TEXT MINING - SENTIMENT ANALYSIS /OPINION MINING

SRI SURYA NULU

EMIS-DATA MINING 8331

SOUTHERN METHODIST UNIVERSITY

SNULU@SMU.EDU

SRISURYANULU@GMAIL.COM

AGENDA

- ✓ Introduction
- ✓ Definition of opinion mining
- ✓ Why opinion mining
- ✓ Applications
- ✓ Elements of opinion mining
- ✓ Entity and aspect
- ✓ Opinion Mathematical model
- ✓ Doc-level/sentence level SA
- ✓ Feature based sentiment analysis
- ✓ Levels of opinion mining
- ✓ Sentiment analysis classification
- ✓ Machine learning approach
- ✓ Challenges
- ✓ Success stories

INTRODUCTION

- Opinion mining is one of the most active research areas in natural language processing and is also widely studied in data mining, Web mining, and text mining.
- The growing importance of sentiment analysis/opinion mining coincides with the growth of social media such as reviews, forum discussions, blogs, micro-blogs, Twitter, and social networks.
- Sentiment analysis systems are being applied in almost every business and social domain.
- Businesses are using opinion mining in many ways to promote products smartly by showcasing opinion analysis to customers.

Source: Sentiment Analysis and Opinion Mining, Bing Liu

DEFINITION

Sentiment analysis or opinion mining is a Computational study of opinions, sentiments, evaluations, attitudes, appraisal, affects, views, emotions, subjectivity, etc. expressed in text.

Reviews, blogs, discussions, news, comments, feedback, or any other documents.

OPINION MINING OR SENTIMENT ANALYSIS?

Sentiment analysis is more widely used in industry, but both of them mean the same.

Source: Sentiment Analysis and Opinion Mining, Bing Liu

OPINION MATTERS !!

(BE PANG AND LILLIAN LEE)

- 81% of Internet users have done online research on a product at least once.
- 20% of the users do so on a typical day.
- Among readers of online reviews about restaurants, hotels, and various services (e.g., travel agencies or doctors), between 73% and 87% report that reviews had a significant influence on their purchase.
- Consumers are willing to pay from 20% to 99% more for a 5-star-rated item than a 4-star-rated item (the variance stems from what type of item or service is considered).
- 32% have provided a rating on a product, service, or person via an online ratings system, and 30% (including 18% of online senior citizens) have posted an online comment or review regarding a product or service

Sentiment without action is the ruin of the soul. — Edward Abbey

APPLICATIONS

Businesses and organization: interested in opinions

a) *Product and service benchmarking*

b) *Market intelligence*

c) *Survey on a topic*

Individuals: interested in others opinion when

a) *Purchasing a product*

b) *Using a service*

c) *Tracking political topics*

d) *Other decision making tasks*

Ads Placements: placing ads in user-generated content

a) *Place an ad when one praises a product*

b) *Place an ad from a competitor if one criticize a product*

Opinion search: providing general search for opinions

Source: Sentiment Analysis and Opinion Mining, Bing Liu

ELEMENTS OF OPINION MINING

"I bought an iPhone a few days ago. It is such a nice phone. The touch screen is really cool. The voice quality is clear too. It is much better than my old blackberry, which was a terrible phone and so difficult to type with its tiny keys. However my mother was mad with me as I did not tell her before I bought the phone. She also thought the phone was too expensive, and wanted me to return it to the shop. …"

Entity	Attributes
Opinion targets	Entities and their features/aspects
Sentiments	positive and negative
Opinion holders	Persons who hold the opinions
Time	when opinions are expressed

EXAMPLE

"I bought an iPhone a few days ago. It is such a nice phone. The touch screen is really cool. The voice quality is clear too. Although the battery life was not long that is ok for me. However my mother was mad with me as I did not tell her before I bought the phone. She also thought the phone was too expensive, and wanted me to return it to the shop. …"

One can look at this review/blog at the

1. document level, i.e., is this review + or -?
2. sentence level, i.e., is each sentence + or -?
3. entity and feature/aspect level
4. It's a regular opinion or a comparative opinion

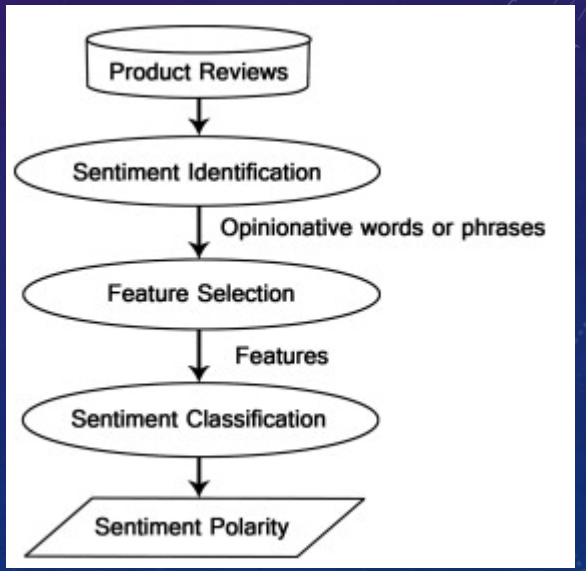

- ✓ Sentiment Analysis can be considered as a classification process.
- ✓ The target of SA is to find opinions, identify the sentiments they express, and then classify their polarity as shown.

Source: Sentiment Analysis Algorithms and Applications
Walaa Medhat Ahmed Hassan Hoda Korashy,2013

ENTITY AND ASPECT

Definition (entity): An *entity e is a product, person, event*, organization, or topic. *e is represented as* a hierarchy of components, sub-components, and so on. Each node represents a component and is associated with a set of attributes of the component.

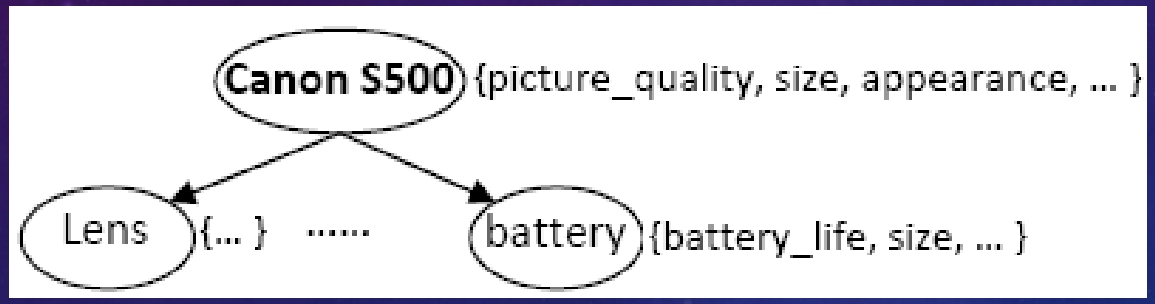

An opinion can be expressed on any node or attribute of the node.

For simplicity, we use the term aspects (features) to represent both components and attributes.

Source: Sentiment Analysis and Opinion Mining, Bing Liu

A REGULAR OPINION

An *opinion has the follow basic components*

$(g_j, so_{ijl}, h_i, t_l)$

Where

g_i *is a target.*

so_{ijl}*, is the sentiment value of the opinion from opinion holder* h_i *on target* g_j *at the time* t_l*.* so_{ijl} *is positive, negative or neutral.*

h_i *is opinion holder.*

t_l *is the time when the opinion is expressed.*

Source: Sentiment Analysis and Opinion Mining, Bing Liu

OPINION MATHEMATICAL MODEL

An *opinion* is a quintuple

$(e_j, a_{jk}, so_{ijkl}, h_i, t_l)$

where

e_j is a target entity.

a_{jk} is an aspect/feature of the entity e_j

so_{ijkl} s the sentiment value of the opinion from the opinion holder *hi* on aspect a_{jk} of entity e_j at a time t_l. so_{ijkl} is +ve, -ve, or neutral, or a more granular rating.

h_i is opinion holder.

t_l *is the time when the opinion was expressed.*

STRUCTURE THE UNSTRUCTURED (HU AND LIU 2004)

Goal: Given an opinionated document
Discover all quintuples (e_j, a_{jk}, so_{ijkl}, h_i, t_l)
E.g., sentiment classification at the document or sentence level.

With the quintuples,
Unstructured Text → Structured Data

- ✓ Traditional data and visualization tools can be used to slice, dice and visualize the results.
- ✓ Enable qualitative and quantitative analysis.

5-1-2008 "*I bought an iPhone a few days ago. It is such a nice phone. The touch screen is really cool. The voice quality is clear too. It is much better than my old Blackberry, which was a terrible phone and so difficult to type with its tiny keys. However, my mother was mad with me as I did not tell her before I bought the phone. She also thought the phone was too expensive, …*"

In quintuples

(iPhone, GENERAL, +, Abc123, 5-1-2008)

(iPhone, touch_screen, +, Abc123, 5-1-2008)

….

SENTIMENT CLASSIFICATION: DOC-LEVEL

(PANG AND LEE, ET AL 2002 AND TURNEY 2002)

Classify a document (e.g., a review) based on the overall sentiment expressed by opinion holder

Classes: Positive, or negative (and neutral)

In the model, $(o_j, f_{jk}, so_{ijkl}, h_i, t_l)$

It assumes

- Each document focuses on a single object and contains opinions from a single opinion holder.
- It considers opinion on the object, o_j (or $o_j = f_{jk}$)

SENTIMENT CLASSIFICATION: SENTENCE-LEVEL

- Sentence-level sentiment analysis has two tasks:
  - Subjectivity classification: Subjective or objective.
    - Objective: e.g., *I bought an iPhone a few days ago.*
    - Subjective: e.g., *It is such a nice phone.*
  - Sentiment classification: For subjective sentences or clauses, classify positive or negative.
    - Positive: *It is such a nice phone.*

- However. (Liu, Chapter in NLP handbook)
  - subjective sentences ≠ +ve or –ve opinions
    - E.g., *I think he came yesterday.*
  - Objective sentence ≠ no opinion
    - Imply –ve opinion: *My phone broke in the second day.*

MINING IN REAL WORLD

FEATURE BASED SENTIMENT ANALYSIS

- ***Feature-Based Opinion Mining*** model, which is now also called ***Aspect-Based Opinion Mining***
- Sentiment classification at both at document and sentence levels are not enough.
- They do not tell what people like or dislike.
- A positive opinion on an object does not mean that the opinion holder like everything.
- A negative opinion on an object does not mean that the opinion holder dislike everything.
- With all quintuples various different analysis become possible and effective.

Source: Sentiment Analysis and Opinion Mining, Bing Liu

SUMMARY OF FEATURE BASED OPINION

"I bought an iPhone a few days ago. It is such a nice phone. The touch screen is really cool. The voice quality is clear too. It is much better than my old blackberry, which was a terrible phone and so difficult to type with its tiny keys. However my mother was mad with me as I did not tell her before I bought the phone. She also thought the phone was too expensive, and wanted me to return it to the shop. …"

Goal: *produce a feature-based opinion summary of multiple reviews*

Task 1: *Identify and extract object features that have been commented on by an opinion holder*
Task 2: *Determine polarity of opinions on features*
Task 3: *Group feature synonyms*

Summary of feature based opinion

Feature 1: Touch screen
Positive : 212
The touch screen was really cool.
The touch screen was so easy and can do amazing things.
…….
…….
Negative: 6
The screen is easily scratched.
I have a lot of difficulty in removing finger marks from the touch screen.
Touch screen sometimes does not respond.
……
Feature 2: voice quality
Positive : 17
……

Google products | sony camera | Search Products

Sony Cyber-shot DSC-W370 14.1 MP Digital Camera (Silver)

Overview - Online stores - Nearby stores - **Reviews** - Technical specifications - Similar items - Accessories

$140 online, $170 nearby

159 reviews +1 0

Reviews

Summary - Based on 159 reviews

1 | 2 | 3 stars | 4 stars | 5 stars

What people are saying

pictures	"We use the product to take quickly photos."
features	"Impressive panoramic feature."
zoom/lens	"It also record better and focus better on sunny days."
design	"It has the slightest grip but it's sufficient."
video	"Video zoom is choppy."
battery life	"Even better, the battery lasts long."
screen	"I Love the Sony's 3" screen which I really wanted."

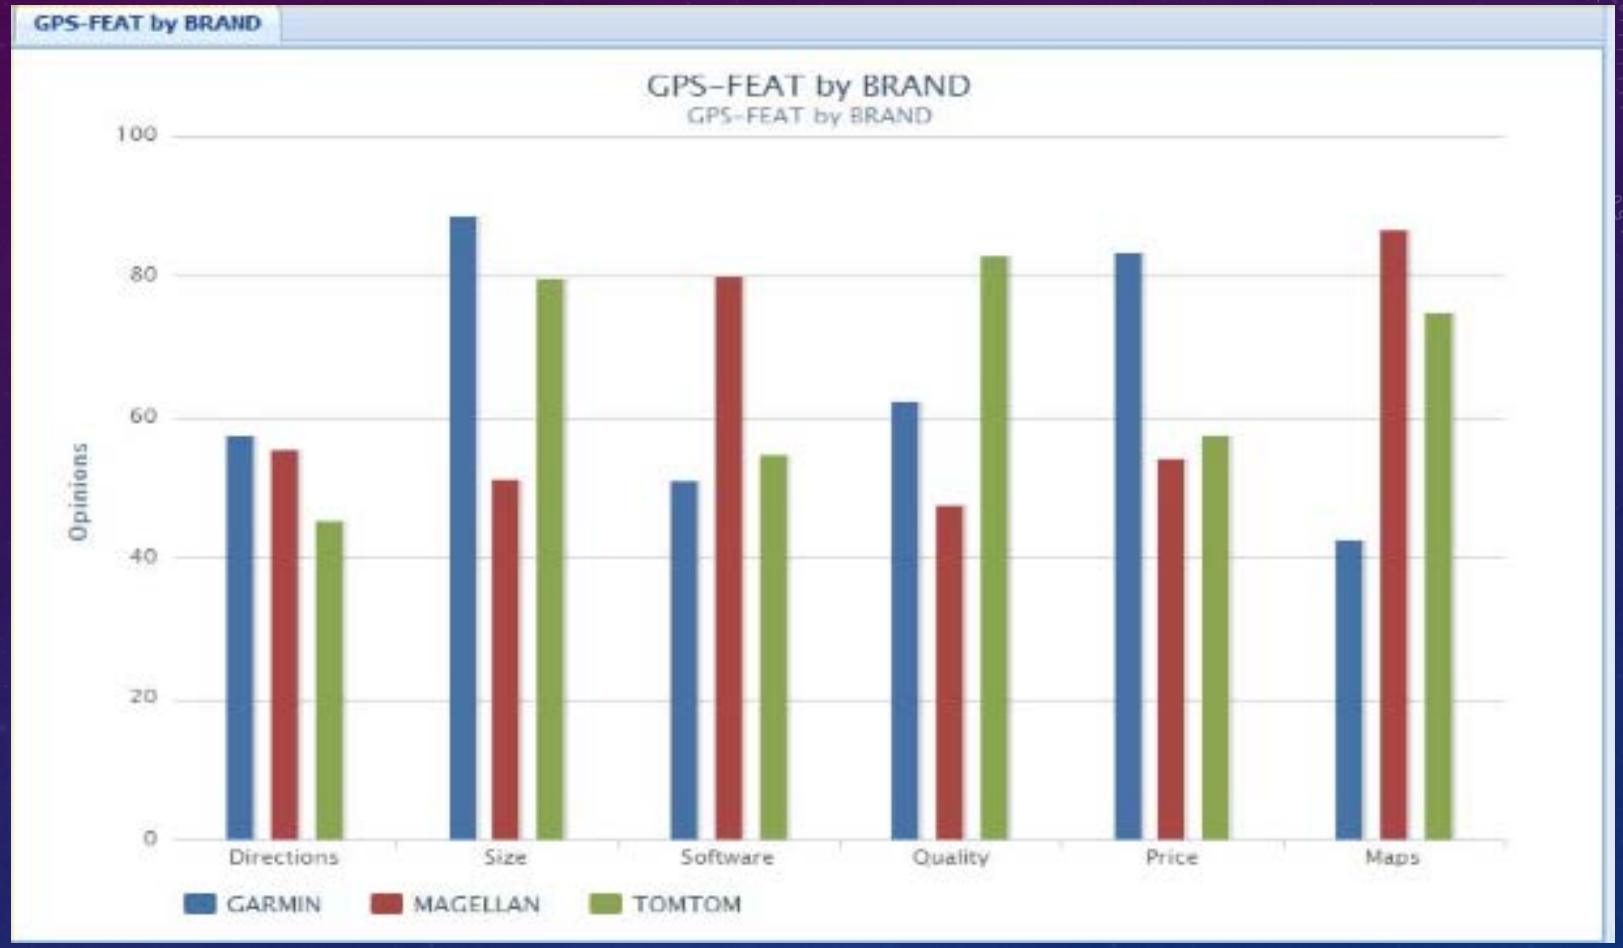
GPS-FEAT by BRAND
GPS-FEAT by BRAND
GPS-FEAT by BRAND
100
80
60
40
20
0
Opinions
Directions
Size
Software
Quality
Price
Maps
GARMIN
MAGELLAN
TOMTOM

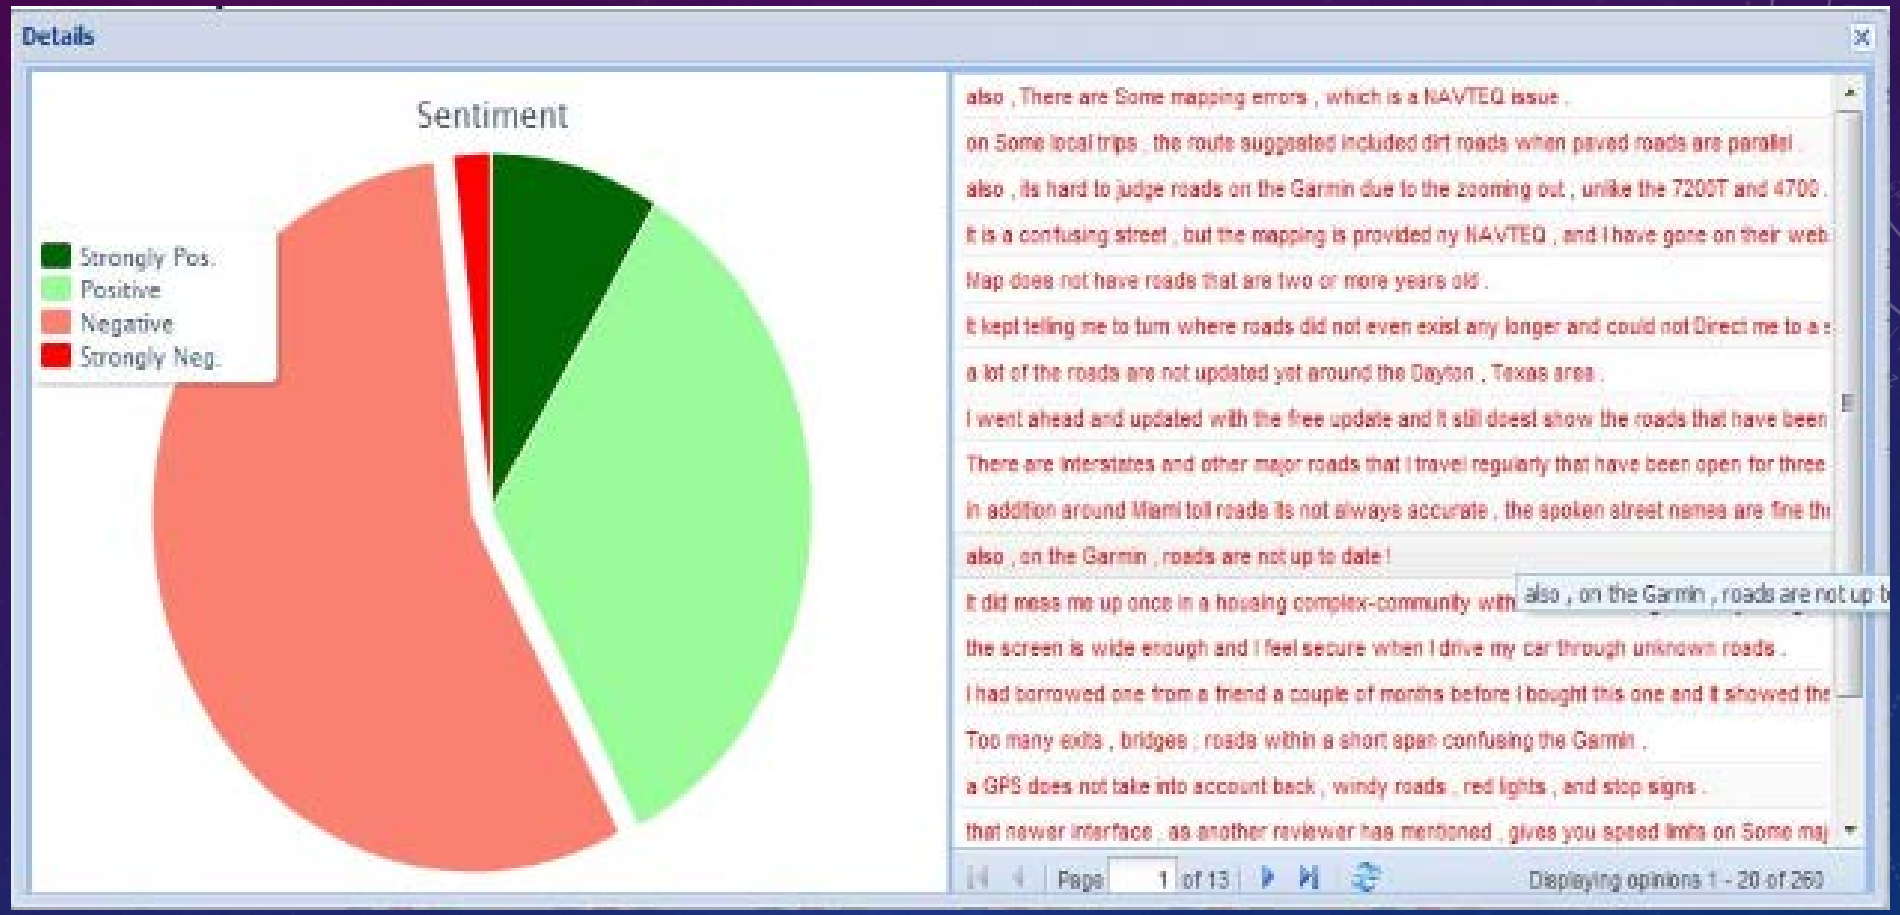
Details
Sentiment
Strongly Pos.
Positive
Negative
Strongly Neg.
also , There are Some mapping errors , which is a NAVTEQ issue .
on Some local trips , the route suggested included dirt roads when paved roads are parallel .
also , its hard to judge roads on the Garmin due to the zooming out , unlike the 7200T and 4700 .
It is a confusing street , but the mapping is provided ny NAVTEQ , and I have gone on their web
Map does not have roads that are two or more years old .
It kept telling me to turn where roads did not even exist any longer and could not Direct me to a
a lot of the roads are not updated yet around the Dayton , Texas area .
I went ahead and updated with the free update and it still doest show the roads that have been
There are Interstates and other major roads that I travel regularly that have been open for three
in addition around Miami toll roads its not always accurate , the spoken street names are fine th
also , on the Garmin , roads are not up to date !
It did mess me up once in a housing complex-community with
also , on the Garmin , roads are not up t
the screen is wide enough and I feel secure when I drive my car through unknown roads .
I had borrowed one from a friend a couple of months before I bought this one and it showed the
Too many exits , bridges , roads within a short span confusing the Garmin .
a GPS does not take into account back , windy roads , red lights , and stop signs .
that newer interface , as another reviewer has mentioned , gives you speed limits on Some maj
Page 1 of 13
Displaying opinions 1 - 20 of 260

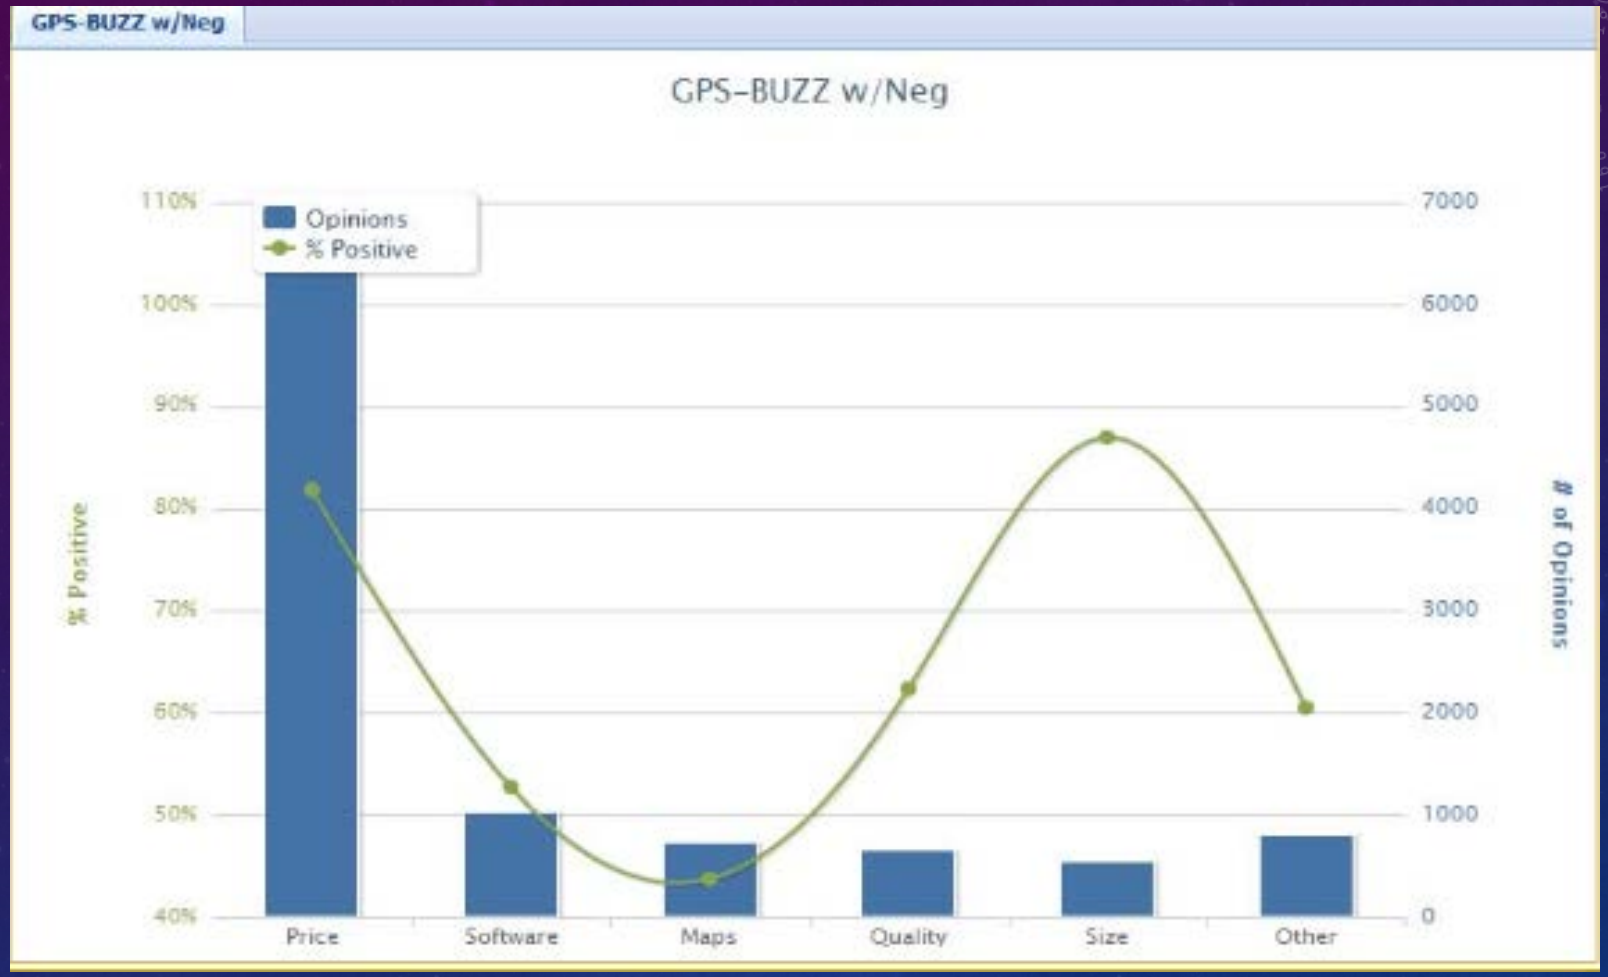
GPS-BUZZ w/Neg
GPS-BUZZ w/Neg
Opinions
% Positive
110%
100%
90%
80%
70%
60%
50%
40%
% Positive
7000
6000
5000
4000
3000
2000
1000
0
# of Opinions
Price
Software
Maps
Quality
Size
Other


GPS-TREND w/Neg
GPS-TREND w/Neg
Opinions
% Negative
32.5%
30%
27.5%
25%
22.5%
20%
% Negative
1000
800
600
400
200
0
# of Opinions
3/1/2009 - 4/1/2009
4/1/2009 - 5/1/2009
5/1/2009 - 6/1/2009
6/1/2009 - 7/1/2009
7/1/2009 - 8/1/2009
8/1/2009 - 9/1/2009
9/1/2009 - 10/1/2009
10/1/2009 - 11/1/2009
11/1/2009 - 12/1/2009
12/1/2009 - 1/1/2010
1/1/2010 - 2/1/2010
2/1/2010 - 2/28/2010

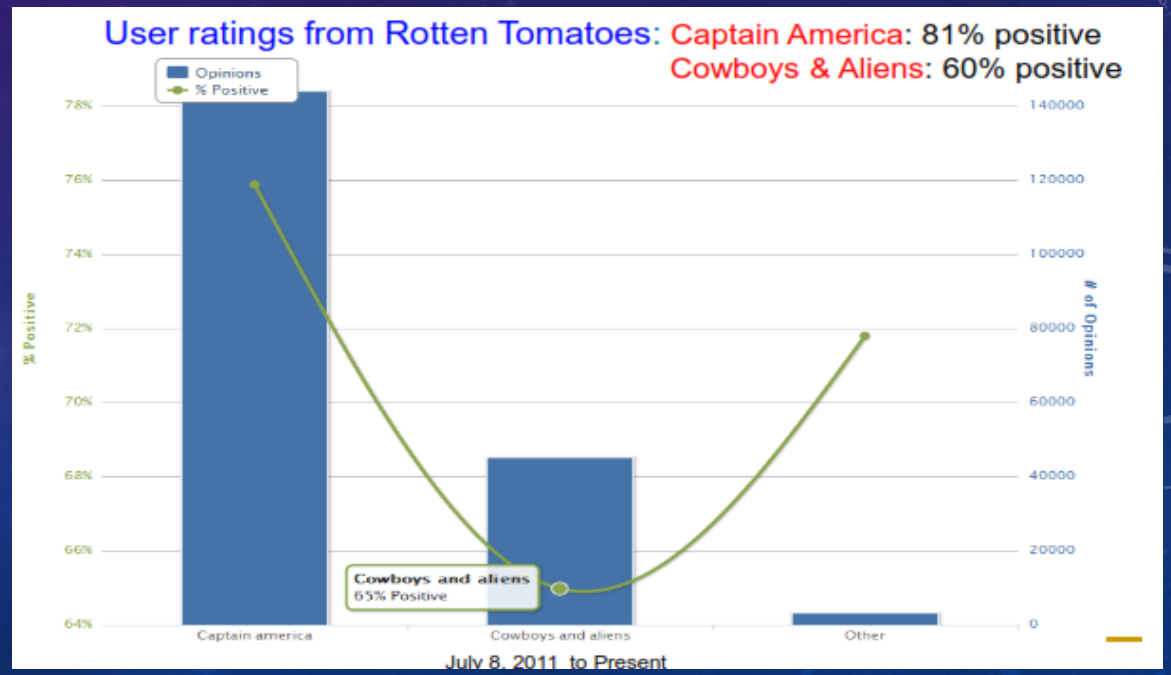
User ratings from Rotten Tomatoes: Captain America: 81% positive
Cowboys & Aliens: 60% positive
Opinions
% Positive
78%
76%
74%
72%
70%
68%
66%
64%
% Positive
140000
120000
100000
80000
60000
40000
20000
0
# of Opinions
Cowboys and aliens
65% Positive
Captain america
Cowboys and aliens
Other
July 8, 2011 to Present

LEVELS OF OPINION MINING

- Tweets from Twitter are probably the easiest, they are short and thus usually straight to the point
- Reviews are next, entities are 90% given and there is little noise
- Discussions, comments, and blogs are hard,Multiple entités, comparaisons, noisy, sarcasm, etc.
- Determining sentiments seems to be easier.
- Extracting entities and aspects is harder.
- Combining them is even harder.

Source: Sentiment Analysis and Opinion Mining, Bing Liu

SENTIMENT ANALYSIS IS DIFFICULT

"This past Saturday, I bought a Nokia phone and my girlfriend bought a Motorola phone with Bluetooth. We called each other when we got home. The voice on my phone was not so clear, worse than my previous Samsung phone. The battery life was short too. My girlfriend was quite happy with her phone. I wanted a phone with good sound quality. So my purchase was a real disappointment. I returned the phone yesterday."

Entity Extraction

Information Extraction

Sentiment Extraction

Data Extraction

Co-reference resolution

Relation extraction

Synonyms match

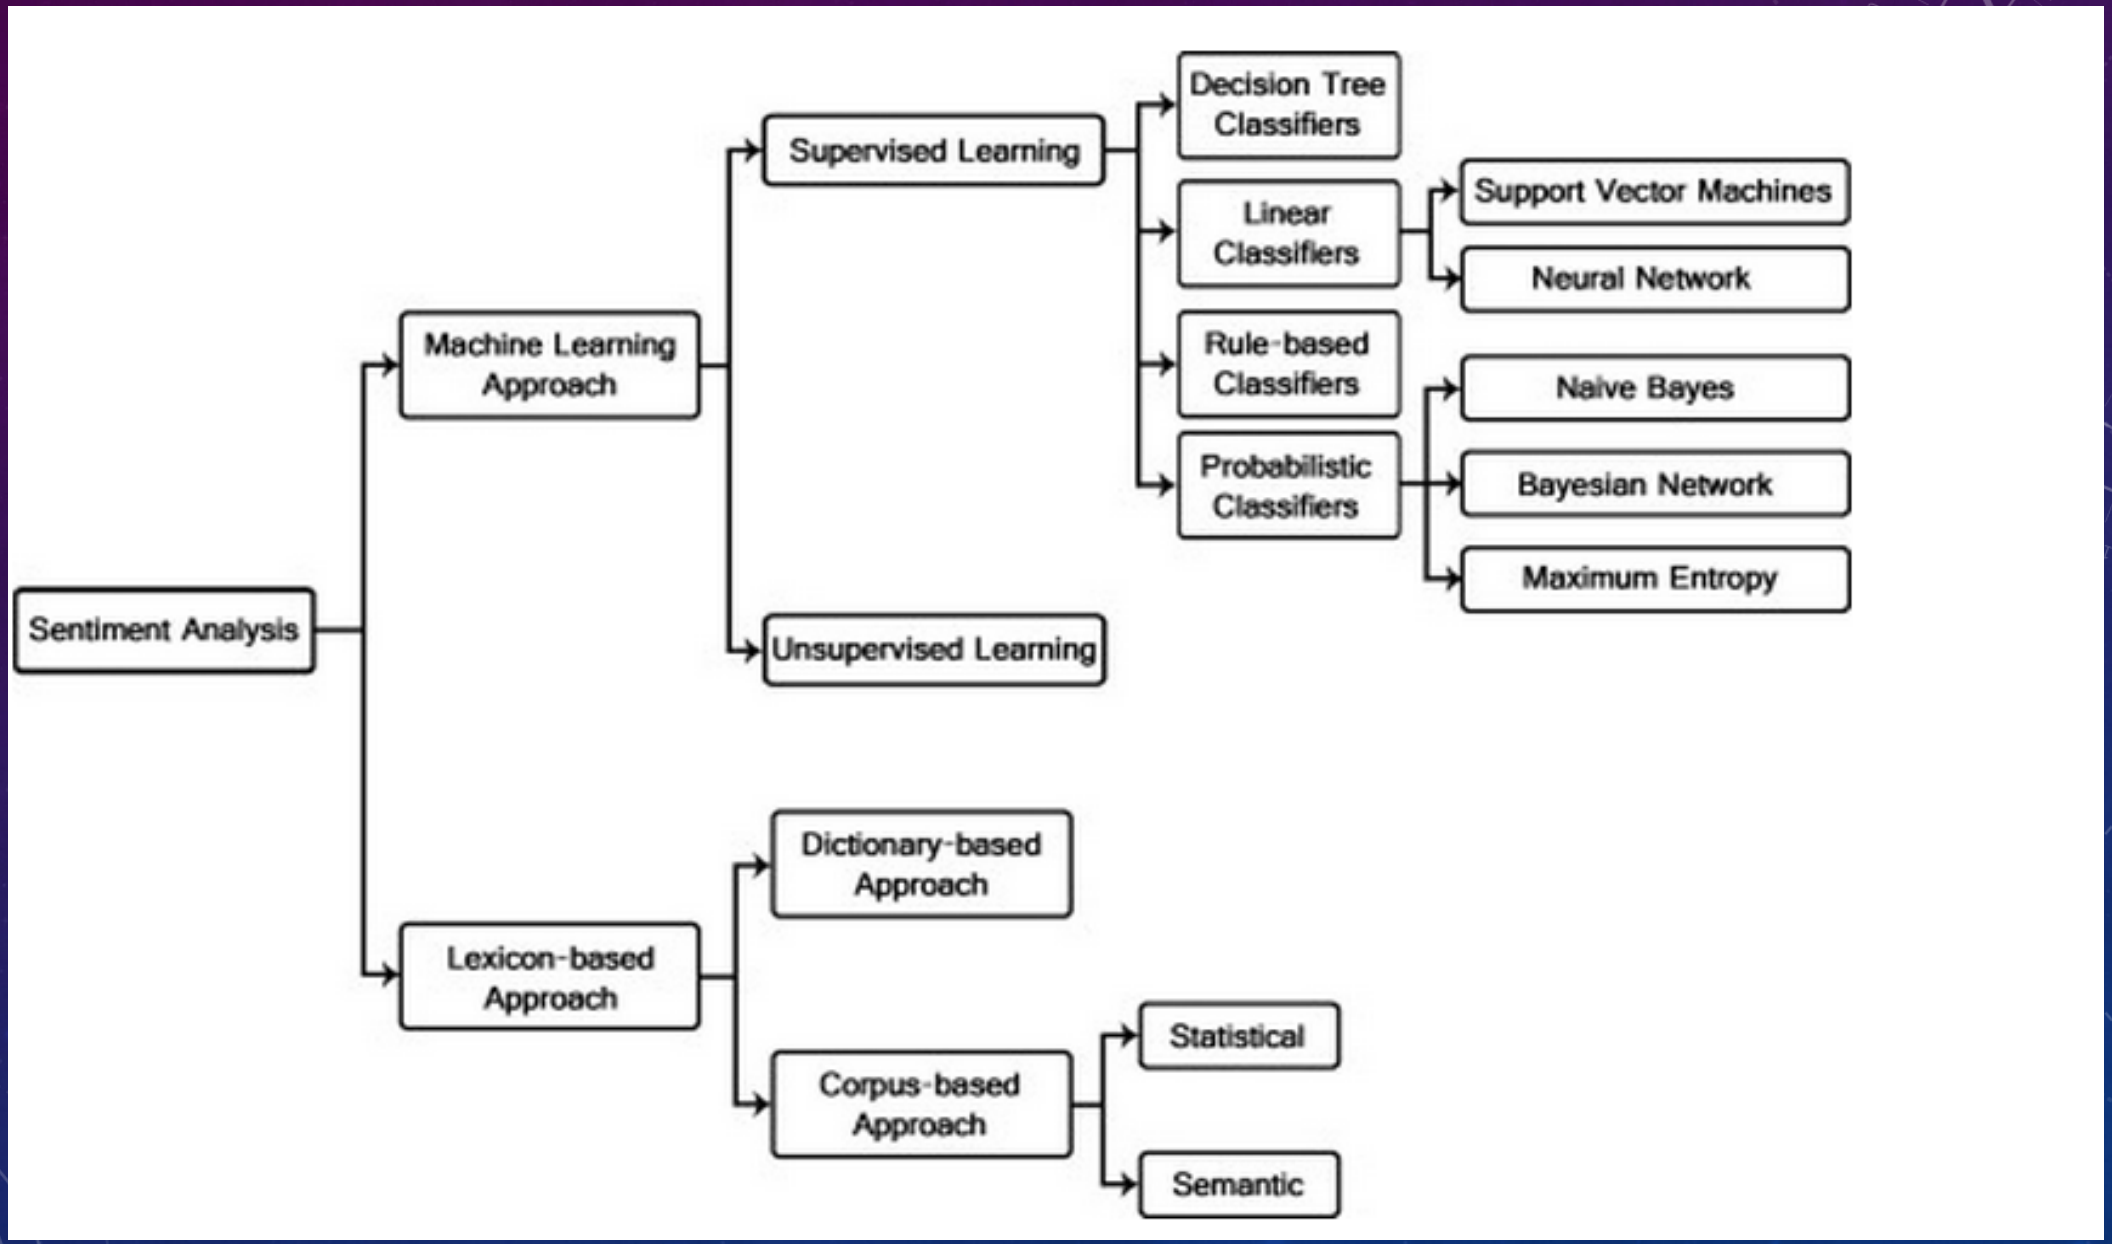

Source: Sentiment Analysis Algorithms and Applications
Walaa Medhat Ahmed Hassan Hoda Korashy,2013

MACHINE LEARNING APPROACH

POSITIVE POLARITY (GOOD)

-a mesmerizing cinematic poem from the first frame to the last ."
-a well-put-together piece of urban satire .
-one can't deny its seriousness and quality .
-hard to resist .
-a naturally funny film , home movie makes you crave chris smith's next movie .
-a true-blue delight .
-a fun ride .
-a surprisingly funny movie .
-the script is smart and dark - hallelujah for small favors .
-a flick about our infantilized culture that isn't entirely infantile .

-unfortunately the story and the actors are served with a hack script .
-too slow for a younger crowd , too shallow for an older one .
-terminally brain dead production .
-one lousy movie .
-this movie . . . doesn't deserve the energy it takes to describe how bad it is .
-a cleverly crafted but ultimately hollow mockumentary .
-it's an 88-minute highlight reel that's 86 minutes too long .
-the whole affair is as predictable as can be .

NEGATIVE POLARITY (BAD)

POSITIVE POLARITY (GOOD)

-b!nftnfsj{joh!djofnbujd!qpfn!gspn!uif!gjstu!gsbnf!up!uif!mbtu!/
-b!xfmm-qvu-uphfuifs!qjfdf!pg!vscbo!tbujsf!/
-pof!dbo(u!efoz!jut!tfsjpvtoftt!boe!rvbmjuz!/
-ibse!up!sftjtu!/
-b!obuvsbmmz!gvooz!gjmn ! ipnf!npwjf!nblft!zpv!dsbwf!disjt!tnjui(t!ofyu!npwjf!/
-b!usvf-cmvf!efmjhiu!/
-b!gvo!sjef!
-b!tvsqsjtjohmz!gvooz!npwjf!/
-uif!tdsjqu!jt!tnbsu!boe!ebsl! - !ibmmfmvkbi!gps!tnbmm!gbwpst!/
-b!gmjdl!bcpvu!pvs!jogboujmj{fe!dvmuvsf!uibu!jto(u!foujsfmz!jogboujmf!/!

-vogpsuvobufmz!uif!tupsz!boe!uif!bdupst!bsf!tfswfe!xjui!b!ibdl!tdsjqu!/
-upp!tmpx!gps!b!zpvohfs!dspxe!-!upp!tibmmpx!gps!bo!pmefs!pof!/
-ufsnjobmmz!csbjo!efbe!qspevdujpo!/
-pof!mpvtz!npwjf!/!
-uijt!npwjf!/!/!/!epfto(u!eftfswf!uif!fofshz!ju!ublft!up!eftdsjcf!ipx!cbe!ju!jt!/!
-b!dmfwfsmz!dsbgufe!cvu!vmujnbufmz!ipmmpx!npdlvnfoubsz!/!
-ju(t!bo!99-njovuf!ijhimjhiu!sffm!uibu(t!97!njovuft!upp!mpoh!/!
-uif!xipmf!bggbjs!jt!bt!qsfejdubcmf!bt!dbo!cf!/

NEGATIVE POLARITY (BAD)

“it's rather like a lifetime special -- pleasant , sweet and forgettable . “

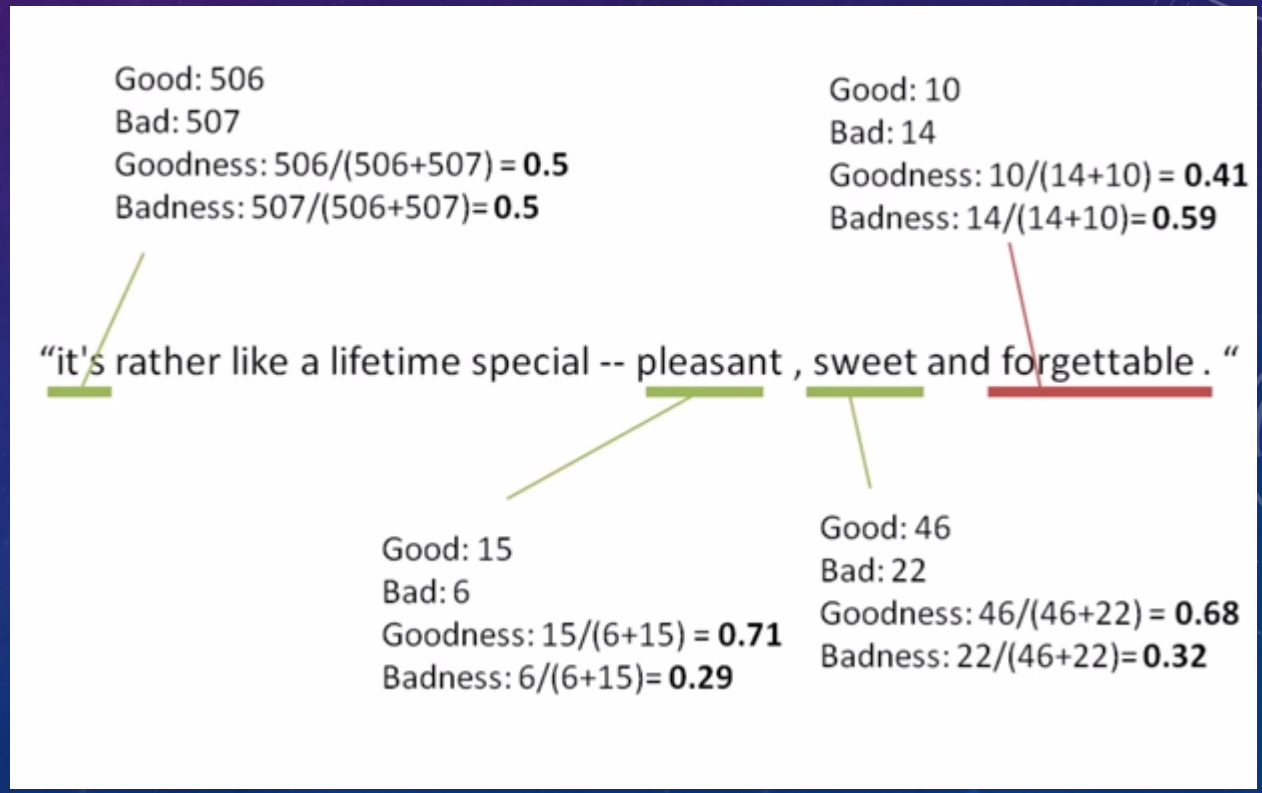

"it's rather like a lifetime special -- pleasant , sweet and forgettable . "

	#GOOD	#BAD	GOODNESS	BADNESS
it's	506	507	0.5	0.5
rather	42	63	0.4	0.6
like	242	396	0.61	0.39
a	3446	3112	0.53	0.47
lifetime	3	5	0.38	0.62
special	29	40	0.42	0.58
pleasant	15	6	0.71	0.29
sweet	46	22	0.68	0.32
and	3198	2371	0.57	0.43
forgettable	10	14	0.42	0.58

SUM: 5.22 4.8

So we should classify
this as a POSITIVE review!

TOP 10 BEST/WORST:

riveting	unfunny
gem	badly
engrossing	poorly
vividly	flat
wonderfully	bore
polished	pointless
lively	offensive
heartwarming	plodding
startling	product
spare	disguise

“this movie makes Catwoman look like a great movie.”

“a terrible movie that some people will nevertheless find moving .”

“well-made but mush-hearted .”

“your children will be occupied for 72 minutes . “

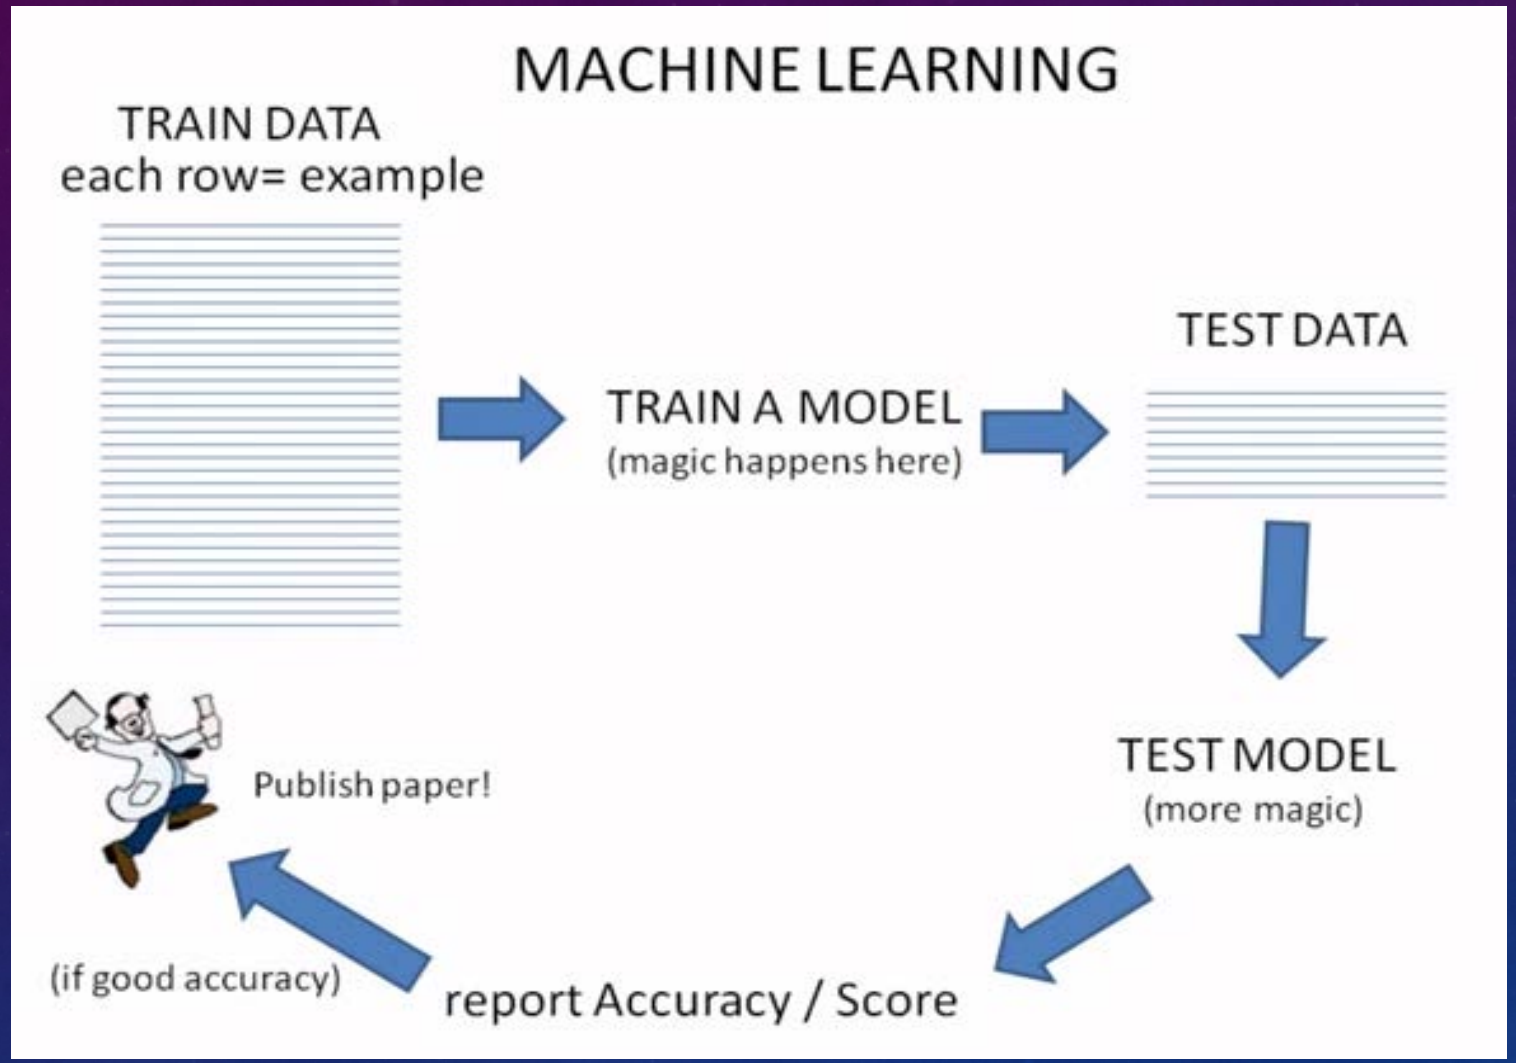
MACHINE LEARNING
TRAIN DATA
each row= example
TRAIN A MODEL
(magic happens here)
TEST DATA
TEST MODEL
(more magic)
report Accuracy / Score
(if good accuracy)
Publish paper!

CHALLENGES IN OPINION MINING

- Polarity terms are context sensitive.
  - Ex. Small can be good for ipod size, but can be bad for LCD monitor size.
  - Even in the same domain, use different words depending on target feature.
    - Ex. Long ‘ipod’ battery life vs. long ‘ipod’ loading time
  - Partially solved (query dependent sentiment classification)
- Implicit and complex opinion expressions
  - Rhetoric expression, metaphor, double negation.
  - Ex. The food was like a stone.
  - Need both good IR and NLP techniques for opinion mining.
- Cannot divide into positive/negative clearly
  - Not all opinions can be classified into two categories
  - Interpretation can be changed based on conditions.
  - Ex. 1) The battery life is ‘long’ if you do not use LCD a lot. (positive)
    2) The battery life is ‘short’ if you use LCD a lot. (negative)

For each battleground state every week, the campaign's call centers conducted 5,000 to 10,000 so-called short-form interviews that quickly gauged a voter's preferences, and 1,000 interviews in a long-form version that was more like a traditional poll.

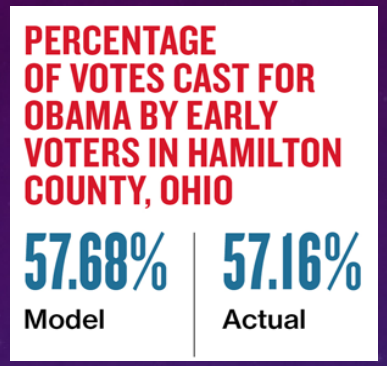

Unilever Dove-brand pro.age campaign Catherine Cardoso, Associate Insights Manager at Unilever, says "We were very pleased with the results and the depth of insight. The results were helpful beyond understanding reactions to our campaign. We also gained an understanding of what motivates people on discussion boards, which issues are most important to women in our target group, and how to create better products and messaging for them.

Cisco "used the sentiment engine to determine which executives have the highest correlation to positively moving the stock price when they deliver positive news. They found that certain executives had a positive influence on the markets, while others actually had a negative influence because of the tone of their delivery."

REFERENCES

1. B. Liu, “Sentiment Analysis and Subjectivity.” A Chapter in Handbook of Natural Language Processing, 2nd Edition, 2010. (An earlier version) B. Liu, “Opinion Mining”, A Chapter in the book: Web Data Mining, Springer, 2006.
2. Sentiment analysis algorithms and applications: A survey, Walaa Medhat, Ahmed Hassan, Hoda Korashy, 2013.
3. Opinion Mining and Sentiment Analysis, Bo Pang and Lillian Lee, 2008.
4. Lecture from Dr. Sreerama K. Murthy Ph.D., Johns Hopkins Univ.
5. A. Agarwal and P. Bhattacharyya, "Sentiment analysis: A new approach for effective use of linguistic knowledge and exploiting similarities in a set of documents to be classified," in *Proceedings of the International Conference on Natural Language Processing (ICON)*, 2005.
6. White paper on Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis by Theresa Wilson, Janyce Wiebe and Paul Hoffmann

QUESTIONS ?

THANK YOU


THE MRS. OLSEN MUG SALE HAS BEEN GREEN-LIGHTED!
SHE SAID YES?
SHE DIDN'T SAY NO.
ARE YOU SURE NOT-NO MEANS YES?
FROM SOMEONE THAT NEGATIVE? UH, YEAH.
BE SURE.
SURE ENOUGH TO GO BUY SUPPLIES.
I MIGHT NOT HAVE NOT ALREADY BOUGHT SOME.