ACL 2012

Multilingual Sentiment and Subjectivity Analysis

Rada Mihalcea, University of North Texas
Carmen Banea, University of North Texas
Janyce Wiebe, University of Pittsburgh

What is subjectivity and sentiment analysis?

- Subjectivity and sentiment analysis focuses on the automatic identification of private states in natural language (Wiebe et al., 2005)

“I love Jeju Island.”

Subjectivity Analysis	Subjective			Objective
Sentiment Analysis	Positive	Negative	Neutral	---

Madagascar - Rotten Tomatoes - Firefox

File Edit View Go Bookmarks Tools Help

http://www.rottentomatoes.com/m/madagascar/ Go

https://www.onerailw... Getting Started Latest BBC Headlines

CLIPS & TRAILERS

View All 14 Clips & Trailers for Madagascar

VIDEO

Kim Williamson
BOXOFFICE MAGAZINE

Dreamworks lets their computers do the talking in a fine cartoon freed of the weight of a significant message.

FULL REVIEW Comments

Ron Wilkinson
MONSTERS AND CRITICS

DreamWorks set the bar awfully low last fall with *Shark Tale*, but like a limbo champion it slips lower still with *Madagascar*.

FULL REVIEW Comments

Jeffrey Westhoff
NORTHWEST HERALD (CRYSTAL LAKE, IL)

Lacking any meat on its bones, ['Madagascar'] is a cartoon about nothing. Not that there's anything wrong with that.

FULL REVIEW Comments

Joe Williams
ST. LOUIS POST-DISPATCH

***Madagascar* is content with being a genial kids' movie, cute, sporadically clever, rarely memorable.**

FULL REVIEW Comments

Glenn Whipp
LOS ANGELES DAILY NEWS

Nearly all of

Philip Wuntch
DALLAS MORNING NEWS

Brilliantly written, animated and performed, it not only bridges th child-adult divide, but also makes its friendly humor smart in its own way

FULL REVIEW Comm

Bruce Westbrook
HOUSTON CHRONICLE

It's a good-humored, pleasant confection that has all kinds of relaxed fun bringing computer-anim savvy to the old-fashioned world of Looney Tunes cartoons.

FULL REVIEW Comm

Kenneth Turan
LOS ANGELES TIMES

The movie sadly runs aground o old gags that no amount of eye-popping color and frenzied action ca spin into gold.

FULL REVIEW Comm

Peter Travers

http://www.rottentomatoes.com/click/movie-1146129/reviews_viewer.php?rid=1437676&fb=no

File Edit View Go Bookmarks Tools Help
http://news.auto.cz/
Go
https://www.onerailw...
Getting Started
Latest BBC Headlines
KTERÉ AUTO JE LEPŠÍ?
zkušenosti majitelů vozů
diskuze
hodnocení
moje.auto.cz
Hledat: Články | Tapety | Videa | Mobil | Autobazar | Zboží | Prodejní akce | Blogy
Najít na AUTO.CZ
Podrobné hledání
Jak hledat
AUTO.CZ > NEWS > Úvodní strana
nepřihlášený | přihlásit
česky slovensky polski
HomePage AUTO.CZ
Nová Škoda Fabia
Přehled dle značky
Autobazar
Akční nabídky
NEWS
Úvodní strana
BRNO 2007
Ženeva 2007
Spy photos
Nové modely
Aktuality
Jízdní dojmy
Tuning
Technika
Bezpečnost
Co se píše v tisku
Spy Photos: Nový Mercedes-Benz třídy E na Grossglockneru
Kde jinde testovat zimní podmínky, když panují třiceti stupňová vedra napříč Evropou, aniž bychom museli za polární kruh? Kde jinde vyzkoušet účinnost, nastavení a zejména vadnutí brzd? Kde jinde testovat chlazení při několikakilometrovém stoupání?
19 příspěvků v diskuzi
Spy photos | 25.06. 2007, 11:43
Přidejte zkušenosti s vaším autem, nebo hodnoťte auta ostatních na moje.auto.cz
již od 339 900 Kč.
Nepřehlédněte
Inzerujte zdarma na AUTO.CZ
Svět počítačů
CD ZDARMA Vyzkoušejte si Visty
Test 22" LCD monitorů
Done
NEWS.auto.cz - Vše o autech na jednom místě - Firefox

File Edit View Go Bookmarks Tools Help

http://vip.bokee.com/ Go blogs china

https://www.onerailw... Getting Started Latest BBC Headlines

博客中国

每天五分钟 为思想加油

www.blogchina.com

每天5分钟 给思想加油

设为首 申请专 收藏本

热点话题 更多

谢东首次回应"侯家私生子"传言
张纪中怒斥大小龄童欲重拍《西游记》
5岁艾滋男童食道堵塞 医院拒绝手术
突发事件应对法删除媒体不得擅发信息
南京水库万斤活鱼遭数百民众哄抢 图

突发事件让媒体说话，天塌不下来

6月24日提交全国人大常委会二审的《突发事件应对法(草案)》，立法信号表明，媒体的监督作用正发挥越来越大的作用，且其中删除有关新闻媒体不得"违规擅自发布"突发事件信息的规定，表明了一个可喜的进步。[全文]

[评论|突发事件报道禁令,形式上进步也不错|禁令是否保障公民基本权利]

高考热过之后议"状元"

高考状元的出炉后成了大街小巷的老百姓嘴边的噱头，那些高考状元有理由也应该被众人推上神坛，成为众人顶礼膜拜的对象。但对高考状元轰轰烈烈地大势宣传会否不妥？网友们各持观点...[全文]

[评论|高考状元是我们这个时代的象征|高考状元，中国教育的"怪胎"]

[相关|填志愿应给学生更多空间|过度关注高考事与愿违|高考落榜后的出路]

[生活] 资深猎头老总：我不忍扼杀失信学生的未来 王常江 19
[社会] "电梯房"民主实验的失败有价值 鲁宁 0
[管理] 小心"当场击毙"击毙程序正义 熊培云 6

焦点追踪

香港回归十周年专题征文

为庆祝香港回归十周年，举《香港回归十年》、《我看港》征文活动。[详细]

网友："香港对祖国的认有了根本转变！"

我们需要怎样的新儒

专访新儒家陈明[专题][博客

一定要尊重日本

互动反馈

blog问答 专栏

欢迎提意见 与读者意见 流[详细]

Waiting for vip.bokee.com...

博客中国——汇聚万名意见领袖，中国第一思维集散地 | 每天5分钟，给思想加油 - Firefox

File Edit View Go Bookmarks Tools Help

http://www.roon.ro/ Go stiri romania

https://www.onerailw... Getting Started Latest BBC Headlines

ROMÂNIA ONLINE www.roon.ro

felicitari.roon.ro

SERVICII

25.6.107, 19:13

Ma	Mi	J	V
26.	27.	28.	29.
33°	28°	20°	25°

Statii meteo

CURS VALUTAR

Cotatiile in lei noi:

EURO	3,218
USD	2,400
100HUF	1,287

Altele... · Convertor

OFERTE

publitim.ro

Toate categoriile

Subiecte modificate

Subiectele pentru bacalaureat au fost modificate cu doar câteva zile înainte de începerea examenelor, deşi au fost declarate definitive încă din februarie. Vreme de patru luni, elevii au învăţat pe subiecte greşite. Bacalaureatul 2007, care a început luni, anunţă şi în acest an un potenţial scandal.

- **Testele naţionale vor fi desfiinţate**

Ameninţat cu moartea!
Jurnalistul BBC răpit, Alan Johnston, este ameninţat cu moartea prin detonare.

Cod portocaliu în 8 judeţe
Temperaturile ar putea depăşi 40 de grade marţi după-amiază în sudul

Focare de gripă aviară!
În Nurnberg, au fost descoperite 8 păsări moarte purtătoare ale virusului H5N1.

DIVERTISMENT

Fa ROON pagina ta de start

Pozele tale!

Blog NOU

GOLD LINK

Blogreview – revista blogosferei sau ce scriu bloggerii români

Încarcă filmele tale sau vizionează clipuri haioase

Sudoku, jocuri arcade. Intră în pagina de jocuri şi relaxează-te!

Vrei să fii auzit? Dezbate subiectele fierbinţi pe Forum

Maşini tunate, cai putere, fete frumoase in pagina de Tuning

Done

Ştiri Romania, portal de stiri din Romania, stiri nationale, informaţii din Romania pe Romania Online - Firefox

Top Ten Languages on the Web

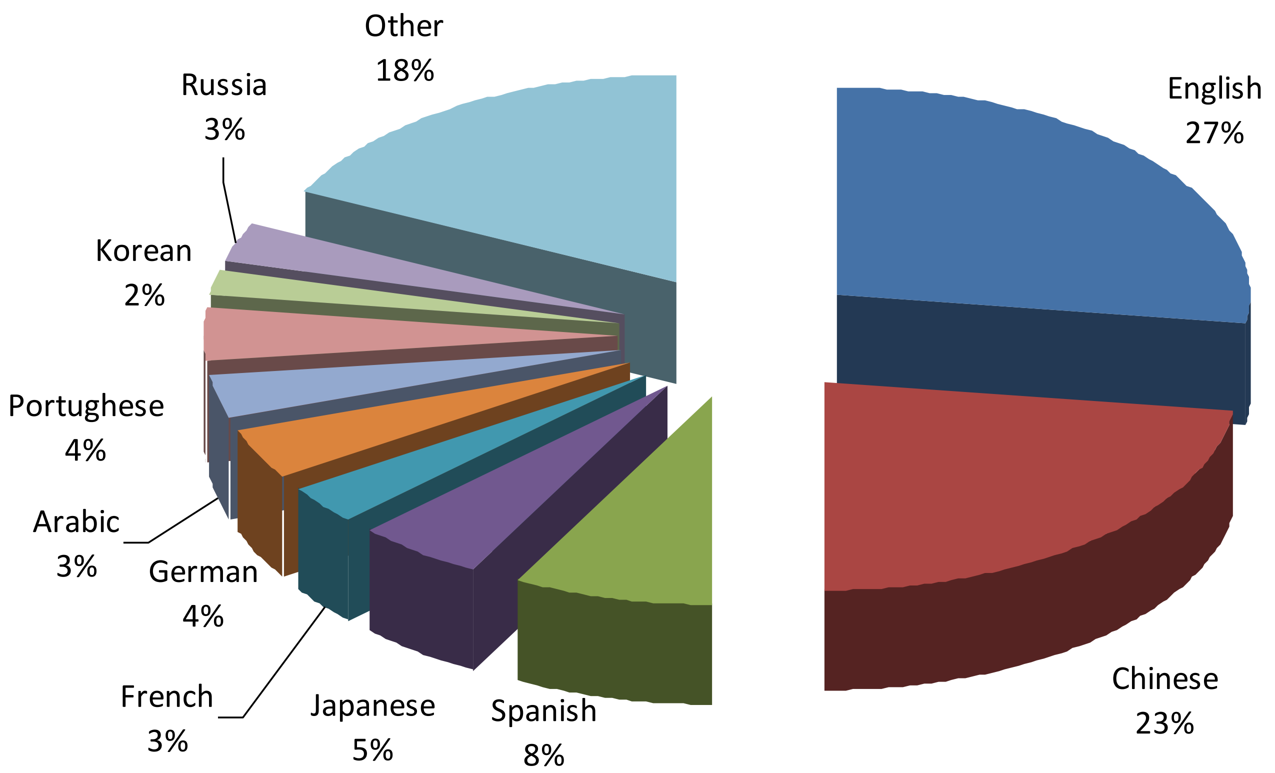

internetworldstats.com, March, 2011

Overview

- I. Sentiment and subjectivity analysis
  - Definitions, Applications
- II. Sentiment and subjectivity analysis on English
  - Lexicons, Corpora, Tools
- III. Word- and phrase-level annotations
- IV. Sentence level annotations
- V. Document level annotations
- VI. What works, what doesn't

Some slides have been adapted from tutorials/lectures given by Carmen Banea, Bing Liu, Janyce Wiebe

I. Sentiment and subjectivity analysis

Definitions & Applications

What is subjectivity?

- The **linguistic** expression of somebody's **opinions**, **sentiments**, emotions, evaluations, beliefs, speculations ***(private states)***

- **Private state:** state that is not open to objective observation or verification

  Quirk, Greenbaum, Leech, Svartvik (1985). *A Comprehensive Grammar of the English Language*.

- Subjectivity analysis classifies content in **objective** or **subjective**

Examples

- The desire to give Broglio as many starts as possible.
- The Pirates have a 9-6 record this year and the Redbirds are 7-9.
- Suppose he did lie beside Lenin, would it be permanent ?
- One of the obstacles to the easy control of a 2-year old child is a lack of verbal communication.

Examples

- It offers a breath of the fresh air of true sophistication.
- This is a thoughtful, provocative, insistently humanizing film.
- The movie is a sentimental mess that never rings true.
- While the performances are often engaging, this loose collection of largely improvised numbers would probably have worked better as a one-hour TV documentary.

Application: Product Review Mining

- Sleek and well designed, the iPhone remains the best touchscreen phone that you can buy. We doubt FaceTime will be a big draw, but the excellent quality photos and videos are impressive, as are the new iOS 4 features.
- I love it. Coming from a 3GS u can see the difference in display pix and games and movies videos the list can go on :)
- After all, it's not a bad phone but hey, it doesn't worth the price tag. Really overrated, it lacks basic features, the platform is very closed and restrictive.
- It costs two times more than models produced by another companies. Also I think that apple phones can′t be tuned well because of lack of settings and huge amount of restrictions.

Application: Opinion Question Answering

Q: What is the international reaction to the reelection of Robert Mugabe as President of Zimbabwe?

A: African observers **generally approved** of his victory while Western Governments **strongly denounced** it.

Opinion QA is more complex
Automatic subjectivity analysis can be helpful
Stoyanov, Cardie, Wiebe EMNLP05
Somasundaran, Wilson, Wiebe, Stoyanov ICWSM07

Application: Information Extraction

"**The Parliament** exploded into fury against the government when word leaked out…"

Observation: subjectivity often causes false hits for IE

Goal: augment the results of IE

Subjectivity filtering strategies to improve IE Riloff, Wiebe, Phillips AAAI05

More applications

- **Product feature review :** What features of the ThinkPad T43 do customers like and which do they dislike?
- **Review classification:** Is a review positive or negative toward the movie?
- **Tracking sentiments toward topics over time:** Is anger ratcheting up or cooling down?
- **Prediction (election outcomes, market trends):** Will Clinton or Obama win?
- **Expressive text-to-speech synthesis**
- **Text semantic analysis** (Wiebe and Mihalcea, 2006) (Esuli and Sebastiani, 2006)
- **Text summarization** (Carenini et al., 2008)

What is sentiment analysis?

- Also known as opinion mining
- Attempts to identify the opinion/sentiment that a person may hold towards an object
- It is a finer grain analysis compared to subjectivity analysis

Sentiment Analysis	Subjectivity analysis
Positive	Subjective
Negative	
Neutral	Objective

Components of an opinion

- Basic components of an opinion:
  - Opinion holder: The person or organization that holds a specific opinion on a particular object.
  - Object: on which an opinion is expressed
  - Opinion: a view, attitude, or appraisal on an object from an opinion holder.

Opinion mining tasks

- At the document (or review) level:
  - Task: sentiment classification of reviews
  - Classes: positive, negative, and neutral
  - Assumption: each document (or review) focuses on a single object (not true in many discussion posts) and contains opinion from a single opinion holder.
- At the sentence level:
  - Task 1: identifying subjective/opinionated sentences
    - Classes: objective and subjective (opinionated)
  - Task 2: sentiment classification of sentences
    - Classes: positive, negative and neutral.
    - Assumption: a sentence contains only one opinion; not true in many cases.
    - Then we can also consider clauses or phrases.

Opinion mining tasks

- At the feature level:
  - *Task 1: Identify and extract object features that have been commented on by an opinion holder (e.g., a reviewer).*
  - *Task 2: Determine whether the opinions on the features are* positive, negative or neutral.
  - *Task 3: Group feature synonyms.*
    - Produce a feature-based opinion summary of multiple reviews.
- Opinion holders: identify holders is also useful, e.g., in news articles, etc, but they are usually known in the user generated content, i.e., authors of the posts.

Facts and Opinions

- Two main types of textual information on the Web.
  - Facts and Opinions
- Current search engines search for facts (assume they are true)
  - Facts can be expressed with topic keywords.
- Search engines do not search for opinions
  - Opinions are hard to express with a few keywords
    - How do people think of Motorola Cell phones?
  - Current search ranking strategy is not appropriate for opinion retrieval/search.

Applications

- Businesses and organizations:
  - product and service benchmarking.
  - market intelligence.
  - Business spends a huge amount of money to find consumer sentiments and opinions.
    - Consultants, surveys and focused groups, etc
- Individuals: interested in other's opinions when
  - purchasing a product or using a service,
  - finding opinions on political topics
- Ads placements: Placing ads in the user-generated content
  - Place an ad when one praises a product.
  - Place an ad from a competitor if one criticizes a product.
- Opinion retrieval/search: providing general search for opinions.

Two types of evaluations

- Direct Opinions: sentiment expressions on some objects, e.g., products, events, topics, persons.
  - E.g., "the picture quality of this camera is great"
  - Subjective
- Comparisons: relations expressing similarities or differences of more than one object. Usually expressing an ordering.
  - E.g., "car x is cheaper than car y."
  - Objective or subjective.

II. Sentiment and subjectivity analysis on English

Lexicons, Corpora, Tools

Main resources

- Lexicons
  - General Inquirer (Stone et al., 1966)
  - OpinionFinder lexicon (Wiebe & Riloff, 2005)
  - SentiWordNet (Esuli & Sebastiani, 2006)

- Annotated corpora
  - MPQA corpus (Wiebe et. al, 2005)
  - Used in statistical approaches (Hu & Liu 2004, Pang & Lee 2004)

- Tools
  - Algorithm based on minimum cuts (Pang & Lee, 2004)
  - OpinionFinder (Wiebe et. al, 2005)

Main resources

- Lexicons
  - General Inquirer (Stone et al., 1966)
  - OpinionFinder lexicon (Wiebe & Riloff, 2005)
  - SentiWordNet (Esuli & Sebastiani, 2006)

- Annotated corpora
  - Used in statistical approaches (Hu & Liu 2004, Pang & Lee 2004)
  - MPQA corpus (Wiebe et. al, 2005)

- Tools
  - Algorithm based on minimum cuts (Pang & Lee, 2004)
  - OpinionFinder (Wiebe et. al, 2005)

Lexicons: who does lexicon development ?

- Humans

- Semi-automatic

- Fully automatic

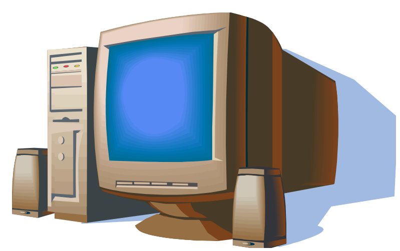

What should be added to a lexicon?

- Find relevant words, phrases, patterns that can be used to express subjectivity
- Determine the polarity of subjective expressions

Words

- Adjectives Hatzivassiloglou & McKeown 1997, Wiebe 2000, Kamps & Marx 2002, Andreevskaia & Bergler 2006
  - positive: honest important mature large patient
    - Ron Paul is the only honest man in Washington.
    - Kitchell's writing is unbelievably mature and is only likely to get better.
    - To humour me my patient father agrees yet again to my choice of film

Words

- Adjectives
  - negative: harmful hypocritical inefficient insecure
    - It was a macabre and hypocritical circus.
    - Why are they being so inefficient ?

Words

- Adjectives
  - Subjective (but not positive or negative sentiment): curious, peculiar, odd, likely, probable
    - He spoke of Sue as his probable successor.
    - The two species are likely to flower at different times.

Words

- Other parts of speech Turney & Littman 2003, Riloff, Wiebe & Wilson 2003, Esuli & Sebastiani 2006
  - Verbs
    - positive: **praise, love**
    - negative: **blame, criticize**
    - subjective: **predict**
  - Nouns
    - positive: **pleasure, enjoyment**
    - negative: **pain, criticism**
    - subjective: **prediction, feeling**

Phrases

- Phrases containing adjectives and adverbs Turney 2002, Takamura, Inui & Okumura 2007
  - positive: high intelligence, low cost
  - negative: little variation, many troubles

How to find them? Using patterns

- Lexico-syntactic patterns Riloff & Wiebe 2003
- way with <np>: … to ever let China use force to have its way with …
- expense of <np>: at the expense of the world's security and stability
- underlined <dobj>: Jiang's subdued tone … underlined his desire to avoid disputes …

How to find them? Using association

- How do we identify subjective items?

- Assume that contexts are coherent

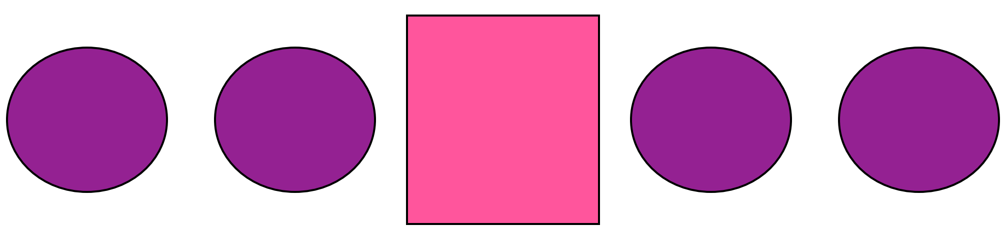

Conjunction

Web Results **1** - **10** of about **762,000** for **"was very nice and"**.

The Homestay Experience - Cultural Kaleidoscope 2006
My host's home **was very nice and** comfortable. I got to try all types of food; Malaysian, Chinese, Indonesian and I loved it all. My host's parents were very **...**
www.gardenschool.edu.my/studentportal/aec/Kaleidoscope06/experience.asp - 10k -
Cached - Similar pages - Note this

PriceGrabber User Rating for Watch Your Budget - PriceGrabber.com
Reviews, Camera I purchased **was very nice and** a bargain. There was a problem with shipping, but was resolved quickly. Buy with confidence from this vendor. **...**
www.pricegrabber.com/rating_getreview.php/retid=5821 - Similar pages - Note this

Testimonials
"Everybody **was very nice and** service was as fast as they possibly could. **...** "Staff member who helped me **was very nice and** easy to talk to." **...**
www.sa.psu.edu/uhs/news/testimonials.cfm - 22k - Cached - Similar pages - Note this

Naxos Villages - Naxos Town or Chora Reviews: Very nice and very **...**
-Did you enjoy the trip to Naxos Town: Yes it **was very nice and** very scenic -In order to get to the village were there enough signs in order to find it: It **...**

Statistical association

- If words of the same orientation likely to co-occur together, then the presence of one makes the other more probable (co-occur within a window, in a particular context, etc.)

- Use statistical measures of association to capture this interdependence
  - E.g., Mutual Information (Church & Hanks 1989)

How to find them? Using similarity

- How do we identify subjective items?

- Assume that contexts are coherent
- Assume that alternatives are similarly subjective ("plug into" subjective contexts)

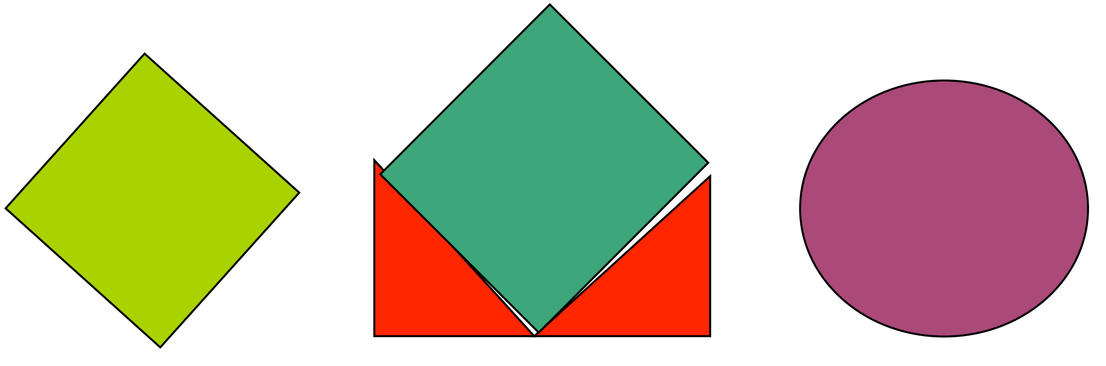

How? Summary

- How do we identify subjective items?

- Assume that contexts are coherent
- Assume that alternatives are similarly subjective
- Take advantage of specific words

Home | Concordance | Word List | Word Sketch | Thesaurus | Sketch-Diff

cause British National Corpus freq = 20207

object	15651	5.8
damage	938	10.09
harm	276	8.91
injury	295	8.38
problem	1014	8.37
trouble	249	8.32
death	383	7.96
delay	146	7.87
confusion	137	7.8
difficulty	223	7.74
disruption	111	7.71
distress	101	7.52
concern	190	7.35
pain	126	7.24
chaos	82	7.22
accident	126	7.22
loss	190	7.14
controversy	81	7.12
pollution	88	7.04
havoc	64	7.01
cancer	93	7.01
stir	62	7.0
suffering	70	6.99
disease	141	6.95
explosion	72	6.93
embarrassment	63	6.87

subject	9100	6.2
negligence	55	7.34
virus	53	7.14
smoking	27	6.36
defect	29	6.32
bacterium	26	6.23
infection	32	6.09
factor	76	6.07
assault	28	6.05
pollution	31	6.04
recession	28	5.99
stress	28	5.88
accident	36	5.8
bomb	26	5.78
disease	50	5.74
fire	45	5.63
lack	37	5.6
organism	18	5.58
deficiency	16	5.57
fault	21	5.57
delay	20	5.55
damage	30	5.51
weather	22	5.39
explosion	15	5.27
drought	12	5.27
parasite	12	5.26

modifier	1971	1.4
reasonable	26	8.72
indirectly	16	7.77
possibly	27	7.67
thereby	26	7.66
mainly	32	7.51
inevitably	20	7.48
partly	22	7.47
probably	51	7.1
thus	36	7.01
recklessly	8	6.97
in part	10	6.94
undoubtedly	11	6.91
deliberately	14	6.81
certainly	26	6.77
intentionally	7	6.77
sometimes	31	6.69
often	72	6.67
directly	24	6.66
reportedly	9	6.64
usually	37	6.61
either	28	6.49
primarily	10	6.45
largely	18	6.39
allegedly	7	6.39
also	196	6.39

and/or	130	0.0
permit	18	6.09
contribute	5	4.21
use	6	0.22

part_intrans	10	0.0
by	9	6.71

unary rels		
np_VPto	2407	25.0
prep_Sing	158	11.3
np_pp	4585	5.1

pp_by-p	3374	16.4
negligence	46	8.14
defect	21	6.81
bacterium	17	6.62
virus	17	6.4
smoking	13	6.4
fault	19	6.15
lack	37	5.98
deficiency	10	5.95
shortage	12	5.75
blockage	6	5.71
breach	14	5.66
parasite	7	5.65
default	7	5.59
error	18	5.58
abnormality	7	5.58
theft	9	5.58
pollution	14	5.57
enteritis	5	5.53
build-up	6	5.51
fall	17	5.5
exposure	11	5.5
warming	6	5.46
drought	6	5.45
failure	24	5.38
recession	11	5.36

Existing lexicons: General Inquirer

- abide,POSITIVE
- able,POSITIVE
- abound,POSITIVE
- absolve,POSITIVE
- absorbent,POSITIVE
- absorption,POSITIVE
- abundance,POSITIVE

- abandon,NEGATIVE
- abandonment,NEGATIVE
- abate,NEGATIVE
- abdicate,NEGATIVE
- abhor,NEGATIVE
- abject,NEGATIVE
- abnormal,NEGATIVE

Existing lexicons: Opinion Finder

- type=weaksubj len=1 word1=able pos1=adj stemmed1=n polarity=positive polannsrc=tw mpqapolarity=weakpos
- type=weaksubj len=1 word1=abnormal pos1=adj stemmed1=n polarity=negative polannsrc=ph mpqapolarity=strongneg
- type=weaksubj len=1 word1=abolish pos1=verb stemmed1=y polannsrc=tw mpqapolarity=weakneg
- type=strongsubj len=1 word1=abominable pos1=adj stemmed1=n intensity=high polannsrc=ph mpqapolarity=strongneg
- type=strongsubj len=1 word1=abominably pos1=anypos stemmed1=n intensity=high polannsrc=ph mpqapolarity=strongneg
- type=strongsubj len=1 word1=abominate pos1=verb stemmed1=y intensity=high polannsrc=ph mpqapolarity=strongneg
- type=strongsubj len=1 word1=abomination pos1=noun stemmed1=n intensity=high polannsrc=ph mpqapolarity=strongneg
- type=weaksubj len=1 word1=above pos1=anypos stemmed1=n polannsrc=tw mpqapolarity=weakpos
- type=weaksubj len=1 word1=above-average pos1=adj stemmed1=n polarity=positive polannsrc=ph mpqapolarity=strongpos

Existing lexicons: SentiWordNet

- P: 0.75 O: 0.25 N: 0 **good**#101123148
  having desirable or positive qualities especially those suitable for a thing specified; "good news from the hospital"; "a good report card"; "when she was good she was very very good"; "a good knife is one good for cutting“
- P: 0 O: 1 N: 0 **good**#2 full#6 00106020
  having the normally expected amount; "gives full measure"; "gives good measure"; "a good mile from here"
- P: 0 O: 1 N: 0 **short**# 201436003
  (primarily spatial sense) having little length or lacking in length; "short skirts"; "short hair"; "the board was a foot short"; "a short toss"
- P: 0.125 O: 0.125 N: 0.75 **short**#3 little#6 02386612
  low in stature; not tall; "he was short and stocky"; "short in stature"; "a short smokestack"; "a little man"

Main resources

- Lexicons
  - General Inquirer (Stone et al., 1966)
  - OpinionFinder lexicon (Wiebe & Riloff, 2005)
  - SentiWordNet (Esuli & Sebastiani, 2006)

- Annotated corpora
  - MPQA corpus (Wiebe et. al, 2005)
  - Used in statistical approaches (Hu & Liu 2004, Pang & Lee 2004)

- Tools
  - Algorithm based on minimum cuts (Pang & Lee, 2004)
  - OpinionFinder (Wiebe et. al, 2005)

MPQA: definitions and annotation scheme

- Manual annotation: human markup of corpora (bodies of text)
- Why?
  - Understand the problem
  - Create gold standards (and training data)

Wiebe, Wilson, Cardie LRE 2005

Wilson & Wiebe ACL-2005 workshop

Somasundaran, Wiebe, Hoffmann, Litman ACL-2006 workshop

Somasundaran, Ruppenhofer, Wiebe SIGdial 2007

Wilson 2008 PhD dissertation

Overview

- Fine-grained: expression-level rather than sentence or document level

- Annotate
  - Subjective expressions
  - material attributed to a source, but presented objectively

Corpus

- MPQA: www.cs.pitt.edu/mqpa/databaserelease (version 2)
- English language versions of articles from the world press *(187 news sources)*
- Also includes contextual polarity annotations (later)
- Themes of the instructions:
  - No rules about how particular words should be annotated.
  - Don't take expressions out of context and think about what they *could* mean, but judge them as they are used in that sentence.

Other gold standards

- Derived from manually annotated data
- Derived from "found" data (examples):
  - Blog tags Balog, Mishne, de Rijke EACL 2006
  - Websites for reviews, complaints, political arguments
    - amazon.com Pang and Lee ACL 2004
    - complaints.com Kim and Hovy ACL 2006
    - bitterlemons.com Lin and Hauptmann ACL 2006

Gold standard data example

- Positive movie reviews

offers a breath of the fresh air of true sophistication .

a thoughtful , provocative , insistently humanizing film .

with a cast that includes some of the top actors working in independent film , lovely & amazing involves us because it is so incisive , so bleakly amusing about how we go about our lives .

a disturbing and frighteningly evocative assembly of imagery and hypnotic music composed by philip glass .

not for everyone , but for those with whom it will connect , it's a nice departure from standard moviegoing fare .

- Negative movie reviews

unfortunately the story and the actors are served with a hack script .

all the more disquieting for its relatively gore-free allusions to the serial murders , but it falls down in its attempts to humanize its subject .

a sentimental mess that never rings true .

while the performances are often engaging , this loose collection of largely improvised numbers would probably have worked better as a one-hour tv documentary .

interesting , but not compelling .

Main resources

- Lexicons
  - General Inquirer (Stone et al., 1966)
  - OpinionFinder lexicon (Wiebe & Riloff, 2005)
  - SentiWordNet (Esuli & Sebastiani, 2006)

- Annotated corpora
  - Used in statistical approaches (Hu & Liu 2004, Pang & Lee 2004)
  - MPQA corpus (Wiebe et. al, 2005)

- Tools
  - Algorithm based on minimum cuts (Pang & Lee, 2004)
  - OpinionFinder (Wiebe et. al, 2005)

Lexicon-based tools

- Use sentiment and subjectivity lexicons
- Rule-based classifier
  - A sentence is subjective if it has at least two words in the lexicon
  - A sentence is objective otherwise

Corpus-based tools

- Use corpora annotated for subjectivity and/or sentiment
- Train machine learning algorithms:
  - Naïve bayes
  - Decision trees
  - SVM
  - …
- Learn to automatically annotate new text

III. Word- and phrase-level annotations

Dictionary-based

Corpus-based

Hybrid

Trends explored so far

- *Manual annotations* involving human judgment of words and phrases
- *Automatic annotations* based on knowledge sources (e.g. dictionary)
- *Automatic annotations* based on information derived from corpora (co-occurrence metrics, patterns)

Dictionary-based: Subjectivity
Mihalcea et al., 2007 - translation

- OpinionFinder lexicon (English)
  - 6,856 entries, 990 multi-word expressions
- Bilingual English-Romanian dictionary
  - Dictionary 1 (authoritative source) 41,500 entries; Dictionary 2 (online, back-up) 4,500 entries
- Resulting lexicon of 4,983 entries (Romanian)

- English lexicon contains inflected words, but lemmatized form is needed to querya dictionary, yet lemmatization can affect subjectivity:
  - ***memories*** (En, pl, subj) → ***memorie*** (Ro, sg, obj)
- Ambiguous entries both in source and target language; 49.6% subjective entries from those correctly translated
  - *fragile* (En, subj) → *fragil* (Ro, obj) [breakable objects vs. delicate]
  - Rely on usage frequency listed by the dictionary
- Multi-word expressions difficult to translate (264/990 translated)
  - If not in the dictionary, word-by-word approach, further validated by counts on search engine: *one-sided* (En, subj) → *cu o latura* (Ro, obj)

Dictionary-based: Polarity
Kim and Hovy, 2006 - bootstrapping

WordNet structure

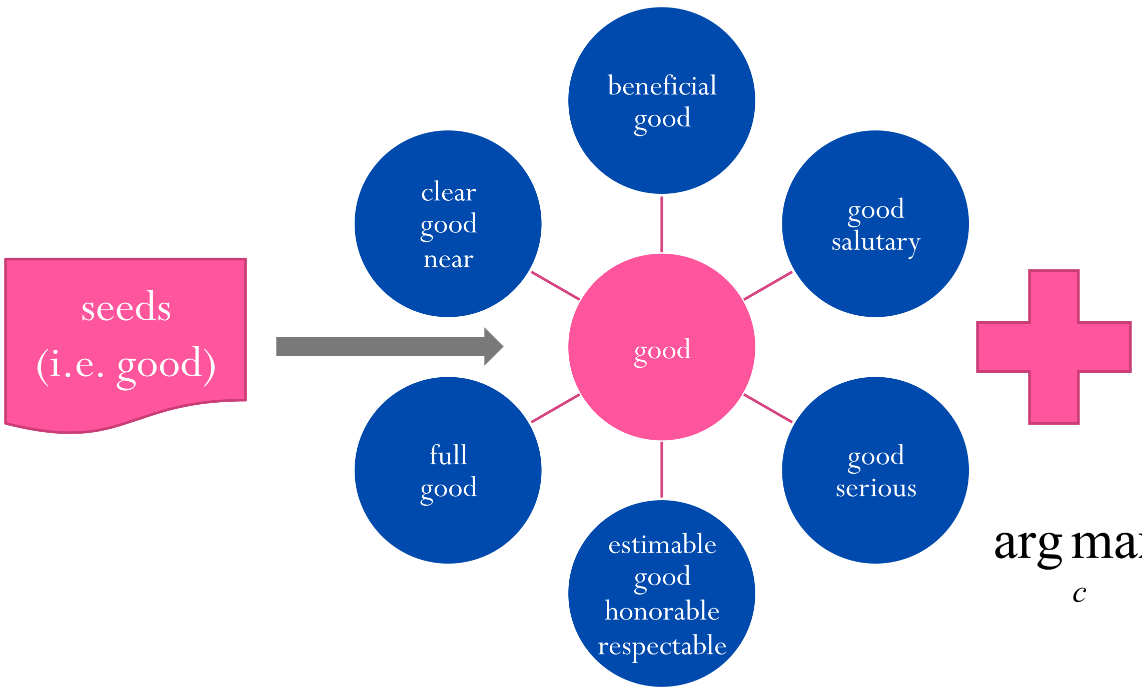

estimated closeness of candidate to positive, negative, and neutral classes

$$\arg\max_{c} P(c)\prod_{k=1}^{n} P(f_k \mid c)^{count(f_k, synset(w))} \; *$$

- Resulted in an English polarity lexicon: 1,600 verbs and 3,600 adjectives
- The lexicon is then translated into German using an automatically generated translation dictionary (obtained from European Parliament corpus via word alignment)
- using a rule based classifier on a document level polarity dataset – avg F-measure=55%

* Note: f_k stands for feature k of class c (who is a synonym of the word), w for word, and c for class.

Dictionary-based: Polarity

Hassan et al., 2011 – multilingual WordNets and Random Walk

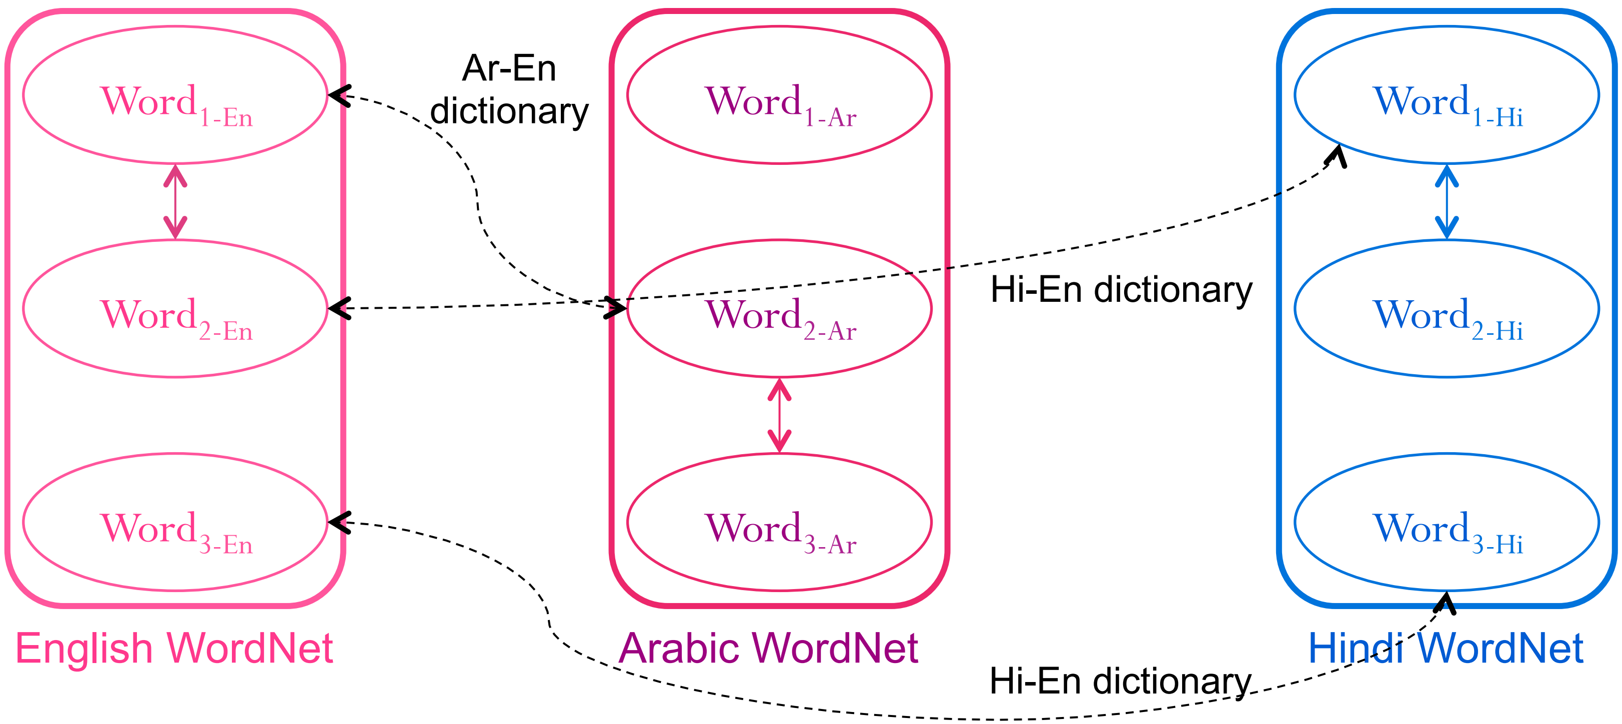

- Predict sentiment orientation based on the mean hitting time to two sets of positive and negative seeds (General Inquirer lexicon – Stone et al., 1966)
- Mean hitting time is the average number of steps a random walker starting at node *i* will take to reach node *j* for the first time (Norris, 1997)
- For Arabic, the accuracy is 92% (approx 30% more than using the SO-PMI method proposed by Turney and Littman, 2003); for Hindi, the accuracy also increases by 20%.

Dictionary-based: Polarity

Pérez-Rosas et al., 2012 – lexicon through WordNet traversal

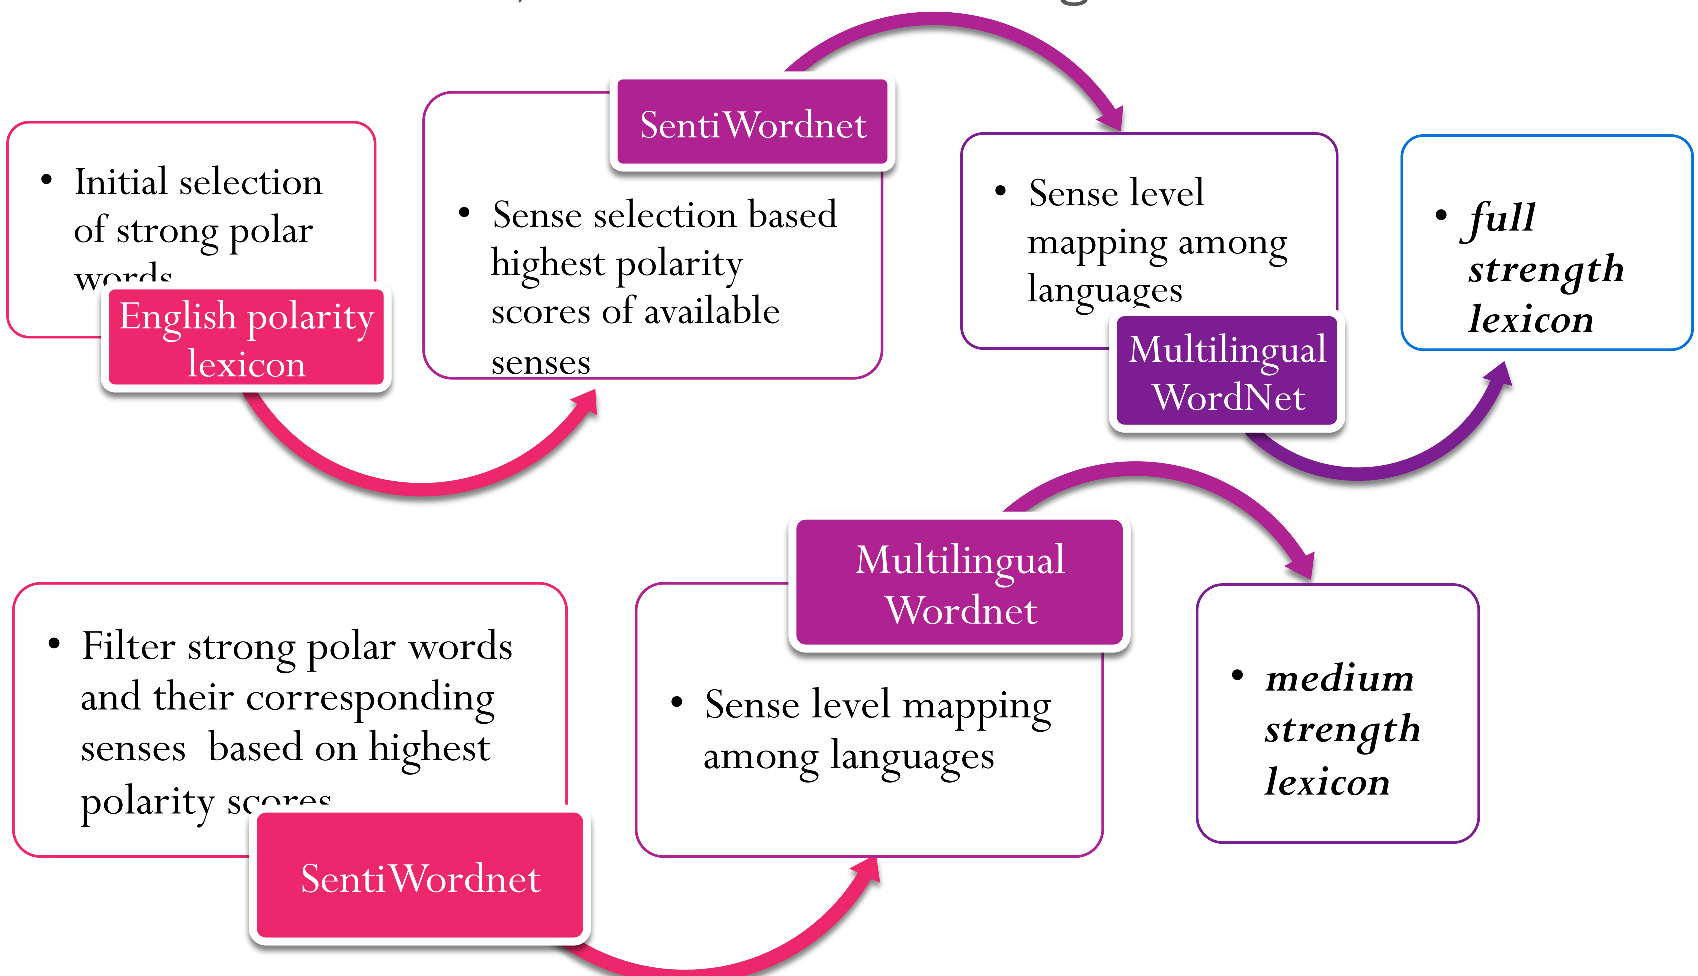

- accuracy values of 90% (full strength lexicon) and 74% (medium strength lexicon) when transferring the sentiment information from English.

Dictionary-based: Subjectivity

Banea et al., 2008 - bootstrapping

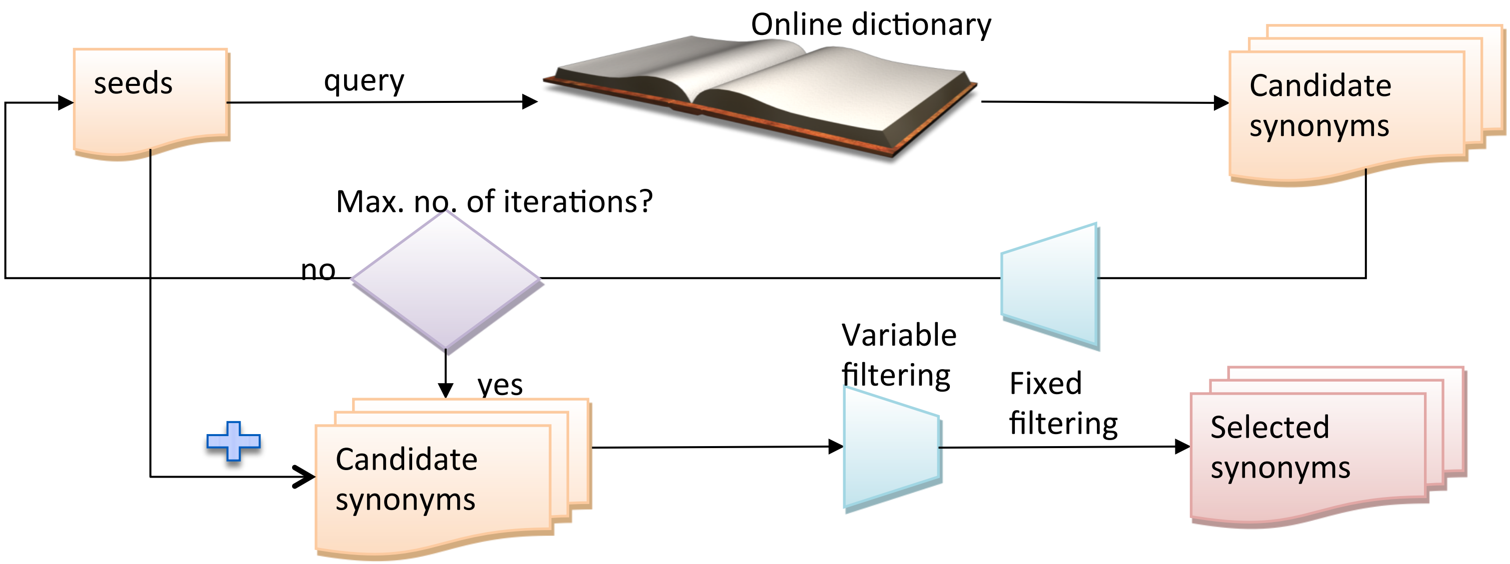

- 60 seeds evenhandedly sampled from nouns, verbs, adjectives, adverbs
- Small training corpus to derive co-occurrence matrix and train LSA to compute the similarity between each candidate and the original seeds
- Online / offline dictionary → extract & parse definition → get candidates → lemmatize → compute similarity scores → accept / discard candidates
- Extracted a subjectivity lexicon of 3,900 entries; using a rule based classifier applied to a sentence level subjectivity dataset – F-measure is 61.7%

Corpus-based: Polarity

Kaji and Kitsuregawa, 2007

1 billion web pages

- HTML layout information (e.g. list markers or tables) that explicitly indicate the evaluation section of a review: pros/cons, minus/plus
- Japanese specific language structure

corpus of polar sentences 220k pos / 280k neg

adjectives & adjectival phrases

Seed data set: 405 pos/neg adjective phrases

polarity_value(w)=PMI(w, pos)-PMI(w,neg)

threshold

polarity lexicon

- Lexicon of 8,166 to 9,670 Japanese entries
- threshold of 0: P_{pos}=76.4%, P_{neg}=68.5%
- threshold of 3: P_{pos}=92.0%, P_{neg}=87.9%

Corpus-based: Polarity
Kanayama and Nasukawa, 2006

- Domain dependent sentiment analysis by using a domain-independent lexicon to extract domain dependent polar atoms.
- **Polar atom**
  - The minimum human-understandable syntactic structures that specify the polarity of clauses
  - **Tuple (polarity, verb/adjective [optional arguments])**
- System uses intra- and inter-sentential coherence to identify polarity shifts (i.e. polarity will not change unless encountering an adversative conjunction)
- Confidence of polar atoms calculated based on its occurrence in positive v. negative contexts

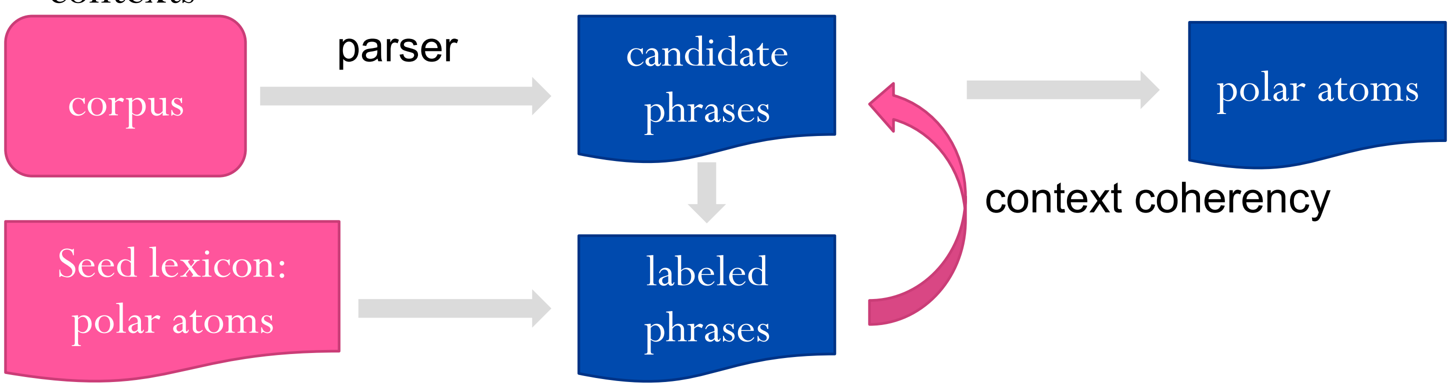

- 4 domains, 200 – 700 polar atoms (in Japanese) per domain with a precision from 54% (phones) to 75% (movies)

Corpus-based: Opinion

Kobayashi et al., 2005 - bootstrapping

- Similar method to Kanayama and Nasukawa's
- Extracts **opinion triplets** = (subject, attribute, value), treated from an anaphora resolution frameset
    - i.e. product is easy to determine, but finding the attribute of a value is similar to finding the antecedent in an anaphora resolution task; attribute may/may not be present

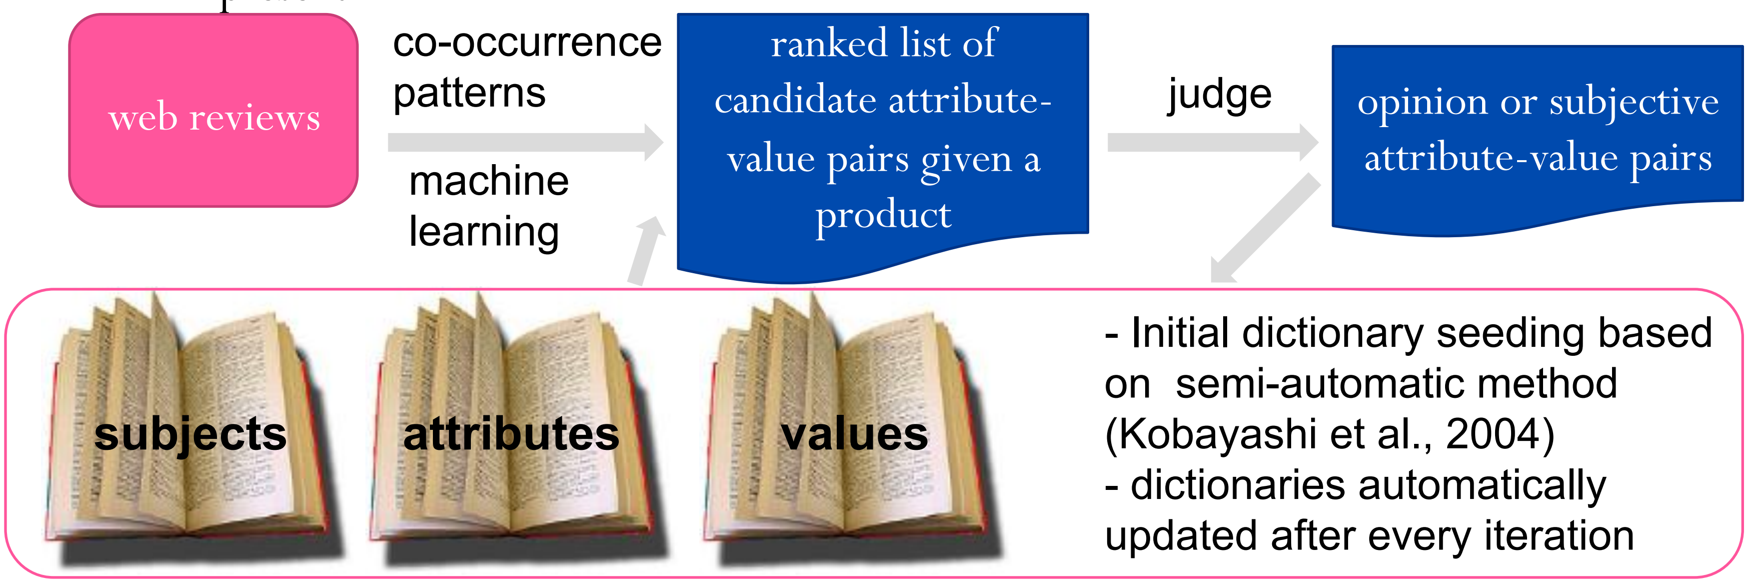

- 3,777 attribute expressions and 3950 value expressions in Japanese
- Coverage of 35% to 45% vis-à-vis manually extracted expressions

Hybrid: Affect

Pitel and Grefenstette, 2008

- Classify words in 44 paired affect classes (e.g., love - hate, courage - fear)
- Each class is associated with a positive/negative orientation

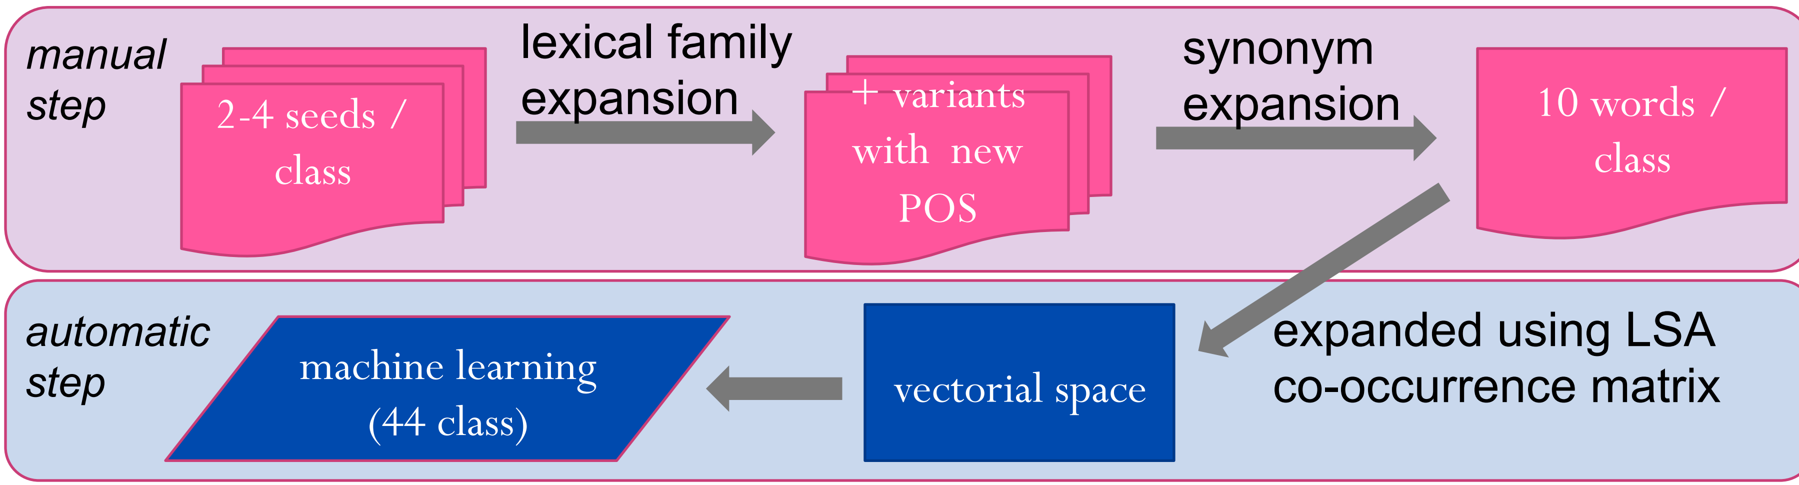

- For LSA – short windows → highly semantic information, large windows → thematic / pragmatic information
- Varied windows is 42 ways, based on no. of words in co-occurrence window and position vis-à-vis reference word → concatenated LSA vectors of 300 dimensions (trained on French EuroParl) →vectorial space of 12,600 dimensions
- Labeled 2632 French words – 54% are correctly classified in the top 10 classes

Other approaches

- Takamura et al., 2006
  - finding the polarity of phrases such as "light laptop" (both words individually are neutral)
  - on a dataset of 12,000 adjective-noun phrases drawn from Japanese newswire → a model based on triangle and "U-shaped" graphical dependencies achieves 81%
- Suzuki et al., 2006
  - focus on evaluative expressions (subjects, attributes and values)
  - use an expectation maximization algorithm and a Naïve Bayes classifier to annotate the polarity of evaluative expressions
  - accuracy of 77% (baseline of 47% - assigning the majority class)
- Bautin et al., 2008
  - Polarity of entities (e.g. George Bush, Vladimir Putin) in 9 languages (Ar, Cn, En, Fr, De, It, Jp, Kr, Es)
  - Translation of documents into English, and calculation of entity polarity using association measures between its occurrence and positive/negative words from a English sentiment lexicon; thus polarity analysis in source language only

IV. Sentence-level annotations

Dictionary-based

Corpus-based

Rule-based classifier

- Use the lexicon to build a classifier
- Rule-based classifier
  - (Riloff & Wiebe, 2003)
  - *Subjective*: two or more (strong) subjective entries
  - *Objective*: at most two (weak) subjective entries in the previous, current, next sentence combined
- Variations are also possible
  - E.g., three or more clues for a subjective sentence
  - Depending on the quality/strength of the classifier

Sentence-level gold standard data set

- Gold standard constructed from SemCor
  - (Mihalcea et al., 2007; Banea et al., 2008,2010)
  - 504 sentences from five English SemCor documents
  - Manually translated in Romanian
  - Labeled by two annotators
  - Agreement 0.83% (κ=0.67)
  - Baseline: 54% (all subjective)
- Also available
  - Spanish (manual translation)
  - Arabic, German, French (automatic translations)

Using the automatically built lexicons

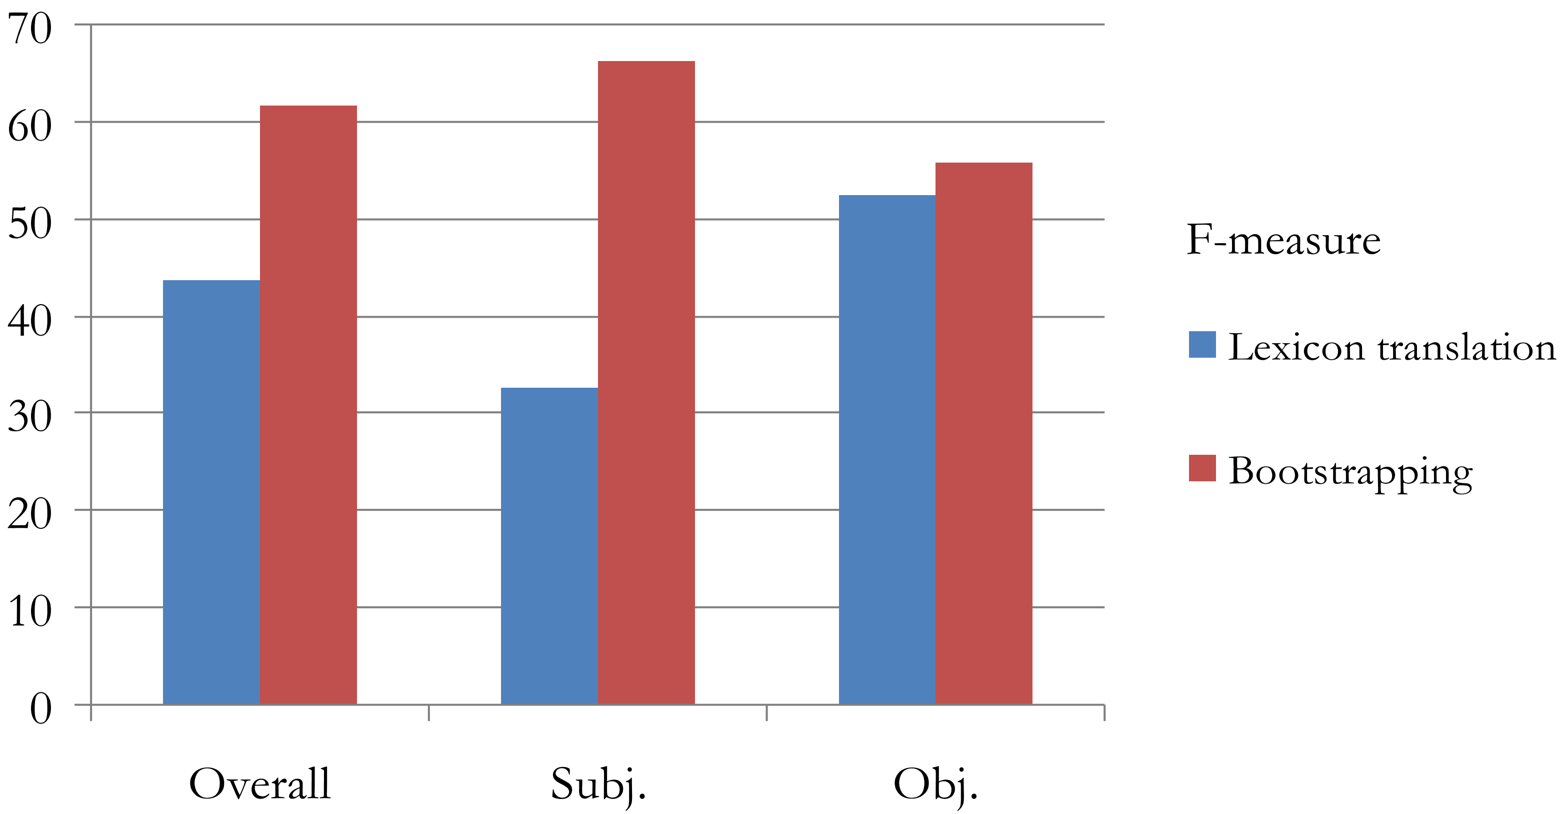

Sentiment units obtained with “deep parsing”

- (Kanayama et. al, 2004)
- Use a machine translation system based on deep parsing to extract “sentiment units” with high precision from Japanese product reviews
- Sentiment unit = a touple between a sentiment label (positive or negative) and a predicate (verb or adjective) with its argument (noun)
- Sentiment analysis system uses the structure of a transfer-based machine translation engine, where the production rules and the bilingual dictionary are replaced by sentiment patterns and a sentiment lexicon, respectively

Sentiment units obtained with "deep parsing"

- Sentiment units derived for Japanese are used to classify the polarity of a sentence, using the information drawn from a full syntactic parser in the target language
- Using about 4,000 sentiment units, when evaluated on 200 sentences, the sentiment annotation system was found to have high precision (89%) at the cost of low recall (44%)

Corpus-based methods

- Collect data in the target language
- Sources:
  - Product reviews
  - Movie reviews
- Extract sentences labeled for subjectivity using min-cut algorithm on graph representation
- Use HTML structure to build large corpus of polar sentences

Extract Subjective Sentences with Min-Cut

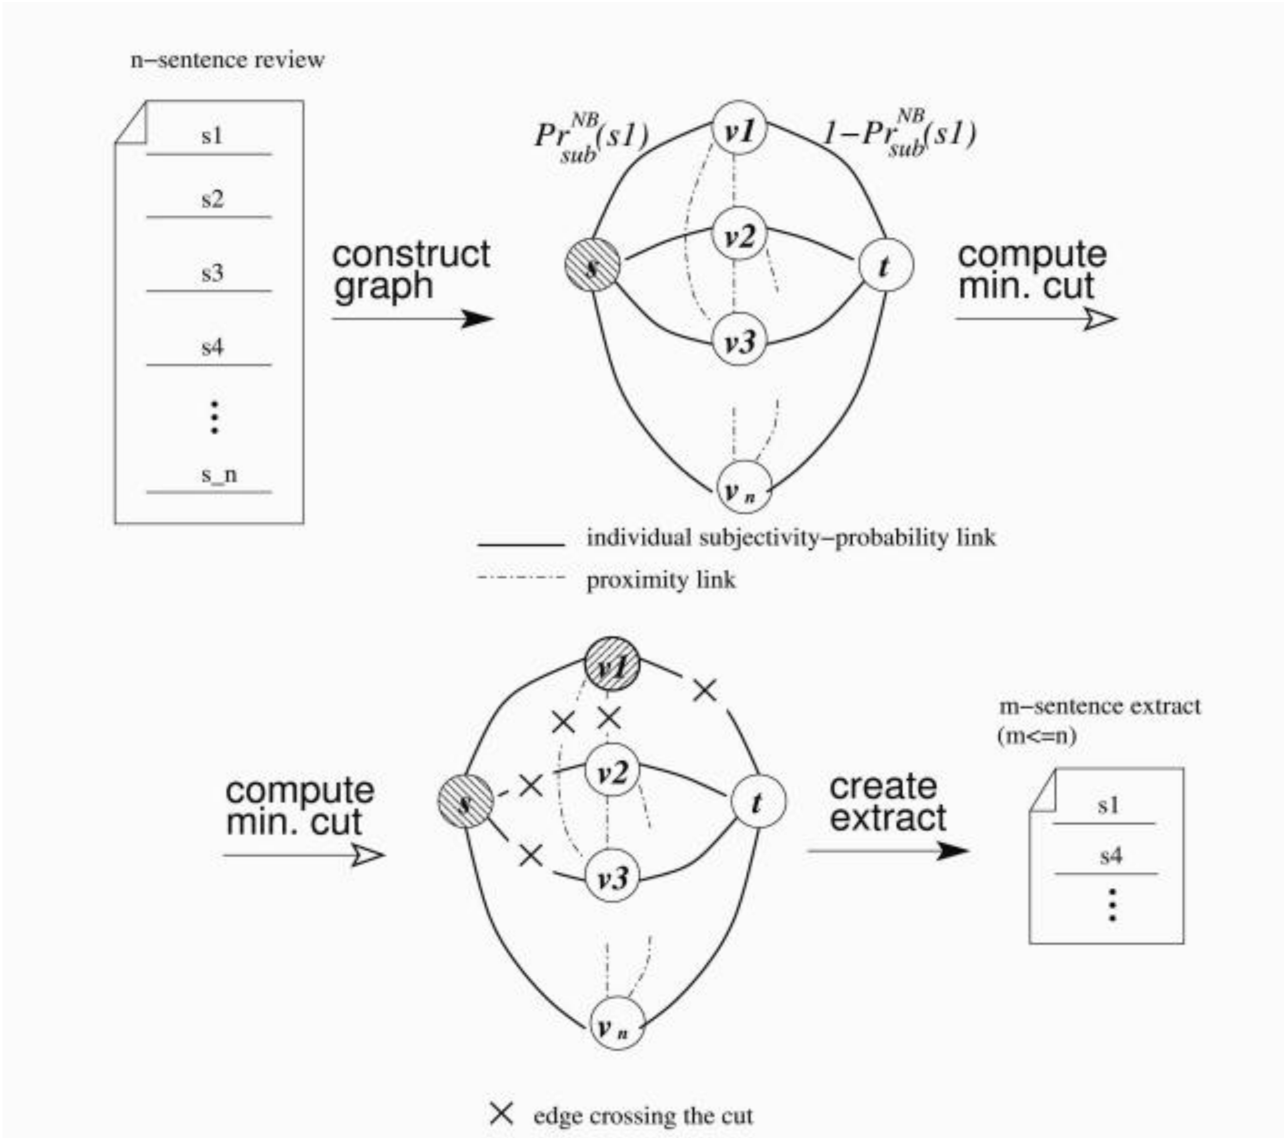

(Pang & Lee, 2004)

Cut-based Algorithm

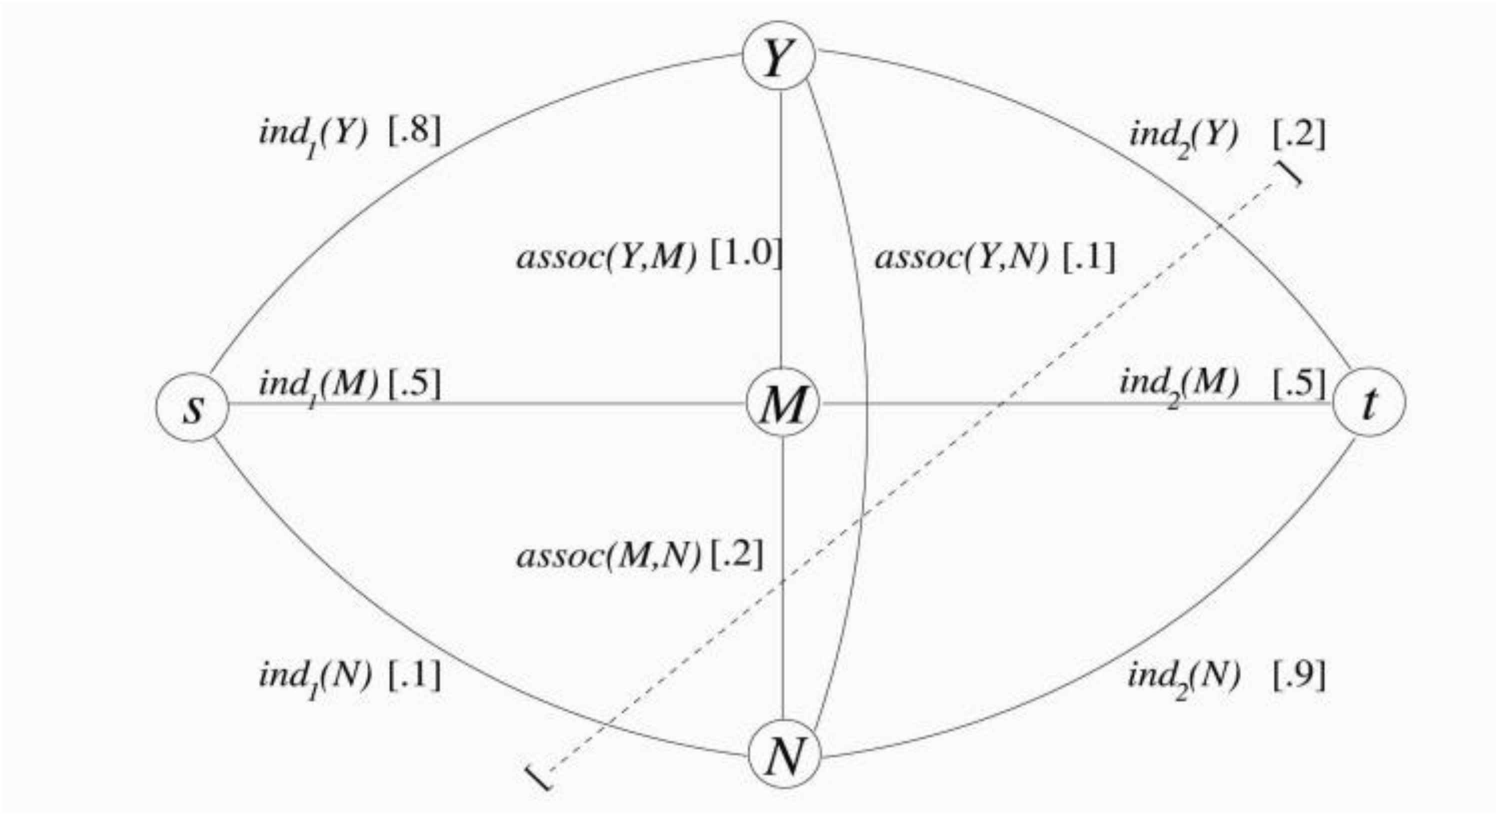

s and t correspond to subjective/objective classification

Extraction of Subjective Sentences

- Assign every individual sentence a subjectivity score
  - e.g. the probability of a sentence being subjective, as assigned by a Naïve Bayes classifier, etc
- Assign every sentence pair a proximity or similarity score
  - e.g. physical proximity = the inverse of the number of sentences between the two entities
- Use the min-cut algorithm to classify the sentences into objective/subjective

Building a labeled corpus from the Web

- (Kaji & Kitsuregawa, 2006, 2007)
- Collect a large corpus of sentiment-annotated sentences from the Web
- Use structural information from the layout of HTML pages (e.g., list markers or tables that explicitly indicate the presence of the evaluation sections of a review, such as "pros"/"cons", "minus"/"plus", etc.), as well as Japanese-specific language structure (e.g., particles used as topic markers)
- Starting with one billion HTML documents, about 500,000 polar sentences are collected, with 220,000 being positive and the rest negative
- Manual verification of 500 sentences, carried out by two human judges, indicated an average precision of 92%

Sentence-level classifiers

- A subset of this corpus, consisting of 126,000 sentences, is used to build a Naive Bayes classifier.
- Using three domain specific data sets (computers, restaurants and cars), the precision of the classifier was found to have an accuracy ranging between 83% (computers) and 85% (restaurants)
- Web data is a viable alternative
- Easily portable across domains

Cross-Language Projections

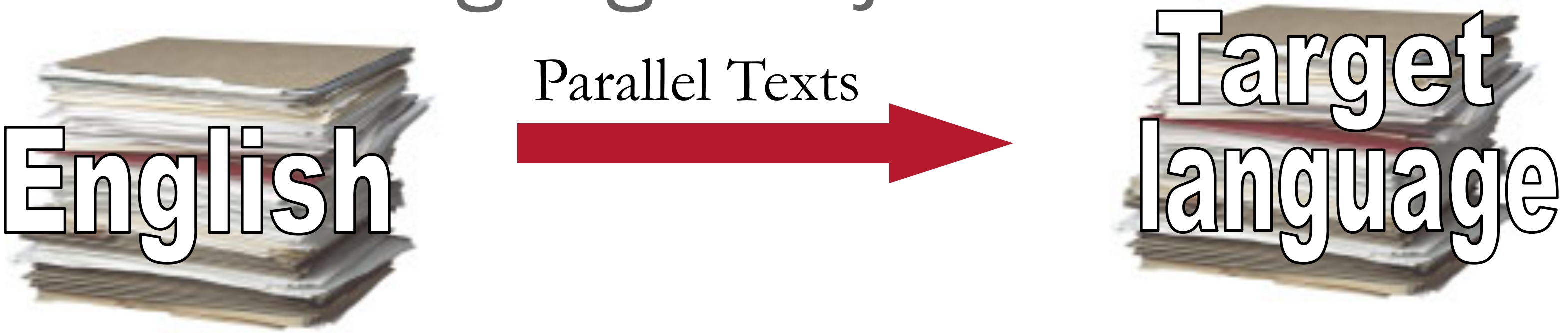

- Eliminate some of the ambiguities in the lexicon by accounting for context
- Subjectivity is transferable across languages – dataset with annotator agreement 83%-90% (kappa .67-.82)

  S: [en] Suppose he did lie beside Lenin, would it be permanent ?

  S: [ro] Sa presupunem ca ar fi asezat alaturi de Lenin, oare va fi pentru totdeauna?
- Solution:
  - Use manually or automatically translated parallel text
  - Use manual or automatic annotations of subjectivity on English data
- (Mihalcea et al., 2007; Banea et al., 2008)

Cross-Language Projections

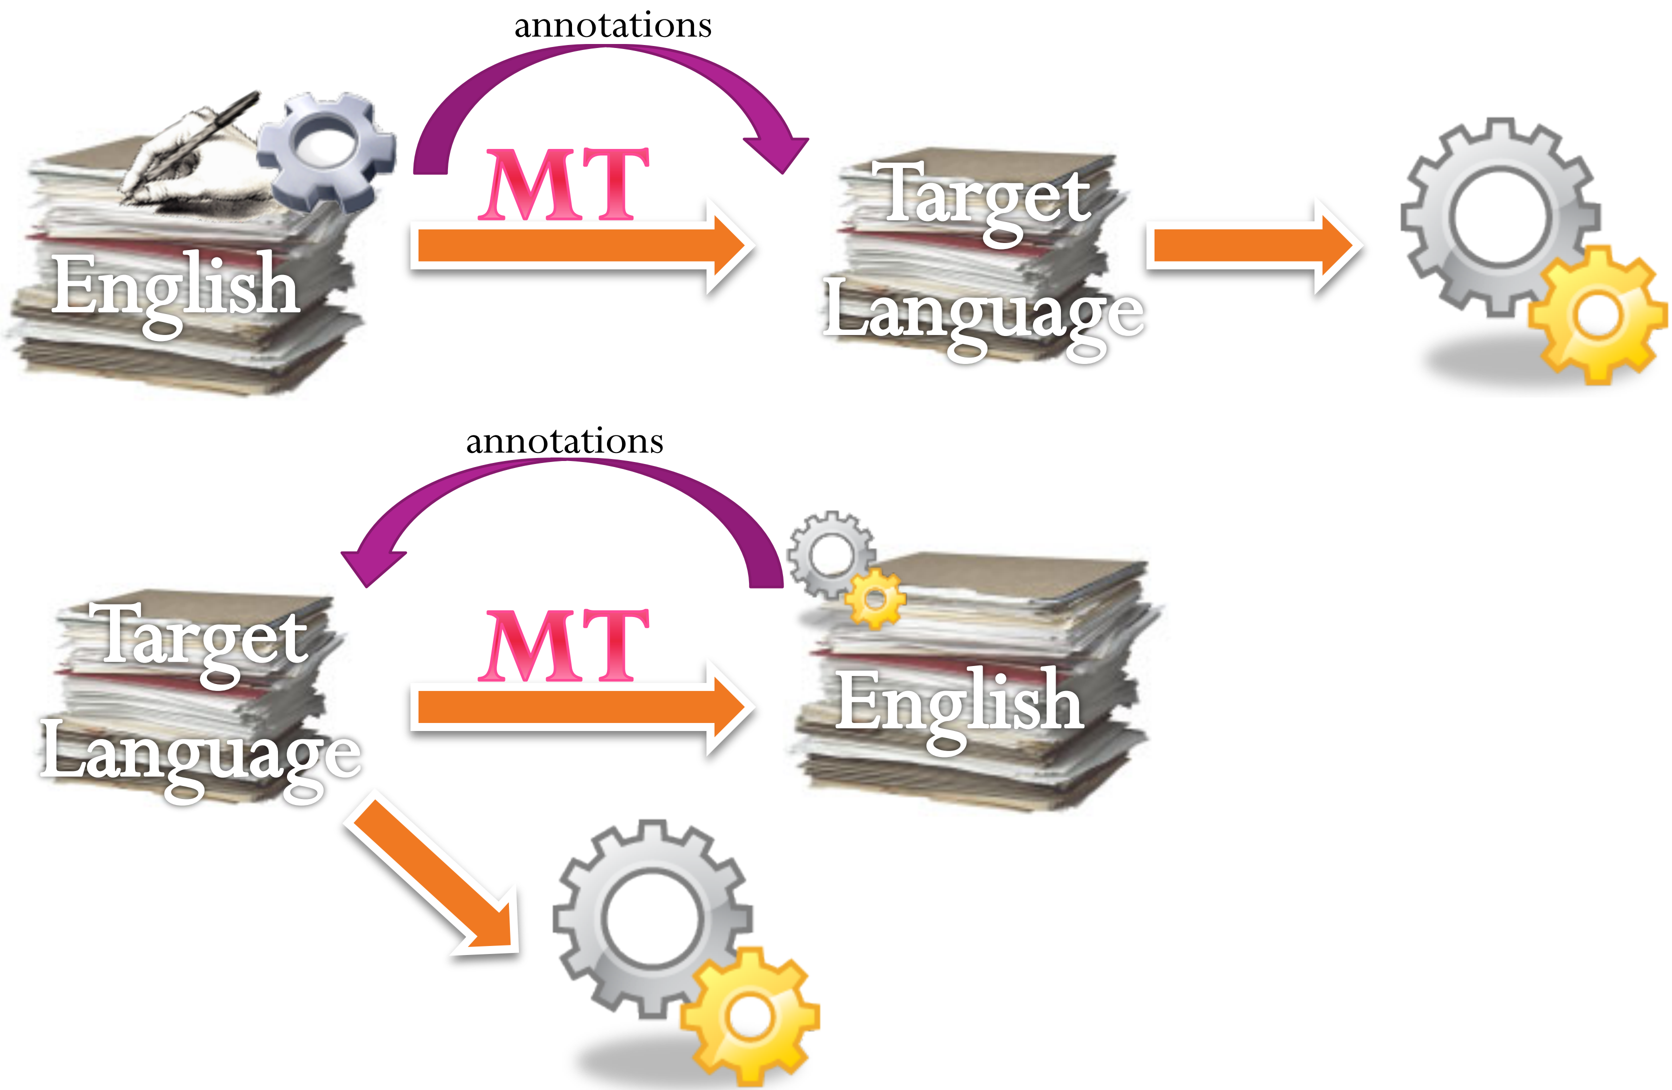

Manual annotation in source language

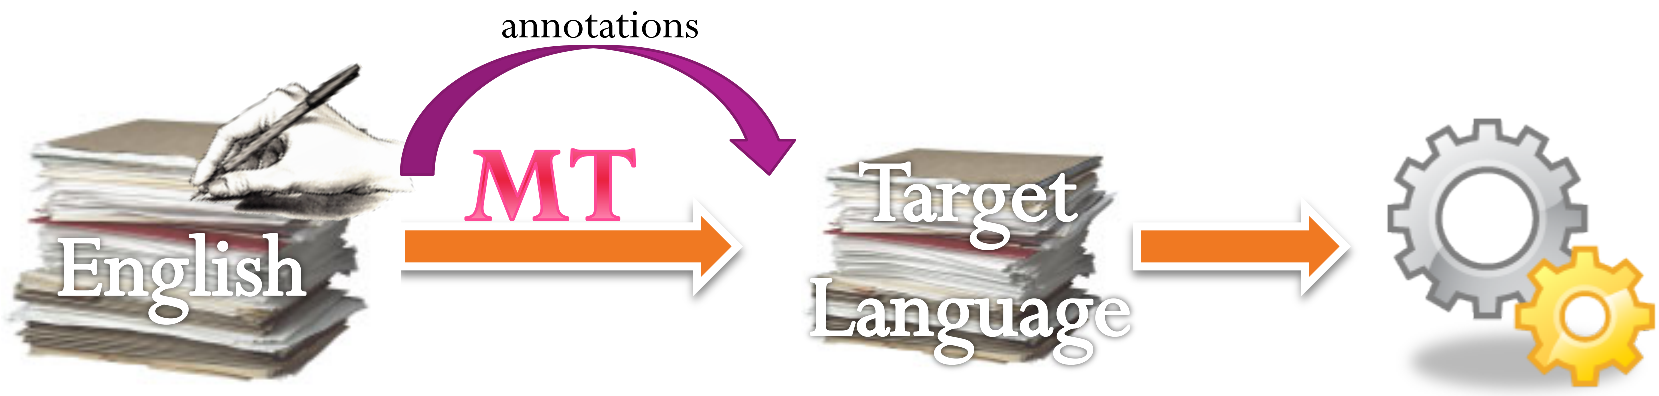

- Manually annotated corpus: MPQA (Wiebe et. al, 2005)
  - A collection of 535 English language news articles
  - 9700 sentences; 55% are subjective & 45% are objective
- Machine translation engine:
  - Language Weaver – Romanian

Source to target language MT

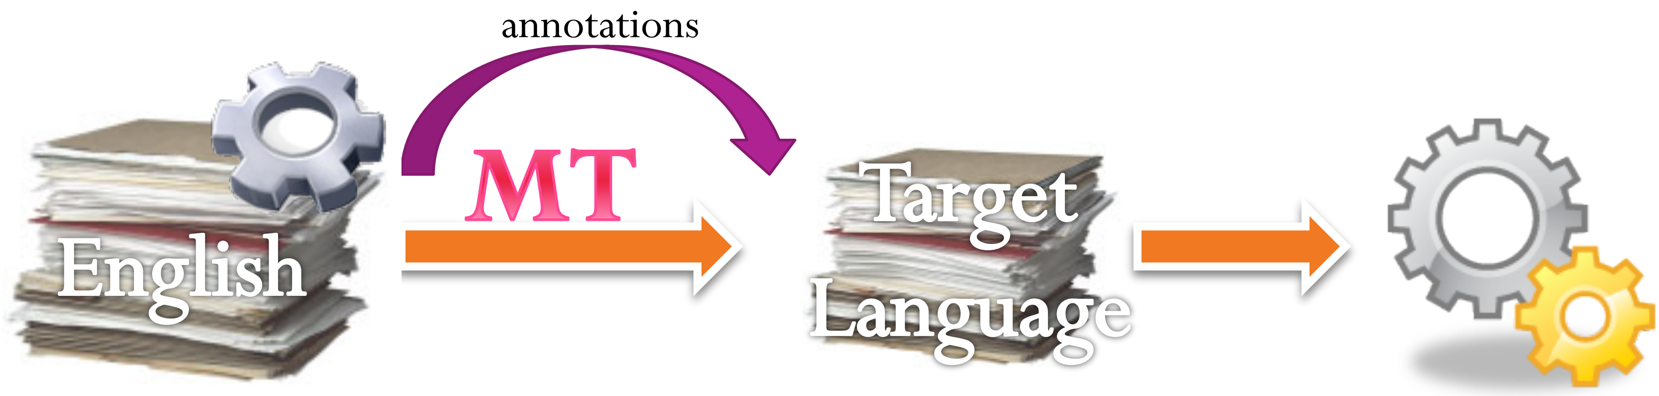

- Raw Corpus: subset of SemCor (Miller et. al, 1993)
  - 107 documents; balanced corpus covering topics such as sports, politics, fashion, education, etc.
  - Roughly 11,000 sentences
- Subjectivity Annotation Tool: OpinionFinder High-Coverage classifier (Wiebe et. al, 2005)
- Machine translation engine:
  - Language Weaver – Romanian

Target to source language MT

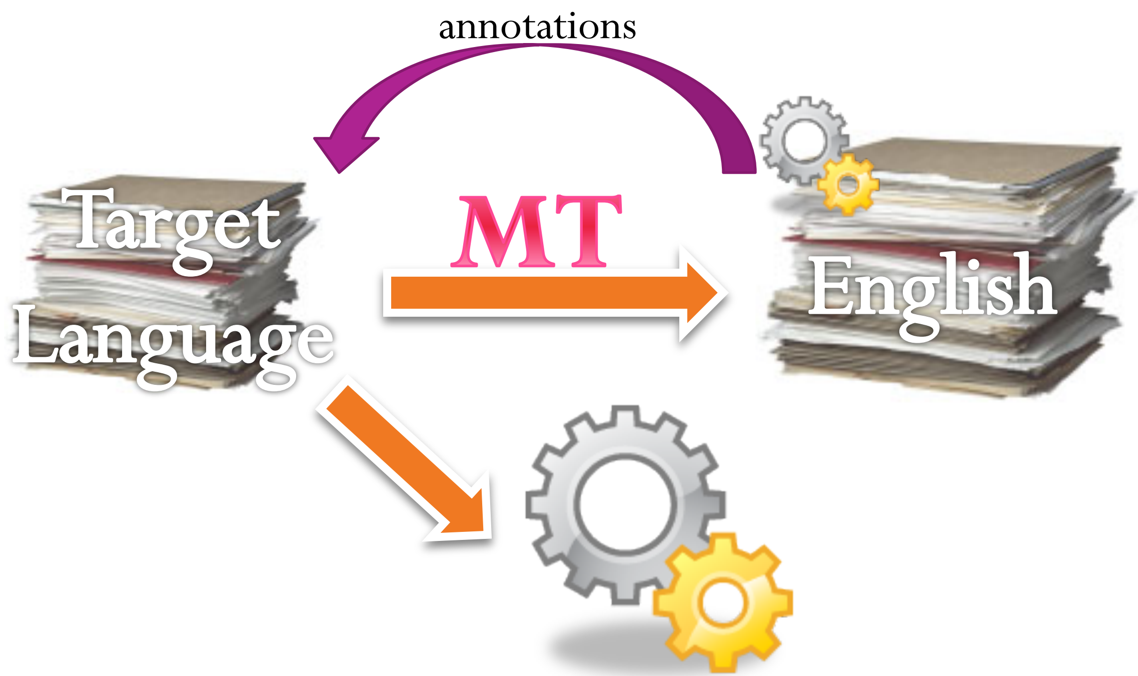

- Same setup as in the automatic annotation experiment
- But the direction of the MT starts from the target language to the source language

Results for cross-lingual projections

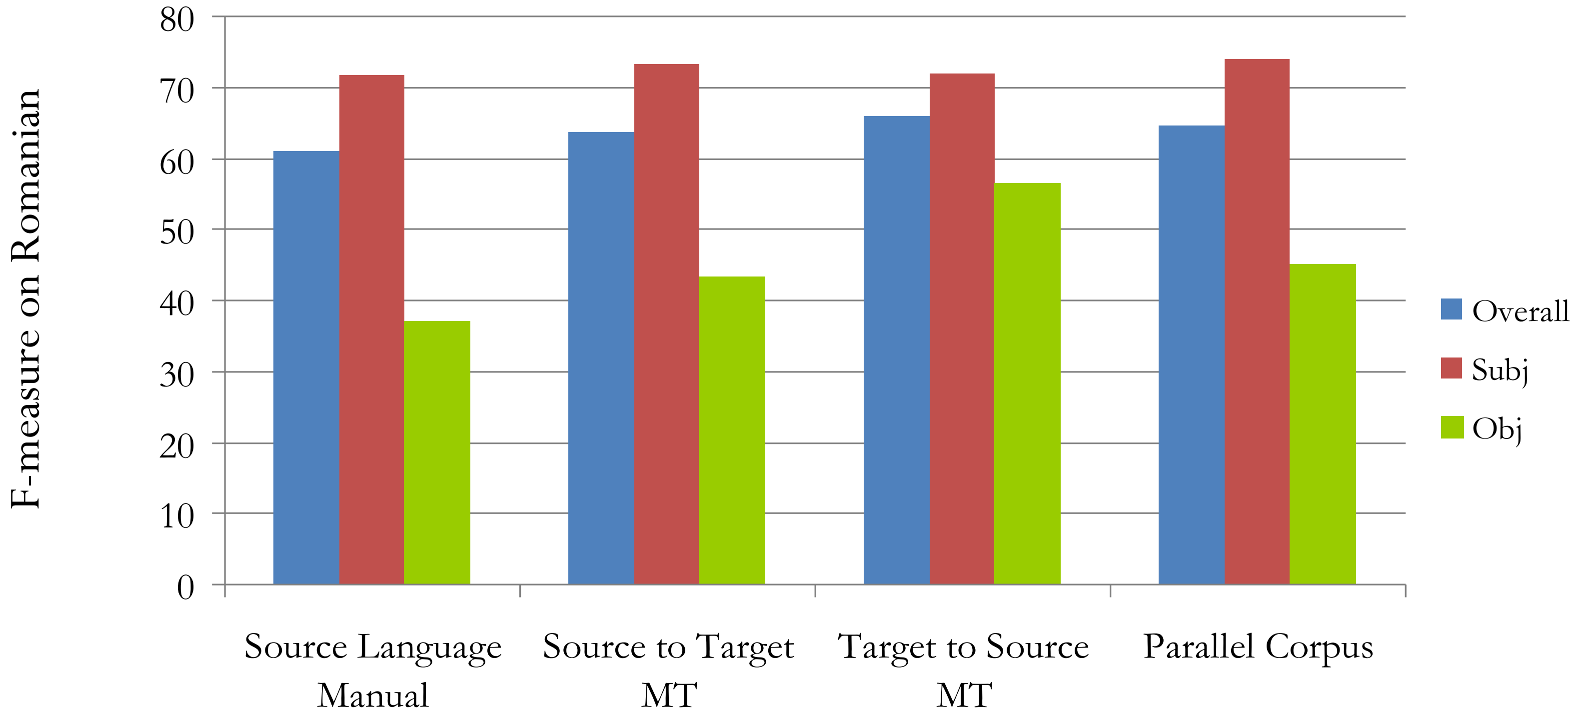

Portability to Spanish

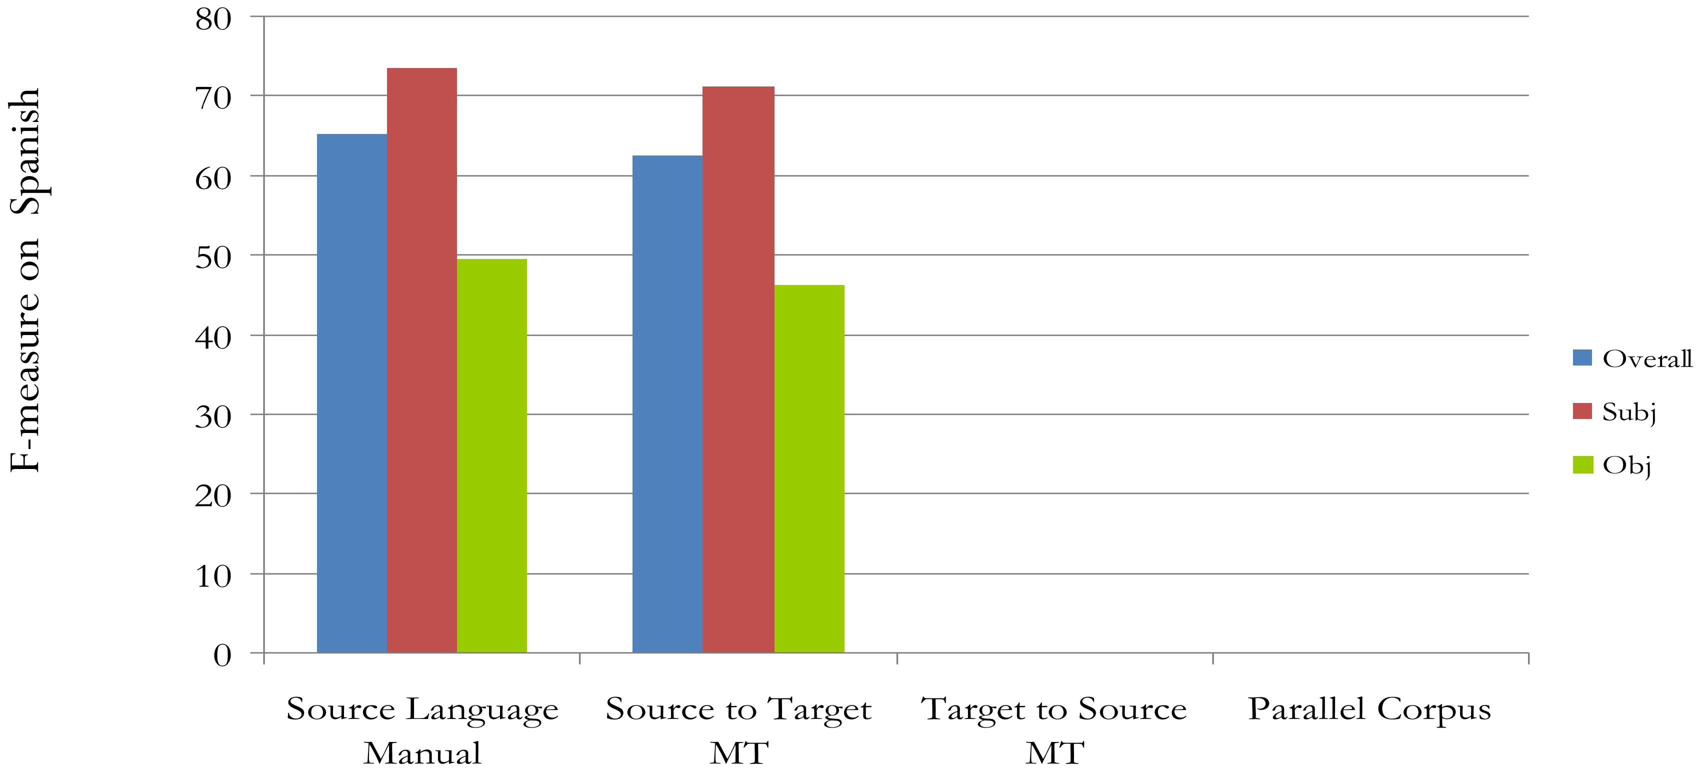

Similar experiments on Asian languages
Kim et al., 2010

- Test set: 859 sentence chunks in Korean, English, Japanese and Chinese.
- Train set: MPQA translated into Korean, Japanese and Chinese using Google Translate.
- Lexicon: translated the OpinionFinder lexicon into the target languages and used a rule based classifier. Strong subj. words – 1; weak subj. words -0.5; if sentence > 1, then subj.

V. Document-level annotations

Dictionary-based

Corpus-based

Dictionary-based: Rule-based polarity

Wan, 2008

- Annotating Chinese reviews using:
  - Method 1:
    - a Chinese polarity lexicon (3,700 pos / 3,100 neg)
    - negation words (13) and intensifiers (148)
  - Method 2:
    - machine translation of Chinese reviews into English
    - OpinionFinder subjectivity / polarity lexicon in English
- Polarity of a document $= \sum \mathit{sentence\ polarity}$
- Sentence polarity $= \sum \mathit{word\ polarity}$
- Evaluations on 886 Chinese reviews:
  - Method 1: accuracy 74.3%
  - Method 2: accuracy 81%; can reach 85% if combining different translations and methods

Dictionary-based: Polarity
Zagibalov and Carroll, 2008 - Bootstrapping

- Identifying "lexical items" (i.e. sequences of Chinese characters that occur between non-character symbols, which include a negation and an adverbial)
- "Zone" – sequence of characters occurring between punctuation marks
- Polarity of a document $= \sum zone\ positive - \sum zone\ negative$
- Zone polarity $= \sum lexical\ item\ polarity$
- Lexical item polarity $\propto length(lexical\ item)^{2} * prev\_polarity\_score / lenght(zone)$ *neg_coeff

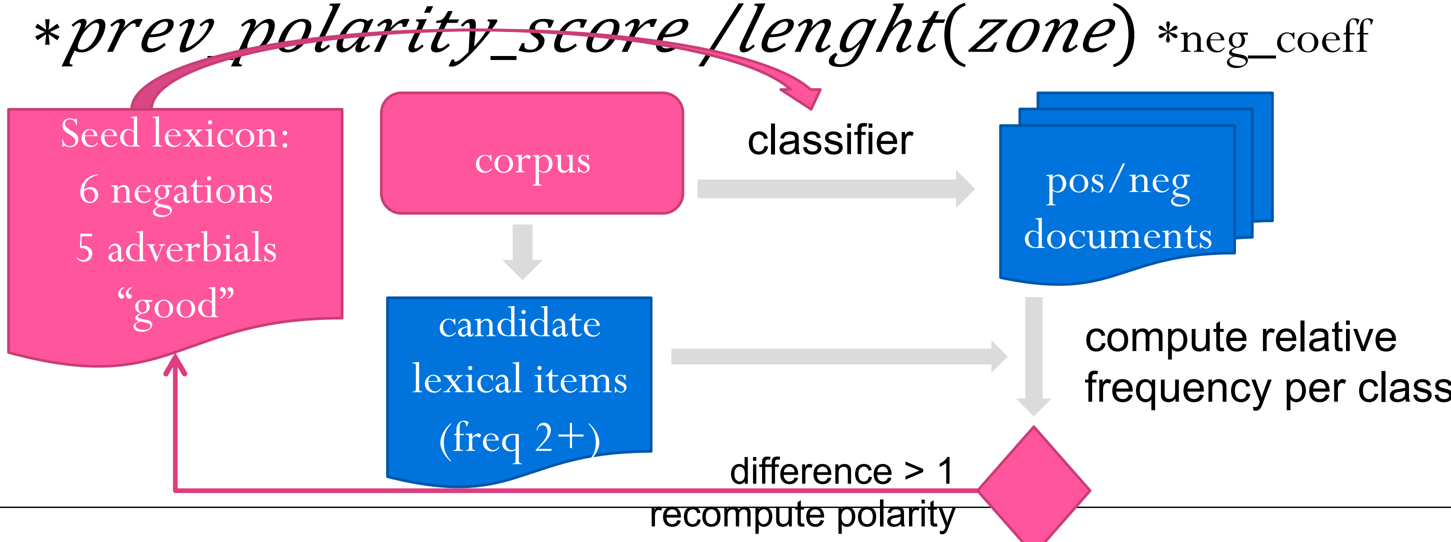

Dictionary-based: Polarity
Kim and Hovy, 2006

- The dictionary-based lexicon construction method using WordNet (discussed previously) generates an English lexicon of 5,000 entries
- Lexicon is translated into German using an automatically generated translation dictionary based on the EuroParl using word alignment
- German lexicon employed in a rule-based system that annotates 70 emails for polarity
- Document polarity:
  - Positive class – a majority of positive words
  - Negative class – count of negative words above threshold
- 60% accuracy for positive polarity, 50% accuracy for negative polarity

Corpus-based: Polarity

Li and Sun, 2007

- Train a machine learning classifier if a set of annotated data exists
  - Experimented with SVM, NB and maximum entropy
- Training set of 6,000 positive / 6,000 negative Chinese hotel reviews, test set of 2,000 positive / 2,000 negative reviews
- Accuracy up to 92% depending on classifier and feature set

Corpus-based: Polarity

Wan, 2009 – Co-training

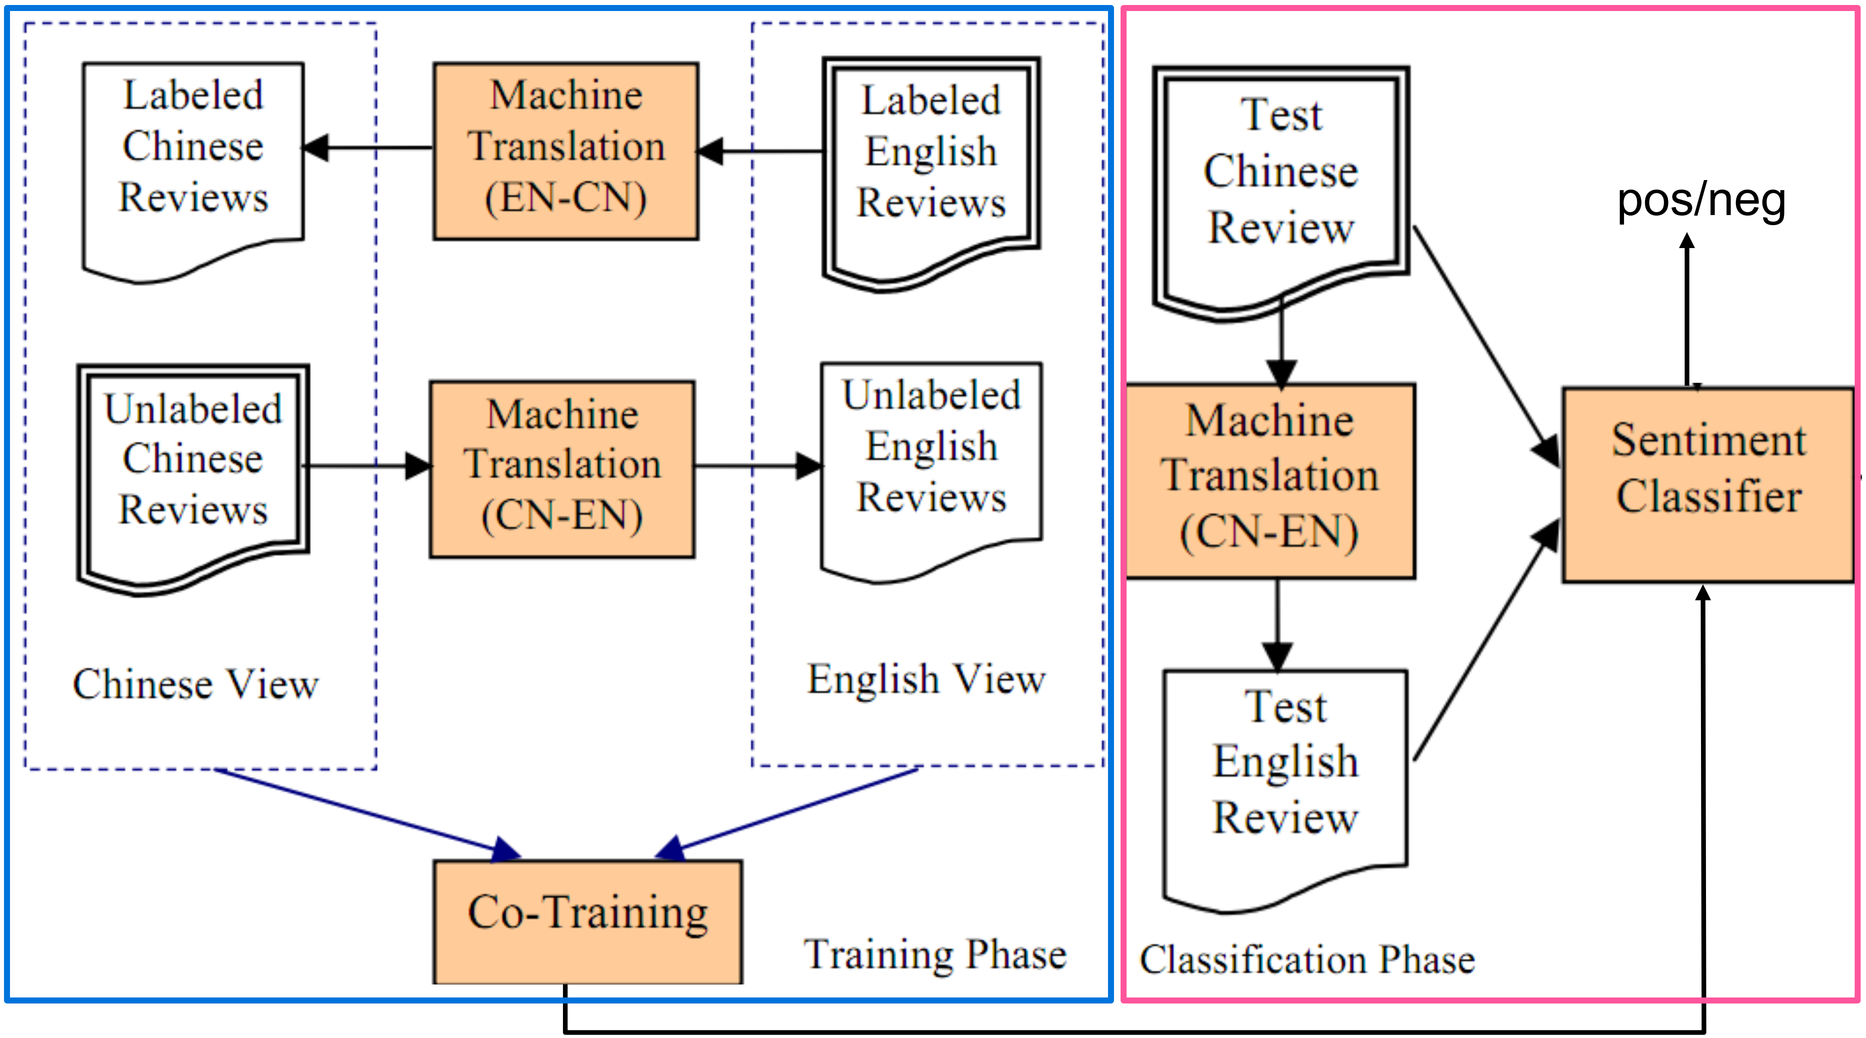

Corpus-based: Polarity
Wan, 2009 – Co-training

- Performance initially increases with the number of iterations
- Degradation after a particular number of iterations
- Best results reported on the 40^{th} iteration, with an overall F-measure of 81%, after adding 5 positive and 5 negative examples at every step
- Method is successful because it uses both cross-language and within-language knowledge

Corpus-based: Polarity

Wei and Pal, 2010 – Structural correspondence learning

- Frame multilingual polarity detection as a special case of domain adaptation, where cross-lingual pivots are used to model the correspondence between features from both domains.
- Instead of using the entire feature set (like Wan, 2009), from the machine translated text only the pivots are maintained (based on method proposed by Blitzer et al., 2007) and appended to the original text; the rest is discarded as MT noise.
- Then apply SCL to find a low dimensional representation shared by both languages.
- They show that using only pivot features outperforms using the entire feature set.
- Improve over Wan, 2009 by 2.2% in overall accuracy.

Hybrid: Polarity

Boyd-Graber and Resnik, 2010 – Multilingual Supervised LDA

- Model for sentiment analysis that learns consistent "topics" from a multilingual corpus.
- Both topics and assignments are probabilistic:
  - Topic = latent concept that is represented through a probabilistic distribution of vocabulary words in multilingual corpora; it displays a consistent meaning and relevance to observed sentiment.
  - Each document is represented as a probability distribution over all the topics and is assigned a sentiment score.
- Alternative to co-training that does not require parallel text or machine translation systems.
- Can use comparable text originating from multiple languages in a holistic framework and provides the best results when it is bridged through a dictionary or a foreign language WordNet aligned with the English WordNet.

Hybrid: Polarity (cont.)

Boyd-Graber and Resnik, 2010 – Multilingual Supervised LDA

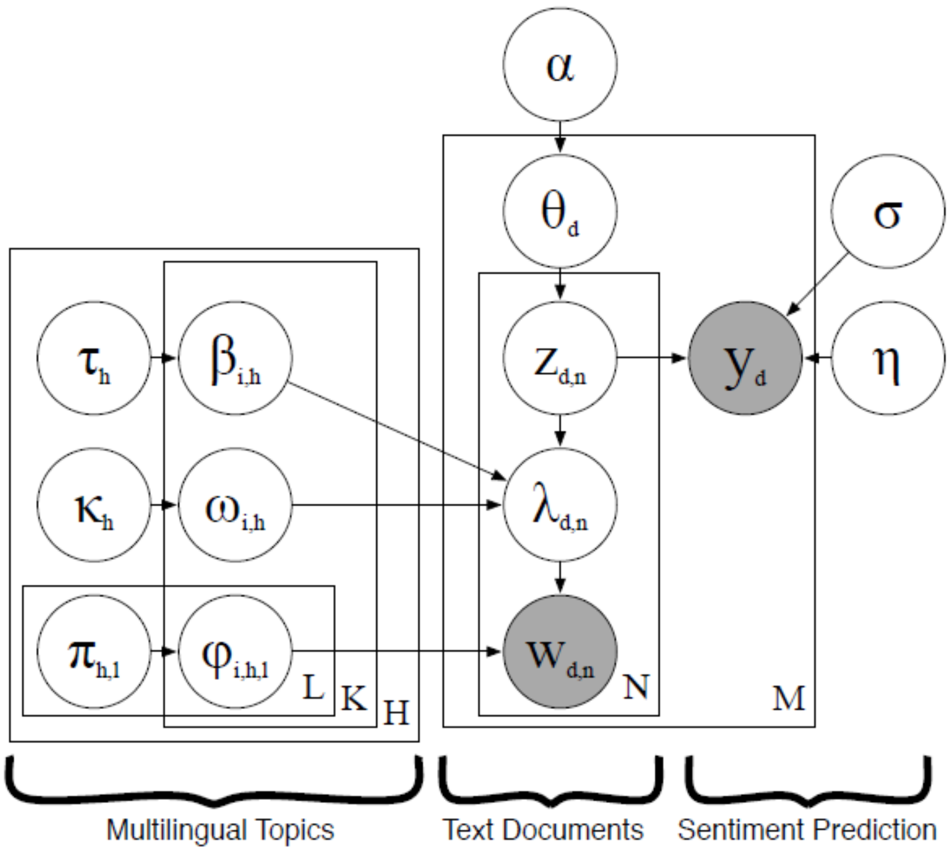

- Model views sentiment across all languages from the perspective imparted by the topics present.
- Better than when porting resources from a source to a target language, when sentiment is viewed from the perspective of the *donor* language.

VI. What works, what doesn’t

Comparative results

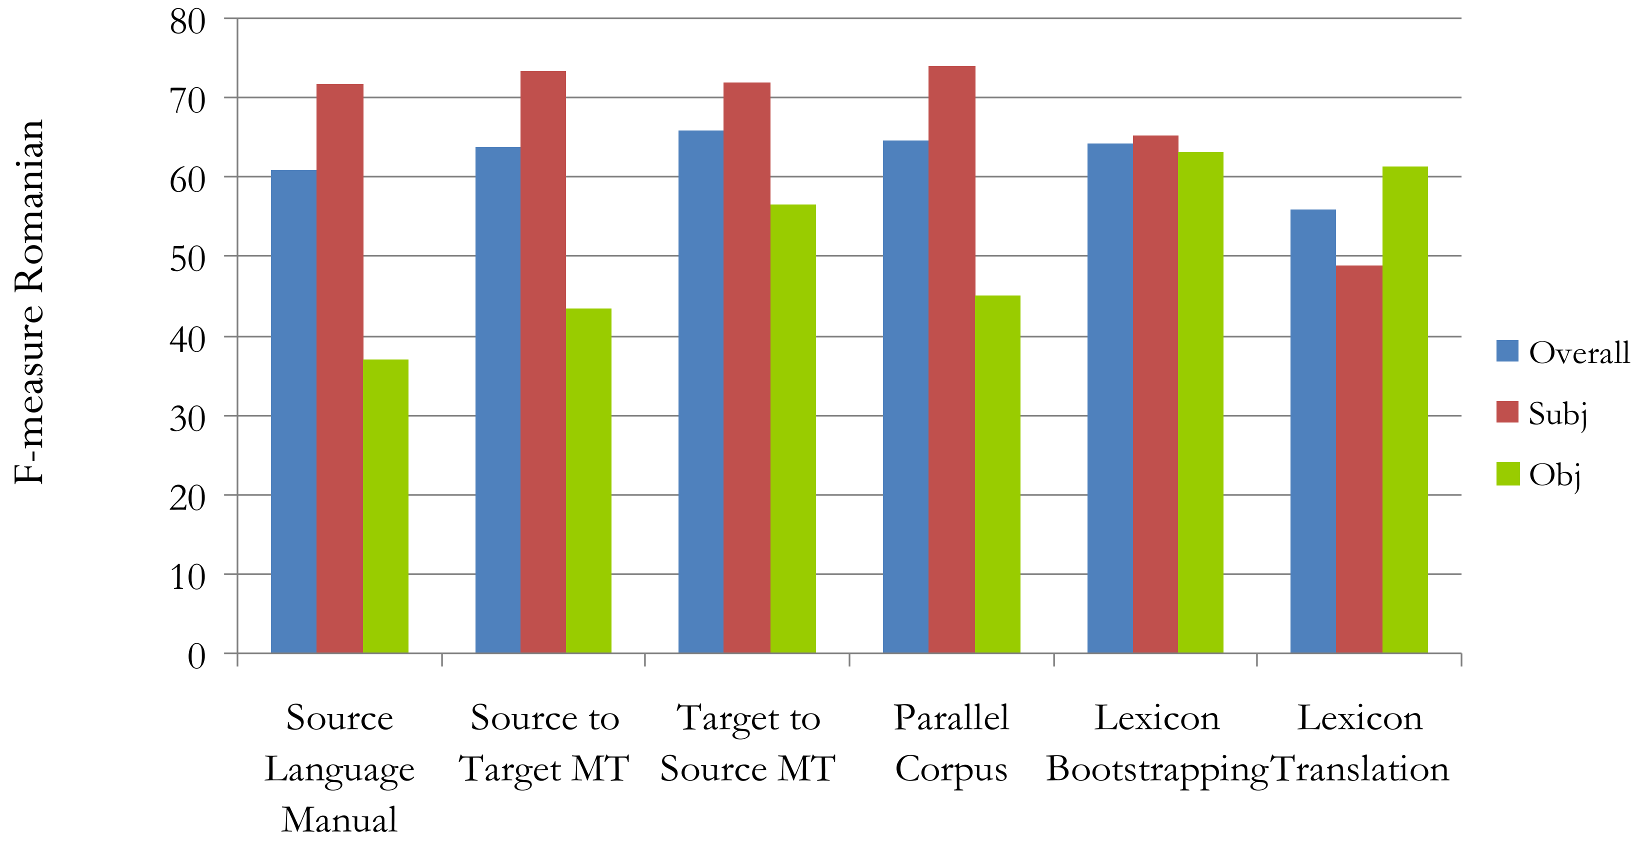

Comparative results

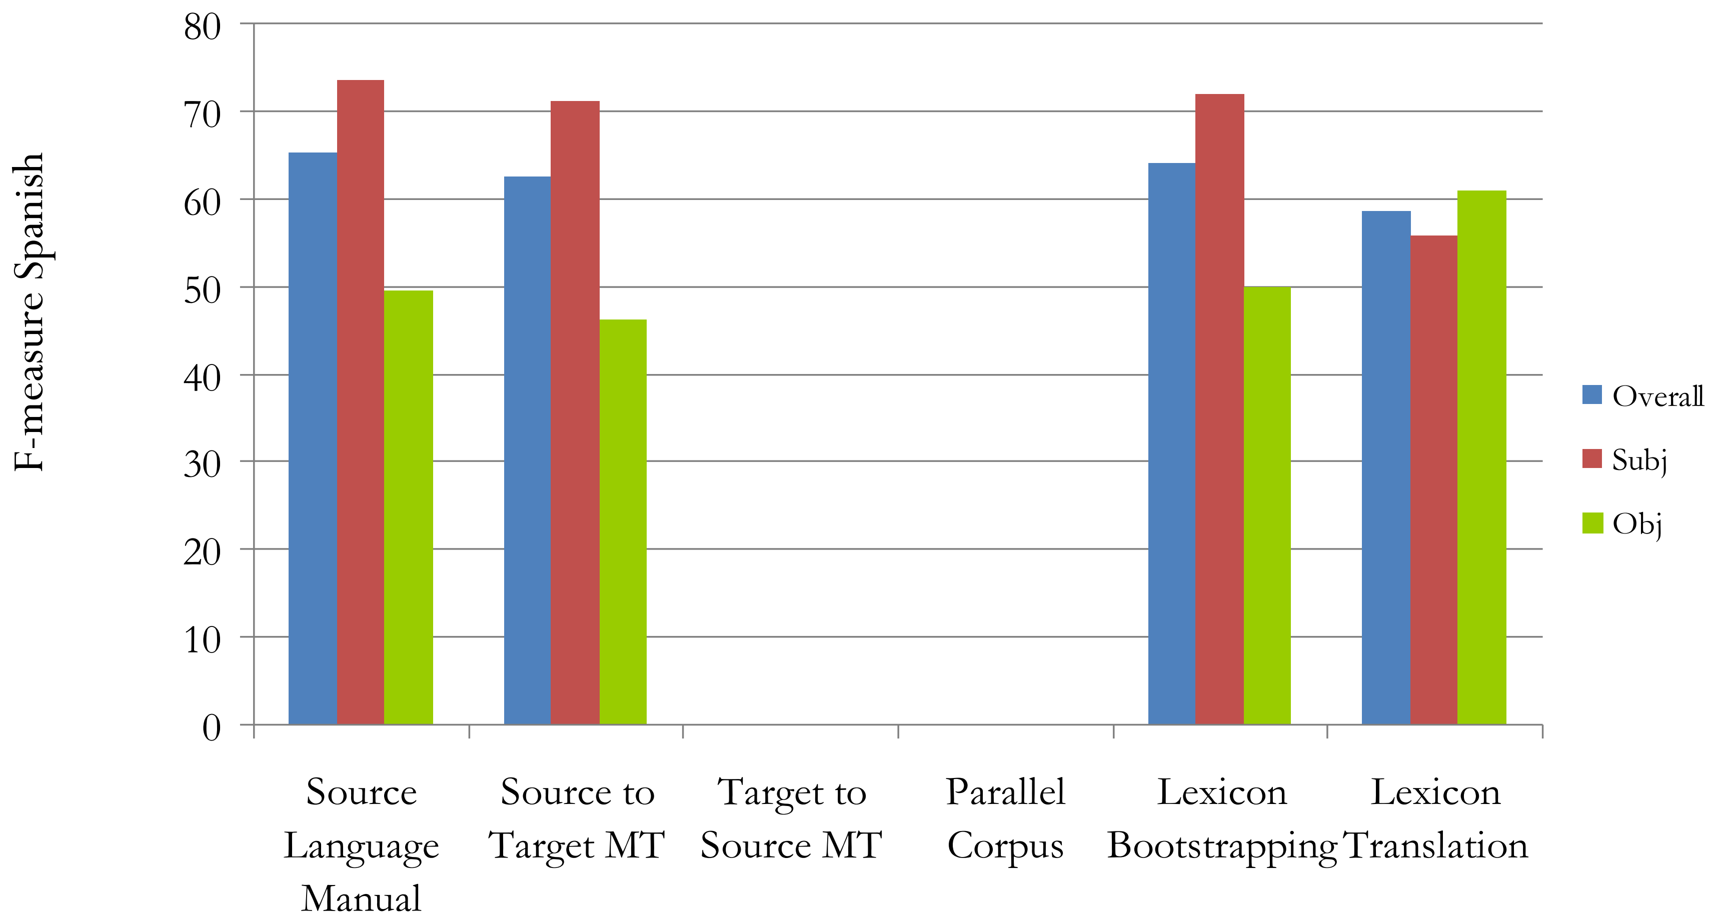

Lessons Learned

- Best Scenario: Manually Annotated Corpora
  - The best scenario is when a corpus manually annotated for subjectivity exists in the target language
  - Unfortunately, this is rarely the case, as large manually annotated corpora exist only for a handful of languages
    - e.g., the English MPQA corpus

Lessons Learned

- Second Best: Corpus-based Cross-Lingual Projections
  - The second best option is to construct an annotated data set by doing cross-lingual projections from a major language
  - This assumes a “bridge” can be created between the target language and a major language such as English, in the form of parallel texts constructed via manual or automatic translations
    - Target language translation tends to outperform source language translation
    - Automatic translation leads to performance comparable to manual translations

Lessons Learned

- Third Best: Bootstrapping a Lexicon
  - The third option is to use bootstrapping starting with a set of seeds
  - No advanced language processing tools are required, only a dictionary in the target language
  - The seed set is expanded using words related found in the dictionary
  - Running the process for several iterations can result in large lexicons with several thousands entries

Lessons Learned

- Fourth best: translating a lexicon
  - If none of the previous methods is applicable, the last resort is to automatically translate an already existing lexicon from a major language
  - The only requirements are a subjectivity lexicon in the source language, and a bilingual dictionary
  - Although very simple and efficient (a lexicon of over 5,000 entries can be created in seconds), the accuracy of the method is rather low, mainly due to the challenges that are typical to a context-free translation process: ambiguity, morphology, phrase translations, etc.

Conclusions

- Sentiment and subjectivity analysis is a very active area in natural language processing
  - Contributions from growing number of research teams
  - Hot commercial applications
    - Understanding social media
- There is growing interest in enabling its application to other languages
  - Continuously increasing number of documents in languages other than English

References

C. O. Alm, D. Roth, and R. Sproat. Emotions from text: Machine learning for text-based emotion prediction. In Proceedings of the Human Language Technologies Conference/Conference on Empirical Methods in Natural Language Processing (HLT/EMNLP-2005), pages 347–354, Vancouver, Canada, 2005.

K. Balog, G. Mishne, and M. de Rijke. Why are they excited? identifying and explaining spikes in blog mood levels. In Proceedings of the 11th Meeting of the European Chapter of the Association for Computational Linguistics (EACL-2006), 2006.

C. Banea, R. Mihalcea, and J. Wiebe. A bootstrapping method for building subjectivity lexicons for languages with scarce resources. In Proceedings of the Learning Resources Evaluation Conference (LREC 2008), Marrakech, Morocco, 2008.

C. Banea, R. Mihalcea, J. Wiebe, and S. Hassan. Multilingual subjectivity analysis using machine translation. In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP 2008), Honolulu, Hawaii, 2008.

M. Bautin, L. Vijayarenu, and S. Skiena. International sentiment analysis for news and blogs. In Proceedings of the International Conference on Weblogs and Social Media, Seattle, WA, 2008.

J. Boyd-Graber and P. Resnik. Holistic sentiment analysis across languages: Multilingual Supervised Latent Dirichlet Allocation. In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP 2010), MIT, Massachusetts, 2010.

G. Carenini, R. Ng, and X. Zhou. Summarizing emails with conversational cohesion and subjectivity. In Proceedings of the Association for Computational Linguistics: Human Language Technologies (ACL-HLT 2008), Columbus, Ohio, 2008.

A. Esuli and F. Sebastiani. Determining term subjectivity and term orientation for opinion mining. In Proceedings the 11th Meeting of the European Chapter of the Association for Computational Linguistics (EACL-2006), pages 193–200, Trento, IT, 2006.

A. Esuli and F. Sebastiani. SentiWordNet: A publicly available lexical resource for opinion mining. In Proceedings of the 5th Conference on Language Resources and Evaluation (LREC 2006), Genova, IT, 2006.

References

A. Hassan, A. Abu-Jbara, R. Jha, D. Radev. Identifying the semantic orientation of foreign words. In Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics, pages 592-597, Portland, Oregon, June 2011.

V. Hatzivassiloglou and K. McKeown. Predicting the semantic orientation of adjectives. In Proceedings of the Conference of the European Chapter of the Association for Computational Linguistics, pages 174–181, 1997.

M. Hu and B. Liu. Mining and summarizing customer reviews. In Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining 2004 (KDD 2004), pages 168–177, Seattle, Washington, 2004.

Y. Hu, J. Duan, X. Chen, B. Pei, and R. Lu. A new method for sentiment classification in text retrieval. In IJCNLP, pages 1–9, 2005.

N. Kaji and M. Kitsuregawa. Automatic construction of polarity-tagged corpus from html documents. In Proceedings of the International Conference on Computational Linguistics / Association for Computational Linguistics, Sydney, Australia, 2006.

N. Kaji and M. Kitsuregawa. Building lexicon for sentiment analysis from massive collection of html documents. In Proceedings of the Conference on Empirical Methods in Natural Language Processing, Prague, Czech Republic, 2007.

H. Kanayama and T. Nasukawa. Fully automatic lexicon expansion for domain-oriented sentiment analysis. In Proceedings of the Conference on Empirical Methods in Natural Language Processing, Sydney, Australia, 2006.

H. Kanayama, T. Nasukawa, and H. Watanabe. Deeper sentiment analysis using machine translation technology. In International Conference on Computational Linguistics, 2004.

N. Kando, T. Mitamura, and T. Sakai. Introduction to the ntcir-6 special issue. ACM Transactions on Asian Language Information Processing (TALIP), 7(2), 2008.

S.-M. Kim and E. Hovy. Identifying and analyzing judgment opinions. In Proceedings of the Human Language Technology Conference - North American chapter of the Association for Computational Linguistics, New York City, NY, 2006.

J. Kim, J.-J. Li, J.-H. Lee. Evaluating multilanguage-comparability of subjectivity analysis systems. In Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, pages 595-603, Uppsala, Sweden, July 2010.

References

N. Kobayashi, K. Inui, K. Tateishi, and T. Fukushima. Collecting evaluative expressions for opinion extraction. In Proceedings of IJCNLP 2004, pages 596–605, 2004.

T. K. Landauer, P. Foltz, and D. Laham. Introduction to latent semantic analysis. Discourse Processes, 25, 1998.

J. Li, , and M. Sun. Experimental study on sentiment classification of chinese review using machine learning techniques. In International Conference on Natural Language Processing and Knowledge Engineering, 2007.

L. Lloyd, D. Kechagias, and S. Skiena. Lydia: A system for large-scale news analysis. In String Processing and Information Retrieval (SPIRE 2005), 2005.

R. Mihalcea, C. Banea, and J. Wiebe. Learning multilingual subjective language via cross-lingual projections. In Proceedings of the Association for Computational Linguistics, Prague, Czech Republic, 2007.

G. Miller. Wordnet: A lexical database. Communication of the ACM, 38(11), 1995.

G. Miller, C. Leacock, T. Randee, and R. Bunker. A semantic concordance. In Proceedings of the 3rd DARPA Workshop on Human Language Technology, Plainsboro, New Jersey, 1993.

F. Och and H. Ney. Improved statistical alignment models. In Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics, Hongkong, October 2000.

B. Pang and L. Lee. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In Proceedings of the 42nd Meeting of the Association for Computational Linguistics, Barcelona, Spain, July 2004.

v. Pérez-Rosas, C. Banea, R. Mihalcea. Learning sentiment lexicons in Spanish. In Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC 2012). Istanbul, Turkey, May 2012.

G. Pitel and G. Grefenstette. Semi-automatic building method for a multidimensional affect dictionary for a new language. In Proceedings of the Sixth International Language Resources and Evaluation (LREC' 08), 2008.

R. Quirk, S. Greenbaum, G. Leech, and J. Svartvik. A Comprehensive Grammar of the English Language. Longman, New York, 1985.

References

E. Riloff and J. Wiebe. Learning extraction patterns for subjective expressions. In Conference on Empirical Methods in Natural Language Processing (EMNLP-03), pages 105–112, 2003.

E. Riloff, J. Wiebe, and T. Wilson. Learning subjective nouns using extraction pattern bootstrapping. In Proceedings of the Seventh Conference on Natural Language Learning (CoNLL-2003), 2003.

P. Stone. General Inquirer: Computer Approach to Content Analysis. MIT Press, 1968.

C. Strapparava and R. Mihalcea. Semeval-2007 task 14: Affective text. In Proceedings of the 4th International Workshop on the Semantic Evaluations (SemEval 2007), Prague, Czech Republic, 2007.

Y. Suzuki, H. Takamura, and M. Okumura. Application of semi-supervised learning to evaluative expression classification. In Proceedings of the 7th International Conference on Intelligent Text Processing and Computational Linguistics, 2006.

H. Takamura, T. Inui, and M. Okumura. Latent variable models for semantic orientations of phrases. In Proceedings of the 11th Meeting of the European Chapter of the Association for Computational Linguistics, 2006.

P. Turney. Thumbs up or thumbs down? semantic orientation applied to unsupervised classification of reviews. In Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL 2002), pages 417–424, Philadelphia, 2002.

V. Vapnik. The Nature of Statistical Learning Theory. Springer, New York, 1995.

X. Wan. Using bilingual knowledge and ensemble techniques for unsupervised chinese sentiment analysis. In Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing, 2008.

X. Wan. Co-training for cross-lingual sentiment classification. In Proceedings of the Joint Conference of the Association of Computational Linguistics and the International Joint Conference on Natural Language Processing, Singapore, August 2009.

B. Wei and C. Pal. Cross Lingual Adaptation: An experiment on sentiment classifications. In Proceedings of the ACL 2010 Conference Short Papers (ACL 2010), pages 258-262., Uppsala, Sweden, July 2010.

References

J. Wiebe, R. Bruce, and T. O'Hara. Development and use of a gold-standard data set for subjectivity classifications. In Proceedings of the 37th annual meeting of the Association for Computational Linguistics on Computational Linguistics, pages 246–253, 1999.

J. Wiebe and R. Mihalcea. Word sense and subjectivity. In Proceedings of the Annual Meeting of the Association for Computational Linguistics, Sydney, Australia, 2006.

J. Wiebe and E. Riloff. Creating subjective and objective sentence classifiers from unannotated texts. In Proceedings of the 6th International Conference on Intelligent Text Processing and Computational Linguistics (CICLing-2005) (invited paper), Mexico City, Mexico, 2005.

J. Wiebe, T. Wilson, and C. Cardie. Annotating expressions of opinions and emotions in language. Language Resources and Evaluation, 39(2-3):165–210, 2005.

T. Wilson. Fine-grained Subjectivity and Sentiment Analysis: Recognizing the Intensity, Polarity, and Attitudes of private states. PhD thesis, Intelligent Systems Program, University of Pittsburgh, 2007.

T. Wilson. Fine-grained Subjectivity and Sentiment Analysis: Recognizing the Intensity, Polarity, and Attitudes of Private States. PhD thesis, University of Pittsburgh, 2008.

T. Wilson, J. Wiebe, and P. Hoffmann. Recognizing contextual polarity in phrase-level sentiment analysis. In Proceedings of Human Language Technologies Conference/Conference on Empirical Methods in Natural Language Processing (HLT/EMNLP 2005), Vancouver, Canada, 2005.

I. Witten and E. Frank. Data Mining: Practical Machine Learning Tools and Techniques. Morgan Kaufmann, 2005.

H. Yu and V. Hatzivassiloglou. Towards answering opinion questions: Separating facts from opinions and identifying the polarity of opinion sentences. In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2003), pages 129–136, Sapporo, Japan, 2003.

T. Zagibalov and J. Carroll. Automatic seed word selection for unsupervised sentiment classification of Chinese text. In Proceedings of the Conference on Computational Linguistics, 2008.