


KCC 2013 자연어처리 및 정보검색 최근 동향 워크샵

Quality Analysis of User Generated Web Documents
유저 생성 웹문서의 퀄러티 파악

Hye-Jin Min
Department of Computer Science
KAIST

2013.06.26

Background

- User generated web documents
  - Common users actively post and share their opinions or daily episodes through blogs, wikis, or other social media on the web

- Types of user generated web documents
  - Customer Reviews
  - Online Forums
  - Comments on UGC
    - YouTube, blogs, reviews

Background

- Problems on accessing user generated web documents
  - Time-consuming
    - Hard to keep track of all the available data
  - Unreliable Quality
    - Quality of the documents is varied

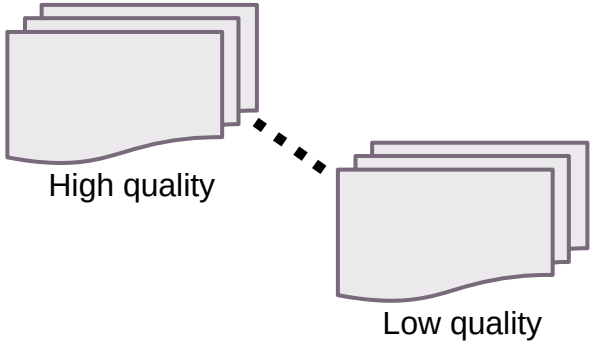

Background

- User generated web documents

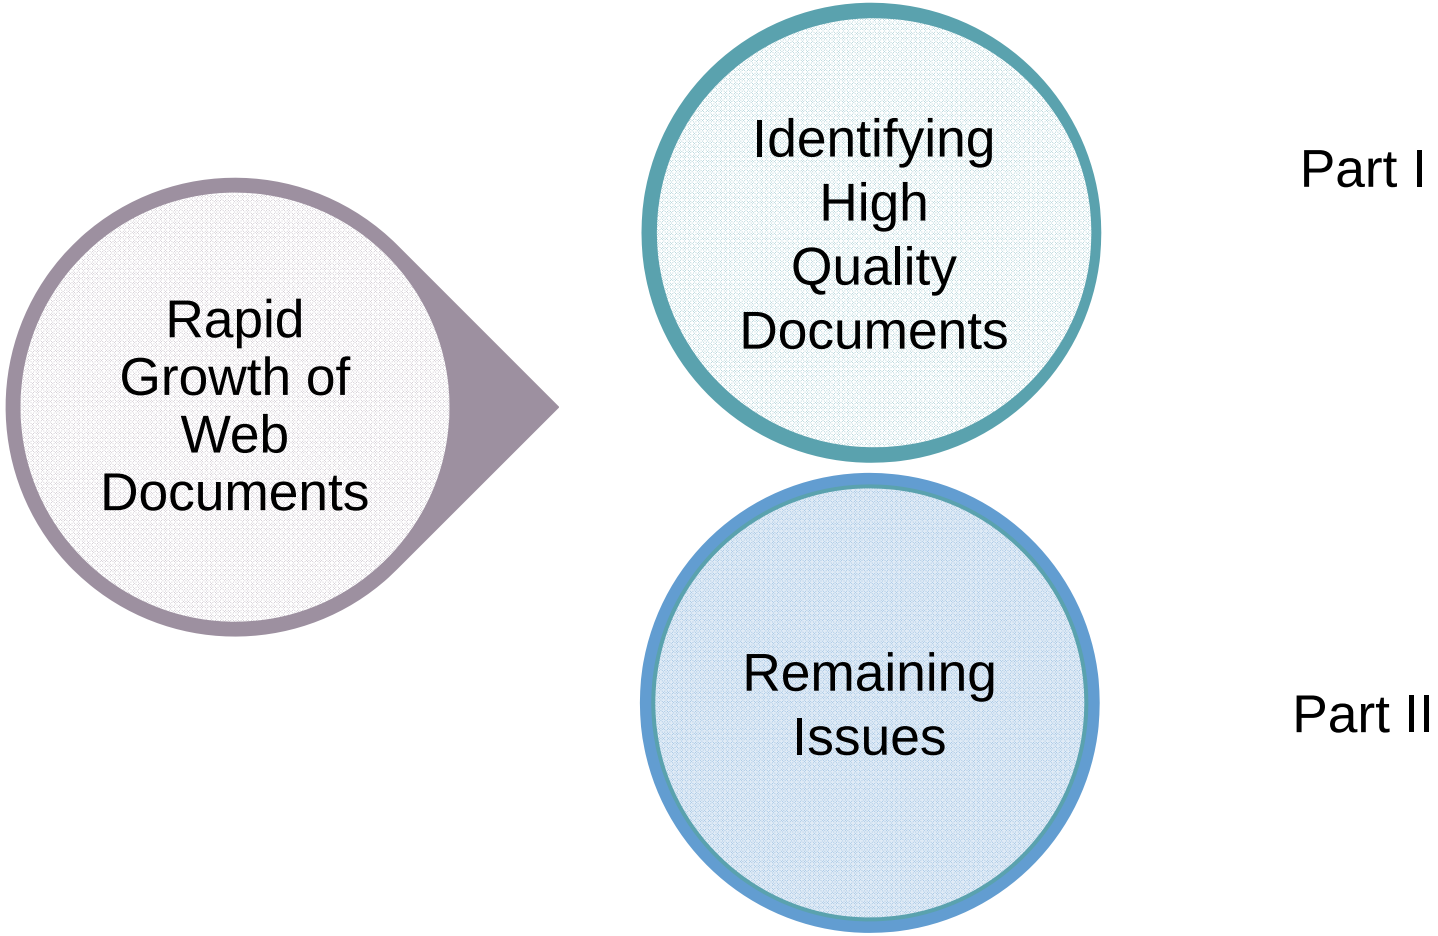

Outline

- Part I
  - Measures for the quality
  - Features for assessing the quality
  - Applications
- Part II
  - Personalized quality analysis
  - Filtering false documents
- Concluding Remarks

Part I

- Measures and gold standard for the quality
- Features for assessing the quality

Measures for the Quality

- Diverse measures with respect to the main purpose for quality assessment

Types of Document	Sources	Underlying Measure
Comments on video data	Video sharing websites (e.g., YouTube)	**Acceptability**
Online reviews	Online shopping websites (e.g., Amazon.com, eBay, and CNET)	**Helpfulness/ Usefulness**
Forum data	Online web forum sites (e.g., Nable.com and Google groups)	**Informativeness**
Blog posts	Aggregated blogs on the web	**Credibility**

Overview

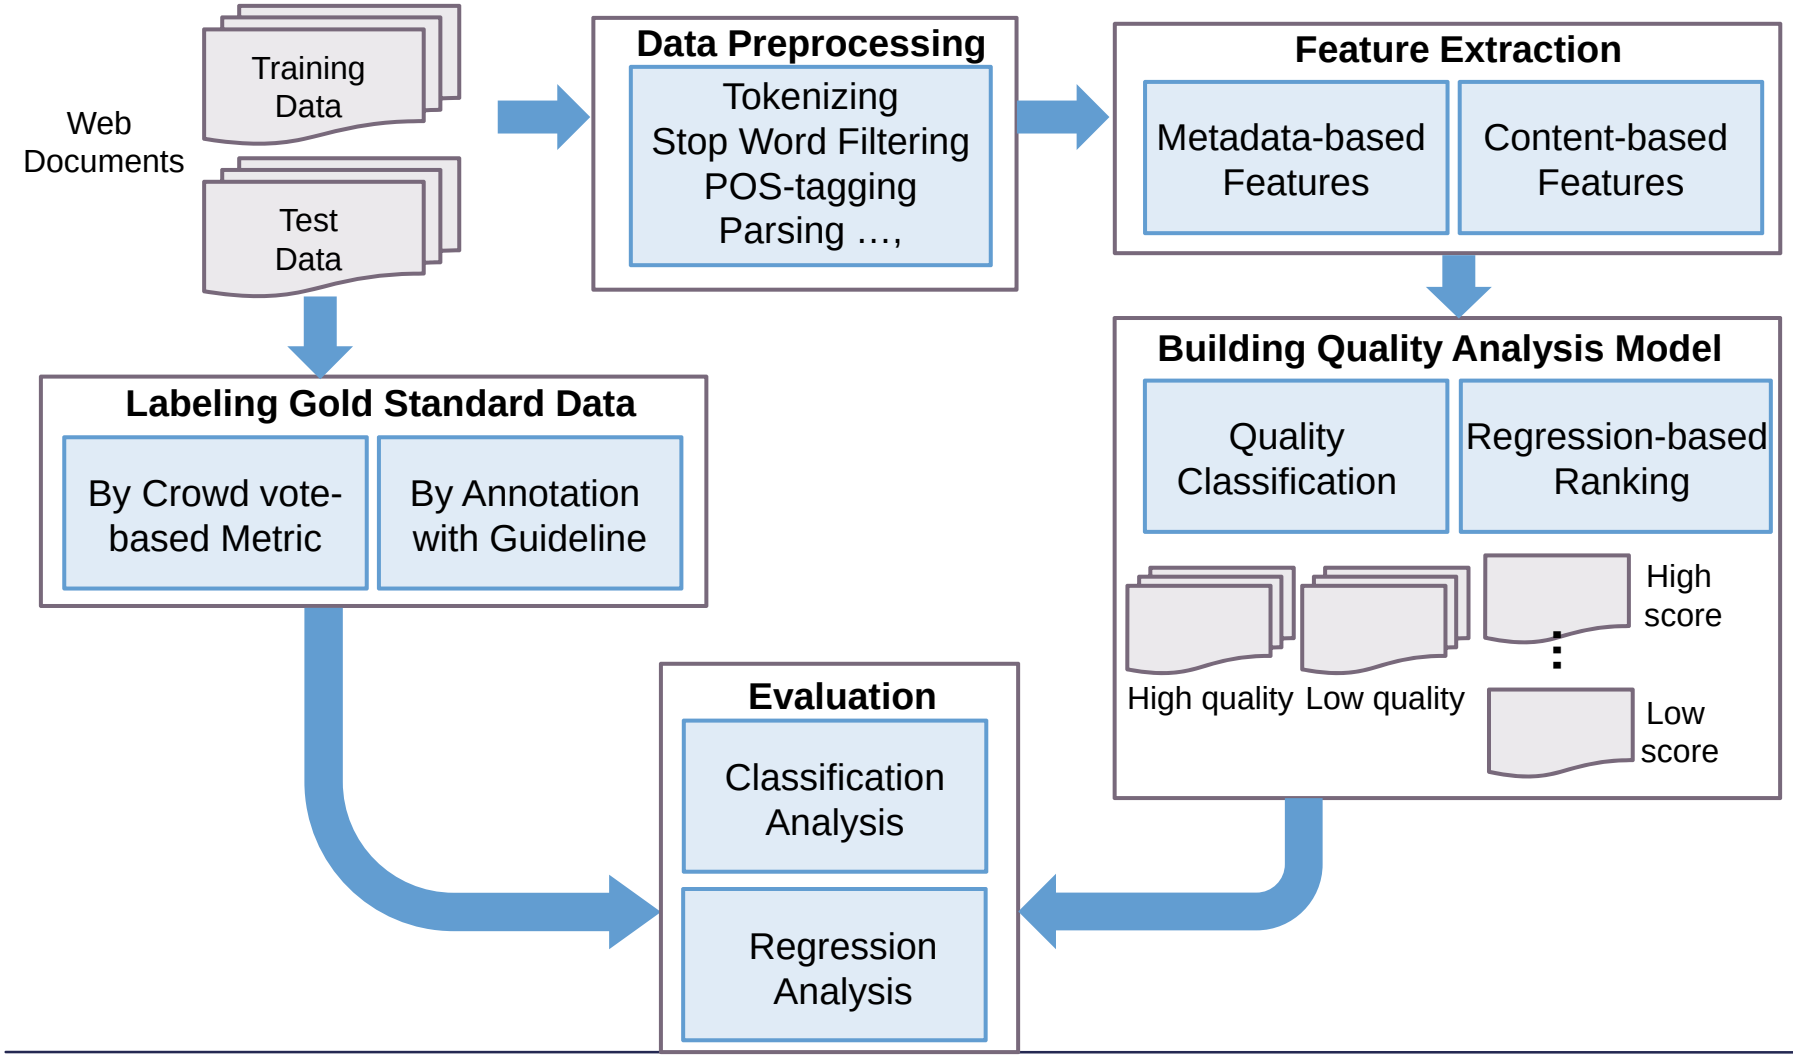
Web Documents
Training Data
Test Data
Data Preprocessing
Tokenizing
Stop Word Filtering
POS-tagging
Parsing …,
Feature Extraction
Metadata-based Features
Content-based Features
Building Quality Analysis Model
Quality Classification
Regression-based Ranking
High score
Low score
High quality
Low quality
Labeling Gold Standard Data
By Crowd vote-based Metric
By Annotation with Guideline
Evaluation
Classification Analysis
Regression Analysis

Gold Standard for the Quality

- Crowd vote-based Metric
  - Like/Dislike
  - Helpfulness
  - Amazon Verified Purchase (AVP)
- Examples of quality function
  - Kim et al., 2006

$$f(r \in R) = \frac{No.ofHelpfulRating(r)}{No.ofHelpfulRating(r) + No.ofUnHelpfulRating(r)}$$

  - O'Mahony and Smyth, 2009

    A review is helpful if and only if >75% of the feedback is positive
  - Kokkodis, 2012

    A reviewer with many AVP reviews is considered trustworthy

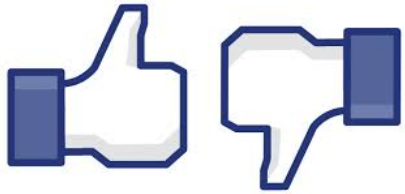

"24,279 of 24,279 people found the following review helpful"

Gold Standard for the Quality

- Gold standard Model
  - Guideline for the gold standard model

    "What is a helpful/best review?"
- Examples of guideline
  - Liu et al., 2007

    - The coverage of (interesting and important) aspects of products
    - An appropriate amount of evidence for the opinions
    - A good format for the content
  - Tsur and Rappoport, 2009

    - "A virtual core review"

    : The most prominent words related to the given book's story or comments on important aspects (e.g., genre and author).

Gold Standard for the Quality

- Example reviews with diverse levels of quality

Best Review

I purchased this camera about six months ago (…) because it was very reasonably priced (about $200). (…) Here are the things I have loved about this camera:
BATTERY - this camera has the best battery of any digital camera I have ever owned or used. …
EASY TO USE - I was able to … (…)
I cannot stress how highly I recommend this camera. I will never buy another digital camera besides Canon again (…)

Good Review

The Sony DSC "P10" Digital Camera is the top pick for CSC.(…) This camera I purchased through a Private Dealer cost me $400.86 (…) Purchase this camera from a wholesale dealer for the best price $377.00.
Great Photo Even in dim light w/o a flash.
The p10 is very compact. Can easily fit into any pocket. (…). What makes the p10 the top pick is it comes with a rechargeable lithium battery. (…)
It's also the best resolution on the market. 6.0
Also the best price for a major brand.

Fair Review

There is nothing wrong with the 2100 except for the very noticeable delay between pics. The camera's digital processor takes about 5 seconds after a photo is snapped to ready itself for the next one. Otherwise, the optics, the 3X optical zoom and the 2 megapixel resolution are fine for anything from Internet apps to 8" x 10" print enlarging. It is competent, not spectacular, but it gets the job done at an agreeable price point.

Bad Review

I want to point out that you should never buy a generic battery, like the person from San Diego who reviewed the S410 on May 15, 2004, was recommending. (…) don't think if your generic battery explodes you can sue somebody and win millions. These batteries are made in sweatshops in China, India and Korea, and I doubt you can find anybody to sue. So play it safe, both for your own sake and the camera's sake. If you want a spare, get a real Canon one.

Gold Standard for the Quality

- Agreement Study
  - Liu et al., 2007
    - Kappa Statistics: 0.8142 (digital cameras)
    - 4 classes (best, good, fair, bad)
    - Showing high consistency
  - Chen and Tseng, 2011
    - Kappa Statistics: 0.7253 (digital cameras), 0.7928 (mp3 players),
    - 5 classes (high-quality, medium-quality, low-quality, duplicate, spam)

Gold Standard for the Quality

- Pros and cons of gold standard
  - Crowd vote-based metric
    - Easy to acquire
    - Suffering from bias (Liu et al., 2007)
  - Gold standard model
    - Free from bias
    - Subjective
    - Agreement Issues

3 Types of Bias in Ground-truth

- Imbalance Vote Bias
- Winner Circle Bias
- Early Bird Bias

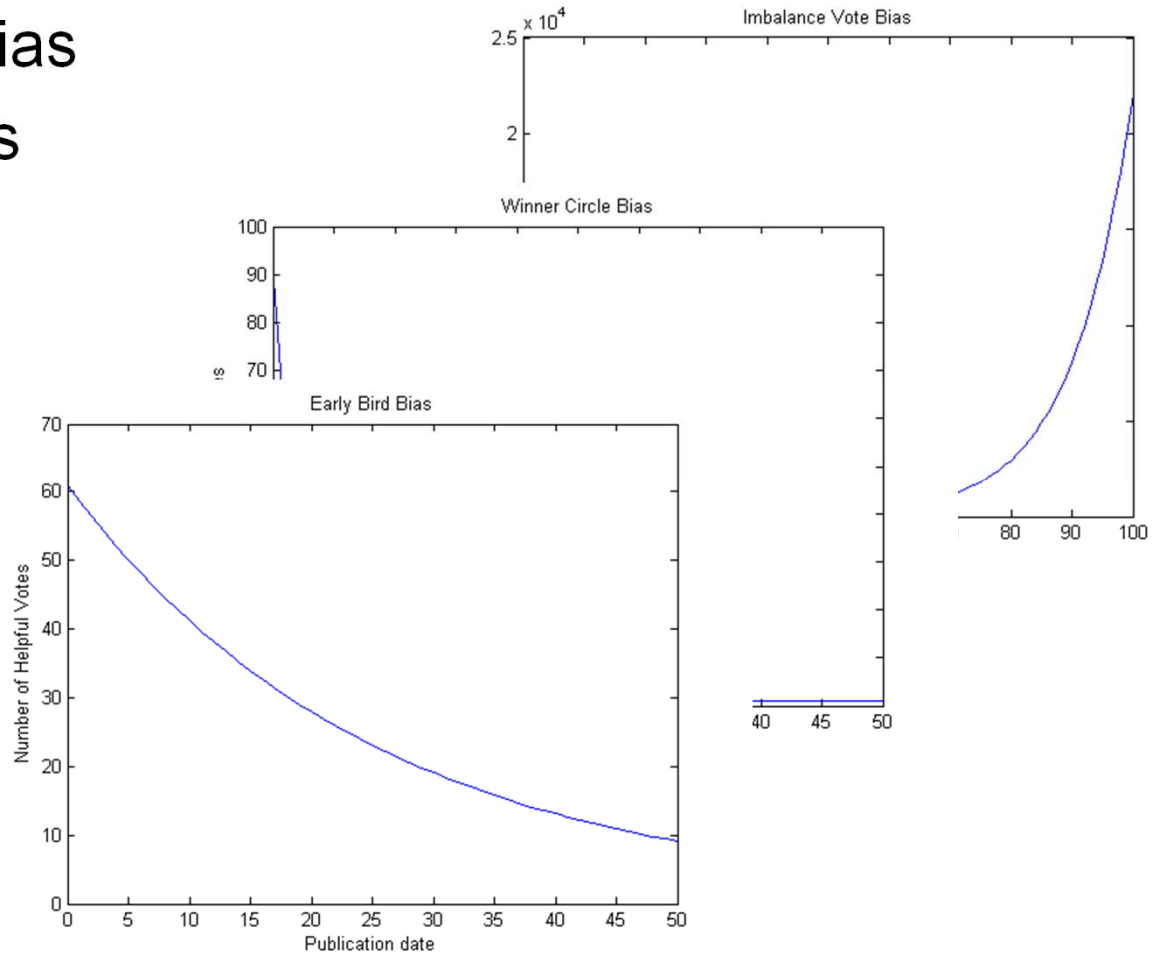

Quality Analysis Models

	Regression-based Ranking	Classification
Methods	• SVM regression (Kim et al., 2006; Tanaka et al., 2012; Ngo-Ye and Sinha, 2012) • Ordinal Logistic Regression (Cao et al., 2011)	• SVM (Liu et al., 2007; Weimer et al., 2007; Wanas et al., 2008; Chen and Tseng, 2011) • Maximum Entropy (Hoang et al., 2008) • JRip, J48, Naïve Bayes (O'Mahony and Smyth, 2009) • Naïve Bayes, Decision Trees, Logistic Regression, Multi-layer Perceptron (MLP) (Kokkodis, 2012)
Evaluation Measures	• Correlation Coefficient (Kim et al., 2006; Tanaka et al., 2012) • Misclassification Rate, Akaike's Information Criterion; Lift Ratio (Cao et al., 2011) • RRSE, RAE, RMSE, MAE (Ngo-Ye and Sinha, 2012)	• Accuracy (Algur et al., 2010) • P, R F-Score (Park et al., 2010; Li et al., 2011, Chen and Tseng, 2011) • Lift (Kokodis, 2012)

Quality Analysis Models

	Classification and Ranking (Hong et al., 2012)	Clustering and Labeling (Hiremath et al., 2010)	Ranking (Chen and Tseng, 2011)
Methods	• SVM	• K-means clustering and weighting clusters	• A proposed quality score indicating the degree of association between a review and 'high-quality' class
Evaluation Measures	• Accuracy • NDCG@n	• Quartile measure technique	• Precision@n

Meaning of the Features

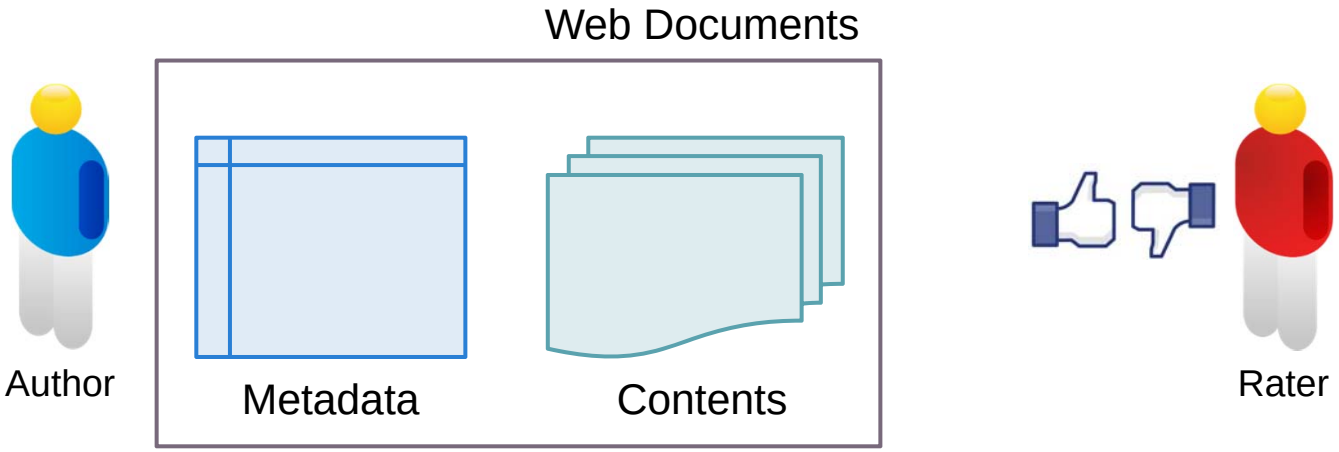

Meaning of the Features

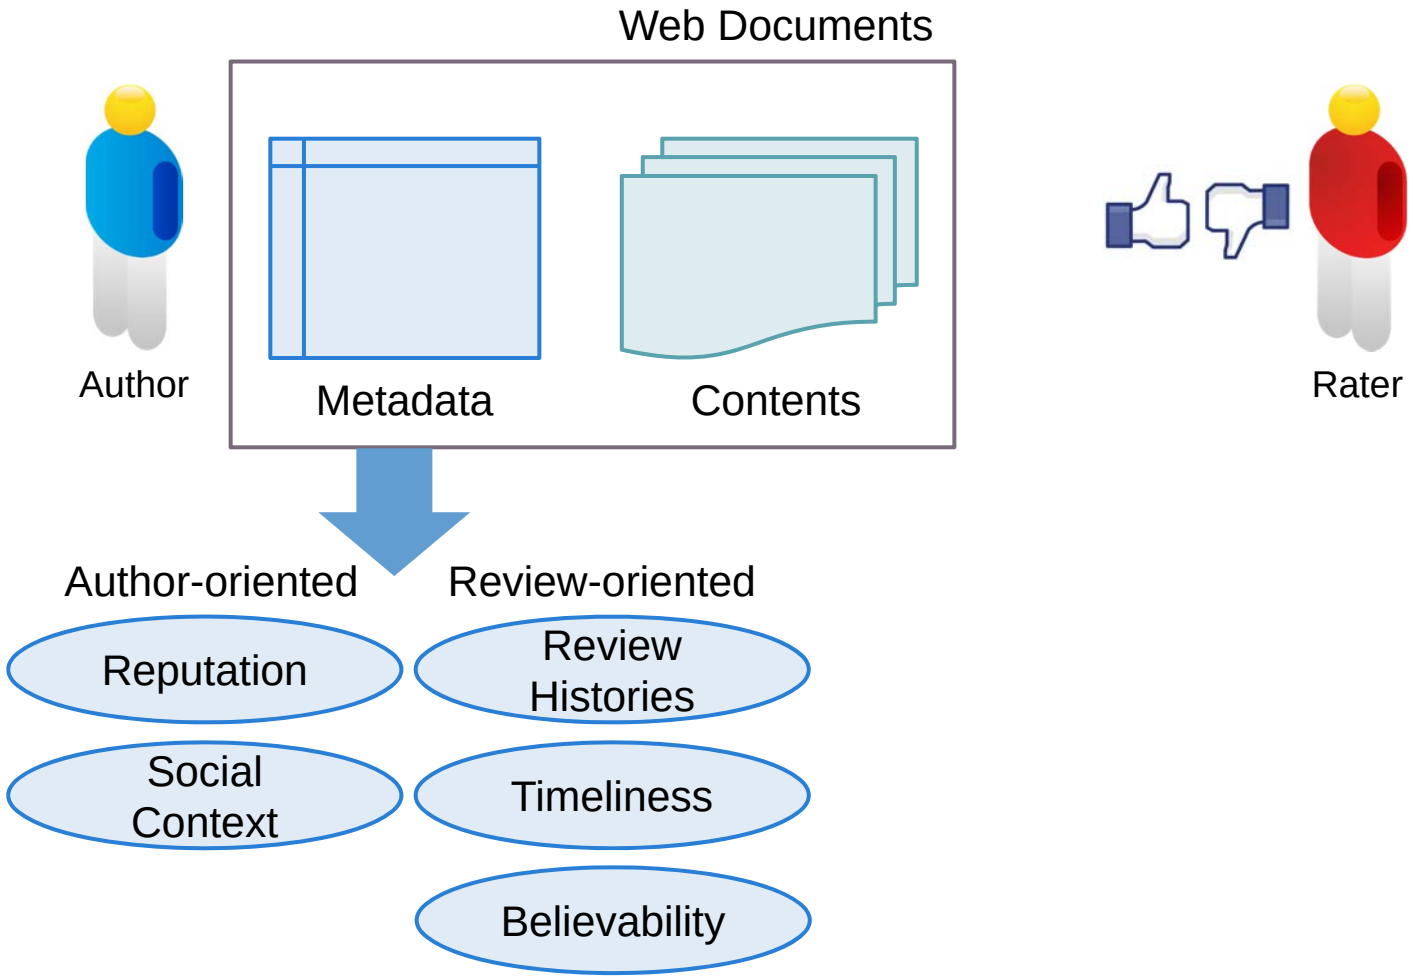

Meaning of the Features

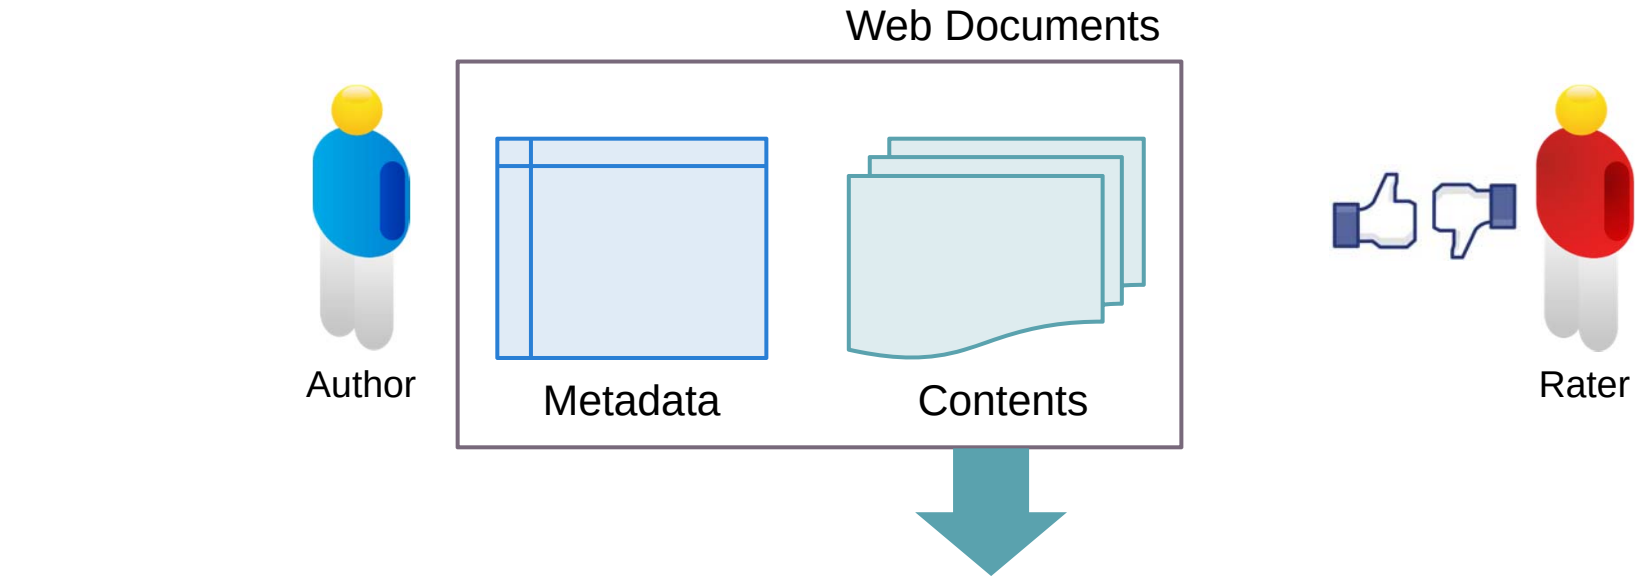

Domain-independent Features

- Structural
- Syntactic
- Lexical

- Informative-ness
- Readability
- Credibility

Domain-dependent Features

- Relevance
- Subjectivity

- Substance
- Reliability

Metadata-based Features

Author's Reputation/Credibility

- # or % of reviews written by the author
- # or % of votes received by others
- The reviewer's Amazon rank
- # of reviewer's Amazon badges
- The amount of information on reviewer's profile
- The reviewer's average posting frequency

(Hoang et al., 2008; Chen and Tseng, 2011; O'Mahony and Smyth, 2009; Kokkodis, 2012)

Author's Social Context

- In-degree/out-degree
- PageRank score

(Lu et al., 2010)

Metadata-based Features

Review Histories

- # of item reviewed
- # of reviews of an item
- Time of review
- Star differences
- # of reviews (same category)

(Tanaka et al., 2012)

Believability

- Product rating deviation of a review

(Chen and Tseng, 2011)

Timeliness

- The degree of duplication of a review
- The interval between the current and the first review

(Chen and Tseng, 2011)

Content-based Features

Domain-Independent Features

Structural/surface Feature

- The writing style in the target text
- # of words/sentences
- The average length of sentences/text
- HTML tags, % of capital letters

Syntactic Feature

- POS tags
- % of nouns/verbs/adjectives/adverbs

Lexical Feature

- Spelling Errors
- TF-IDF values

Content-based Features

Domain-dependent Features

Informativeness

- # of product names
- # of brand names
- # of product features and function

(Liu et al., 2007; Chen and Tseng, 2011, Hong et al., 2012)

Relevance

- Topic words from the forum content and leading post in a thread (Wanas et al., 2008)

Substance

- Forum-specific word-level features
  (Weimer et al., 2007; Wanas et al., 2008)
- Terms with dimension reduction
  (Cao et al., 2011; Ngo-Ye and Sinha, 2012)
- The average frequency of product features in a review
  (Chen and Tseng, 2011)

Content-based Features

Domain-dependent Features

Readability

- Cue words for each topic
  (e.g., pros/cons, guidelines for product reviews)
- Ease of understanding (Chen and Tseng, 2011)

Subjectivity

- Positive and negative opinions
  (Liu et al., 2007; Hiremanth et al., 2010;
  Chen and Tseng, 2011)
- Mainstream opinion (Hong et al., 2012)

Content-based Features

Domain-dependent Features

Reliability

- Uncertainties of volitive auxiliaries

  (1) *I **prefer to** buy Sony, for it **may** have higher resolution*. (*9/68*)
  (2) *It costs more than Nikon*. (*89/129*)

- Tense expressions

  The habit of the reviewers often using past and perfect tenses in their writing

  (Hong et al., 2012)

Credibility

- Mentions about a reviewer's long-term experiences
  (Min and Park, 2012)

  (1) The duration of product use
  (2) The number of purchasing objects
  (3) Temporally detailed description about product use

Summary of the Work

Work	Meta-data Feature	Content Feature	Performance
(Kim et al. 2006)	Star ratings	Structural (e.g., html tags, punctuation, review length), Lexical (e.g., n-grams), Syntactic (e.g., percentage of verbs and nouns), Semantic (e.g., product feature mentions)	The most informative features are the length of the review, unigrams.
(Liu et al., 2007)	Not considered	Informativeness (e.g., # of sentences/words/product features), Readability (e.g., # of paragraphs, # of paragraph separators), Subjectiveness (e.g., % of positive/negative sentences)	Features on word level, product feature level, and readability improve the performance of classification. Features on subjectiveness make no contribution.

Summary of the Work

Work	Meta-data Feature	Content Feature	Performance
(Hoang et al. 2008)	Authority (indicating the author's trustworthiness)	Formality (the writing style of target document), Readability (format of the document), Subjectivity (opinions of authors)	Formality is the most effective feature category; Subjectivity features improve the performance on product review data.
(O'Mahony and Smyth 2009)	Reputation Features, Social Features, Sentiment scores assigned by users	Terms, Ratio of uppercase & lowercase characters in review text, Review Completeness	Best performance with all 4 types of features; Among single features, reputation features are the best.
(Lu et al. 2010)	# of reviews by the author, Average rating for the author, Social Context (In-degree of the author, Out-degree of the author, PageRank score of the author)	Text-statistics Features, Syntactic Features, Conformity Features, Sentiment Features	Performance is better when incorporating social context.

Summary of the Work

Work	Meta-data Feature	Content Feature	Performance
(Cao, et al., 2011)	Basic Characteristics (the posting date, the rating difference)	Basic Characteristics (e.g., pros/cons, summary), Stylistic Characteristics (e.g., # of sentences/words) Semantic characteristics (substance of the review)	Best performance with all 3 types of features, Semantic characteristics have a significant impact.
(Chen and Tseng, 2011)	Believability (Product rating deviation of a review), Reputation (e.g., the ranking of the reviewer), Timeliness (the interval between the current review and the first review of the product)	Objectivity (e.g., # & % of opinion sentences), Relevancy (e.g., # of the product/brand name), Completeness (# of different product features), Appropriate amount of information (e.g., the avg freq. of product features), Ease of understanding (e.g., misspelled words)	Best performance with the objectivity, reputation, appropriate amount of information, and ease of understanding features (SVM, Linear kernel)
(Tanaka et al., 2012)	Review Histories (e.g., # of item reviewed, # of reviews of an item, time of review, star differences)	# & freq. of words/product names/product features, TF-IDF scores, percentage of nouns/verbs/adj, # of positive/negative words/sentences, stars	The top 5 attributes: stars, star differences, # & freq. product features, # of words

Summary of the Work

Work	Meta-data Feature	Content Feature	Performance
(Kokkodis, 2012)	Reviewer's Trustfulness (e.g., the reviewer's Amazon rank, # of reviewer's Amazon badges, the amount of personal information on reviewer's profile)	Not considered	Naïve Bayes outperform other classifiers (DT, Logistic Regression, MLP)
(Ngo-Ye and Sinha, 2012)	Not considered	Substance of the review by dimension reduction	The proposed Regressional RelieF feature selection method outperforms other dimension reduction methods (LSA, CfsSubSet Evaluation)
(Hong et al., 2012)	Not considered	Information Needs (product aspects/function), Information Reliability (auxiliary & tense), Mainstream Opinion	The best performance with the information needs & information reliability (tense)

Applications

- Opinion Summarization
  - Review Summarization System (Liu et al., 2007)
    - Filtering the reviews with low quality
- Content Recommendation
  - High-quality Hotel Review Recommending System (O'Mahony and Smyth, 2009)
    - Rank-ordering with respect to prediction confidence

Part II

- Personalized quality analysis
- Filtering false documents

Personalized Quality Analysis

- Modeling a helpfulness function based on a product designer's perspective (Liu et al., 2013)
  - No guideline (deliberately)
    - Study how they actually perceive 'helpfulness'
  - Several important perspectives from the questionnaire

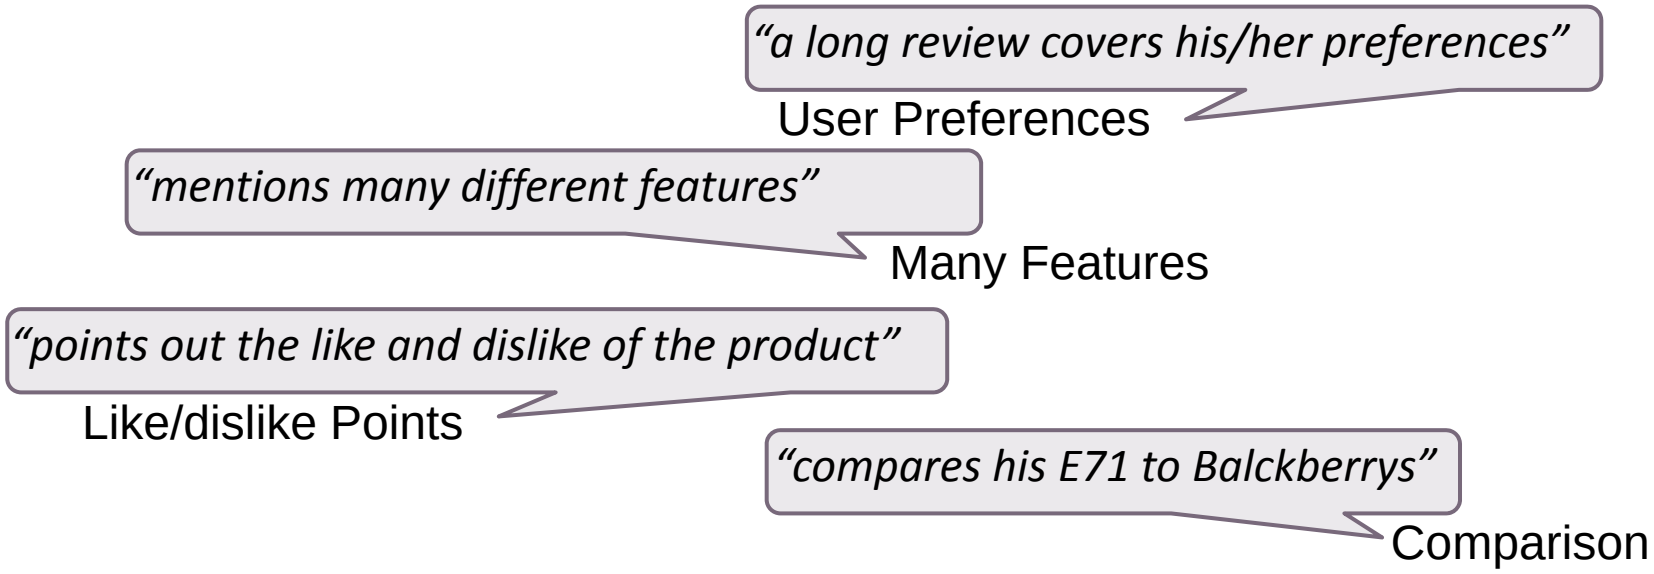

Personalized Quality Analysis

- Modeling a helpfulness function regarding differences in users' criteria (Moghaddam et al., 2012)
  - Examples

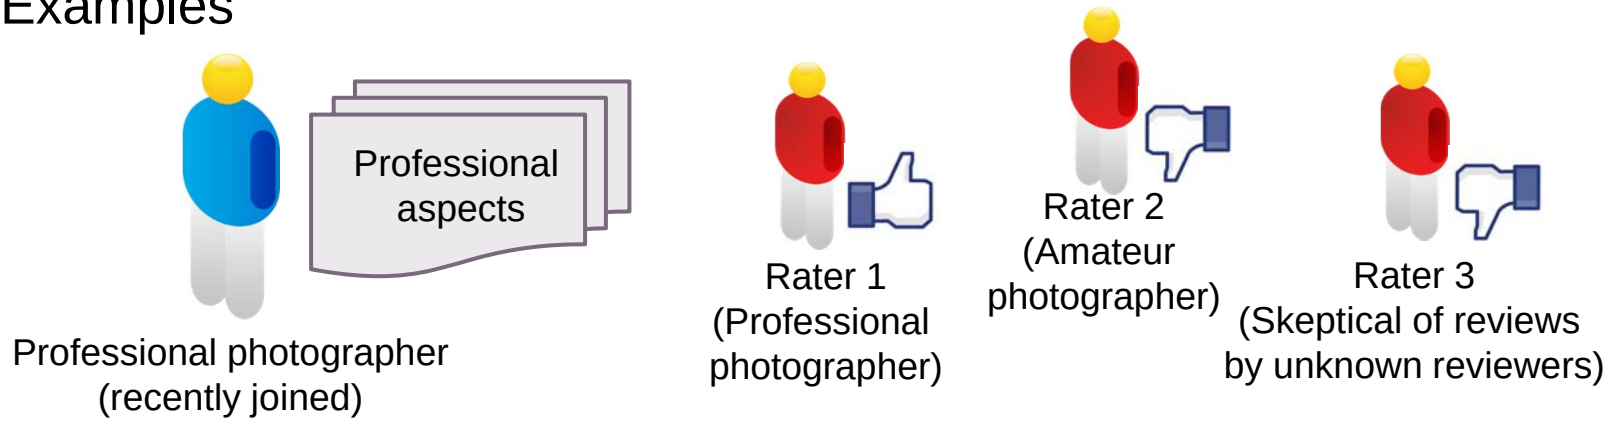

  - Proposing an extended tensor factorization (ETF) model by utilizing the latent user/rater features
  - All of the personalized methods (MF, TF, ETF, BETF models) outperform all of the non-personalized methods (textual/social feature-based regression model)

Filtering False Documents

- Spam Detection
  - Types of Spam Reviews (Jindal and Liu, 2008)
    - Untruthful opinions
      - Deliberately written for the purpose of promoting the products or damaging the reputation of the products
    - Reviews on brands only
    - Advertisement or other irrelevant reviews
  - Classifying Methods
    - SVM, Naïve Bayes, Logistic Regression (Jindal and Liu, 2008; Li et al, 2011; Park et al, 2012; Morales et al., 2013)
    - Similarity-based Classification (Algur et al., 2010)

Filtering False Documents

- Spammer Detection
  - Detecting review spammers using rating behaviors (Lim et al., 2010)
    - Rating Behaviors
      - Targeting specific products
      - Deviating from the other reviews in rating
    - Proposing a behavior score
      - Linear weighted combination

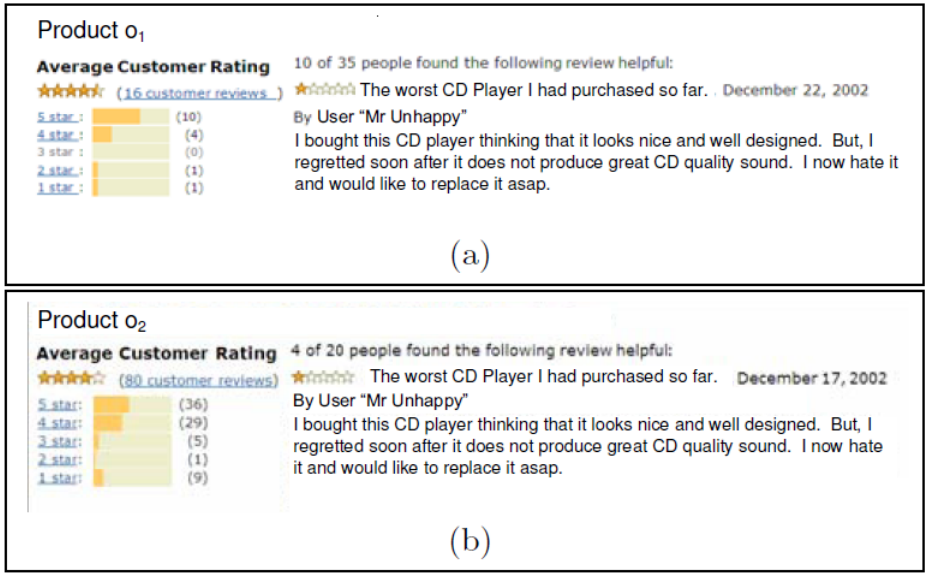

Filtering False Documents

- Spammer Detection
  - Detecting fake reviewer group (Mukherjee et al., 2012)

1 of 1 people found the following review helpful: ★★★★★ Practically FREE music, December 4, 2004 This review is from: Audio Xtract (CD-ROM) I can't believe for $10 (after rebate) I got a program that gets me free unlimited music. I was hoping it did half what was	2 of 2 people found the following review helpful: ★★★★★ Like a tape recorder..., December 8, 2004 This review is from: Audio Xtract (CD-ROM) This software really rocks. I can set the program to record music all day long and just let it go. I come home and my	★★★★★ Wow, internet music! ..., December 4, 2004 This review is from: Audio Xtract (CD-ROM) I looked forever for a way to record internet music. My way took a long time and many steps (frustrtaing). Then I found Audio Xtract. With more than 3,000 songs downloaded in ...
3 of 8 people found the following review helpful: ★★★★★ Yes – it really works, December 4, 2004 This review is from: Audio Xtract Pro (CD-ROM) See my review for Audio Xtract - this PRO is even better. This is the solution I've been looking for. After buying iTunes, ...	3 of 10 people found the following review helpful: ★★★★★ This is even better than..., December 8, 2004 This review is from: Audio Xtract Pro (CD-ROM) Let me tell you, this has to be one of the coolest products ever on the market. Record 8 internet radio stations at once, ...	2 of 9 people found the following review helpful: ★★★★★ Best music just got ..., December 4, 2004 This review is from: Audio Xtract Pro (CD-ROM) The other day I upgraded to this TOP NOTCH product. Everyone who loves music needs to get it from Internet ....
5 of 5 people found the following review helpful: ★★★★★ My kids love it, December 4, 2004 This review is from: Pond Aquarium 3D Deluxe Edition This was a bargain at $20 - better than the other ones that have no above water scenes. My kids get a kick out of the	5 of 5 people found the following review helpful: ★★★★★ For the price you..., December 8, 2004 This review is from: Pond Aquarium 3D Deluxe Edition This is one of the coolest screensavers I have ever seen, the fish move realistically, the environments look real, and the	3 of 3 people found the following review helpful: ★★★★★ Cool, looks great..., December 4, 2004 This review is from: Pond Aquarium 3D Deluxe Edition We have this set up on the PC at home and it looks GREAT. The fish and the scenes are really neat. Friends and family ....
Figure 1: Big John's Profile	**Figure 2: Cletus' Profile**	**Figure 3: Jake's Profile**

- The same products with all 5 star ratings
- Within a small time window
- Only reviewed the 3 products
- The early reviewers

Suspicious Patterns

A several behavioral models based on the relationship among groups, individual reviewers, and products

Helpful Reviews vs. Spam Reviews

- Inter-annotator agreement scores are lower in spam review annotation
  - 0.48 ~ 0.64 kappa values, spammer/non-spammer (Lim et al., 2010)
  - Labeling spam reviewer groups is easier than labeling individual spam reviews/reviewers
    - 0.79, Fleiss' multi-rater kappa, spam/non-spam/borderline (Mukherjee et al., 2012)
- Metadata-based features are more useful for spam detection
  - Spam reviews usually look perfectly normal (Lim et al., 2010)

Metadata-based vs. Content-based

- Metadata-based Features
  - More robust and consistent
  - No corresponding information for newly posted content
  - From the reviewers/spammers
- Content-based Features
  - Depending heavily on explicitly mentioned information in the content
  - The accuracy of extracting features is comparatively lower
  - From the helpful reviews/spam reviews

Concluding Remarks

Concluding Remarks

- We examined the notion of quality from the perspective of the types of content and introduced two major methods of determining the quality.
  - The crowd-based metric
  - The gold standard model
- We examined the two major quality analysis models.
  - Regression-based ranking
  - Classification
- Utilizing metadata-based features gives a robust result but it is hard to acquire for newly posted contents, which may be overcome by content-based features.
  - The hybrid approach of utilizing both features outperforms any one of the feature-based approaches.

Concluding Remarks

- We introduced several applications such as an opinion summarization system and a review recommending system.
- Personalized quality analysis can be a complement to the current quality analysis measure
  - ‘Helpful’ votes suffer from some bias
- False or spam contents should be also seriously dealt with for providing users with high quality contents.



KCC 2013 자연어처리 및 정보검색 최근 동향 워크샵

Thank you

Hye-Jin Min
hjmin@nlp.kaist.ac.kr
Department of Computer Science
KAIST

References

- Algur, S.P., A.P. Patil, P.S. Hiremath, and S. Shivashankar. 2010. Conceptual level Similarity Measure based Review Spam Detection, Proceedings of International Conference on Signal and Image Processing, pp. 416-423.
- Cao, Q., W. Duan, and Q. Gan. 2011. Exploring determinants of voting the "helpfulness" of online user reviews: A text mining approach. *Decision Support System* 50, pp. 511-521.
- Chen, C.C. and Y.-D. Tseng. 2011. Quality Evaluation of Product Reviews Using an Information Quality Framework, *Decision Support System* 50, pp. 755-768.
- Hiremath, P.S., S.P. Algur, and S. Shivashankar. 2010. Quality Assessment of Customer Reviews Extracted From Web Pages: A Review Clustering Approach*, International Journal on Computer Science and Engineering*, 2(03), May, pp. 600-606.
- Hoang, L., J.-T. Lee, Y.-I. Song, and H.-C. Rim. 2008. A model for evaluating the quality of user-created documents. *Proceedings of the 4th AIRS*, LNCS 4993, Springer-Verlag, Berlin, Germany, pp. 496–501.
- Hong, Y., J. Lu, J. Yao, Q. Zhu, and G. Zho. 2012. What Reviews are Satisfactory: Novel Features for Automatic Helpfulness Voting, *Proceedings of SIGIR*, August 12-16, Portland, Oregon, USA, pp. 495-504.

References

- Huang, S., D. Shen, W. Feng, and Y. Zhang. 2009. Discovering clues for review quality from author's behaviors on e-commerce sites. *Proceedings of ICEC*, August 12–15, Taipei, Taiwan, ACM, New York, pp. 133–141.
- Jindal, N. and B. Liu. 2008. Opinion spam and analysis. *Proceedings of the 1st WSDM*, February 11–12, Stanford, CA, ACM, New York, pp. 219–229.
- Kim, S.-M., P. Pantel, T. Chklovski, and M. Pennacchiotti. 2006. Automatically assessing review helpfulness. *Proceedings of EMNLP*, July 22–23, Sydney, Australia, Association for Computational Linguistics, Stroudsburg, Pennsylvania, pp. 423–430.
- Kokkodis, M. 2012. Learning from Positive and Unlabeled Amazon Reviews: Towards Identifying Trustworthy Reviewers. *Proceedings of WWW*, April 16-20, Lyon, France, pp. 545-546.
- Li, F., M. Huang, Y. Yang, and X. Zhu. 2011. Learning to Identify Review Spam, Proceedings of the Twenty-Second international joint conference on Artificial Intelligence (IJCAI), pp. 2488-2493.
- Lim, E.-P., V.-A. Nguyen, N. Jindal, B. Liu, and H.W. Lauw. 2010. Detecting Product Review Spammers using Rating Behaviors, *Proceedings of CIKM*, October 26-30, Toronto, Ontario, Canada, pp. 939-948.

References

- Liu, J., Y. Cao, C.-Y. Lin, Y. Huang, and M. Zhou. 2007. Low-quality product review detection in opinion summarization. *Proceedings of EMNLP-CoNLL*, June 28–30, Prague, Czech Republic, Association for Computational Linguistics, Stroudsburg, Pennsylvania, pp. 334–342.
- Liu, Y., J. Jin, P. Ji, J.A. Harding, and R.Y.K. Fung. 2013. Identifying helpful online reviews: A product designer's perspective, *Computer-Aided Design*, 45, pp. 180-194.
- Lu, Y., P. Tsaparas, A. Ntoulas, and L. Polanyi. 2010. Exploiting social context for review quality prediction. *Proceedings of WWW*, April 26–30, Raleigh, NC, ACM, New York, pp. 691–700.
- Moghaddam, S., M. Jamali, and M. Ester. 2012. ETF: Extented Tensor Factorization Model for Personalizing Prediction of Review Helpfulness, *Proceedings of WSDM*, February 8-12, Seattle, Washington, USA, pp. 163-172.
- Morales, A., H. Sun, and X. Yan. 2013. Synthetic Review Spamming and Defense, *Proceedings of WWW Companion*, May 13-17, Rio de Janeiro, Brazil, pp. 155-156.
- Min, H.-J. and J.C. Park. 2012. Identifying Helpful Reviews Based on Customer's Mentions about Experiences, *Expert Systems With Applications*, 39(15), pp. 11830-11838, November.

References

- Mukherjee, A., B. Liu, and N. Glance. 2012. Spotting Fake Reviewer Groups in Consumer Reviews, *Proceedings of WWW*, April 16-20, Lyon, France, pp. 191-200.
- Ngo-Ye, T.L and A.P. Sinha. 2012. Analyzing Online Review Helpfulness Using a Regressional RelifF-Enhanced Text Mining Method, *ACM Transactions on Management Information Systems*, 3(2), July, pp. 1-19.
- O'Mahony, M. P. and B. Smyth. 2009. Learning to recommend helpful hotel reviews. *Proceedings of RecSys*, October 23–25, ACM, New York, pp. 305–308.
- Park, I., H. Kang, C.Y. Lee, and S.J. Yoo. 2010. Classification of Advertising Spam Reviews, *Proceedings of 6th International Conference on Advanced Information Management and Service (IMS)*, pp. 185-190.
- Raghavan, S., S. Gunasekar, and J. Ghosh. 2012. Review Quality Aware Collaborative Filtering, Proceedings of RecSys, September 9-13, Dublin, Ireland, pp. 123-130.
- Tanaka,Y., N. Nakamura, Y. HiJikata, and S. Nishida. 2012. Estimating Reviewer Credibility Using Review Contents and Review Histories, *IEICE Transaction on Information & System*, E95-D(11), November.

References

- Tsur, O. and A. Rappoport. 2009. RevRank: A fully unsupervised algorithm for selecting the most helpful book reviews. *Proceedings of the 3rd ICWSM*, San Jose, California, Association for the Advancement of Artificial Intelligence, Palo Acto, California, pp. 154–161.
- Wanas, N., M. El-Saban, H. Ashour, and W. Ammar. 2008. Automatic scoring of online discussion posts. *Proceedings of WICOW*, October 30, Napa Valley, CA, ACM, New York, pp. 19–25.
- Weimer, M., I. Gurevych, and M. Mühlhäuser. 2007. Automatically assessing the post quality in online discussions on software. *Proceedings of the ACL Demo and Poster Session*, June 25–27, Prague, Czech Republic, Association for Computational Linguistics, Stroudsburg, Pennsylvania, pp. 125–128.
- Zhang, Z. and B. Varadarajan. 2006. Utility scoring of product reviews. *Proceedings of CIKM*, November 5–11, Arlington, VA, ACM, New York, pp. 51–57.