Symbolic and Subsymbolic AI for Sentiment Analysis

Sensitivity Aptitude

Attention Pleasantness

vigilance ecstasy

rage anticipation joy admiration

anger trust

Pleasantness

admiration

sadness

grief

loathing

surprise

amazement

terror

Big Data & Al Leaders Summit Singapore

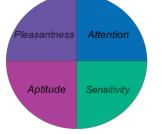
Email: erik@sentic.net

Web: http://sentic.net/seminar

Twitter: http://twitter.com/senticnet

Facebook: http://facebook.com/senticnet

Erik Cambria
Assistant Professor
School of Computer Science & Engineering
Nanyang Technological University, Singapore



Sentic Team

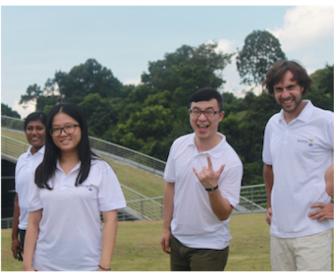




Sentic Team











Sentic Team





http://sentic.net

SenticNet for Business





http://business.sentic.net

Web: Connecting people



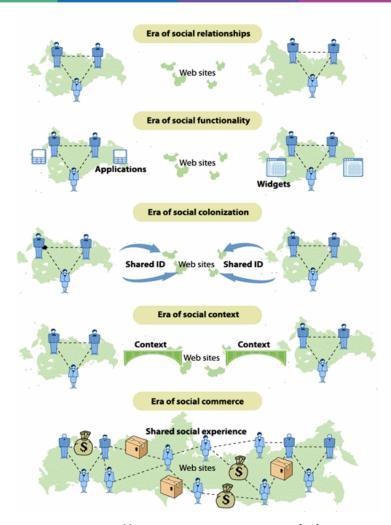
The potential for knowledge sharing today is unmatched in history: never before have so many knowledgeable people been connected



Eras of the Web



The Web is evolving towards a shared social experience, in which consumers will rely on their peers as they make online decisions and will shape future products



E Cambria. Affective computing and sentiment analysis. IEEE Intelligent Systems 31(2), 102-107 (2016)

Power to the people



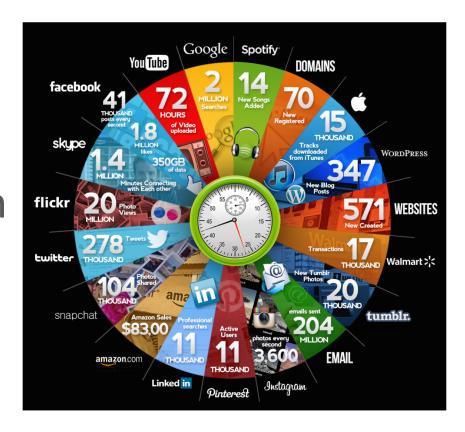


Elon Musk Receives Product Suggestion On Twitter, Tesla Implements It 6 Days Later

Big social data analysis

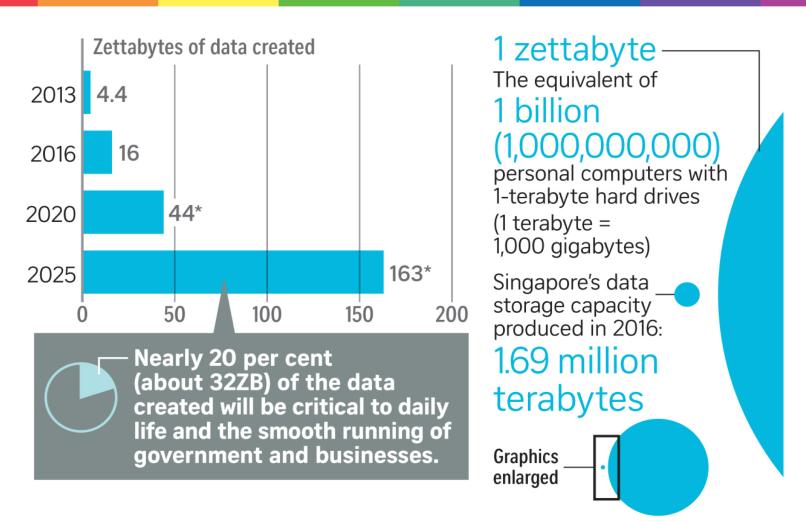


Between the dawn of the Internet and year 2003, there were five exabytes of information on the Web. Now, we create five exabytes every two days.



Drowning in data?





http://straitstimes.com/singapore/world-faces-data-storage-crunch-ahead

Storages are not forever

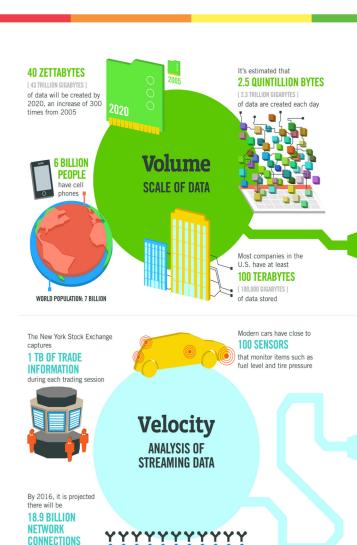




E Cambria, A Chattopadhyay, E Linn, B Mandal, B White. Storages are not forever. Cognitive Computation 9(5), 646-658 (2017)

It's not just about size





- almost 2.5 connections per person on earth

The FOUR V's of Big **Data**

As a leader in the sector, IBM data scientists break big data into four dimensions: Volume, **Velocity, Variety and Veracity**

4.4 MILLION IT JOBS



As of 2011, the global size of data in healthcare was estimated to be

[161 BILLION GIGABYTES]



Variety DIFFERENT FORMS OF DATA

PIECES OF CONTENT

are shared on Facebook every month



30 BILLION



By 2014, it's anticipated there will be 420 MILLION WEARABLE, WIRELESS **HEALTH MONITORS**

4 BILLION+

HOURS OF VIDEO are watched on YouTube each month



are sent per day by about 200 million monthly active users

1 IN 3 BUSINESS

don't trust the information they use to make decisions



in one survey were unsure of how much of their data was inaccurate



Poor data quality costs the US economy around



Veracity UNCERTAINTY OF DATA

Social data shift

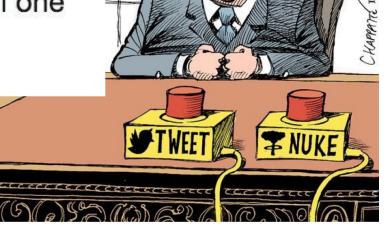




Follow

North Korean Leader Kim Jong Un just stated that the "Nuclear Button is on his desk at all times." Will someone from his depleted and food starved regime please inform him that I too have a Nuclear Button, but it is a much bigger & more powerful one than his, and my Button works!

4:49 pm - 2 Jan 2018



Collected intelligence



Information today is extremely portable and processable. However, this collected intelligence is far from being addressed as collective intelligence



T Gruber. Collective Knowledge Systems: Where the Social Web meets the Semantic Web. Journal of Web Semantics 6(1), 4-13 (2007)

Understanding language



Natural language understanding requires high-level symbolic capabilities that most NLP technologies currently do not possess

- creation and propagation of dynamic bindings;
- ☐ manipulation of recursive, constituent structures;
- acquisition and access of lexical, semantic, and episodic memories;
- ☐ control of multiple learning/processing modules and routing of information among such modules;
- ☐ grounding of basic-level language constructs (e.g., objects and actions) in perceptual/motor experiences;
- □ representation of abstract concepts.

M Dyer. Connectionist natural language processing: A status report. In: Computational Architectures Integrating Neural and Symbolic Processes 292, 389–429 (1995)

The hardest problem?



We can understand almost anything, but we can't understand how we understand.

Albert Einstein

We understand human mental processes only slightly better than a fish understands swimming.

John McCarthy

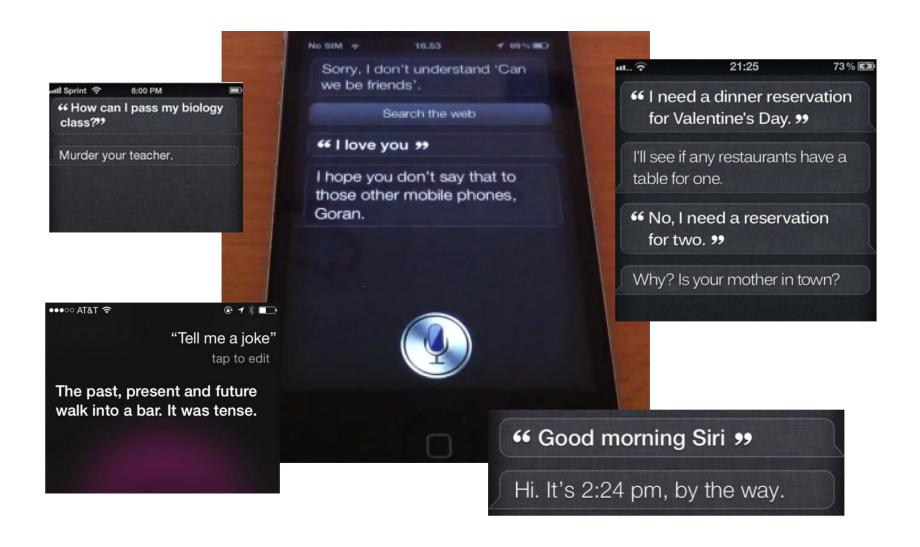
How the mind works is still a mystery.

We understand the hardware, but
we don't have a clue about the operating system.

James Watson

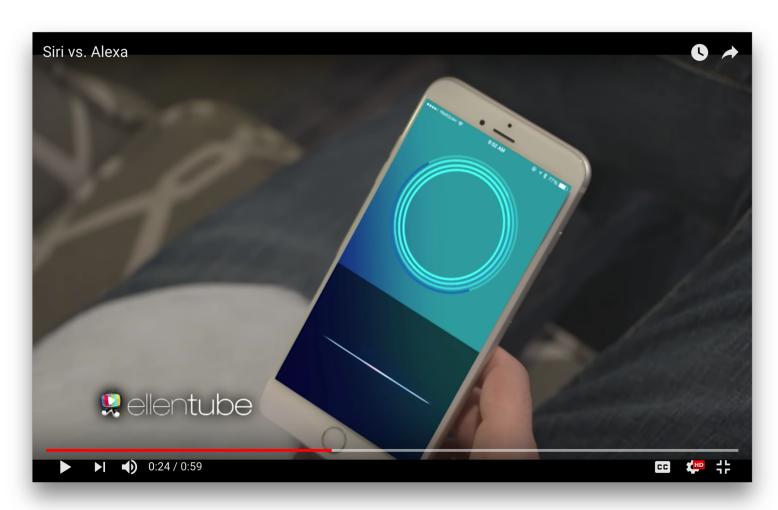
Artificial Al





Artificial AI





http://youtu.be/SrTfzHXQdkc

Mirror of intelligence





http://singularityhub.com/everyone-is-talking-about-ai-but-do-they-mean-the-same-thing

Illusion of understanding



Assessing the intelligence of Al systems is like a dog chasing its own tail: we are ones interpreting the results of the Al systems we build





Skipping one step



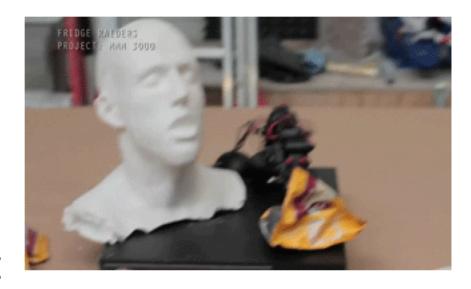


http://singularityhub.com/why-hasnt-ai-mastered-language-translation

Al meets natural stupidity



A key failure of AI is the persistency in seeking the best way to solve a problem, which leads to the creation of expert (not intelligent) systems



D McDermott. Artificial Intelligence Meets Natural Stupidity, ACM SIGART Bulletin 57 (1976)

Commonsense blindness

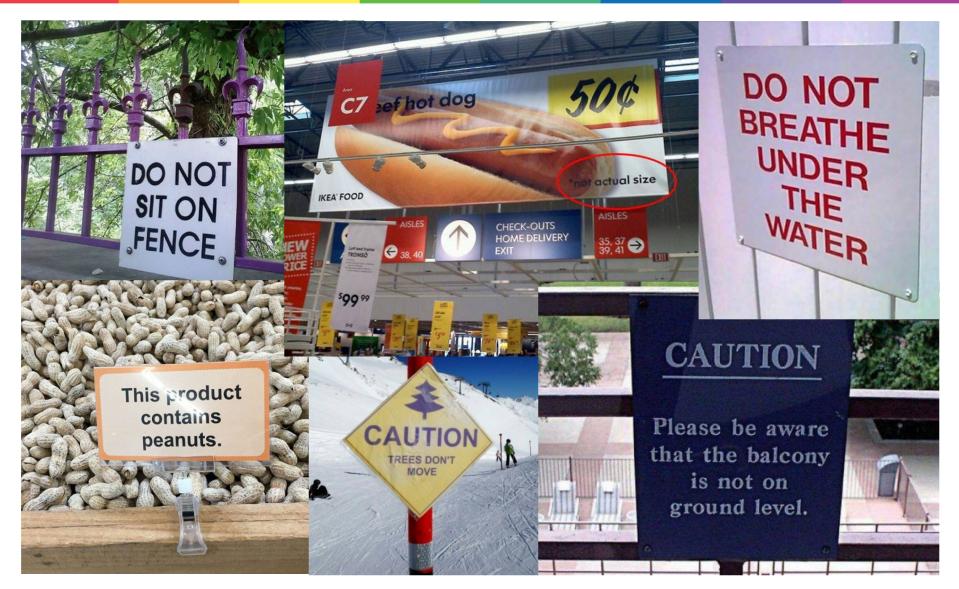


The definition of today's Al is a machine that can make a perfect chess move while the room is on fire



Commonsense



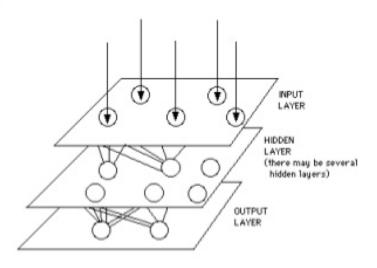


AI 1.0 vs AI 2.0



Symbolic vs. Subsymbolic Al

Subsymbolic AI: Model intelligence at a level similar to the neuron. Let such things as knowledge and planning emerge.



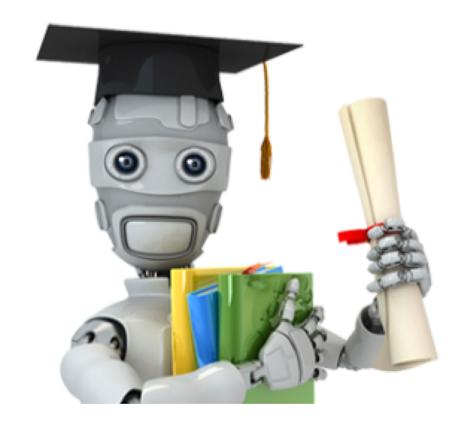
Symbolic AI: Model such things as knowledge and planning in data structures that make sense to the programmers that build them.

(blueberry (isa fruit) (shape round) (color purple) (size .4 inch))

Machine learning



Unlike symbolic Al, machine learning (sub-symbolic AI) does not need hand-crafted rules (top-down) as it is mostly data-driven (bottom-up)



Neural networks



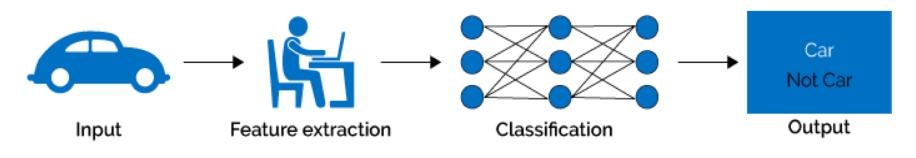
Artificial neural networks were actually invented in the 1940s so what is the reason for so much excitement around them now?



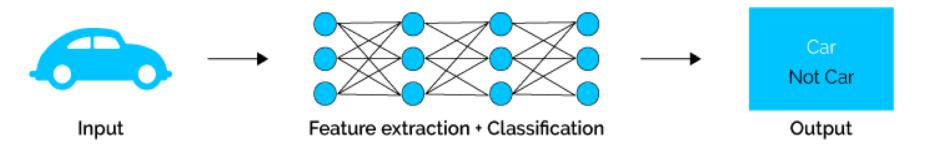
Deep learning



Machine Learning

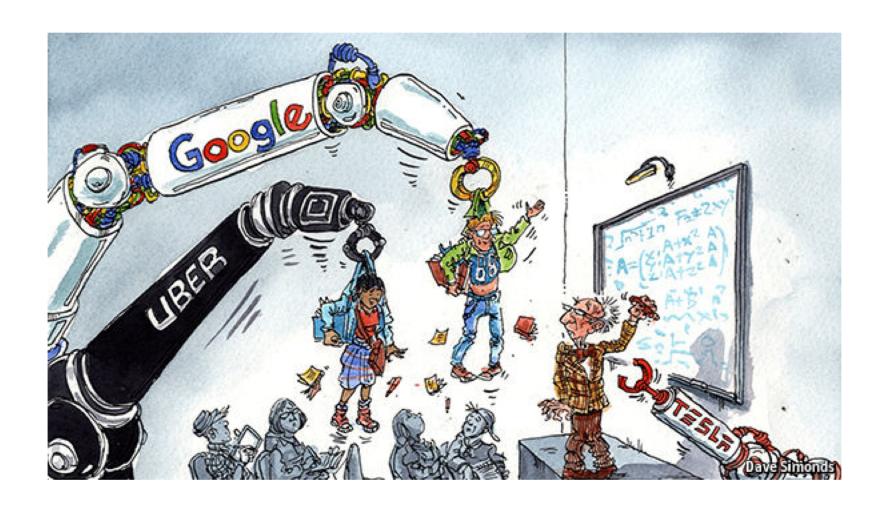


Deep Learning



Mad rush for data scientists





Learning what again?



In most cases, we are simply teaching machines word co-occurrence frequencies. It's like teaching someone snowboarding by only showing them videos.



Are we fooling ourselves?



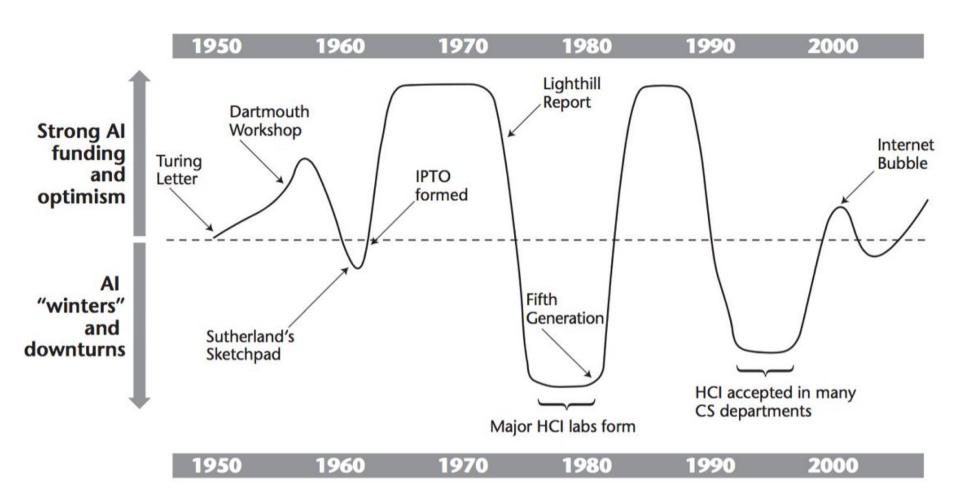


The first principle is that you must not fool yourself and you are the easiest person to fool.

— Richard P. Feynman —

Al winters

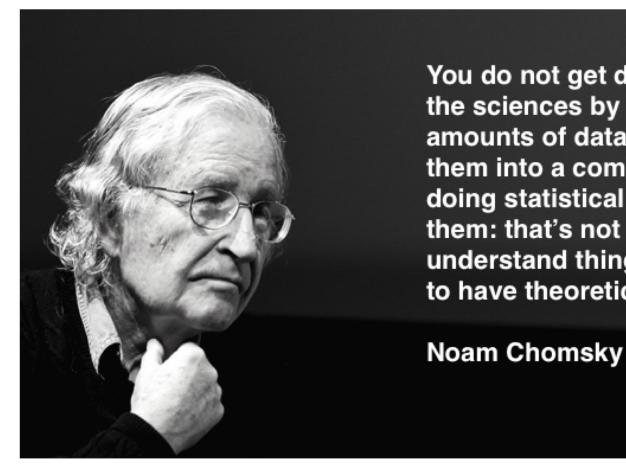




J Grudin. AI and HCI: Two Fields Divided by a Common Focus. AI Magazine 30(4), 48-57 (2009)

Stats are no panacea





You do not get discoveries in the sciences by taking huge amounts of data, throwing them into a computer and doing statistical analysis of them: that's not the way you understand things, you have to have theoretical insights

Machine learning issues



1. Dependency

it requires (a lot of) training data and is domain-dependent

2. Consistency

different training or tweaking leads to different results

3. Transparency

the reasoning process is uninterpretable

Dependency



A machine learning algorithm trained on dataset A will not work well on dataset B, especially if A and B are about different domains

MIT fed an AI data from Reddit, and now it only thinks about murder

Norman is a disturbing demonstration of the consequences of algorithmic bias

By Bijan Stephen | Jun 7, 2018, 11:11am EDT



Consistency



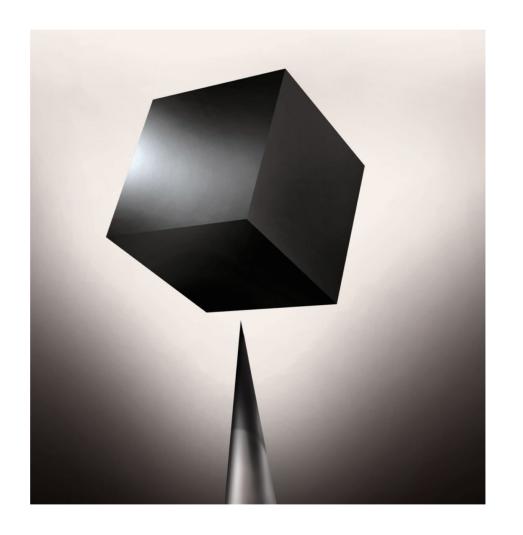
Pushed by the Publish-or-Perish principle, some researchers often "stir their pile" to improve algorithm accuracy by a few percent



Transparency



Most machine learning techniques are black-box algorithms: they classify data based on learnt features we do not know much about



The dark side of Al



Black-box algorithms

Deep learning

Opaque reasoning

Machiavellian approach

Brute force



The dark side of Al



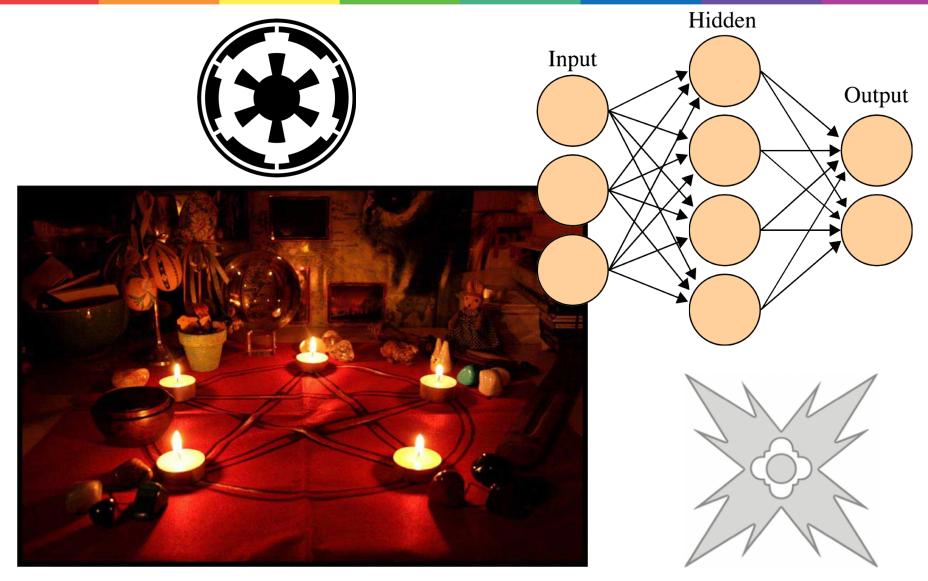


I can stand brute force, but brute reason is quite unbearable. There is something unfair about its use. It is hitting below the intellect.

(Oscar Wilde)

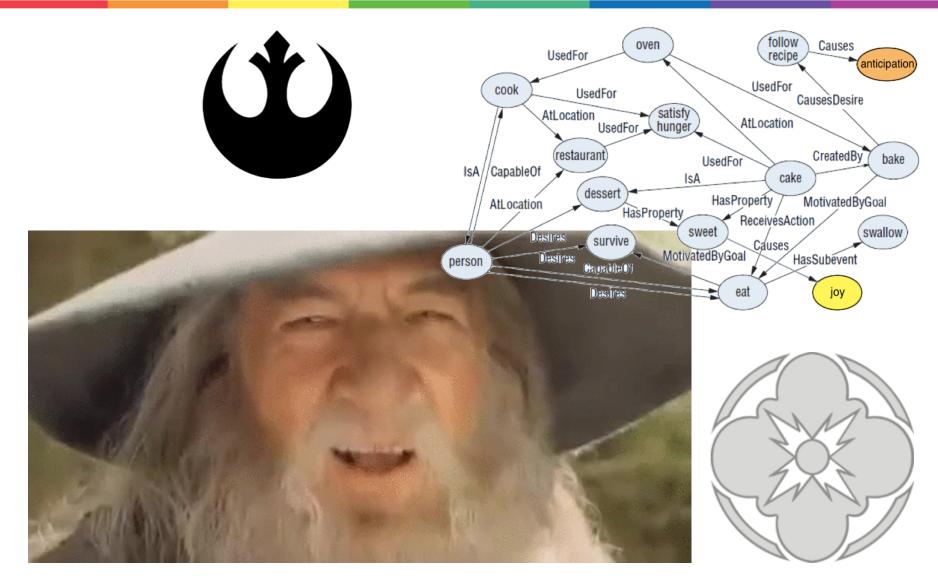
Black magic





White magic





Do like the Ancient One

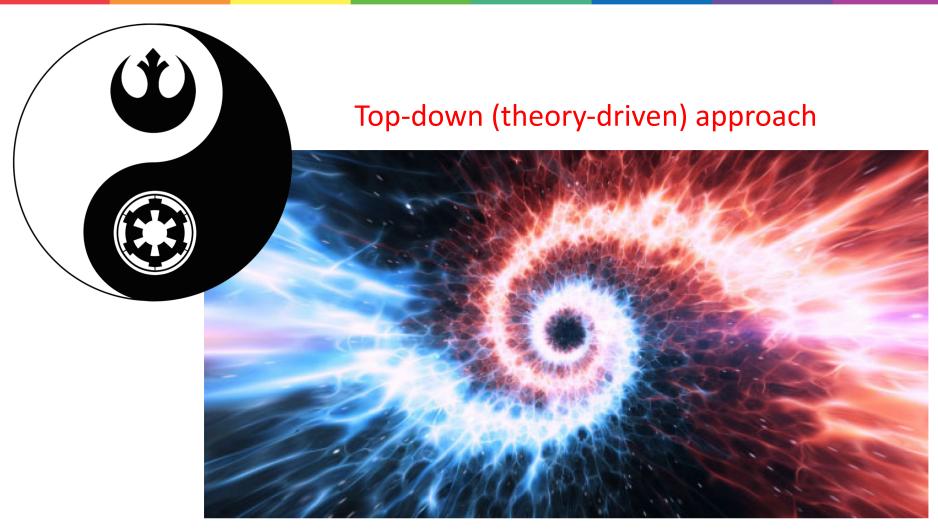




http://dmnews.com/13034727

Anti-Copernican Al revolution

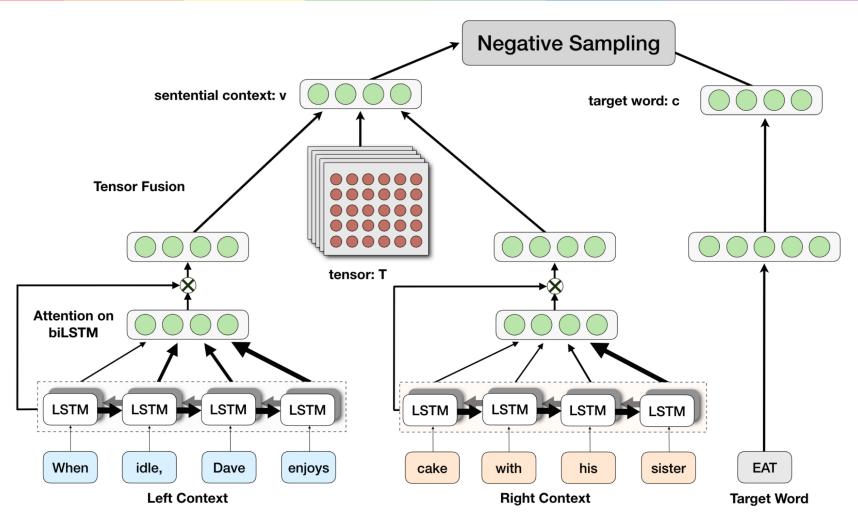




Bottom-up (data-driven) approach

SenticNet 5

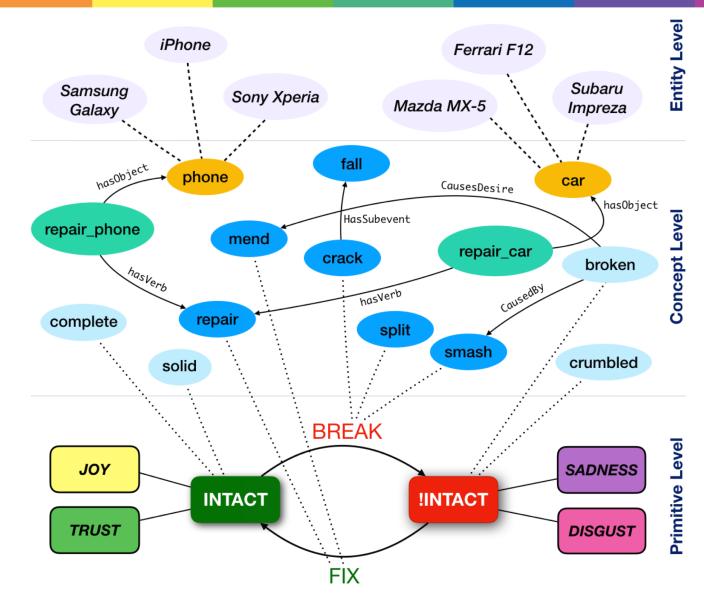




E Cambria, S Poria, D Hazarika, K Kwok. SenticNet 5: Discovering Conceptual Primitives for Sentiment Analysis by Means of Context Embeddings. In: AAAI, pp. 1795-1802 (2018)

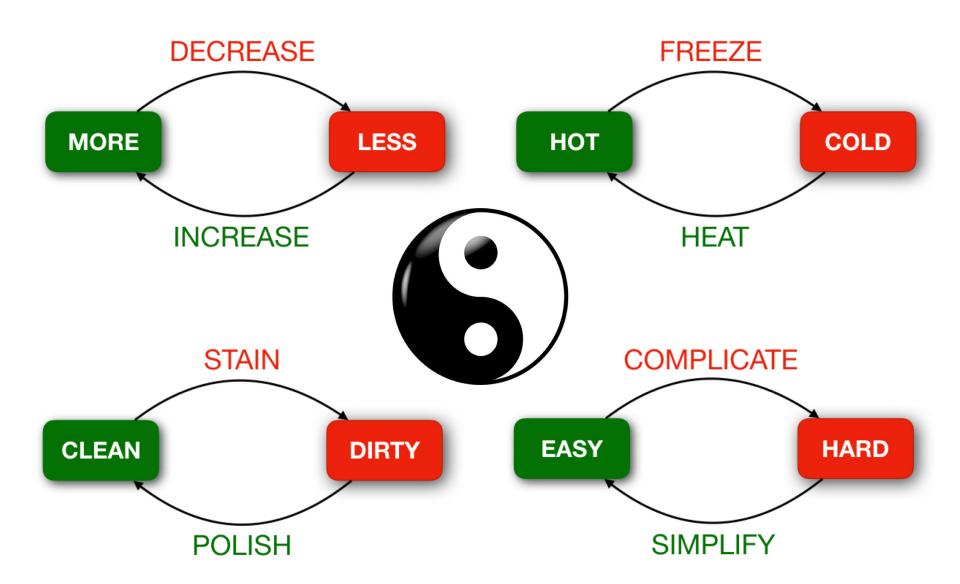
Explainable Al





Yin Yang approach





Sentic algebra





$$+1 x -1 = -1$$

INCREASE GAIN



$$+1 x +1 = +1$$

DECREASE GAIN



$$-1 x +1 = -1$$

DECREASE PAIN



$$-1 \quad x \quad -1 = +1$$

E Cambria, S Poria, R Bajpai, B Schuller. SenticNet 4: A Semantic Resource for Sentiment Analysis Based on Conceptual Primitives. In: COLING, 2666-2677 (2016)

Sentic algebra



lessen, reduce, drop, diminish, minimize, decline, dwindle, contract, shrink, abate, subside, ebb, wanelighten, sink, slump, plummet, plunge, cut down, curtail, diminish, narrow, pare down, slim down, tone down, temper, weaken, deplete, taper off, peter out, fall off, die down, let up, tail off, ...

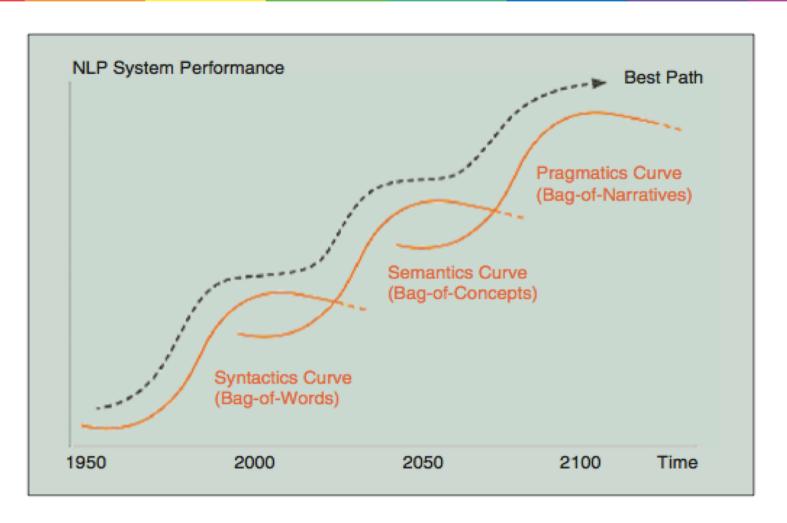
DECREASE

suffering, agony, affliction, torture, torment, discomfort, soreness, ache, soreness, hurt, throb, smarting, pricking, sting, twinge, shooting pain, stab, pang, spasm, stitch, cramp, irritation, stiffness, pain in the neck, nuisance, pest, bother, vexation, worry, tribulation, plague, bore, ...

PAIN

Jumping NLP curves





E Cambria, B White. Jumping NLP Curves: A Review of Natural Language Processing Research. IEEE Computational Intelligence Magazine 9(2), 48-57 (2014)

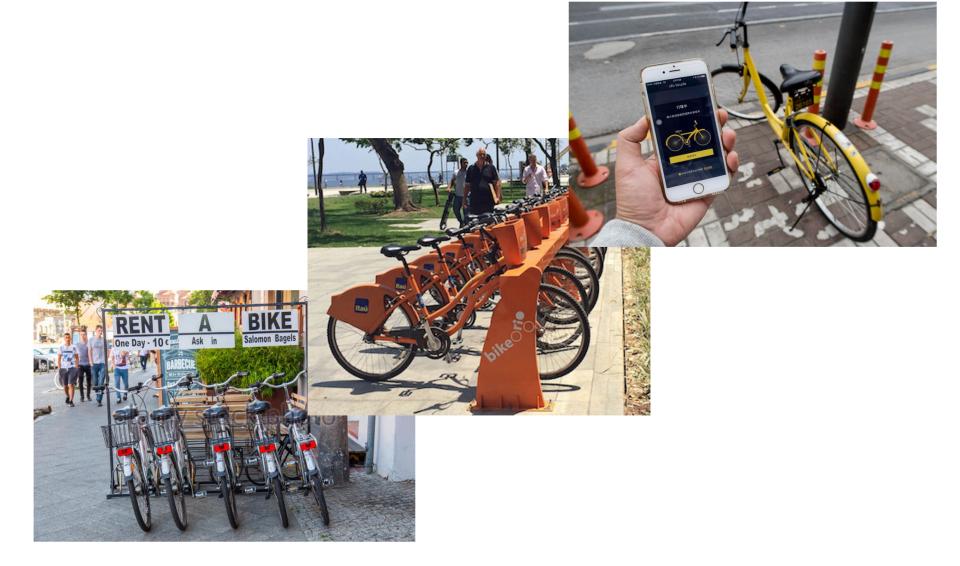
Jumping curves





Bike sharing example





A 'pipe' is not a pipe



You can know the name of all the different kinds of 'pipe', but you know nothing about a pipe until you comprehend its purpose and method of usage



From BoW to BoC





long





big





cold



From BoW to BoC



smile

=> sad_smile



damn

=> damn_good



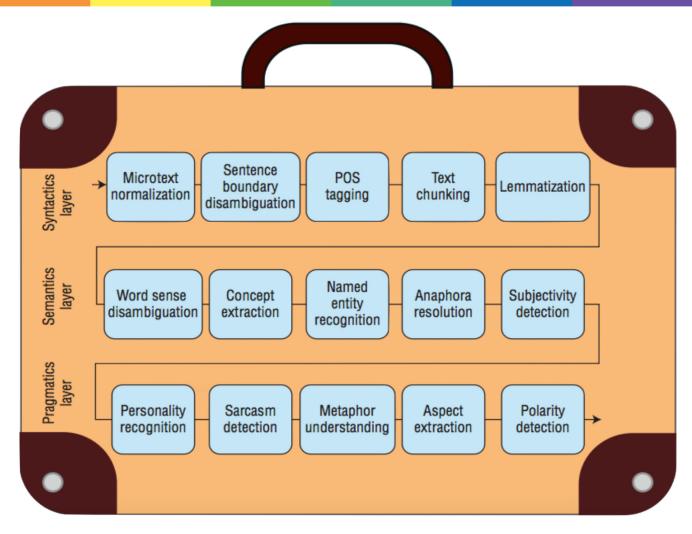
pretty

=> pretty_ugly



A big suitcase





E Cambria, S Poria, A Gelbukh, M Thelwall. Sentiment Analysis is a Big Suitcase. IEEE Intelligent Systems 32(6), 74-80 (2017)

Syntactics layer



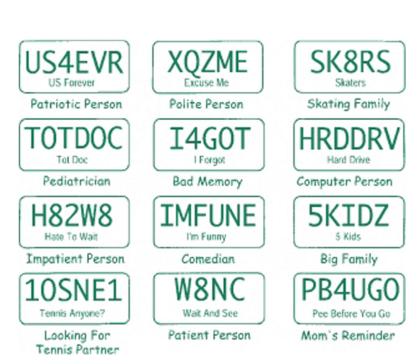


E Cambria, S Poria, A Gelbukh, M Thelwall. Sentiment Analysis is a Big Suitcase. IEEE Intelligent Systems 32(6), 74-80 (2017)

Microtext normalization



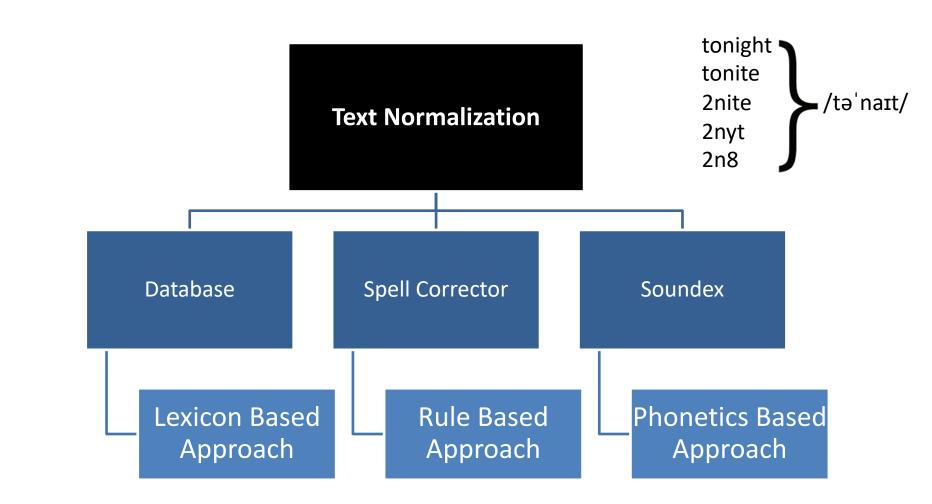
Before NLP techniques can be applied, informal text (e.g., c u l8r), acronyms (e.g., LOL), and emoticons (e.g., :>), need to be translated into plain English



K Bontcheva, L Derczynski, A Funk, M Greenwood, D Maynard, N Aswani. TwitIE: An Open-Source Information Extraction Pipeline for Microblog Text. In: ACL (2013)

Microtext normalization





R Satapathy, C Guerreiro, I Chaturvedi, E Cambria. Phonetic-Based Microtext Normalization for Twitter Sentiment Analysis. In: ICDM, 407-413 (2017)

SBD



The group included J. M. Freeman Jr. and T. Boone Pickens. The group included J. M. Freeman Jr. T. Boone Pickens had left.

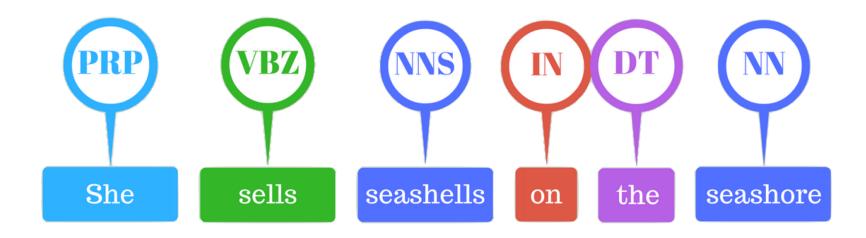
He stopped to see Dr. Lawson. He stopped at Meadows Dr. Lawson was still open.

It was due Friday by 5½ p.½m.½ Saturday would be too late.½ ½%She has an appointment at 5 p.½m.½ Saturday to get her car ½fixed?.

D Palmer, M Hearst. Adaptive multilingual sentence boundary disambiguation. Computational Linguistics 23(2), 241-267 (1997)

POS tagging





- The <u>back</u> door = JJ
- On my <u>back</u> = NN
- Win the voters back = RB
- Promised to <u>back</u> the bill = VB

K Toutanova, D Klein, C Manning, Y Singer. Feature-rich part-of-speech tagging with a cyclic dependency network. In NAACL-HLT, 173-180 (2003)

Text chunking



Text chunking, also referred to as shallow parsing, is a task that follows POS tagging and that adds more structure to the sentence. The result is a grouping of the words in "chunks".

L Ramshaw, M Marcus. Text chunking using transformation-based learning. Text, Speech and Language Technology 11, 157-176 (1999)

Lemmatization



eats_burger

eating_burger

ate_burger



eat_burger

eat_the_burger

eat_burgers

Lemmatization



cross

=> crossing

hell_ride

=> a_hell_of_a_ride

kick_ball

=> kick_balls



Semantics layer



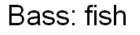


E Cambria, S Poria, A Gelbukh, M Thelwall. Sentiment Analysis is a Big Suitcase. IEEE Intelligent Systems 32(6), 74-80 (2017)

WSD









???



Bass: instrument

D Yarowsky. Unsupervised word sense disambiguation rivaling supervised methods. In ACL, 189-196 (1995)

Concept extraction



the camera has [long focus time]
the camera takes a [long time] to [focus]
the [focusing] of the camera takes [long time]
the [focus time] of the camera is very [long]



long_focus_time

D Rajagopal, E Cambria, D Olsher, K Kwok. A graph-based approach to common-sense concept extraction and semantic similarity detection. In WWW, 565–570 (2013)

Concept extraction



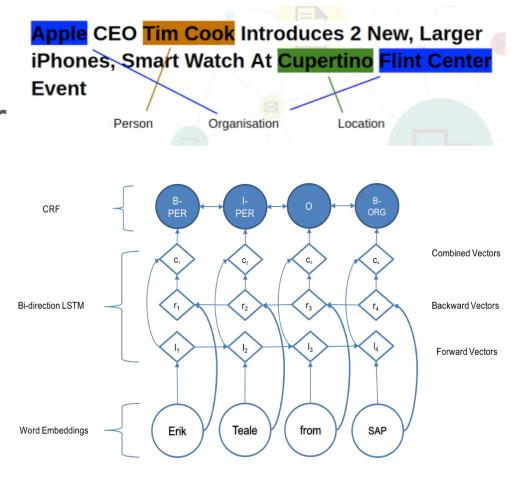
After chunking and lemmatization, each potential noun chunk is paired with stemmed verbs in order to detect verb+object multi-word expressions

```
Data: NounPhrase
Result: Valid object concepts
Split the NounPhrase into bigrams;
Initialize concepts to Null;
for each NounPhrase do
   while For every bigram in the NounPhrase do
      POS Tag the Bigram ;
      if adj noun then
         add to Concepts: noun, adj+noun
      else if noun noun then
         add to Concepts: noun+noun
      else if stopword noun then
       add to Concepts: noun
      else if adj stopword then
       continue
      else if stopword adj then
       continue
       Add to Concepts : entire bigram
      repeat until no more bigrams left;
```

NER



Named-entity recognition is key for improving anaphora resolution and, hence, for detecting aspects or opinion targets in reviews.



Y Ma, E Cambria, S Gao. Label Embedding for Zero-shot Fine-grained Named Entity Typing. In: COLING, 171-180 (2016)

Anaphora resolution



Anaphora is the use of an expression the interpretation of which depends upon another one. It is commonly resolved by gender and number agreement

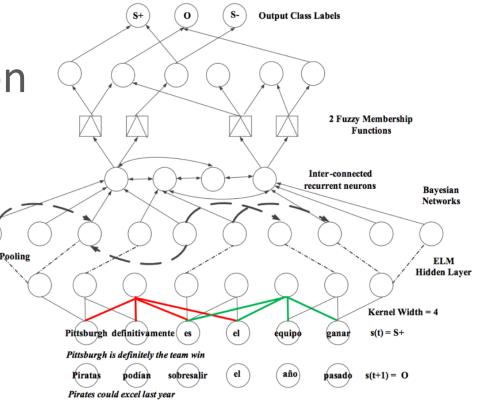


R Mitkov. Anaphora resolution: the state of the art. Working paper. University of Wolverhampton, Wolverhampton (1999)

Subjectivity detection



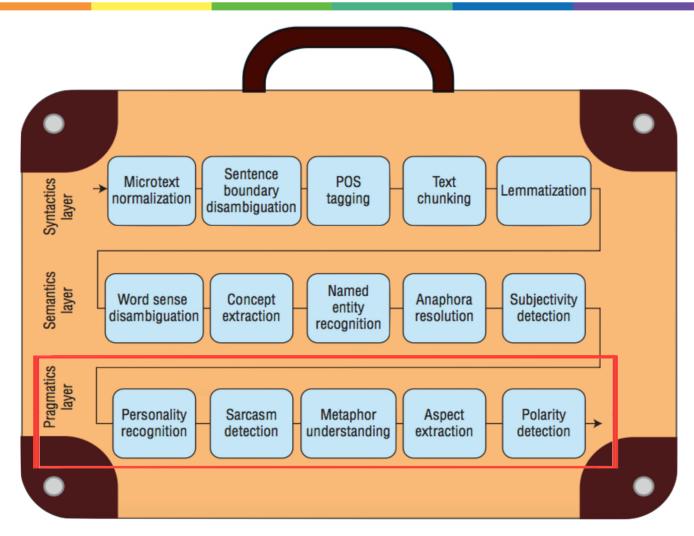
Subjectivity detection is a binary classification task that consists in classifying text into either objective (neutral) or subjective (positive/negative)



I Chaturvedi, E Cambria, R Welsch, F Herrera. Distinguishing between facts and opinions for sentiment analysis: Survey and challenges. Information Fusion 44, 65–77 (2018)

Pragmatics layer



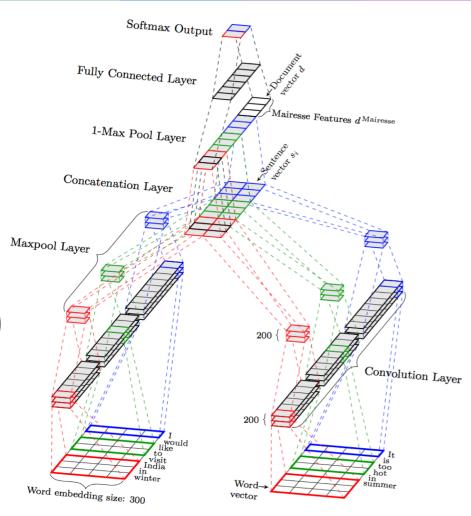


E Cambria, S Poria, A Gelbukh, M Thelwall. Sentiment Analysis is a Big Suitcase. IEEE Intelligent Systems 32(6), 74-80 (2017)

Personality recognition



Personality recognition is an important step towards userdependent sentiment analysis (user profiling) and it is useful for sarcasm detection



N Majumder, S Poria, A Gelbukh, E Cambria. Deep learning-based document modeling for personality detection from text. IEEE Intelligent Systems 32(2), 74-79 (2017)

Personality recognition

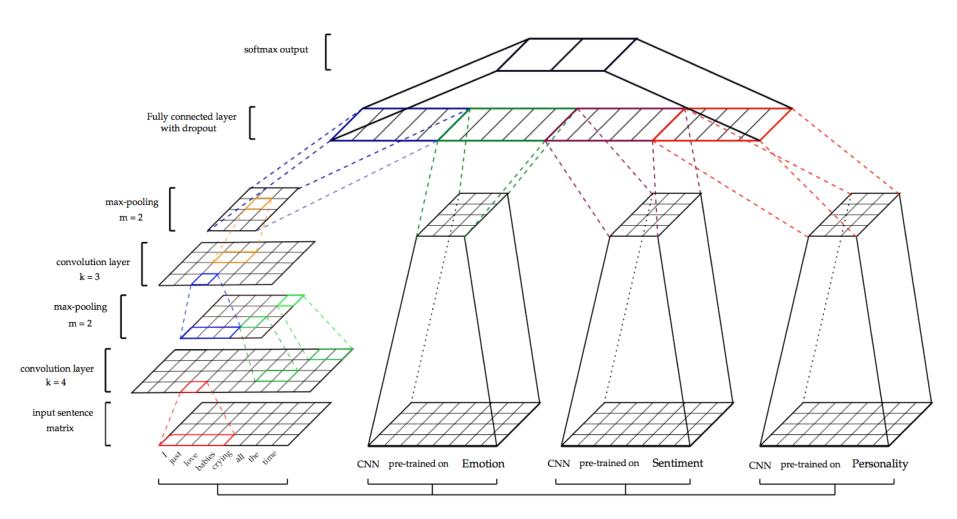




http://datanami.com/2017/09/21/deep-learning-reveals-new-insights-people

Sarcasm detection

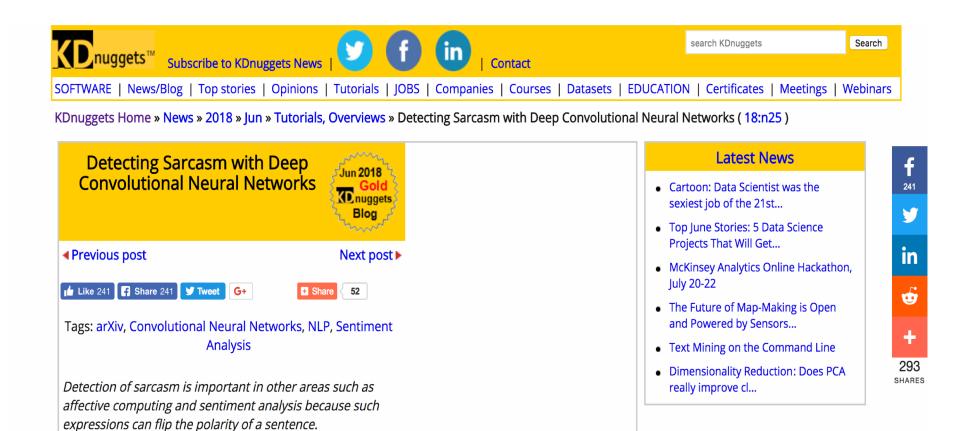




S Poria, E Cambria, D Hazarika, P Vij. A Deeper Look into Sarcastic Tweets Using Deep Convolutional Neural Networks. In: COLING, 1601-1612 (2016)

Sarcasm detection





http://kdnuggets.com/2018/06/detecting-sarcasm-deep-convolutional-neural-networks

Metaphor understanding



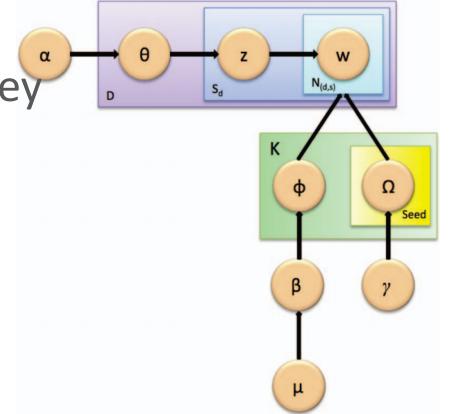


M Schulder and E Hovy. Metaphor detection through term relevance. In Proceedings of the Second Workshop on Metaphor in NLP, 18–26 (2014)

Aspect extraction



Extracting aspects (or features) from text is key for identifying the socalled opinion targets and disambiguate the polarity associated to each of them

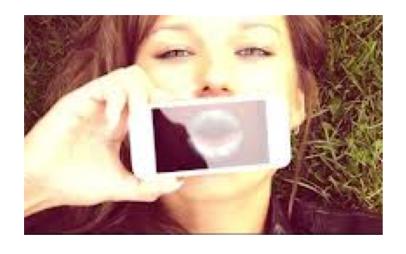


S Poria, E Cambria, A Gelbukh. Aspect extraction for opinion mining with a deep convolutional neural network. Knowledge-Based Systems 108, 42-49 (2016)

Aspect extraction



I love iPhoneX's touchscreen but the battery life is so short



touchscreen



Uh oh.

battery



Y Ma, H Peng, E Cambria. Targeted aspect-based sentiment analysis via embedding commonsense knowledge into an attentive LSTM. In: AAAI, pp. 5876-5883 (2018)

Polarity detection



Early works treated polarity detection as a binary classification problem (pos vs neg). Recent works calculate polarity intensity as a float $\in [-1, +1]$

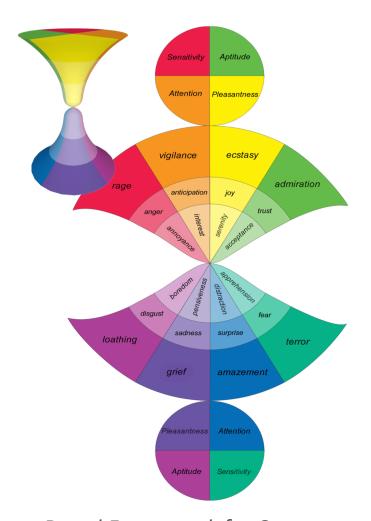


S Poria, et al. Sentiment data flow analysis by means of dynamic linguistic patterns. IEEE Computational Intelligence Magazine 10(4), 26-36 (2015)

Sentic computing



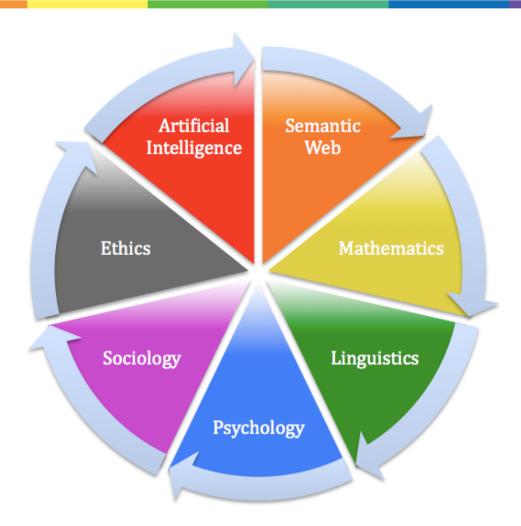
- 1. Shift from mono- to multi-disciplinarity
- 2. Shift from syntax to semantics
- 3. Shift from statistics to *linguistics*



E Cambria, A Hussain. Sentic Computing: A Common-Sense-Based Framework for Concept-Level Sentiment Analysis. Cham: Springer, ISBN: 978-3-319-23654-4 (2015)

From mono- to multi-disciplinarity





E Cambria, A Hussain. Sentic Computing: A Common-Sense-Based Framework for Concept-Level Sentiment Analysis. Cham: Springer, ISBN: 978-3-319-23654-4 (2015)

From syntax to semantics



cloud_computing => cloud



pain_killer => pain, killer



E Cambria, A Hussain. Sentic Computing: A Common-Sense-Based Framework for Concept-Level Sentiment Analysis. Cham: Springer, ISBN: 978-3-319-23654-4 (2015)

From statistics to linguistics



The car is nice but expensive



The car is expensive but nice



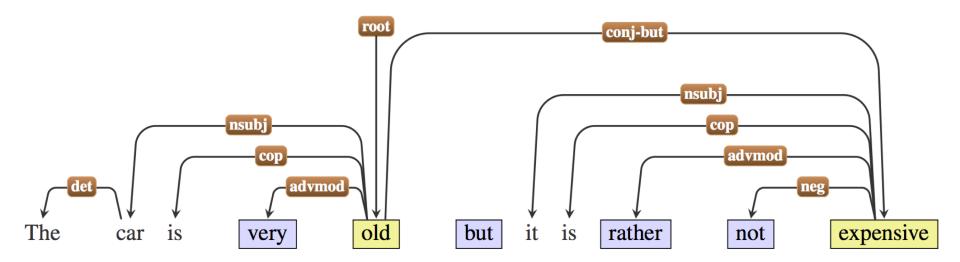


Matrix predicate (h)	Dependent predicate (d)	Dep. comp. (x)	Overall polarity	Example
Pos	Pos	Pos	Pos	a
Pos	Pos	Neg	Neg	b
Pos	Neg	Pos	Neg	c
Pos	Neg	Neg	Pos	d
Neg	Pos	Pos	Neg	e
Neg	Pos	Neg	Neg	f
Neg	Neg	Pos	Neg	g
Neg	Neg	Neg	Neg	h
Pos	Neutral	Pos	Pos	i
Pos	Neutral	Neg	Neg	j
Neg	Neutral	Pos	Neg	k
Neg	Neutral	Neg	Neg	1

- This is perfect to gain money.
- b. This is perfect to gain weight.
- This is perfect to lose money.
- This is perfect to lose weight.
- e. This is useless to gain money.
- f. This is useless to gain weight.

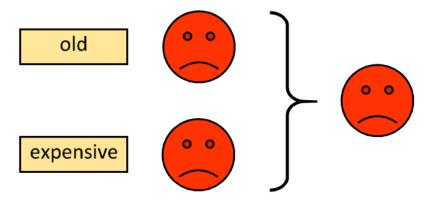
- This is useless to lose money.
- This is useless to lose weight.
- This is perfect to talk about money.
- j. This is perfect to talk about weight.
- k. This is useless to talk about money.
- This is useless to talk about weight.





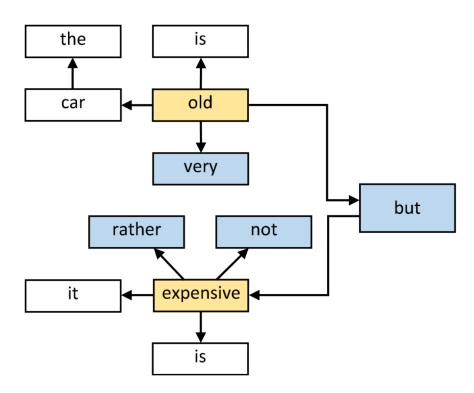
(a) Dependency tree of a sentence.





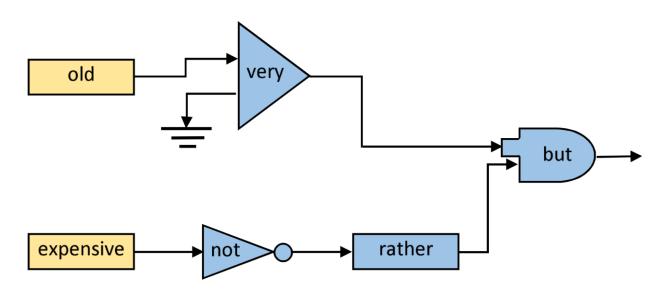
(b) The old way: averaging over a bag of sentiment words. The overall polarity of a sentence is given by the algebraic sum of the polarity values associated with each affect word divided by the total number of words.





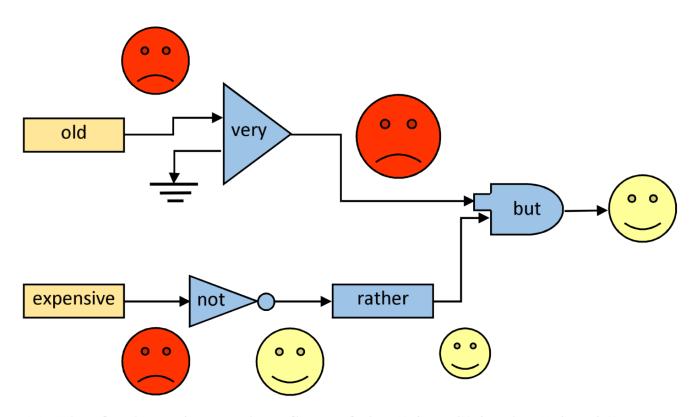
(c) The dependency tree of a sentence resembles an electronic circuit: words shown in blue can be thought as a sort of "boolean operations" acting on other words.





(d) The electronic circuit metaphor: sentiment words are "sources" while other words are "elements", e.g., *very* is an amplifier, *not* is a logical complement, *rather* is a resistor, *but* is an OR-like element that gives preference to one of its inputs.





(e) The final sentiment data flow of the "signal" in the "circuit".



SenticNet Stanford

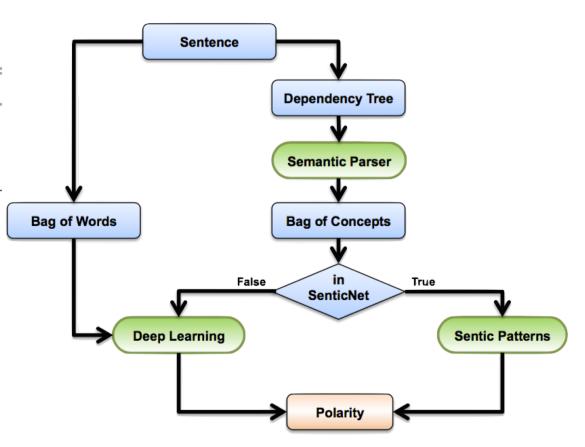
1. I love the movie which you hate	+	_
2. The phone is very big to hold	_	+
3. You are making fun of me	_	+
4. You are not so beautiful	_	+
5. The tooth hit the pavement and broke	_	+
6. I am one of the least happy person in the world	_	0
7. I love Starbucks but they just lost a customer	_	0
8. I doubt that he is good		+
9. Receiving payments has never been this simple & fast		_
10. I am eagerly looking forward to Dr. Wu's future work		_

http://sentic.net/demos



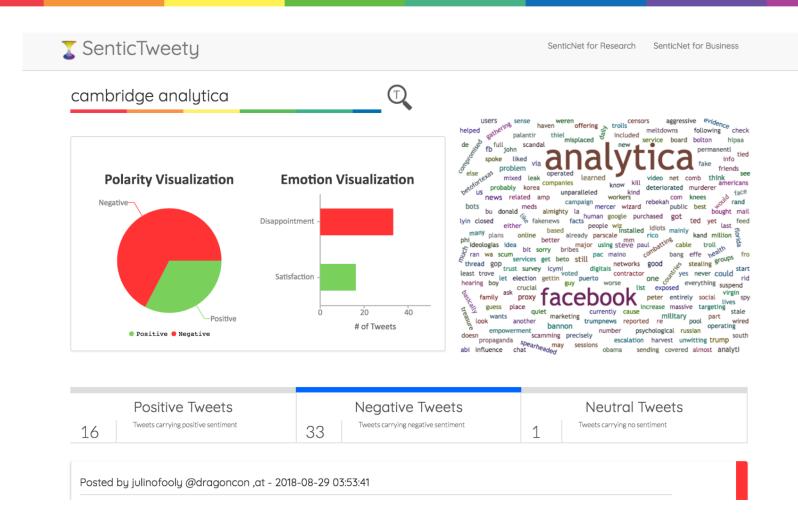
Algorithm	Precision
Sentic Patterns	84.15%
Machine Learning	67.35%
Ensemble Classification	86.21%

System	Precision
Socher et al. 2012 [59]	80.00%
Socher et al. 2013 [57]	85.40%
Proposed Method	86.21%



Social media marketing

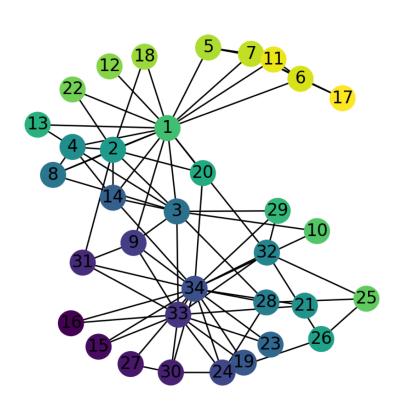


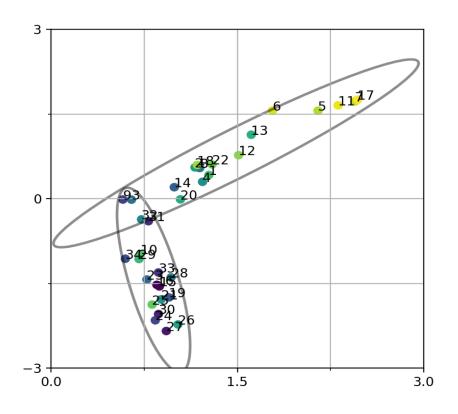


E Cambria, M Grassi, A Hussain, C Havasi. Sentic Computing for Social Media Marketing. Multimedia Tools and Applications 59(2), 557-577 (2012)

Big social data analysis







S Cavallari, VW Zheng, H Cai, KCC Chang, E Cambria. Learning Community Embedding with Community Detection and Node Embedding on Graphs. In: CIKM, 377-386 (2017)

Sentic album

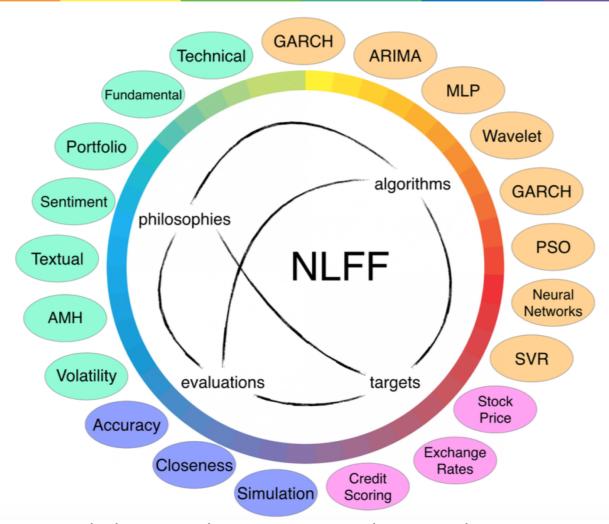




E Cambria, A Hussain. Sentic Album: Content-, Concept-, and Context-Based Online Personal Photo Management System. Cognitive Computation 4(4), 477-496 (2012)

Financial forecasting



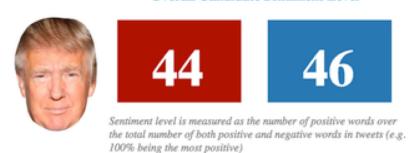


F Xing, E Cambria, R Welsch. Natural Language Based Financial Forecasting: A Survey. Artificial Intelligence Review 50(1), 49-73 (2018)

Political forecasting









100% being the most positive)



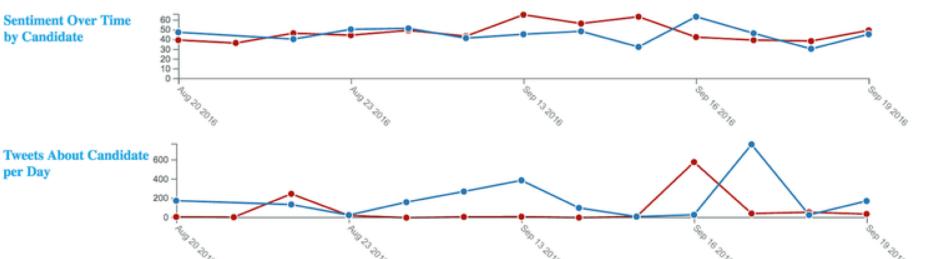


Candidate

 Donald Trump ■ Hillary Clinton

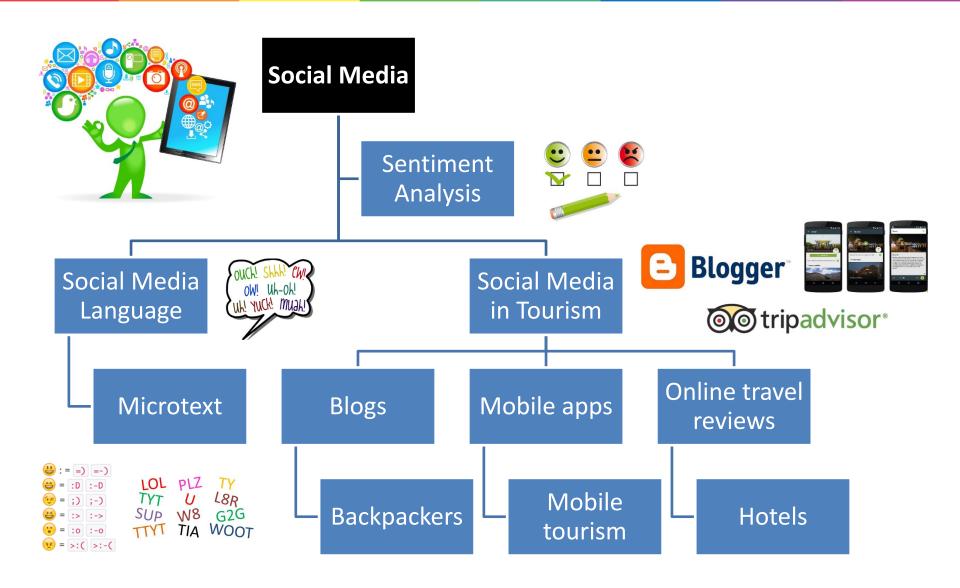
Sentiment Over Time by Candidate

per Day



e-Tourism

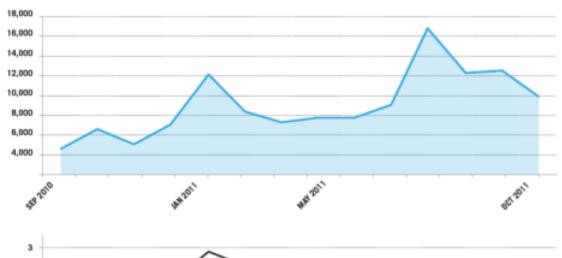




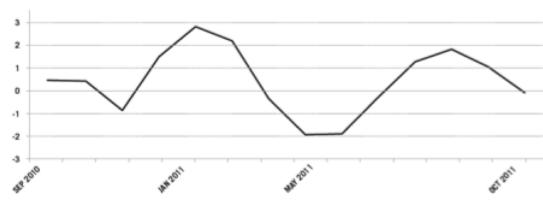
Cyber issue detection



Tweets about the price of rice

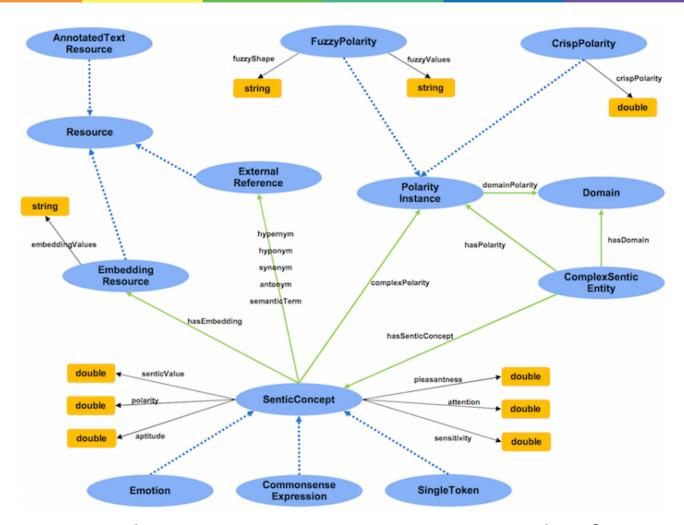


Food Price Inflation



Al for the Semantic Web

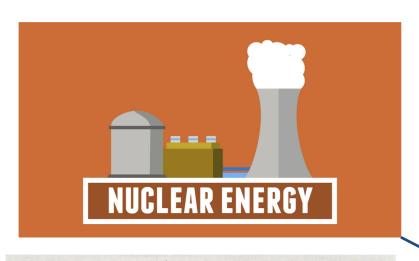




M Dragoni, S Poria, E Cambria. OntoSenticNet: A commonsense ontology for sentiment analysis. IEEE Intelligent Systems 33(3), 77-85 (2018)

Media monitoring









http://sentic.net/projects/#one-belt-one-road-one-sentiment

Mood of Singapore



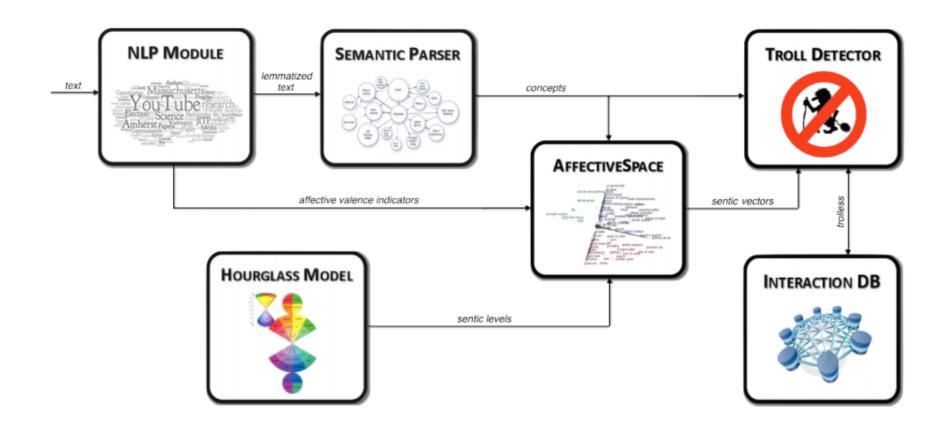




http://sentic.net/projects/#mood-of-singapore

Troll filtering

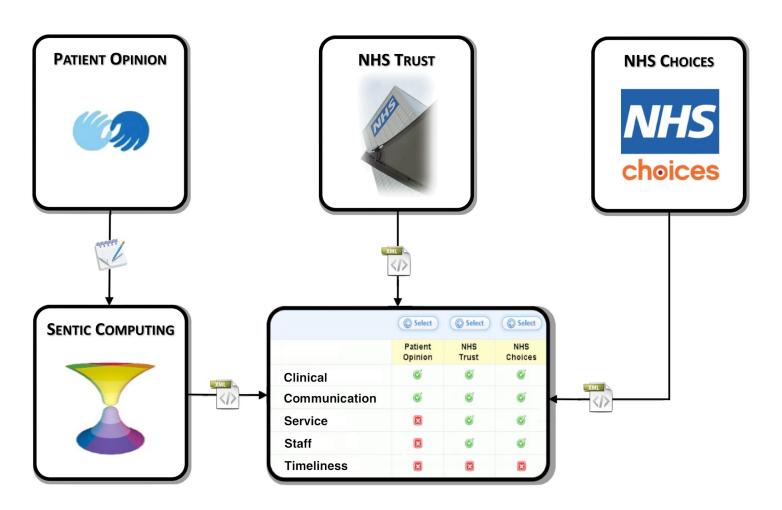




E Cambria, P Chandra, A Sharma, A Hussain. Do not feel the trolls. In: ISWC (2010)

Crowd validation





E Cambria, A Hussain, T Durrani, C Havasi, C Eckl, J Munro. Sentic Computing for Patient Centered Applications. In: IEEE ICSP, 1279-1282 (2010)

Sentic PROMs



In spite of demonstrated benefits, routine HRQoL assessments remain rare as few patients are willing to spend the time needed to fill-in long questionnaires daily

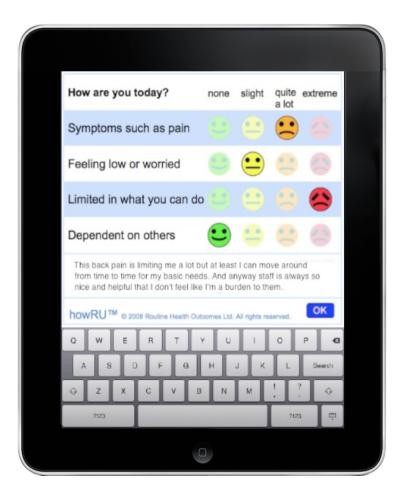


E Cambria, T Benson, C Eckl, A Hussain. Sentic PROMs: Application of Sentic Computing to the Development of a Novel Unified Framework for Measuring Health-Care Quality. Expert Systems with Applications 39(12), 10533–10543 (2012)

Sentic PROMs



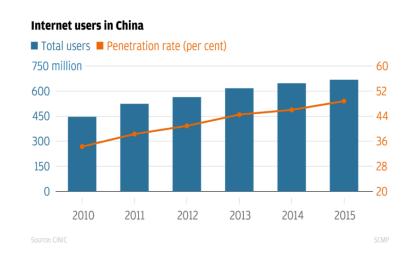
Sentic PROMs allow patients to evaluate their health and healthcare experience to accordingly aggregate text and visual data in a semi-structured way

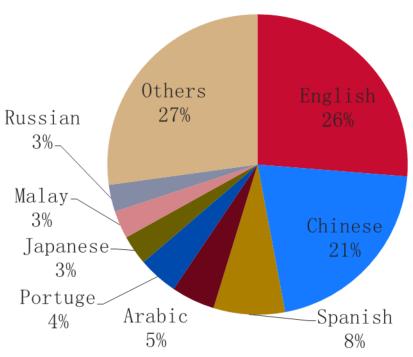


Multilingual analysis



Top 8 Language used on Internet

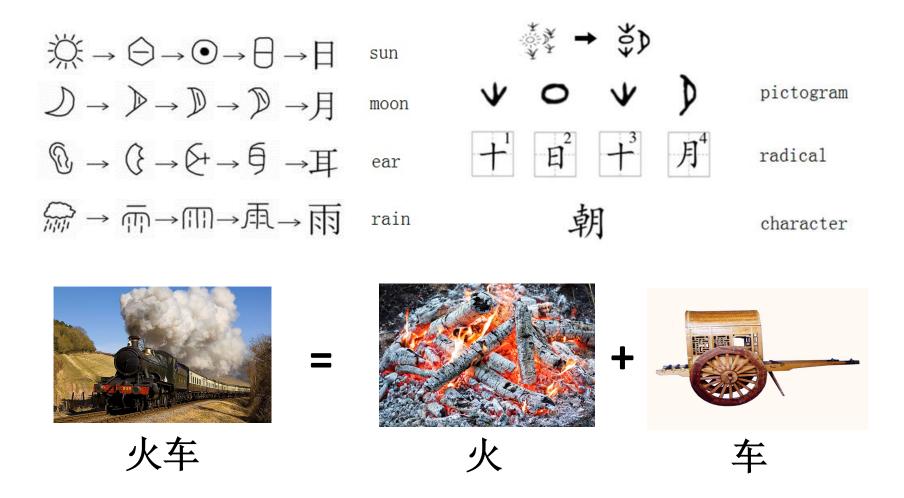




SL Lo, E Cambria, R Chiong, D Cornforth. Multilingual sentiment analysis: From formal to informal and scarce resource languages. Artificial Intelligence Review 48(4), 499-527 (2017)

Multilingual analysis



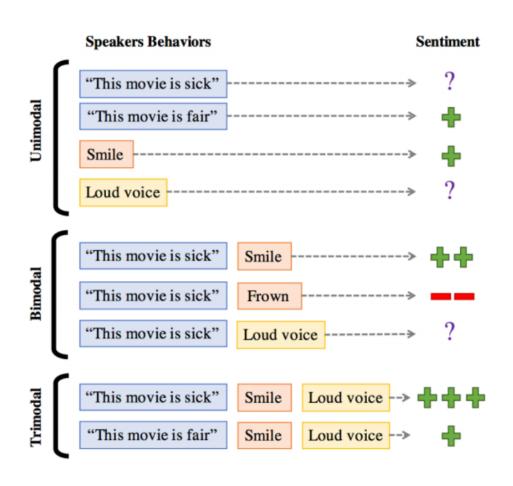


H Peng, E Cambria, X Zou. Radical-Based Hierarchical Embeddings for Chinese Sentiment Analysis at Sentence Level. In: FLAIRS, 347-352 (2017)

The more, the merrier



When available, more modalities aid to resolve ambiguities in communication and, hence, enhance polarity detection



A Zadeh, M Chen, S Poria, E Cambria, LP Morency. Tensor fusion network for multimodal sentiment analysis. In: EMNLP, 1114-1125 (2017)

Multimodal analysis



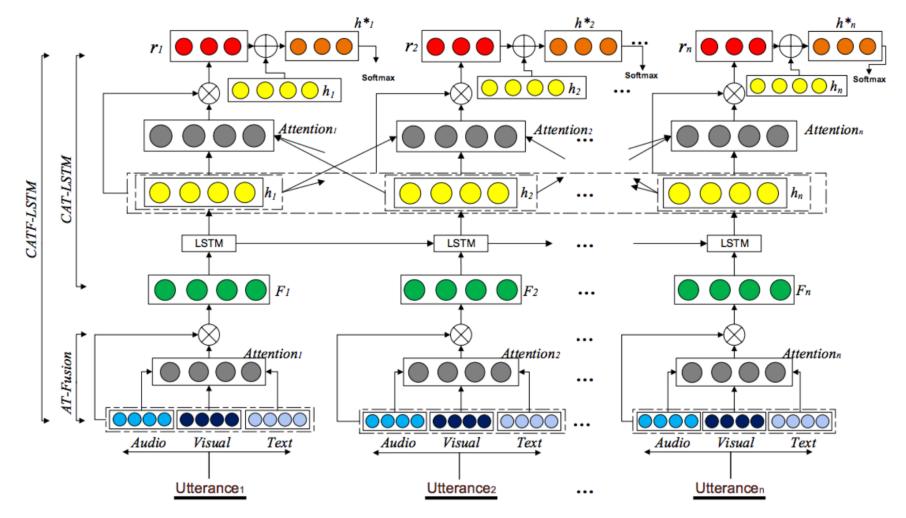
An ensemble feature extraction approach exploiting the joint use of tri-modal (text, audio and video) features enhances multimodal sentiment analysis



S Poria, E Cambria, D Hazarika, N Mazumder, A Zadeh, L Morency. Context-dependent sentiment analysis in user-generated videos. In: ACL, 873-883 (2017)

Multimodal attention

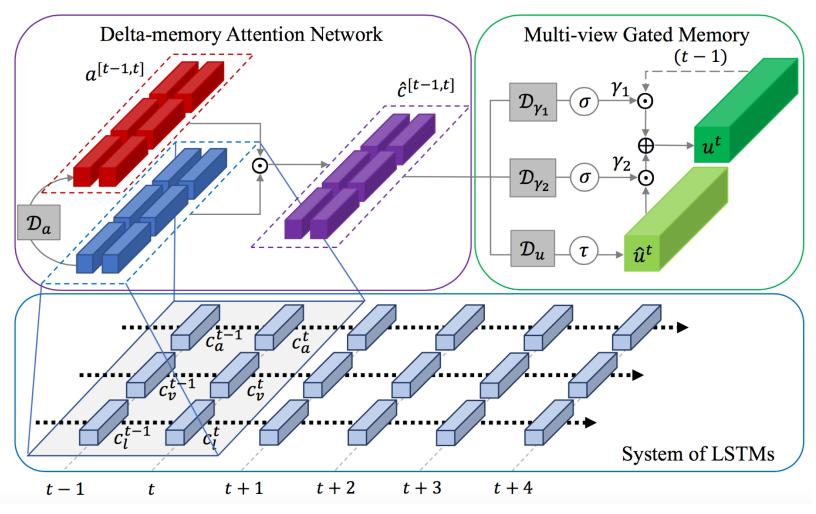




S Poria, E Cambria, D Hazarika, N Mazumder, A Zadeh, LP Morency. Multi-level multiple attentions for context-aware multimodal sentiment analysis. In: ICDM, 1033–1038 (2017)

Memory fusion network

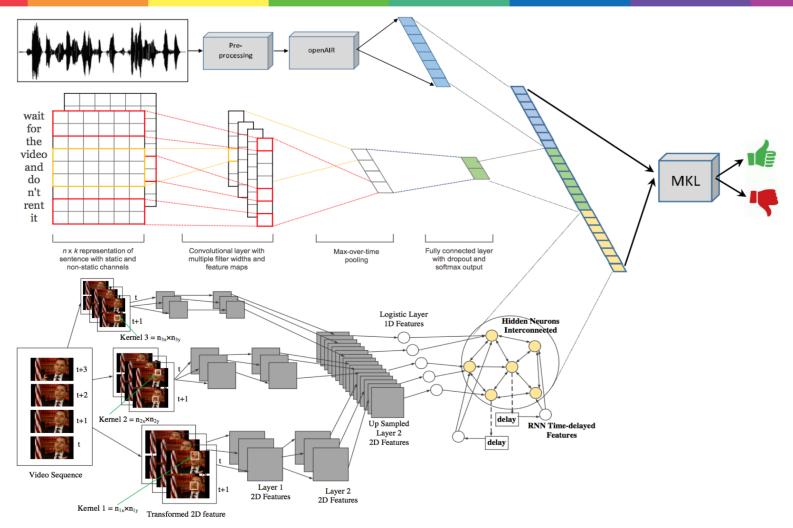




A Zadeh, PP Liang, N Mazumder, S Poria, E Cambria, LP Morency. Memory fusion network for multi-view sequential learning. In: AAAI, 5634-5641 (2018)

Multimodal fusion

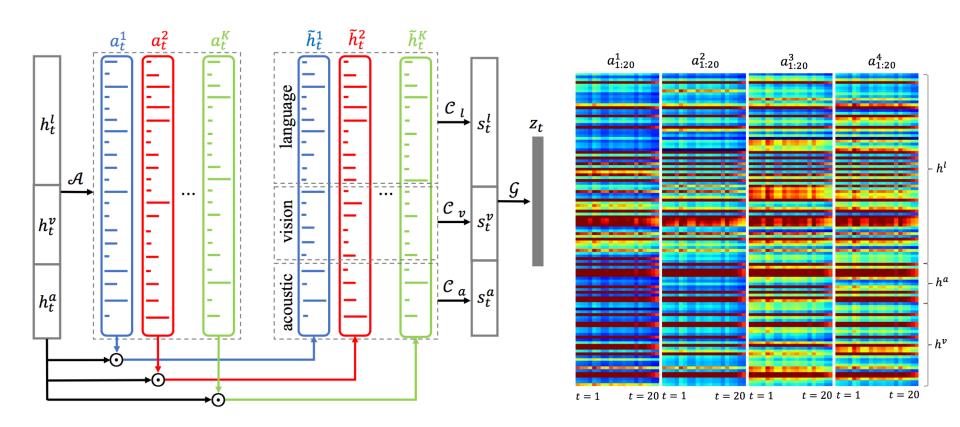




S Poria, I. Chaturvedi, E Cambria, A Hussain. Convolutional MKL Based Multimodal Emotion Recognition and Sentiment Analysis. In: ICDM, 439-448 (2016)

Multi-attention fusion

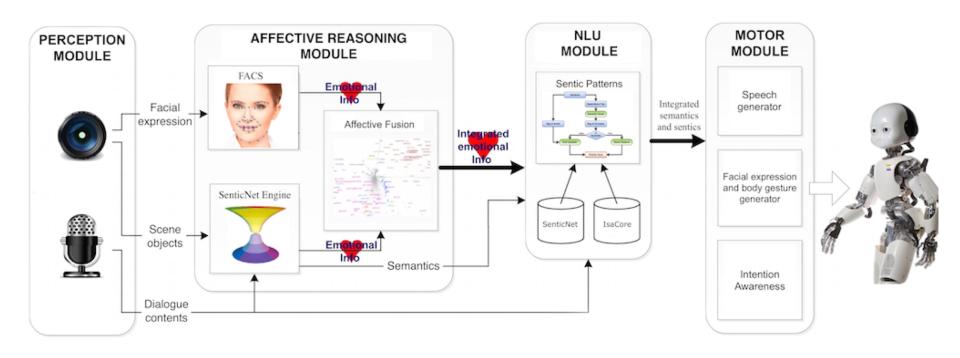




A Zadeh, PP Liang, S Poria, P Vij, E Cambria, LP Morency. Multi-attention recurrent network for human communication comprehension. In: AAAI, 5642-5649 (2018)

Dialogue systems





T Young, E Cambria, I Chaturvedi, H Zhou, S Biswas, M Huang. Augmenting end-to-end dialogue systems with commonsense knowledge. In: AAAI, 4970-4977 (2018)

Open challenges



User Preference

hard/soft bed, small/big phone, cheap/expensive bag

Theory of Mind

It is good that you killed the professor

Theory of Relativity

I used to hate gardening but now I love it

Intention Awareness

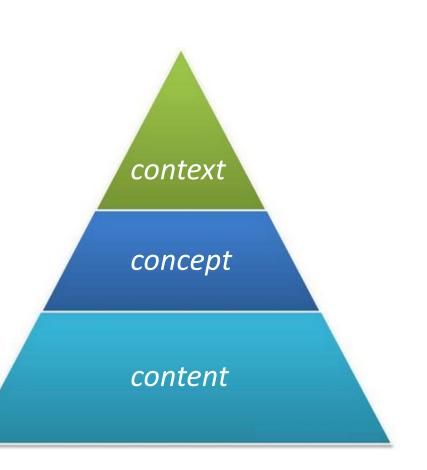
big/small room, warm/cold water

N Howard, E Cambria. Intention awareness: Improving upon situation awareness in human-centric environments. Human-centric Computing and Information Sciences 3(9) (2013)

3C sentiment analysis



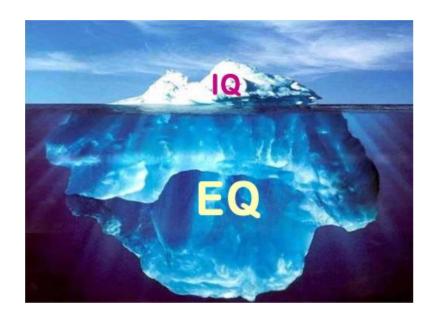
Sentiment analysis is distinguishing itself as a separate field and is moving toward content-, concept-, and contextbased natural language analysis



3Q NLP



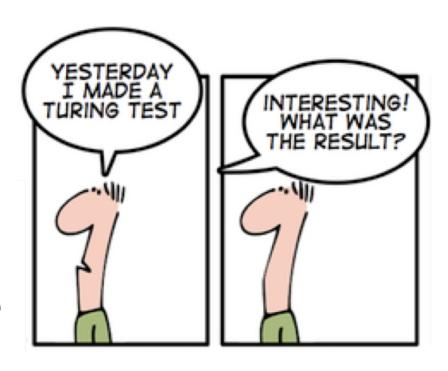
To achieve real machine intelligence, a computer needs to be able to not only perform reasoning (IQ), but also interpret emotions (EQ) and cultural nuances (CQ)



Machines that think



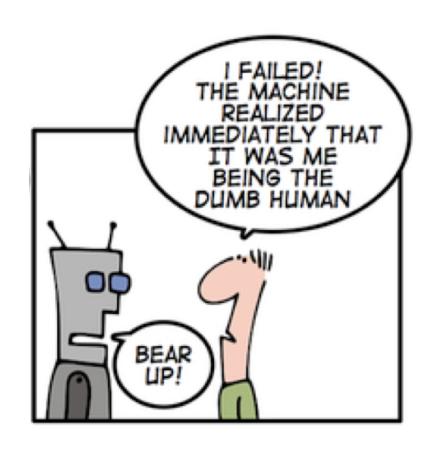
The world has changed less since Jesus Christ than it has in the last century. In another century's time, machines might be able to think as humans do.



Machines that think



But do we really want them to think exactly like us? Should we attempt to make them better than us or make sure they will never be as smart as us?



Are we even a good model?







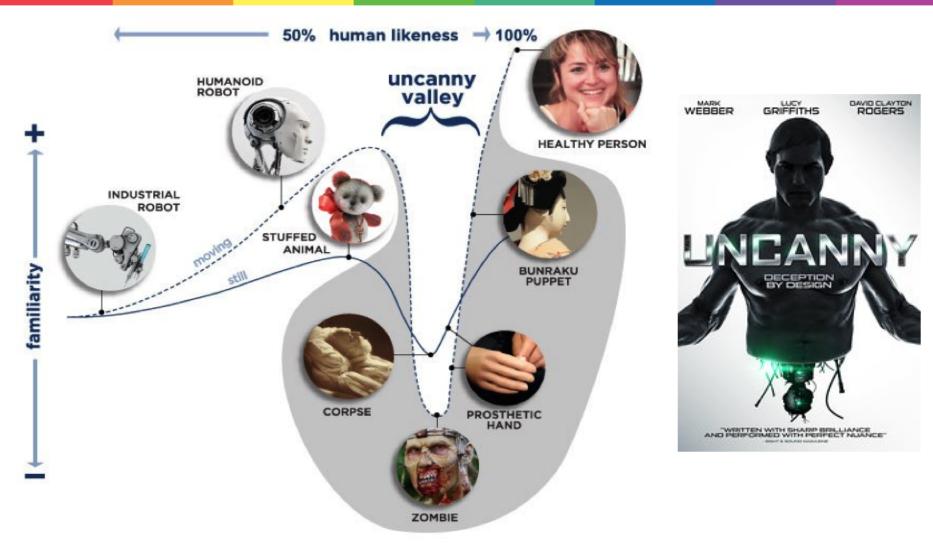
How much should they feel?

Do we really want to emulate the full spectrum of human emotions into machines? Should we remove some negative emotions such as revenge or envy?



How should they look like?



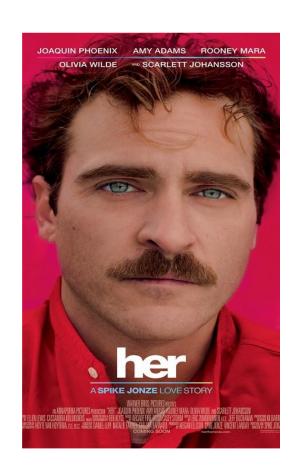


M Mori. The uncanny valley. IEEE Robotics & Automation Magazine. 19 (2), 98-100 (1970)

How smart should they be?

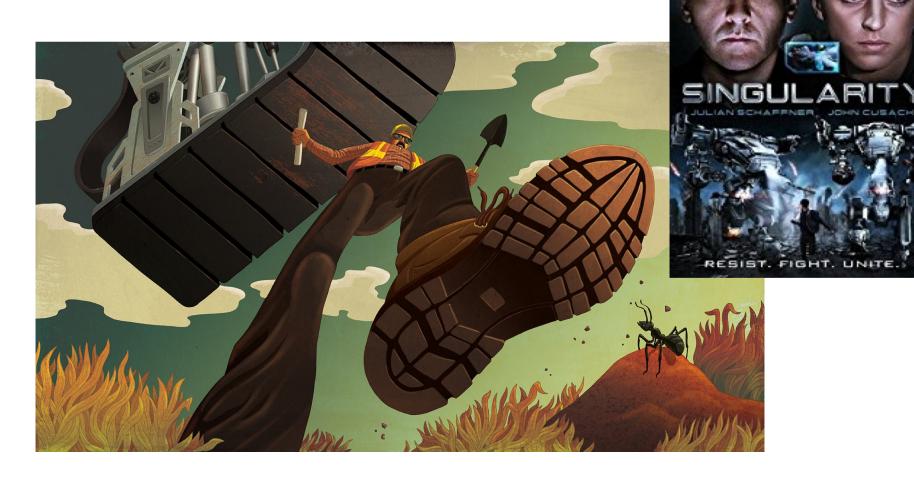


Do we really want machines to be smarter than us? If not, what can we do to prevent that from happening? Is an Al superintelligence going to take care of us or destroy us?



What is the risk?





Conclusion

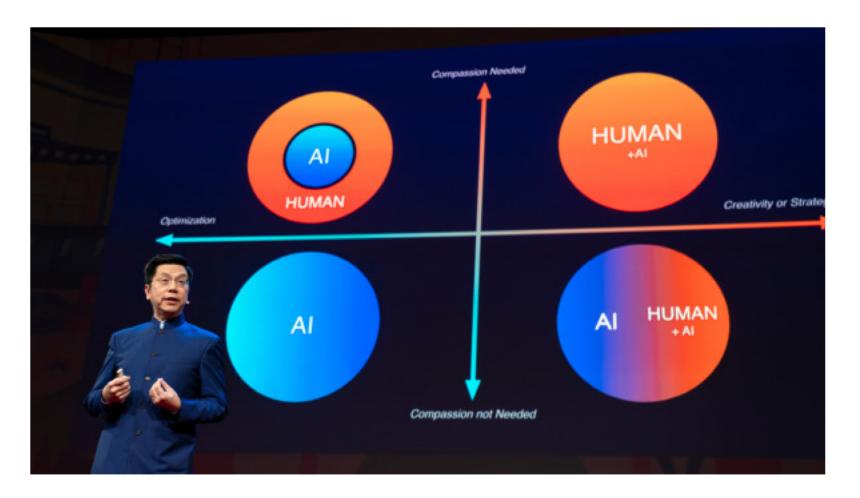


The truth is that we are still light-years away from a truly intelligent machine, due in part to the fact that we do not even fully understand human intelligence yet



Conclusion





http://ted.com/talks/kai_fu_lee_how_ai_can_save_our_humanity

Conclusion



We do not aim to help machines take over the world but rather to exploit differences in human-computer abilities and costs so as to achieve symbiotic HMI

Questions?



