

EECS 498-004: Introduction to Natural Language Processing

Instructor: Prof. Lu Wang

Computer Science and Engineering

University of Michigan

<https://web.eecs.umich.edu/~wangluxy/>

Logistics

- March 17: Dr. Tershia Pinder-Grover (Director, Center for Research on Learning and Teaching in Engineering) will attend our lecture.
- Feel free to reach out to discuss course projects! Many of you have done so, others are also encouraged!
- Last homework will be out this week.

Outline



- Sentiment analysis tasks
- Sentiment lexicons
- Semi-supervised learning of lexicons

Positive or negative movie review?



- unbelievably disappointing



- Full of zany characters and richly applied satire, and some great plot twists



- this is the greatest screwball comedy ever filmed



- It was pathetic. The worst part about it was the boxing scenes.

Apple iPhone 12 Pro Max - 128 GB - Silver - AT&T

★★★★★ (8,093)



- 6.7-Inch All-Screen OLED Display
- Super Retina XDR Display
- 2778-By-1284-Pixel Resolution At 460 Ppi
- 128GB Capacity
- A14 Bionic Chip

Color: Silver



Capacity

128 GB

Connectivity

AT&T

Product details

#3 in Mobile Phones

Apple · iPhone · iPhone 12 Pro · iOS · 6.7" · Facial Recognition · 12 MP front camera · 12 MP rear camera · Smartphone · With Wireless Charging

5G goes Pro. A14 Bionic rockets past every other smartphone chip. The iPhone 12 Pro Max features a 6.7" Super Retina XDR display, LiDAR scanner for ultrafast and accurate depth maps of whatever space you're in and MagSafe wireless charging. And Ceramic Shield delivers four times ... [More](#)

Reviews

4.7




8,093 reviews




Twitter sentiment versus Gallup Poll of Consumer Confidence

Brendan O'Connor, Ramnath Balasubramanyan, Bryan R. Routledge, and Noah A. Smith. 2010. From Tweets to Polls: Linking Text Sentiment to Public Opinion Time Series. In ICWSM-2010



Twitter sentiment:

Johan Bollen, Huina Mao, Xiaojun Zeng. 2011.
[Twitter mood predicts the stock market](#),
Journal of Computational Science 2:1, 1-8.
10.1016/j.jocs.2010.12.007.



Sentiment analysis has many other names

- Opinion extraction
- Opinion mining
- Sentiment mining
- Subjectivity analysis

Why sentiment analysis?

- *Movie*: is this review positive or negative?
- *Products*: what do people think about the new iPhone?
- *Public sentiment*: how is consumer confidence? Is despair increasing?
- *Politics*: what do people think about this candidate or issue?
- *Prediction*: predict election outcomes or market trends from sentiment

Scherer Typology of Affective States

- **Emotion:** brief organically synchronized, evaluation of a major event
 - *angry, sad, joyful, fearful, ashamed, proud, elated*
- **Mood:** diffuse non-caused low-intensity long-duration change in subjective feeling
 - *cheerful, gloomy, irritable, listless, depressed, buoyant*
- **Interpersonal stances:** affective stance toward another person in a specific interaction
 - *friendly, flirtatious, distant, cold, warm, supportive, contemptuous*
- **Attitudes:** enduring, affectively colored beliefs, dispositions towards objects or persons
 - *liking, loving, hating, valuing, desiring*
- **Personality traits:** stable personality dispositions and typical behavior tendencies
 - *nervous, anxious, reckless, morose, hostile, jealous*

Scherer Typology of Affective States

- Emotion and Mood
 - Annoyance in talking to dialog systems
 - Uncertainty of students in tutoring
 - Detecting trauma or depression in conversations (in person or online)
- Interpersonal Stance
 - Romantic interest, flirtation, friendliness
 - Alignment/accommodation/entrainment
- Attitudes = Sentiment (positive or negative)
 - Movie or Products or Politics: is a text positive or negative?
 - “Twitter mood predicts the stock market.”
- Personality Traits
 - Open, Consciencious, Extroverted, Anxious

Scherer Typology of Affective States

- **Emotion:** brief organically synchronized, evaluation of a major event
 - *angry, sad, joyful, fearful, ashamed, proud, elated*
- **Mood:** diffuse non-caused low-intensity long-duration change in subjective feeling
 - *cheerful, gloomy, irritable, listless, depressed, buoyant*
- **Interpersonal stances:** affective stance toward another person in a specific interaction
 - *friendly, flirtatious, distant, cold, warm, supportive, contemptuous*
- **Attitudes: enduring, affectively colored beliefs, dispositions towards objects or persons**
 - *liking, loving, hating, valuing, desiring*
- **Personality traits:** stable personality dispositions and typical behavior tendencies
 - *nervous, anxious, reckless, morose, hostile, jealous*

Sentiment Analysis

- Extraction of **opinions** and **attitudes** from text and speech
- When we say “sentiment analysis”
 - often meaning a binary or an ordinal task
 - like X/ dislike X
 - one-star to 5-stars

Sentiment Analysis (broader view)

- Sentiment analysis is the detection of **attitudes**
“enduring, affectively colored beliefs, dispositions towards objects or persons”

Emily told Charlie that the new movie is disappointing.

1. **Holder (source)** of attitude
2. **Target (aspect)** of attitude
3. **Type** of attitude, e.g. **polarity**:
 - *positive, negative, neutral*, often together with *strength*
4. **Text** containing the attitude
 - Sentence or entire document

Sentiment Analysis

- Simplest task:
 - Is the attitude of this text positive or negative?
- More complex:
 - Rank the attitude of this text from 1 to 5
- Advanced:
 - Detect the target, source, or complex attitude types

Sentiment Analysis

- Simplest task:
 - Is the attitude of this text positive or negative?
- More complex:
 - Rank the attitude of this text from 1 to 5
- Advanced:
 - Detect the target, source, or complex attitude types

Sentiment Classification in Movie Reviews

Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment Classification using Machine Learning Techniques. EMNLP-2002, 79—86.

Bo Pang and Lillian Lee. 2004. A Sentimental Education: Sentiment Analysis Using Subjectivity Summarization Based on Minimum Cuts. ACL, 271-278

- Polarity detection:
 - Is an IMDB movie review positive or negative?
- Data: *Polarity Data 2.0*:
 - <http://www.cs.cornell.edu/people/pabo/movie-review-data>

IMDB data in the Pang and Lee database



when `_star wars_` came out some twenty years ago , the image of traveling throughout the stars has become a commonplace image . [...]

when han solo goes light speed , the stars change to bright lines , going towards the viewer in lines that converge at an invisible point .

cool .

`_october sky_` offers a much simpler image—that of a single white dot , traveling horizontally across the night sky . [...]



“ snake eyes ” is the most aggravating kind of movie : the kind that shows so much potential then becomes unbelievably disappointing .

it’s not just because this is a brian depalma film , and since he’s a great director and one who’s films are always greeted with at least some fanfare .

and it’s not even because this was a film starring nicolas cage and since he gives a brauvara performance , this film is hardly worth his talents .

Baseline Algorithm (adapted from Pang and Lee)

- Tokenization
- Feature Extraction (unigrams, bigrams, POS tags)
- Classification using different classifiers
 - Naïve Bayes
 - Logistic regression
 - SVM

Negation in Sentiment Analysis

They have not succeeded, and will never succeed, in breaking the will of this valiant people.

Negation in Sentiment Analysis

They have not **succeeded**, and will never succeed, in breaking the will of this valiant people.

Negation in Sentiment Analysis

They have **not succeeded**, and will never succeed, in breaking the will of this valiant people.

Negation in Sentiment Analysis

They **have not succeeded, and will never succeed, in breaking the will of this valiant people.**

Negation

Das, Sanjiv and Mike Chen. 2001. Yahoo! for Amazon: Extracting market sentiment from stock message boards. In Proceedings of the Asia Pacific Finance Association Annual Conference (APFA).
Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment Classification using Machine Learning Techniques. EMNLP-2002, 79—86.

Add NOT_ to every word between negation and following punctuation:

didn't like this movie , but I



didn't NOT_like NOT_this NOT_movie but I

Problems: What makes reviews hard to classify?

- Subtlety:

- Perfume review in *Perfumes: the Guide*:

- “If you are reading this because it is your darling fragrance, please wear it at home exclusively, and tape the windows shut.”

- Dorothy Parker on Katherine Hepburn

- “She runs the gamut of emotions from A to B”

Thwarted Expectations and Ordering Effects

- “This film should be **brilliant**. It sounds like a **great** plot, the actors are **first grade**, and the supporting cast is **good** as well, and Stallone is attempting to deliver a good performance. However, it **can't hold up.**”
- Well as usual Keanu Reeves is nothing special, but surprisingly, the **very talented** Laurence Fishbourne is **not so good** either, I was surprised.

Outline

- Sentiment analysis tasks
- • Sentiment lexicons
- Semi-supervised learning of lexicons

What is a Lexicon?

- A (usually hand-built) list of words that correspond to some meaning or class
- Possibly with numeric values
- Commonly used as simple classifiers, or as features to more complex classifiers

- Adjectives

- positive: honest important mature large patient

- He is the only honest man in Washington.

- Her 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

- negative: harmful hypocritical inefficient insecure

- It was a macabre and hypocritical circus.

- Why are they being so inefficient ?

- Verbs
 - positive: **praise, love**
 - negative: **blame, criticize**
- Nouns
 - positive: **pleasure, enjoyment**
 - negative: **pain, criticism**

Phrases

- Phrases containing adjectives and adverbs
 - positive: high intelligence, low cost
 - negative: little variation, many troubles

The General Inquirer

Philip J. Stone, Dexter C Dunphy, Marshall S. Smith,
Daniel M. Ogilvie. 1966. The General Inquirer: A
Computer Approach to Content Analysis. MIT Press

Positiv (1915 words)

Negativ (2291 words)

1	Entry
2586	DAKOTA
2587	DAMAGE#1
2588	DAMAGE#2
2589	DAMN
2590	DAMNABLE
2591	DAMNED
2592	DAMP
2593	DANCE#1
2594	DANCE#2
2595	DANCE#3
2596	DANCER
2597	DANGER
2598	DANGEROUS
2599	DANISH
2600	DARE
2601	DARING
2602	DARK
2603	DARKEN
2604	DARKNESS
2605	DARLING

MPQA Subjectivity Cues Lexicon

Theresa Wilson, Janyce Wiebe, and Paul Hoffmann (2005). Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis. Proc. of HLT-EMNLP-2005.

Riloff and Wiebe (2003). Learning extraction patterns for subjective expressions. EMNLP-2003.

- 6885 words
- Is a subjective word positive or negative?
 - Strongly or weakly?
- <http://mpqa.cs.pitt.edu/lexicons/>
- GNU GPL

type=**weaksubj** len=1 word1=**abandoned** pos1=adj stemmed1=n priorpolarity=**negative**
type=weaksubj len=1 word1=**abandonment** pos1=noun stemmed1=n priorpolarity=negative
type=weaksubj len=1 word1=**abandon** pos1=verb stemmed1=y priorpolarity=negative
type=strongsubj len=1 word1=**abase** pos1=verb stemmed1=y priorpolarity=negative
type=strongsubj len=1 word1=**abacement** pos1=anypos stemmed1=y priorpolarity=negative
type=strongsubj len=1 word1=**abash** pos1=verb stemmed1=y priorpolarity=negative
type=weaksubj len=1 word1=**abate** pos1=verb stemmed1=y priorpolarity=negative

LIWC (Linguistic Inquiry and Word Count)

Pennebaker, J.W., Booth, R.J., & Francis, M.E. (2007). Linguistic Inquiry and Word Count: LIWC 2007. Austin, TX

- Home page: <http://www.liwc.net/>
- 2300 words, >70 classes
- **Affective Processes**
 - negative emotion (*bad, weird, hate, problem, tough*)
 - positive emotion (*love, nice, sweet*)
- **Cognitive Processes**
 - Tentative (*maybe, perhaps, guess*), Inhibition (*block, constraint*)
- **Pronouns, Negation** (*no, never*), **Quantifiers** (*few, many*)
- Not free though!

LIWC: Linguistic Inquiry and Word Count

Positive Emotion	Negative Emotion	Insight	Inhibition	Family	Negate
appreciat*	anger*	aware*	avoid*	brother*	aren't
comfort*	bore*	believe	careful*	cousin*	cannot
great	cry	decid*	hesitat*	daughter*	didn't
happy	despair*	feel	limit*	family	neither
interest	fail*	figur*	oppos*	father*	never
joy*	fear	know	prevent*	grandf*	no
perfect*	griev*	knew	reluctan*	grandm*	nobod*
please*	hate*	means	safe*	husband	none
safe*	panic*	notice*	stop	mom	nor
terrific	suffers	recogni*	stubborn*	mother	nothing
value	terrify	sense	wait	niece*	nowhere
wow*	violent*	think	wary	wife	without

Bing Liu Opinion Lexicon

Minqing Hu and Bing Liu. Mining and Summarizing Customer Reviews. ACM SIGKDD-2004.

- 6786 words
 - 2006 positive
 - 4783 negative

SentiWordNet

Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010 SENTIWORDNET 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining. LREC-2010

- Home page: <http://sentiwordnet.isti.cnr.it/>
- All WordNet synsets automatically annotated for degrees of positivity, negativity, and neutrality/objectiveness
- [estimable(J,3)] “may be computed or estimated”
Pos 0 Neg 0 Obj 1
- [estimable(J,1)] “deserving of respect or high regard”
Pos .75 Neg 0 Obj .25

Disagreements between polarity lexicons

	Opinion Lexicon	General Inquirer	SentiWordNet	LIWC
MPQA	33/5402 (0.6%)	49/2867 (2%)	1127/4214 (27%)	12/363 (3%)
Opinion Lexicon		32/2411 (1%)	1004/3994 (25%)	9/403 (2%)
General Inquirer			520/2306 (23%)	1/204 (0.5%)
SentiWordNet				174/694 (25%)
LIWC				

Online review data

Movie review excerpts (IMDb)

- 10** A great movie. This film is just a wonderful experience. It's surreal, zany, witty and slapstick all at the same time. And terrific performances too.
- 1** This was probably the worst movie I have ever seen. The story went nowhere even though they could have done some interesting stuff with it.

Restaurant review excerpts (Yelp)

- 5** The service was impeccable. The food was cooked and seasoned perfectly... The watermelon was perfectly square ... The grilled octopus was ... mouthwatering...
- 2** ...it took a while to get our waters, we got our entree before our starter, and we never received silverware or napkins until we requested them...

Book review excerpts (GoodReads)

- 1** I am going to try and stop being deceived by eye-catching titles. I so wanted to like this book and was so disappointed by it.
- 5** This book is hilarious. I would recommend it to anyone looking for a satirical read with a romantic twist and a narrator that keeps butting in

Product review excerpts (Amazon)

- 5** The lid on this blender though is probably what I like the best about it... enables you to pour into something without even taking the lid off! ... the perfect pitcher! ... works fantastic.
- 1** I hate this blender... It is nearly impossible to get frozen fruit and ice to turn into a smoothie... You have to add a TON of liquid. I also wish it had a spout ...

Analyzing the polarity of each word in IMDB


Potts, Christopher. 2011. On the negativity of negation. SALT 20, 636-659.

- How likely is each word to appear in each sentiment class?
- Count(“bad”) in 1-star, 2-star, 3-star, etc.
- But can't use raw counts:

- Instead, **likelihood**:
$$P(w | c) = \frac{f(w, c)}{\sum_{w \in C} f(w, c)}$$

- Make them comparable between words
 - **Scaled likelihood**:

$$\frac{P(w | c)}{P(w)}$$




“Potts diagrams”

Potts, Christopher. 2011. NSF workshop on restructuring adjectives.


Positive scalars

good




Negative scalars

disappointing




Emphatics

totally




Attenuators


somewhat




great




bad




absolutely




fairly




excellent




terrible



utterly



pretty




Other sentiment feature: Logical negation

Potts, Christopher. 2011. On the negativity of negation. *SALT* 20, 636-659.


- Is logical negation (*no, not*) associated with negative sentiment?
- Potts experiment:
 - Count negation (*not, n't, no, never*) in online reviews

Potts 2011 Results: More negation in negative sentiment

IMDB (4,073,228 tokens)



Five-star reviews (846,444 tokens)



Outline

- Sentiment analysis tasks
- Sentiment lexicons
- • Semi-supervised learning of lexicons

Outline

- Sentiment analysis tasks
- Sentiment lexicons
- ➔ • Semi-supervised learning of lexicons
 - Sentiment axis
 - Label propagation

Semi-supervised learning of lexicons

- Use a small amount of information
 - A few labeled examples
 - A few hand-built patterns or relations with the examples
- To bootstrap a lexicon

Semantic Axis Methods

(An et al., 2018, Turney and Littman 2003)

- General idea:
- Start with seed words like *good* or *bad* for the two poles
- For each word to be added to lexicon
 - Compute a word representation
 - Use this to measure its distance from the poles
 - Assign it to the pole it is closer to

Initial seeds for different domains

- (1) Start with a single large seed lexicon
- (2) Choose different seed words for different genres

Domain	Positive seeds	Negative seeds
General	good, lovely, excellent, fortunate, pleasant, delightful, perfect, loved, love, happy	bad, horrible, poor, unfortunate, unpleasant, disgusting, evil, hated, hate, unhappy
Twitter	love, loved, loves, awesome, nice, amazing, best, fantastic, correct, happy	hate, hated, hates, terrible, nasty, awful, worst, horrible, wrong, sad
Finance	successful, excellent, profit, beneficial, improving, improved, success, gains, positive	negligent, loss, volatile, wrong, losses, damages, bad, litigation, failure, down, negative

Compute representation

- Can just use off-the-shelf static embeddings
 - word2vec, GloVe, etc.
- Or train word embeddings in a given corpus

Represent each pole

Start with embeddings for seed words:

$$S^+ = \{E(w_1^+), E(w_2^+), \dots, E(w_n^+)\}$$

$$S^- = \{E(w_1^-), \hat{E}(w_2^-), \dots, E(w_m^-)\}$$

Pole centroids are:

$$\mathbf{V}^+ = \frac{1}{n} \sum_1^n E(w_i^+)$$

$$\mathbf{V}^- = \frac{1}{m} \sum_1^m E(w_i^-)$$

Semantic axis is:

$$\mathbf{V}_{axis} = \mathbf{V}^+ - \mathbf{V}^-$$

Word score is cosine with axis

$$\begin{aligned} \text{score}(w) &= \cos(E(w), \mathbf{V}_{axis}) \\ &= \frac{E(w) \cdot \mathbf{V}_{axis}}{\|E(w)\| \|\mathbf{V}_{axis}\|} \end{aligned}$$

Label Propagation Methods

- Alternative to axis methods: propagate sentiment labels on word graphs
- First proposed by Hatzivassiloglou and McKeown (1997)

Hatzivassiloglou and McKeown intuition for identifying word polarity

- Adjectives conjoined by “*and*” have same polarity
 - Fair **and** legitimate, corrupt **and** brutal
 - *fair **and** brutal, *corrupt **and** legitimate
- Adjectives conjoined by “*but*” do not have the same polarity
 - fair **but** brutal

Step 1: seed set construction

- Label **seed set** of 1336 adjectives
 - 657 positive
 - adequate central clever famous intelligent remarkable
reputed sensitive slender thriving...
 - 679 negative
 - contagious drunken ignorant lanky listless primitive
strident troublesome unresolved unsuspecting...

Step 2: expand candidates

- Expand seed set to conjoined adjectives

Google

"was nice and"

[Nice location in Porto and the front desk staff **was nice and helpful** ...](#)

www.tripadvisor.com/ShowUserReviews-g189180-d206904-r12068... +1

Mercure Porto Centro: Nice location in Porto and the front desk staff **was nice and helpful** - See traveler reviews, 77 candid photos, and great deals for Porto, ...

nice, helpful

[If a girl **was nice and classy**, but had some vibrant purple dye in ...](#)

answers.yahoo.com > Home > All Categories > Beauty & Style > Hair +1

4 answers - Sep 21


Question: Your personal opinion or what you think other people's opinions might ...

Top answer: I think she would be cool and confident like katy perry :)

nice, classy


Step 3: graph construction

- Supervised classifier/heuristic rule assigns “polarity similarity” to each word pair, resulting in graph:



Step 4: graph partitioning

- Clustering for partitioning the graph into two



Output polarity lexicon

- Positive

- bold decisive disturbing generous good honest important large mature patient peaceful positive proud sound stimulating straightforward strange talented vigorous witty...

- Negative

- ambiguous cautious cynical evasive harmful hypocritical inefficient insecure irrational irresponsible minor outspoken pleasant reckless risky selfish tedious unsupported vulnerable wasteful...

Output polarity lexicon

- Positive

- bold decisive **disturbing** generous good honest important large mature patient peaceful positive proud sound stimulating straightforward **strange** talented vigorous witty...

- Negative


- ambiguous **cautious** cynical evasive harmful hypocritical inefficient insecure irrational irresponsible minor **outspoken pleasant** reckless risky selfish tedious unsupported vulnerable wasteful...

Another Label Propagation Method (Hamilton et al., 2016)

- 1. Define a graph: connecting each word with k nearest neighbor

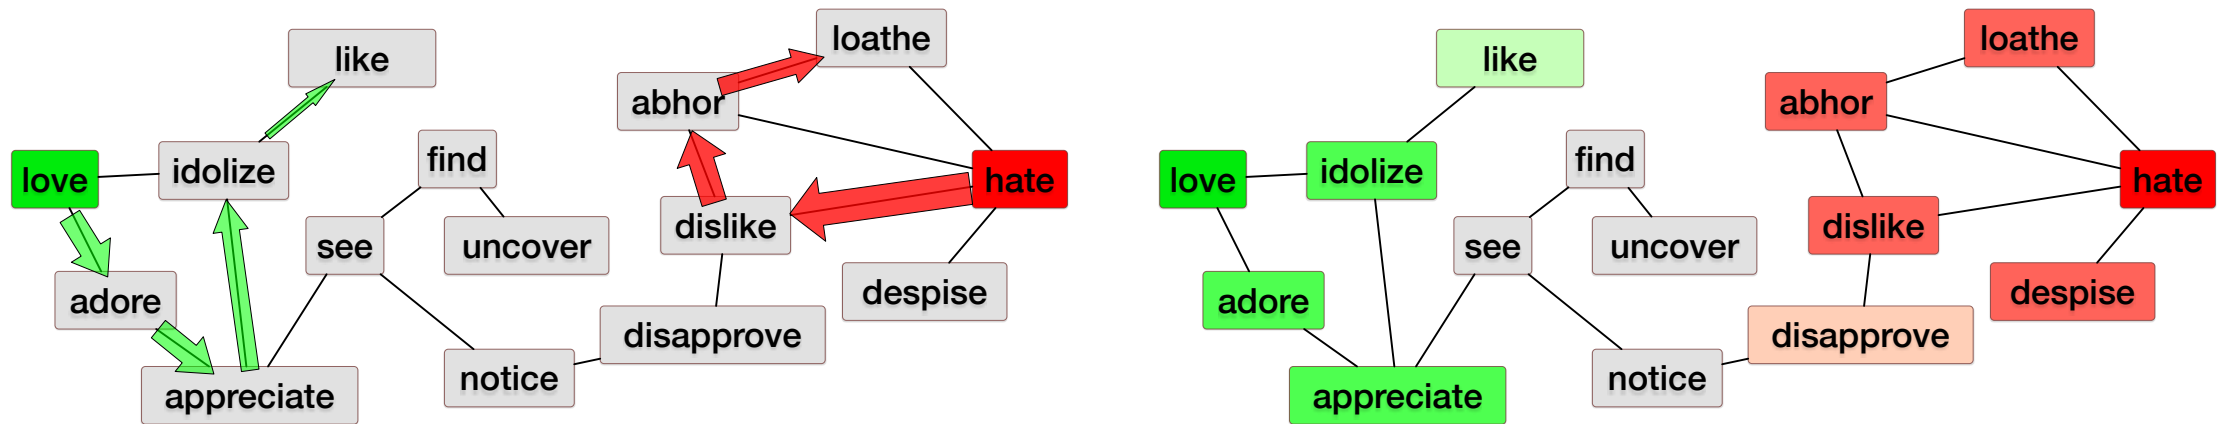
$$\mathbf{E}_{i,j} = \arccos \left(-\frac{\mathbf{w}_i^\top \mathbf{w}_j}{\|\mathbf{w}_i\| \|\mathbf{w}_j\|} \right)$$

- 2. Define a seed set (pos and neg words)



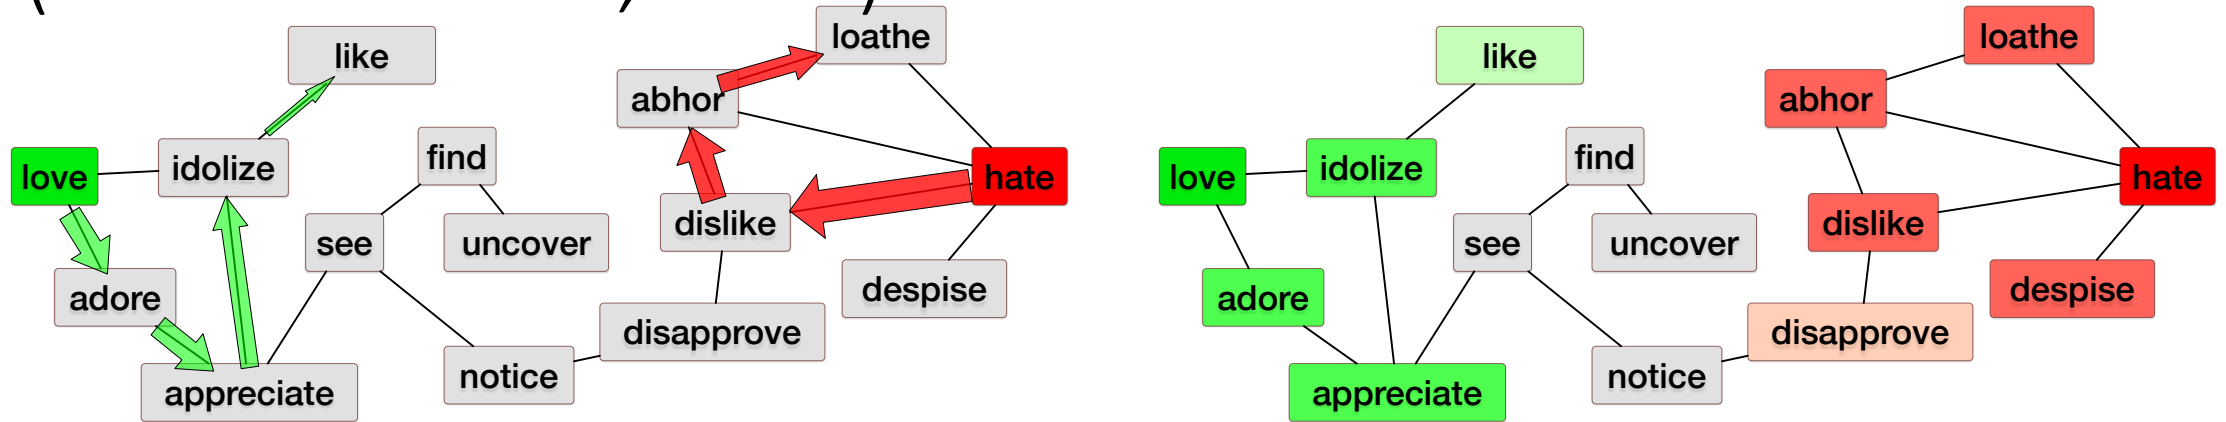
Another Label Propagation Method (Hamilton et al., 2016)

- 3. Propagate polarities from the seed set: randomly walk on the graph



- Random walk: start at a node and then choose a node to move to with probability proportional to the edge probability
- A word's polarity score for a seed set is proportional to the probability of a random walk from the seed set landing on that word

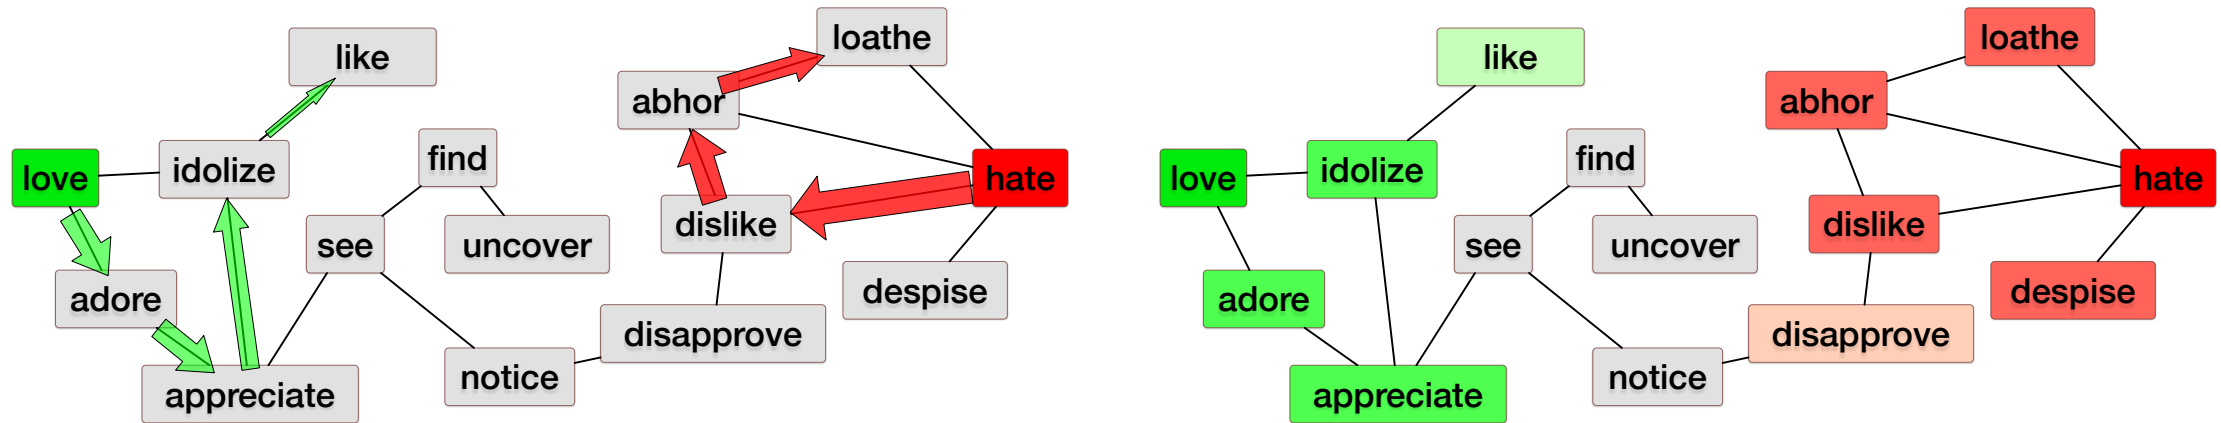
Another Label Propagation Method (Hamilton et al., 2016)



- 4. Create word scores:
 - Walking from positive and negative seed sets
 - Gives $\text{rawscore}^+(w_i)$ and $\text{rawscore}^-(w_i)$
 - Combine into one score:

$$\text{score}^+(w_i) = \frac{\text{rawscore}^+(w_i)}{\text{rawscore}^+(w_i) + \text{rawscore}^-(w_i)}$$

Another Label Propagation Method (Hamilton et al., 2016)



- 5. Assign confidence via bootstrap sampling:
 - Compute the propagation B times over random subsets of the positive and negative seed sets
 - The standard deviation of the bootstrap sampled polarity scores gives a confidence measure.