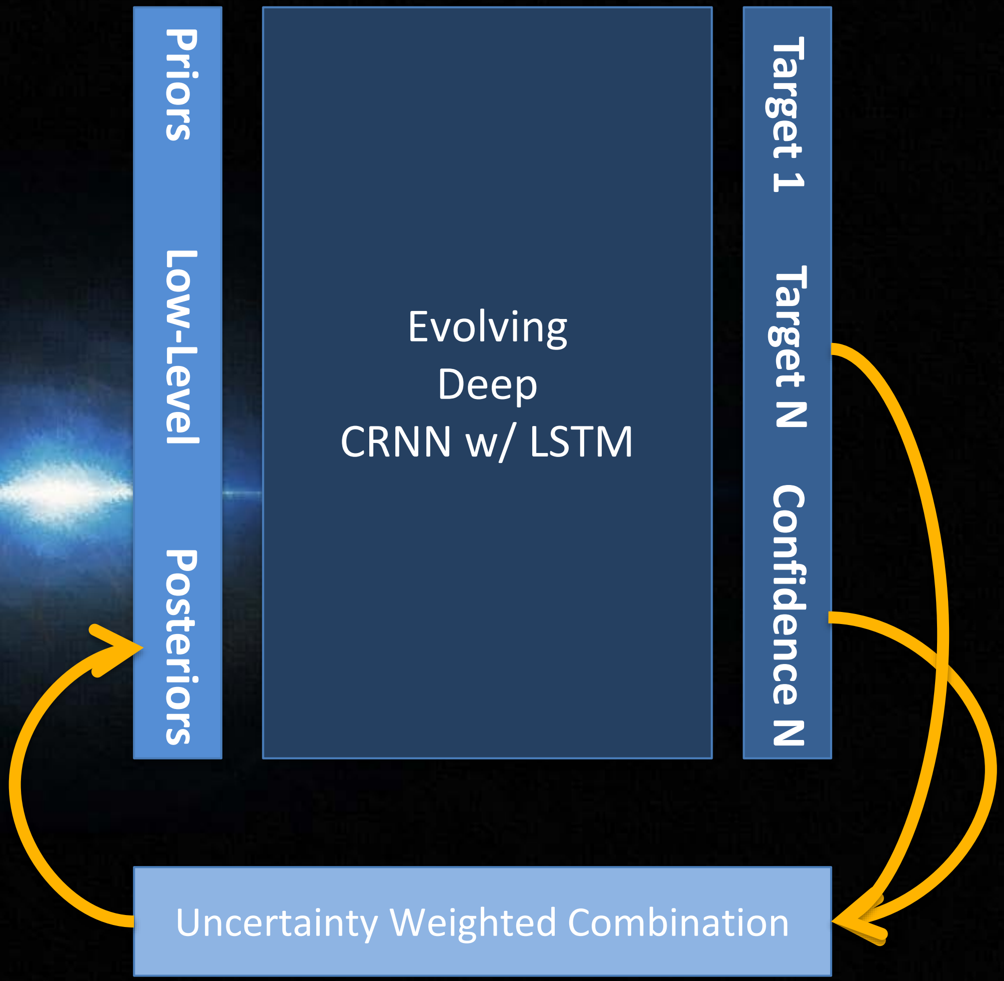

5. Emotion and affect detection

Language and emotion

- Certain languages put more emphasis on some emotions
  - Many more words to express the "same" emotion
  - No equivalence to emotions expressed in certain languages (saudade (PT), dor (RO), Schadenfreude (DE), etc.)

- The manner in which emotions are expressed in language conditions the way in which they are perceived

- Identifying and labeling through language the emotion felt can help to relieve it (Lieberman et al., 2007)

- Can emotions be translated?
  - Studies on the translation of the Bible

Models of emotion

- Two types:
  - Categorical – a certain number of limited emotion "categories" are defined
  - Dimensional – organized in affective dimensions
    - Valence-pleasantness + activity-arousal (Russell)
    - Semantic differentials (Osgood)
    - Three-dimensional model based on levels of presence of hormones (Lövheim)

Models of emotion

Categorical models of emotion:

Ekman (facial expressions):
- 6 basic emotions: joy, anger, fear, sadness, disgust, surprise

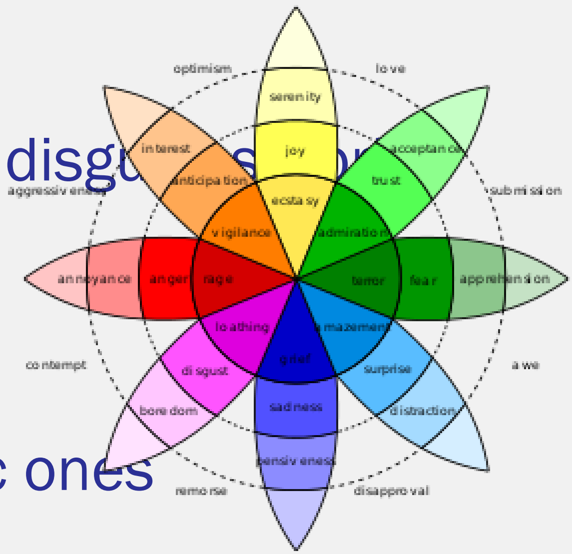

Plutchik's "Wheel of emotions":
- 8 basic emotions
- 8 derivative emotions, combination of basic ones

Shaver (1987)/Parrot (2001) Tree-structured list of emotions:
- 6 basic emotions (instead of disgust, love)
- Secondary and tertiary emotions

Models of emotion

- **"Secondary" Emotions**

  So far not in systems?

  Difficult to label?

  2nd emotion neutral…

  Difficult to classify

  Difficult to evaluate

Minor

Major

[%]	AMU	COL	DEC	ENE	INQ	IRO	JOI	NEG	NEU	PEU	POS	SAT	SED	STR	SUR	TRI
AMU			3	1	3	5	24	1	47		1	3	3	2	8	
COL			15	25	2	3			34					3	13	5
DEC		11		16	6	1			44					2	7	12
ENE	1	23	15		12	5			29		1	1		3	6	4
INQ	1	1	7	12					44			1	2	14	11	8
IRO			3	5					76			4	5		7	
JOI	9			1	1	2			21			4	2	1	58	
NEG			17						42					4	38	
NEU						5			77						18	
PEU				15					34					23	7	21
POS			12	5	17	3			47			1		6	9	
SAT	6		2	1	4	10	4	1	42				6	6	15	2
SED	15		1		15	11			33			18		4	1	1
STR	1	9	5	7	42				18	3		6	1		5	4
SUR	4		4	7	30	4	3		34	2		10		1		2
TRI	1	7	10	8	13				49	1		3	1	6	2	

Models of emotion

Primary emotion	Secondary emotion	Tertiary emotions
Love	Affection	Adoration, affection, love, fondness, liking, attraction, caring, tenderness, compassion, sentimentality
	Lust	Arousal, desire, lust, passion, infatuation
	Longing	Longing
Joy	Cheerfulness	Amusement, bliss, cheerfulness, gaiety, glee, jolliness, joviality, joy, delight, enjoyment, gladness, happiness, jubilation, elation, satisfaction, ecstasy, euphoria
	Zest	Enthusiasm, zeal, zest, excitement, thrill, exhilaration
	Contentment	Contentment, pleasure
	Pride	Pride, triumph
	Optimism	Eagerness, hope, optimism
	Enthrallment	Enthrallment, rapture
	Relief	Relief
Surprise	Surprise	Amazement, surprise, astonishment
Anger	Irritation	Aggravation, irritation, agitation, annoyance, grouchiness, grumpiness
	Exasperation	Exasperation, frustration
	Rage	Anger, rage, outrage, fury, wrath, hostility, ferocity, bitterness, hate, loathing, scorn, spite, vengefulness, dislike, resentment
	Disgust	Disgust, revulsion, contempt
	Envy	Envy, jealousy
	Torment	Torment
Sadness	Suffering	Agony, suffering, hurt, anguish
	Sadness	Depression, despair, hopelessness, gloom, glumness, sadness, unhappiness, grief, sorrow, woe, misery, melancholy
	Disappointment	Dismay, disappointment, displeasure
	Shame	Guilt, shame, regret, remorse
	Neglect	Alienation, isolation, neglect, loneliness, rejection, homesickness, defeat, dejection, insecurity, embarrassment, humiliation, insult
	Sympathy	Pity, sympathy
Fear	Horror	Alarm, shock, fear, fright, horror, terror, panic, hysteria, mortification
	Nervousness	Anxiety, nervousness, tenseness, uneasiness, apprehension, worry, distress, dread

Models of emotion

- Class Hierarchies
  - **Hierarchical**: from main categories (positive, neutral, negative), or the 6 emotions, to sub-categories.
  - **Fixed/scaled categories**: e.g., weak, medium, strong; pure vs. mixed, or antagonistic e.g., a mixture of anger and joy can be irony.

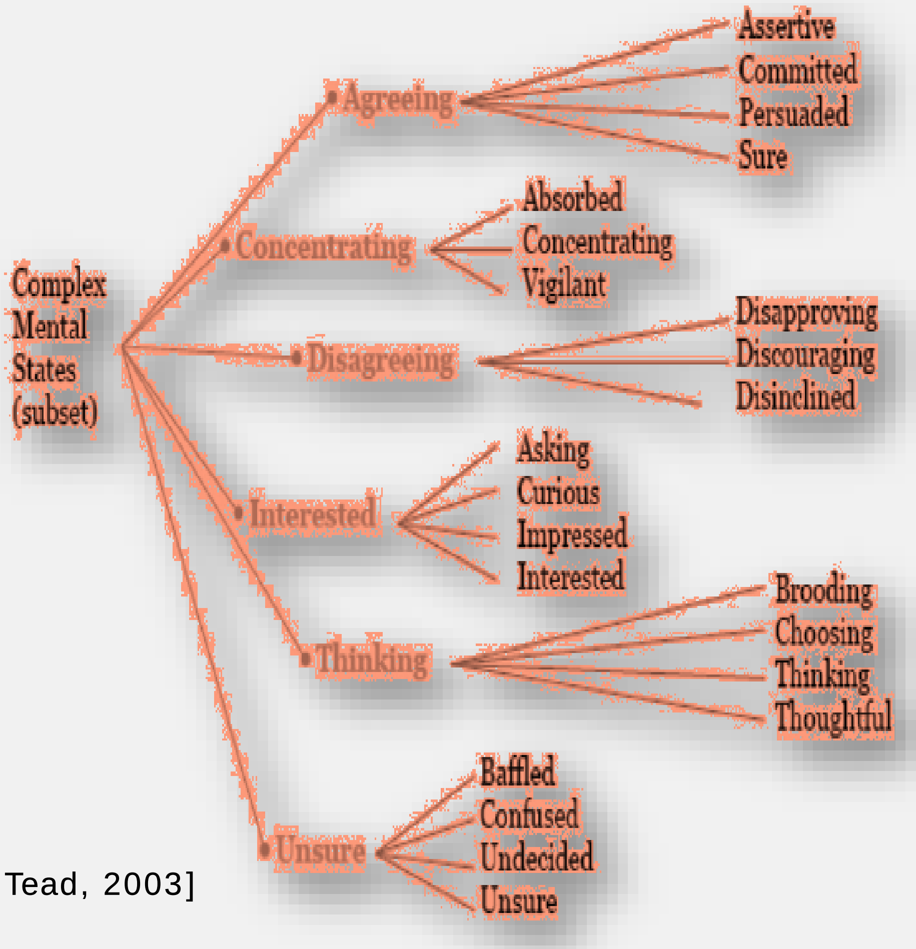

[Baron-Cohen & Tead, 2003]

Models of emotion

- **Soft Emotion Profiles**

  N labelers lead to profile

  Recognition with score

  Thresholding can be used

  Natural extension to detection

  Most difficult to evaluate?

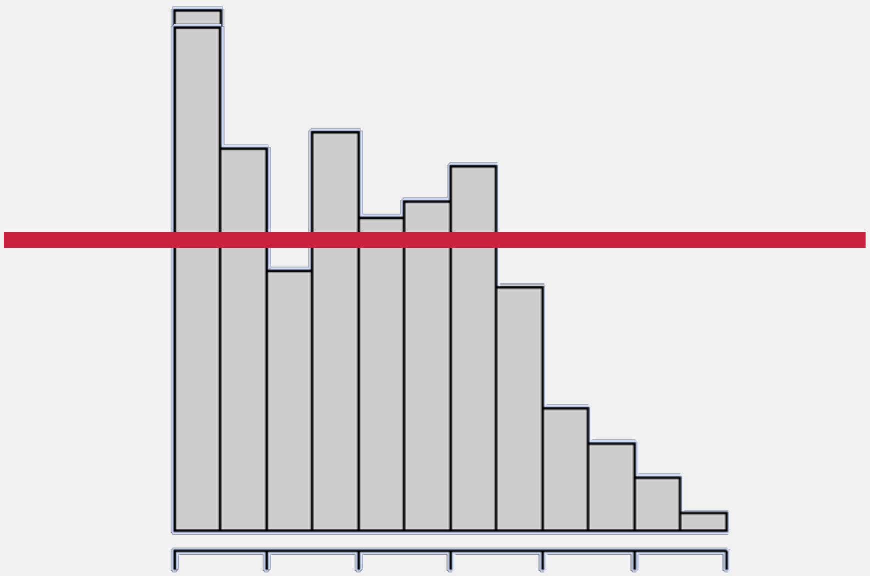

Models of emotion

Dimensional models of emotion:

- Valence/pleasantness + activity/arousal (Russell, 1980)

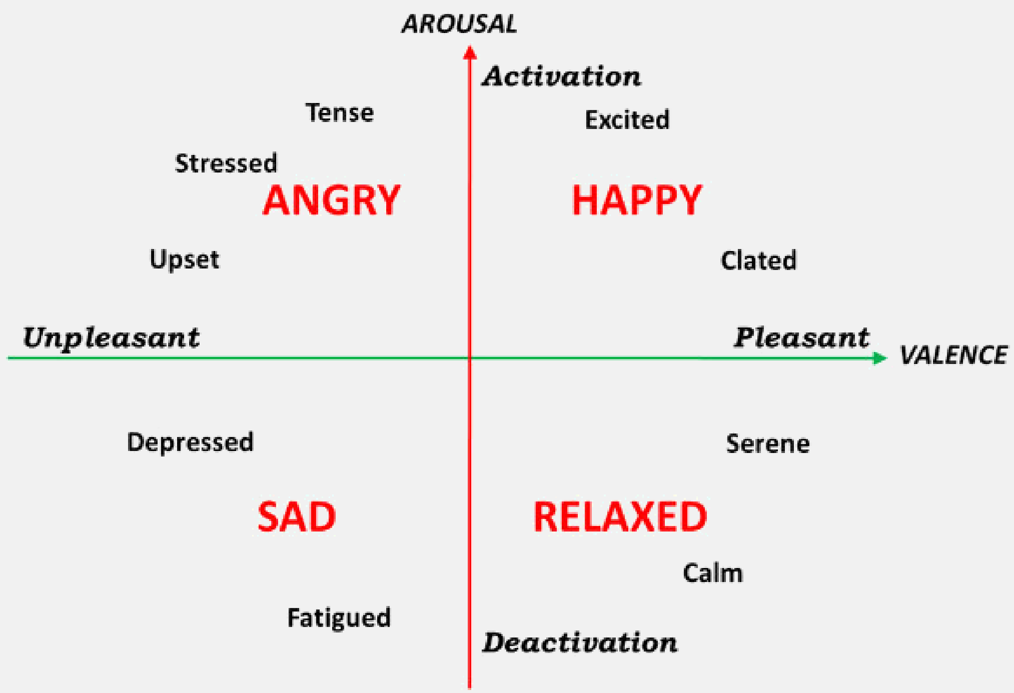

Models of emotion

Dimensional models of emotion:

- Semantic differentials (Osgood, 1957)

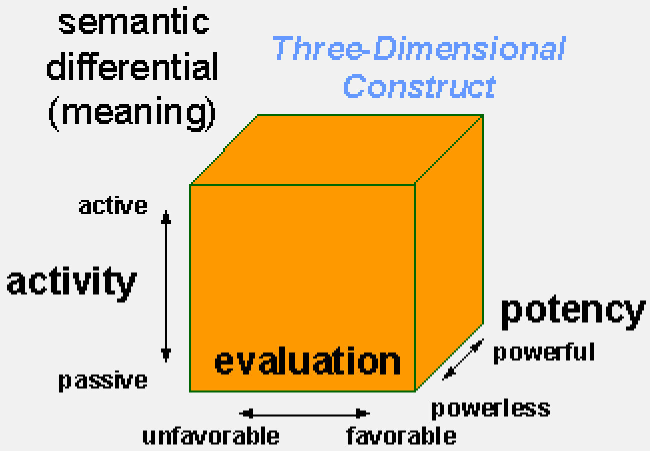

Models of emotion

- Lövheim's (2001) cube of emotion based on Tomkin's 8 basic emotions ("Affect theory"):

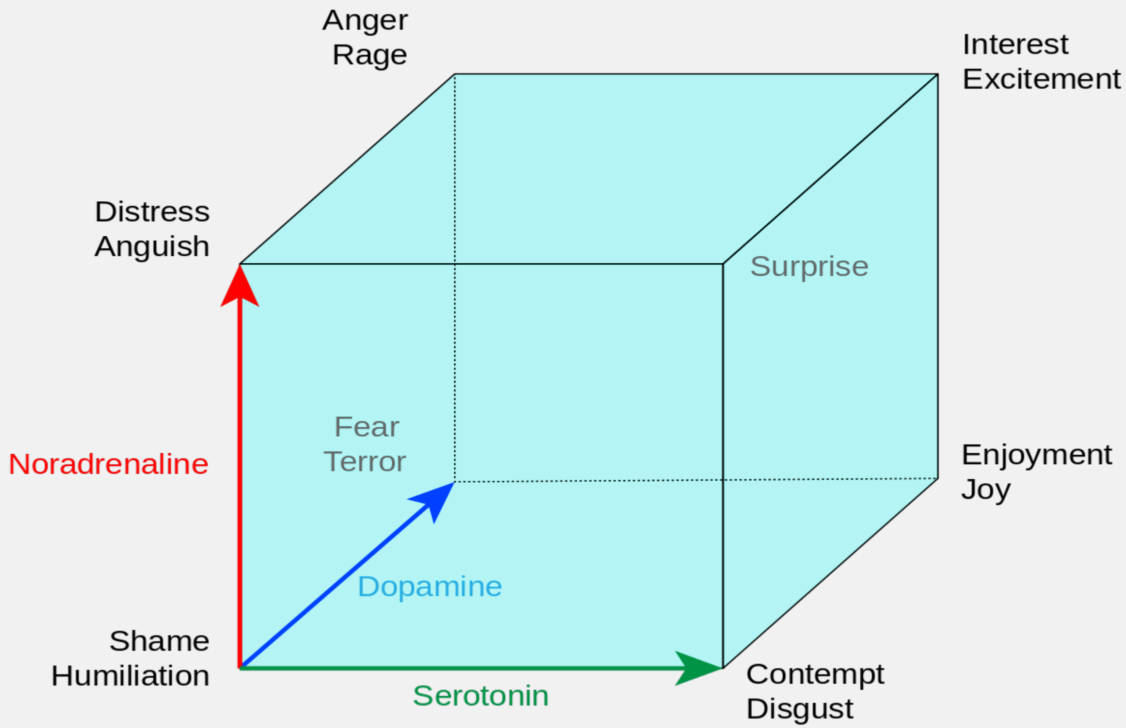

Time

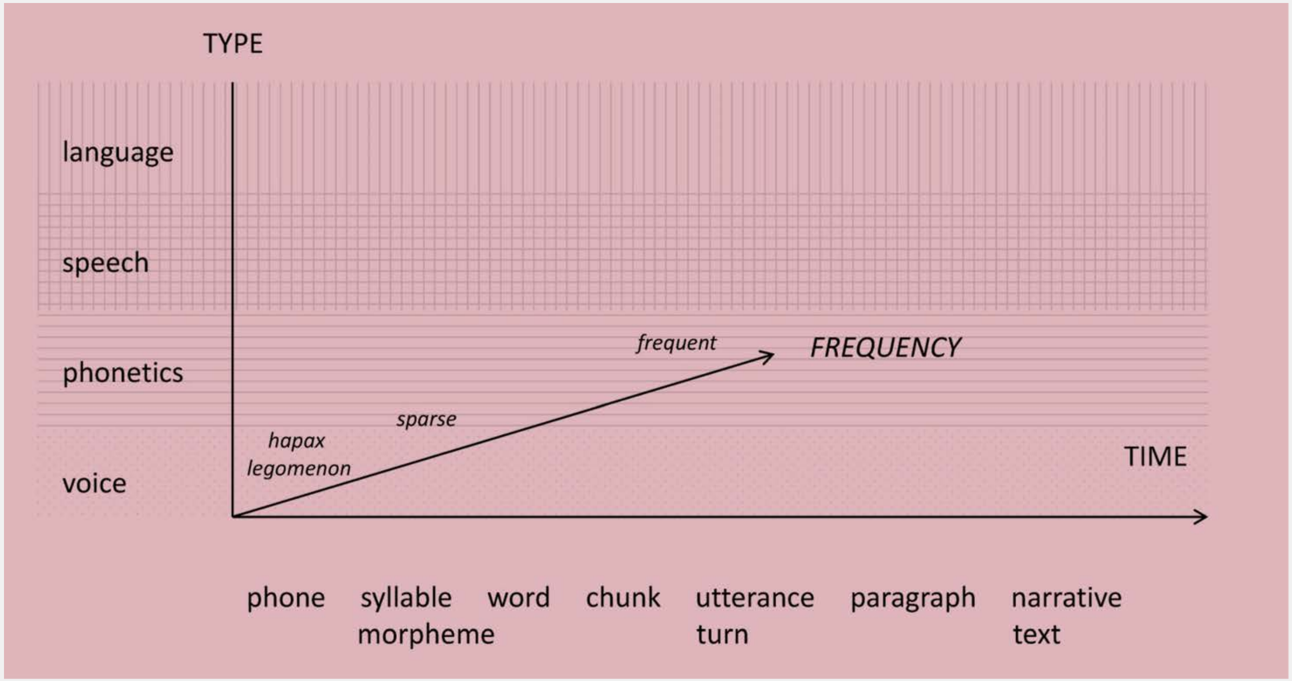

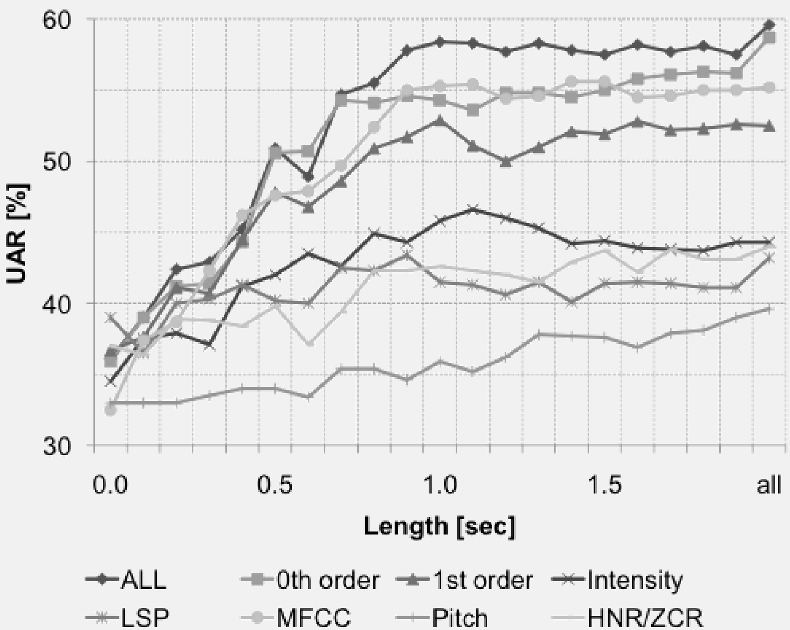

(a) LLD: unweighted average recall.

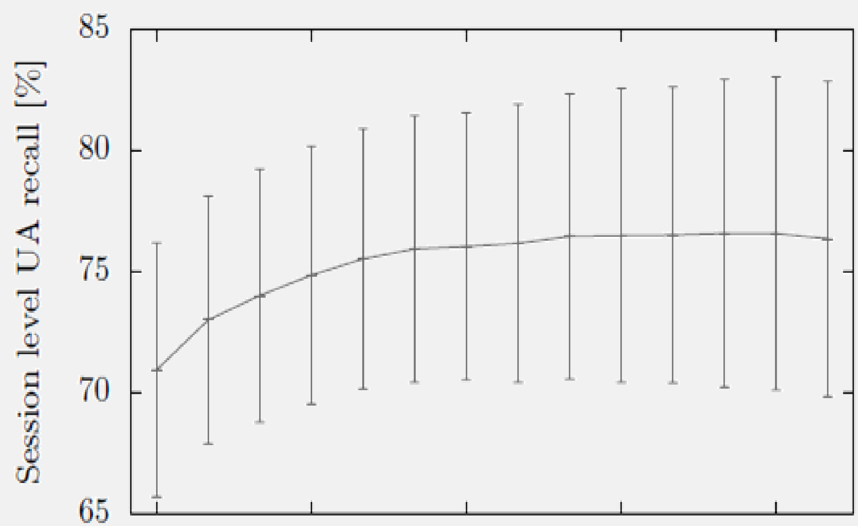

Standards for Emotion

- **Existing Markup Languages**

  EARL (HUMAINE Emotion Annotation and Representation Language)

  SMIL (W3C Synchronized Multimedia Integration Language)

  SSML (W3C Speech Synthesis ML)

  EMMA (W3C Extensible Multi-Modal Annotation ML)

  VHML (Virtual Human ML), APML, RPL, EmoTV Coding Scheme

  …
- **W3C Emotion Markup Language**

```
<emotion>
    <category set="everydayEmotions" name="fear"/>
        <intensity>
            <trace freq="10Hz" samples="0.1 0.1 0.15 0.2 …"/>
        </intensity>
</emotion>
```

Categories of emotion

Adapted from Gabrielsson (2002) – emotions in music

- Expressed emotion: emotion the performer tries to communicate to the (readers)
- Perceived emotion: emotion the reader perceives as being expressed
- Felt (evoked) emotion: emotion felt by the reader, in response to text

- And we can add: emotion directly present in the text

AVEC – The Emotion Challenge

- **AVEC 2011 & 2012: Emotion**
  - AVEC database: 50k Turns
- **AVEC 2013 & 2014: Depression**
  - AVDLC database: 240 h, ~300 subjs.

  Beck Depression Inventory: 0 – 63
- **AV+EC 2015: Emotion**
  - RECOLA database: 6 h, ~30 subjs.

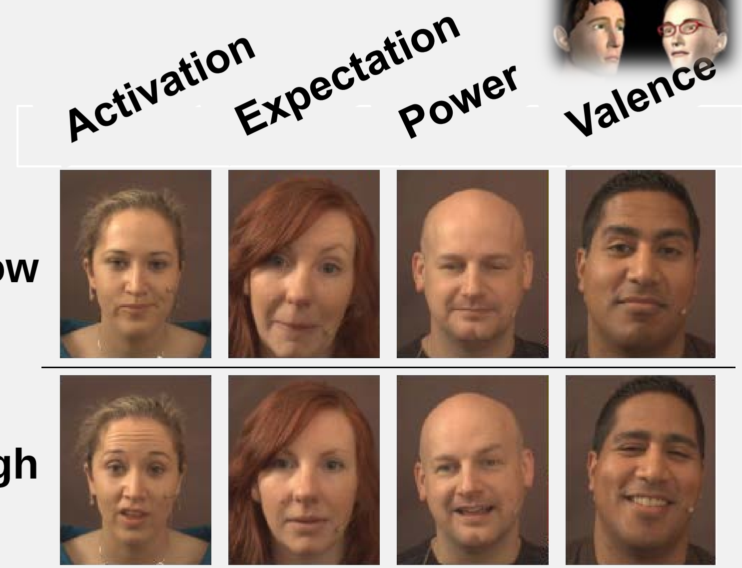

%Acc.	71.2	68.6	59.6	70.2
MAErr	6.3			
%CCC	67.8			

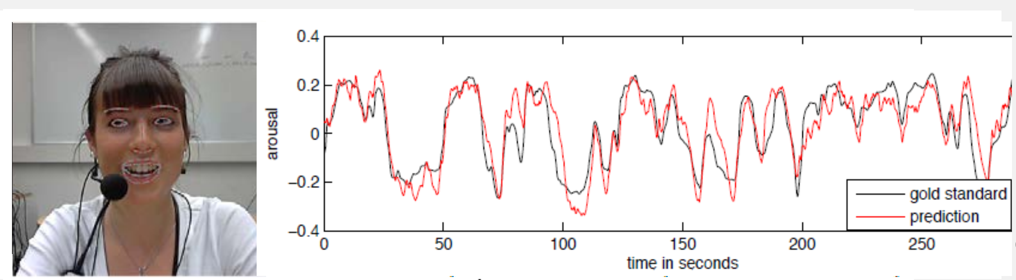

"AVEC 2014 - The Continuous Audio/Visual Emotion and Depression Recognition Challenge", ACM Multimedia, 2014.

"Prediction of Asynchronous Dimensional Emotion Ratings from Audiovisual and Physiological Data", Pattern Recognition Letters, 2014.

AVEC – The Emotion Challenge

- AVEC 2016

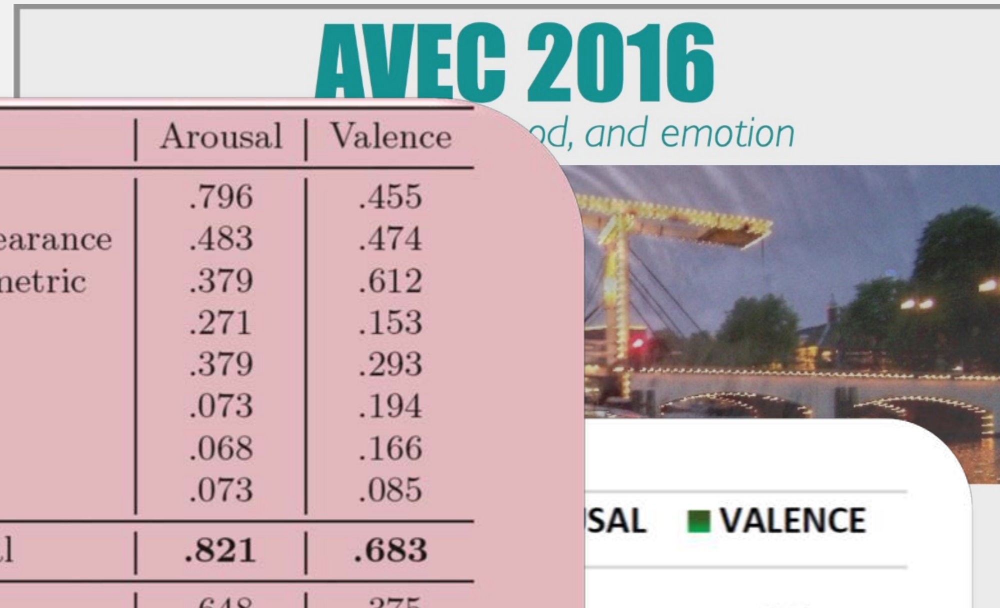

Modality	Arousal	Valence
D-Audio	.796	.455
D-Video-appearance	.483	.474
D-Video-geometric	.379	.612
D-ECG	.271	.153
D-HRHRV	.379	.293
D-EDA	.073	.194
D-SCL	.068	.166
D-SCR	.073	.085
D-Multimodal	**.821**	**.683**
T-Audio	.648	.375
T-Video-appearance	.343	.486
T-Video-geometric	.272	.507
T-ECG	.158	.121
T-HRHRV	.334	.198
T-EDA	.075	.228
T-SCL	.066	.216
T-SCR	.065	.145
T-Multimodal	**.683**	**.639**

"AVEC 2016 - VEC 2016 – Depression, Mood, and Emotion Recognition Workshop and Challenge", ACM Multimedia, 2016.

6. Conclusions

Evaluative Language Beyond Bags of Words: Linguistic Insights and Computational Applications

Computational Linguistics

Farah Benamara Zitoune*
IRIT-Université de Toulouse

Maite Taboada**
Simon Fraser University

Yannick Mathieu†
LLF-CNRS

Submission received: 30 October 2015; Revised version received: 19 February 2016; Accepted for publication: 8 June 2016

Announcement

- ACM Trans. on Multimedia Computing, Communications and Applications Special Section on Multimedia Computing and Applications of Socioaffective Behaviors in the Wild
- ACM Multimedia 6th Audio/Visual Emotion Challenge (AVEC 2016)
- INTERSPEECH 2016 Computational Paralinguistics Challenge (ComParE)
- Knowledge-Based Systems Special Issue on New Avenues in Knowledge Bases for Natural Language Processing
- Image & Vision Computing Special Issue on Multimodal Sentiment Analysis and Mining in the Wild

- W3C Linked Data Models for Emotion and Sentiment Analysis Community Group

Reading

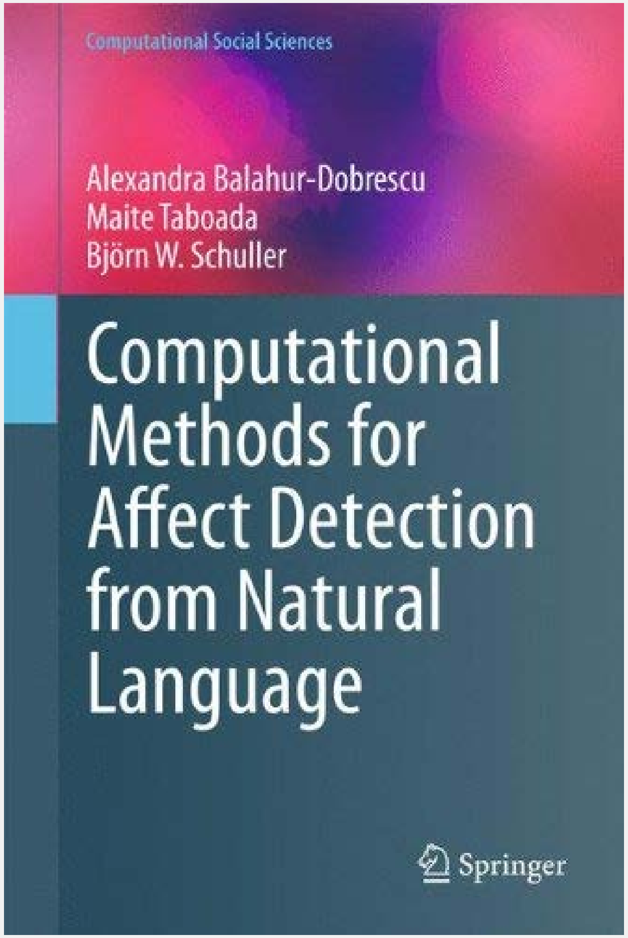
Computational Social Sciences
Alexandra Balahur-Dobrescu
Maite Taboada
Björn W. Schuller
Computational Methods for Affect Detection from Natural Language
Springer

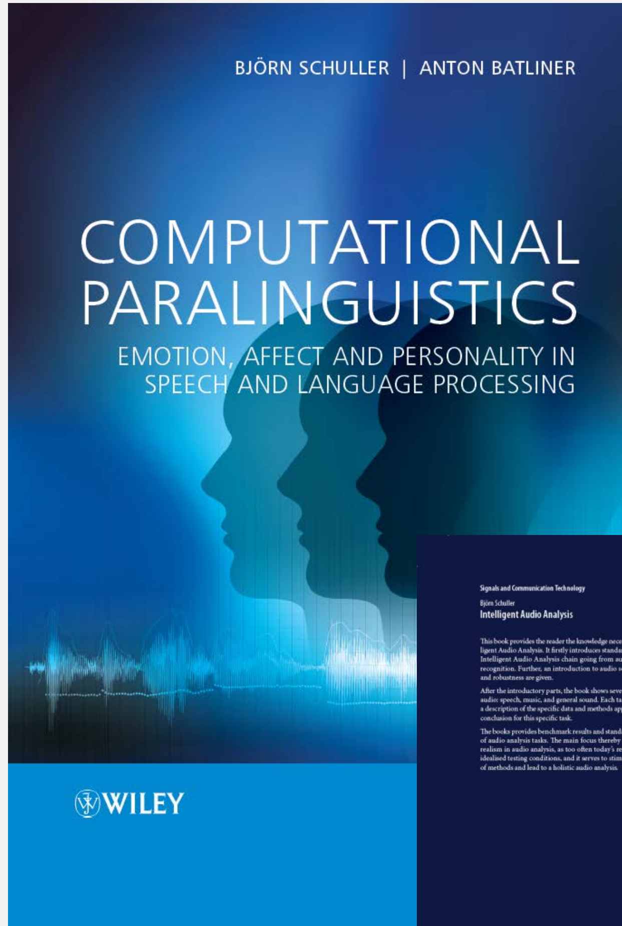
BJÖRN SCHULLER | ANTON BATLINER
COMPUTATIONAL PARALINGUISTICS
EMOTION, AFFECT AND PERSONALITY IN SPEECH AND LANGUAGE PROCESSING
WILEY

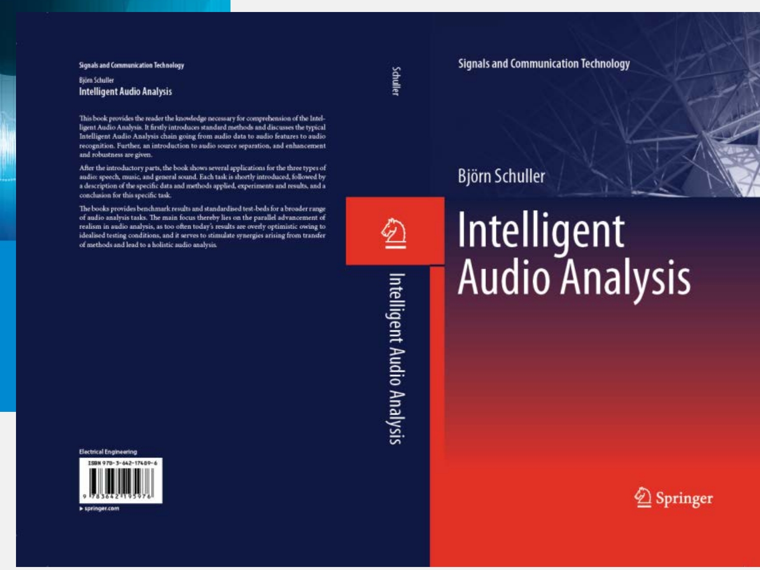
Signals and Communication Technology
Björn Schuller
Intelligent Audio Analysis
Springer
Intelligent Audio Analysis

References

Baccianella, Stefano, Andrea Esuli & Fabrizio Sebastiani. 2009. Multi-facet rating of product reviews. In Mohand Boughanem, Catherine Berrut, Josiane Mothe & Chantal Soule-Dupuy (eds.), *Advances in Information Retrieval: Proceedings of the European Conference on Information Retrieval* (pp. 461-472). Toulouse, France.

Baccianella, Stefano, Andrea Esuli & Fabrizio Sebastiani. 2010. SentiWordNet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. *Proceedings of the Seventh Conference on International Language Resources and Evaluation (LREC'10)* (pp. 2200-2204). Valletta, Malta.

Benamara, Farah, Maite Taboada & Yannick Mathieu. to appear. Evaluative language beyond bags of words: Linguistic insights and computational applications. *Computational Linguistics*.

Biber, Douglas. 2004. Historical patterns for the grammatical marking of stance: A cross-register comparison. *Journal of Historical Pragmatics* 5(1). 107-135.

Choi, Yejin & Claire Cardie. 2008. Learning with compositional semantics as structural inference for subsentential sentiment analysis. *Proceedings of the Conference on Empirical Methods in Natural Language Processing* (pp. 793-801). Honolulu, Hawaii.

Clough, Patricia Ticineto & Jean O'Malley Halley (eds.). 2007. *The Affective Turn: Theorizing the Social*. Durham, NC: Duke UP.

Davidson, Richard J., Klaus R. Scherer & H. Hill Goldsmith (eds.). 2003. *Handbook of Affective Sciences*. Oxford: Oxford University Press.

Esuli, Andrea. 2008. *Automatic generation of lexical resources for opinion mining: models, algorithms and applications.* Pisa, Italy: Università di PisaThesis Type.

Esuli, Andrea & Fabrizio Sebastiani. 2005. Determining the semantic orientation of terms through gloss classification. *Proceedings of the 14th ACM International Conference on Information and Knowledge Management*. Bremen, Germany.

Esuli, Andrea & Fabrizio Sebastiani. 2006a. Determining term subjectivity and term orientation for opinion mining. *Proceedings of EACL-06, the 11th Conference of the European Chapter of the Association for Computational Linguistics*. Trento, Italy.

Esuli, Andrea & Fabrizio Sebastiani. 2006b. SentiWordNet: A publicly available lexical resource for opinion mining. *Proceedings of 5th International Conference on Language Resources and Evaluation (LREC)* (pp. 417-422). Genoa, Italy.

Hatzivassiloglou, Vasileios & Kathleen McKeown. 1997. Predicting the semantic orientation of adjectives. *Proceedings of 35th Meeting of the Association for Computational Linguistics* (pp. 174-181). Madrid, Spain.

Hu, Minqing & Bing Liu. 2004. Mining and summarizing customer reviews. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-2004)* (pp. 168-177). Seattle, WA.

Kennedy, Alistair & Diana Inkpen. 2006. Sentiment classification of movie and product reviews using contextual valence shifters. *Computational Intelligence* 22(2). 110-125.

Kim, Soo-Min & Eduard Hovy. 2004. Determining the sentiment of opinions. *Proceedings of COLING 2004* (pp. 1367-1373). Geneva, Switzerland.

Kiritchenko, Svetlana & Saif Mohammad. 2016. Sentiment composition of words with opposing polarities. *Proceedings of the 7th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis (WASSA)*. San Diego, CA.

Kiritchenko, Svetlana, Xiaodan Zhu & Saif Mohammad. 2014. Sentiment analysis of short informal texts. *Journal of Artificial Intelligence Research* 50. 723-762.

Langacker, Ronald W. 1990. Subjectification. *Cognitive Linguistics* 1(1). 5-38.

Liu, Bing. 2012. *Sentiment Analysis and Opinion Mining*: Morgan & Claypool.

Liu, Bing. 2015. *Sentiment Analysis: Mining Opinions, Sentiments, and Emotions*. Cambridge: Cambridge University Press.

Lyons, John. 1994. Subjecthood and subjectivity. In Marina Yaguello (ed.), *Subjecthhod and Subjectivity: The Status of the Subject in Linguistic Theory*. Paris: Ophrys, Institut français du Royaume-Uni.

Mohammad, Saif. 2015. Sentiment analysis: Detecting valence, emotions, and other affectual states from text. In *Emotion Measurement*.

Mohammad, Saif. 2016. Sentiment analysis: Detecting valence, emotions, and other affectual states from text. In Herbert L. Meiselman (ed.), *Emotion Measurement*. Duxford: Woodhead.

Mohammad, Saif, Bonnie Dorr & Cody Dunne. 2009. Generating high-coverage semantic orientation lexicons from overtly marked words and a thesaurus. *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2009)* (pp. 599-608). Singapore.

Mohammad, Saif, Svetlana Kiritchenko & Xiaodan Zhu. 2013. NRC-Canada: Building the state-of-the-art in sentiment analysis of tweets. *Proceedings of the seventh international workshop on Semantic Evaluation Exercises (SemEval-2013)*. Atlanta, GA.

Mohammad, Saif & Peter Turney. 2010. Emotions evoked by common words and phrases: Using Mechanical Turk to create an emotion lexicon. *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text* (pp. 26-34). Los Angeles, CA.

Moilanen, Karo & Stephen Pulman. 2007. Sentiment composition. *Proceedings of Recent Advances in Natural Language Processing* (pp. 27-29). Borovets, Bulgaria.

Ortony, Andrew, Gerald L. Clore & Allan Collins. 1988. *The Cognitive Structure of Emotions*. Cambridge: Cambridge University Press.

Pang, Bo & Lillian Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. *Proceedings of 42nd Meeting of the Association for Computational Linguistics* (pp. 271-278). Barcelona, Spain.

Pang, Bo & Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. *Proceedings of ACL 2005* (pp. 115-124). Ann Arbor, MI.
Pang, Bo & Lillian Lee. 2008. Opinion Mining and Sentiment Analysis. *Foundations and Trends in Information Retrieval* 2(1-2). 1-135.
Pang, Bo, Lillian Lee & Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment classification using Machine Learning techniques. *Proceedings of Conference on Empirical Methods in NLP* (pp. 79-86). Philadelphia, PA.
Picard, Rosemary W. 1997. *Affective Computing*. Cambridge, MA: MIT Press.
Polanyi, Livia & Annie Zaenen. 2004. Contextual valence shifters. *Proceedings of AAAI Spring Symposium on Exploring Attitude and Affect in Text (AAAI Technical Report SS-04-07)* (pp. 106-111). Stanford University, CA.
Polanyi, Livia & Annie Zaenen. 2006. Contextual valence shifters. In James G. Shanahan, Yan Qu & Janyce Wiebe (eds.), *Computing Attitude and Affect in Text: Theory and Applications* (pp. 1-10). Dordrecht: Springer.
Scherer, Klaus R., Tanja Banziger & Etienne Roesch. 2010. *A Blueprint for Affective Computing*. Oxford: Oxford University Press.
Somasundaran, Swapna & Janyce Wiebe. 2009. Recognizing stances in online debates. *Proceedings of the 47th Meeting of the Association for Computational Linguistics* (pp. 226-234). Singapore.
Stab, Christian & Iryna Gurevych. 2014. Identifying argumentative discourse structures in persuasive essays. *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*. Doha, Qatar.
Stab, Christian & Iryna Gurevych. 2016. Parsing argumentation structures in persuasive essays. *Computational Linguistics*. in press.
Strappavara, Carlo & Oliviero Stock. 2004. WordNet-Affect: An affective extension of WordNet. *Proceedings of 4th International Conference on Language Resources and Evaluation (LREC)* (pp. 1083-1086). Lisbon, Portugal.
Taboada, Maite. 2016. Sentiment analysis: An overview from linguistics. *Annual Review of Linguistics* 2. 325-347.
Taboada, Maite, Julian Brooke, Milan Tofiloski, Kimberly Voll & Manfred Stede. 2011. Lexicon-based methods for sentiment analysis. *Computational Linguistics* 37(2). 267-307.
Thomas, Matt, Bo Pang & Lillian Lee. 2006. Get out the vote: Determining support or opposition from Congressional floor-debate transcripts. *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing* (pp. 327-335). Sydney, Australia.
Turney, Peter. 2002. Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. *Proceedings of 40th Meeting of the Association for Computational Linguistics* (pp. 417-424). Philadelphia, PA.
Turney, Peter & Michael Littman. 2002. *Unsupervised learning of semantic orientation from a hundred-billion-word corpus* (No. ERB-1094, NRC 44929): National Research Council of Canada.
Vis, Kirsten, José Sanders & Wilbert Spooren. 2012. Diachronic changes in subjectivity and stance–A corpus linguistic study of Dutch news texts. *Discourse, Context and Media* 1(2-3). 95-102.

Whitelaw, Casey, Jon Patrick & Maria Herke-Couchman. 2006. Identifying interpersonal distance using systemic features. In James G. Shanahan, Yan Qu & Janyce Wiebe (eds.), *Computing Attitude and Affect in Text: Theory and Applications* (pp. 199-214). Dordrecht: Springer.

Wiebe, Janyce, Eric Breck, Chris Buckley, Claire Cardie, Paul Davis, Bruce Fraser, Diane J. Litman, David R. Pierce, Ellen Riloff, Theresa Wilson, David Day & Mark Maybury. 2003. Recognizing and organizing opinions expressed in the world press. *Working Notes of the AAAI Spring Symposium in New Directions in Question Answering* (pp. 12-19). Stanford, CA.

Wiegand, Michael, Alexandra Balahur, Benjamin Roth, Dietrich Klakow & Andrés Montoyo. 2010. A survey on the role of negation in sentiment analysis. *Proceedings of the Workshop on Negation and Speculation in Natural Language Processing* (pp. 60-68). Uppsala, Sweden.

Wilson, Theresa, Janyce Wiebe & Paul Hoffmann. 2009. Recognizing contextual polarity: An exploration of features for phrase-level sentiment analysis. *Computational Linguistics* 35(3). 399-433.

Wilson, Theresa, Janyce Wiebe & Rebecca Hwa. 2006. Recognizing strong and weak opinion clauses. *Computational Intelligence* 22(2). 73-99.