UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

Time-aware online reputation analysis

Peetz, M.-H.

Publication date
2015
Document Version
Final published version

Link to publication

Citation for published version (APA):
Peetz, M.-H. (2015). *Time-aware online reputation analysis*. [Thesis, fully internal, Universiteit van Amsterdam].

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (https://dare.uva.nl)

Download date:14 Nov 2026

Social media has become an integral part of society. Omnipresent mobile devices allow for immediate sharing of experiences. Experiences can be about brands and other entities. For social media analysts a collection of posts mentioning a brand can serve as a magnifying glass on the prevalent opinion towards a brand: The overall estimation of a its reputation is increasingly based on the aggregation of a brand's reputation polarity in social media posts. This polarity of reputation is currently annotated manually. However, with the dramatic increase of social media, this is no longer feasible.

This thesis aims to facilitate and automate parts of the process to estimate the reputation of a brand. We motivate this by performing user studies with expert social media analysts. We analyse three resulting datasets: a questionnaire, log data of a manual annotation interface, and videos of annotating experts following the think-aloud protocol. We find the online and offline authority of the user posting online influences the annotation decision most. This online and offline authority is therefore a strong indicator for reputation polarity of this posting. Additionally, experts welcome automation of information retrieval and filtering tasks. For both, information retrieval and filtering, as well as for several indicators, the reputation analysts' background information proves vital.

Based on the indicators used for manual annotation, we proceed with the development of algorithms for the automatic estimation of reputation polarity. Unlike earlier, static evaluation scenarios, we follow a dynamic scenario, which mimics the daily workflow of social media analysts. Our algorithms are successful because we distinguish between reputation and sentiment.

The second part of this thesis is motivated by the analysts' desire for automation of retrieval and filtering of new media. For information retrieval, we present two improvements to existing algorithms. The first improvement is based on burst identification in time series of pseudo-relevant documents. We sample terms in those bursts for query modeling, thus improving effectiveness on news and blog corpora. Secondly, recency is an important aspect of relevance in social media. Inspired by memory models from cognitive science, we point out new (cognitive) document priors based on those models. We show that those priors are more effective, efficient, and plausible than commonly used temporal priors.

Background knowledge is essential for information filtering. Additionally, topics around an entity are dynamic. We find that for a consistently strong performance of filtering algorithms, the expertise of social media analysts is needed. The algorithms for filtering presented in this thesis are based on active learning: For documents where the algorithm cannot classify with high certainty, manual labels are requested from the analyst. Using intuitions about bursts and cognitive priors for sampling difficult to classify documents, we need very little help from annotators to reach high effectiveness.

We conclude that many aspects of the annotation of reputation can be automated – using in particular time series analysis, memory models, and low-impact help from expert social media analysts.

ISBN 978-94-6182-540-7

MARIA-HENDRIKE PEETZ TIME-AWARE ONLINE REPUTATION ANALYSIS

TIME-AWARE ONLINE REPUTATION ANALYSIS

MARIA-HENDRIKE PEETZ

Time-Aware Online Reputation Analysis

Maria-Hendrike Peetz

Time-Aware Online Reputation Analysis

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de
Universiteit van Amsterdam
op gezag van de Rector Magnificus
Prof. dr. D.C. van den Boom
ten overstaan van een door het college voor promoties ingestelde
commissie, in het openbaar te verdedigen in
de Agnietenkapel
op dinsdag 24 maart 2015, te 10:00 uur

door

Maria-Hendrike Peetz

geboren te Soltau, Duitsland

Promotiecommissie

Promotor:
Prof. dr. M. de Rijke

Co-Promotor:
Prof. dr. W.M. van Dolen

Overige leden:
Prof. dr. F.M.G. de Jong
Prof. dr. M. Welling
Dr. F. Diaz
Dr. J. Gonzalo

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

SIKS Dissertation Series No. 2015-07
The research reported in this thesis has been carried out under the auspices of SIKS, the Dutch Research School for Information and Knowledge Systems.

COMMIT/

The research was partially supported by the Dutch national program COMMIT.

The research was partially supported by the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement nr 288024.

Copyright © 2015 Maria-Hendrike Peetz, Amsterdam, The Netherlands
Cover by David Graus
Printed by Off Page, Amsterdam

ISBN: 978-94-6182-540-7

Für meinen Vater, Heiner Peetz

Dankwoord

Maarten, hi. Without you, this book wouldn't exist. You knew when to push me and when to stop me. Unlike your own prognosis, I never hated you. Thank you.

Edgar, Wouter, and Manos. You taught me the nitty bits at the beginning, with more patience than I had myself.

Daan, David, Richard, Evgeny, Anne, and Zhaochun. We grew together as PhD students, what a ride. There is much I would like to take from each of you: Open arms from Daan, humor from David, calmness and deep problem solving from Richard, elephant skin from Evgeny, *yes, but* from Anne, and motivation from Zhaochun.

Willemijn. You inspired me to think differently, not from a computer science perspective. You are a great role model.

Martin. I will always remember *my spot*, hacking and drinking away. Thank you for telling me I could program, it took me years to believe you.

Sara. You are the smartest and kindest woman I know. Spending time with you goes easy.

Shinbukan, my beloved dojo. Thank you for throwing me, spending wonderful saturday afternoons in cafés, and (training) holidays at Balaton. You all taught me something. Ferenç-sensei showed me my force-power. Joost-sensei taught me how to fly. Gerald-sensei convinced me that there is a sweet-spot between fear and suicidal bravery. Edo, you believed in me, in academia and aikido. Grainne, you found my grounding. And Desmond, soulmate, you found my mirror.

Jessie und Wendy. Danke für die Stunden die wir lachend, weinend, diskutierend, organisierend, streitend (über likelihoods und andere semantischen Differenzen), einander stützend verbracht haben. Danke Dir besonders, Jessie, für die Inspiration zu Kapitel 7.

Mama. Danke dass Du mir gezeigt hast, dass man alles erreichen kann, aber die Liebe immer das wichtigste im Leben bleibt.

Thomas. In diesem Buch steckt mehr von Dir als Du denkst. Ich liebe Dich.

Contents

1 Introduction **1**
- 1.1 Research Outline and Questions . . . 3
- 1.2 Main Contributions . . . 5
- 1.3 Thesis Overview . . . 6
- 1.4 Origins . . . 8

2 Background **9**
- 2.1 Social Media . . . 9
- 2.2 Reputation . . . 12
  - 2.2.1 Measuring Reputation . . . 12
  - 2.2.2 Reputation Polarity . . . 13
- 2.3 Information Retrieval . . . 16
  - 2.3.1 Temporal Information Retrieval . . . 19
- 2.4 Entity Filtering . . . 21

3 Datasets **23**
- 3.1 TREC News . . . 23
- 3.2 TREC Blog . . . 24
- 3.3 TREC Microblog . . . 24
- 3.4 RepLab . . . 25
  - 3.4.1 RepLab 2012 . . . 25
  - 3.4.2 RepLab 2013 . . . 27
- 3.A Query Sets Used . . . 29
- 3.B Domains . . . 30

4 Analyzing Annotation Procedures and Indicators for Reputation Polarity **31**
- 4.1 Methodological Background . . . 33
  - 4.1.1 Understanding Social Media Users . . . 33
  - 4.1.2 Log Analysis . . . 34
  - 4.1.3 Think Aloud . . . 34
  - 4.1.4 Annotation Interfaces . . . 34
  - 4.1.5 Related Annotation Tasks . . . 35
- 4.2 Datasets . . . 36
  - 4.2.1 Questionnaires . . . 36
  - 4.2.2 Annotation System Logs . . . 38
  - 4.2.3 Think Aloud . . . 42
  - 4.2.4 Summary . . . 42
- 4.3 Annotation Procedures for Reputation Polarity . . . 44
  - 4.3.1 Analysis of Log Data . . . 44
  - 4.3.2 Annotation Difficulty . . . 48
  - 4.3.3 Analysis of Questionnaire . . . 50
  - 4.3.4 Summary . . . 52
- 4.4 Indicators for Reputation Polarity . . . 53
  - 4.4.1 Analysis of Questionnaire . . . 53

4.4.2 Analysis of Think Aloud Videos . . . 56
4.4.3 Summary . . . 59
4.5 Conclusion . . . 61

5 Estimating Reputation Polarity on Microblog Posts 63
5.1 Task Definition . . . 65
5.1.1 Reputation Polarity vs. Sentiment . . . 65
5.2 Modeling Reputation Polarity . . . 66
5.2.1 Baseline: Sentiment features . . . 67
5.2.2 Sender features . . . 69
5.2.3 Message features . . . 69
5.2.4 Reception features . . . 70
5.2.5 Classification . . . 72
5.3 Experimental Setup . . . 72
5.3.1 Preprocessing . . . 73
5.3.2 Sampling . . . 73
5.3.3 Experiments . . . 73
5.3.4 Parameters . . . 75
5.3.5 Runs . . . 75
5.3.6 Evaluation and Significance Testing . . . 75
5.4 Results and Analysis . . . 76
5.4.1 Preliminary Experiment: Sampling . . . 76
5.4.2 Overall Performance . . . 77
5.4.3 Entity-Independent vs. Entity-Dependent vs. Domain-Dependent 79
5.4.4 Feature Analysis . . . 83
5.5 Conclusion . . . 91

6 Using Temporal Bursts for Query Modeling 93
6.1 Temporal Query Models . . . 95
6.1.1 Bursts . . . 97
6.1.2 Term Reweighting . . . 99
6.1.3 Generative Probability of a Document in a Burst . . . 100
6.1.4 Burst Normalization . . . 102
6.1.5 Document Score . . . 103
6.2 Experimental Setup . . . 103
6.2.1 Data . . . 103
6.2.2 Baselines . . . 103
6.2.3 Parameter Settings . . . 103
6.2.4 Evaluation . . . 104
6.3 Results and Discussion . . . 104
6.3.1 Selection of Relevant Documents . . . 104
6.3.2 Document Priors . . . 110
6.3.3 Parameter Optimisation . . . 114
6.3.4 Burst Quality . . . 119
6.4 Conclusion . . . 120
6.A Additional Graphs . . . 121

6.B Vendler Classes of the Queries . . . 122

7 Cognitive Temporal Document Priors 125
7.1 Memory Models . . . 126
7.2 Methods . . . 127
7.2.1 Baselines . . . 127
7.2.2 Retention Functions . . . 128
7.3 Experimental Setup . . . 130
7.3.1 A Framework for Assessing Temporal Document Priors . . . 130
7.3.2 Experiments . . . 131
7.4 Analysis . . . 132
7.4.1 Retrieval Effectiveness . . . 133
7.4.2 Assessing the Document Priors . . . 136
7.4.3 Discussion . . . 138
7.5 Conclusion . . . 139

8 Active Learning for Filtering of Streaming Documents 141
8.1 Active Learning . . . 143
8.2 An Active Learning Approach to Entity Filtering . . . 143
8.2.1 Candidate selection . . . 145
8.2.2 Candidate Annotation . . . 147
8.2.3 Model Updating . . . 147
8.3 Experimental Setup . . . 148
8.3.1 Datasets . . . 148
8.3.2 Document Representation . . . 148
8.3.3 Streaming Scenario . . . 148
8.3.4 Settings . . . 148
8.3.5 Evaluation . . . 149
8.4 Results and Analysis . . . 150
8.4.1 Margin Sampling . . . 150
8.4.2 Recency Priors . . . 152
8.4.3 Temporal Reranking . . . 154
8.4.4 Is There a Universal Best Approach to Sampling Candidates? . 157
8.5 Conclusion . . . 159

9 Conclusions 161
9.1 Answers to Research Questions . . . 161
9.2 Future Research Directions . . . 165

Bibliography 167

Samenvatting 179

Zusammenfassung 181