

Computational Exploration of the Origin of Mood in Literary Texts

Emily Öhman
Waseda University
ohman@waseda.jp

Riikka Rossi
University of Helsinki
riikka.rossi@helsinki.fi

Abstract

This paper is a methodological exploration of the origin of mood in early modern and modern Finnish literary texts using computational methods. We discuss the pre-processing steps as well as the various natural language processing tools used to try to pinpoint where mood can be best detected in text. We also share several tools and resources developed during this process. Our early attempts suggest that overall mood can be computationally detected in the first three paragraphs of a book.

1 Introduction

This paper aims to provide a short exploratory study into a very specific literary concept: *mood*. *Mood* is the general atmosphere that the author creates through their word choices, style and use of imagery and can sometimes even include *tone*. *Mood* is about how the reader feels about the text, and the related concept of *tone* is about how the implied author feels about it and uses words to convey their attitude towards a topic or subject¹ (Turco, 2020).

We use various natural language processing (NLP) tools and methods to attempt to identify the origins and location of mood. For this purpose we have collected a corpus of 975 literary works written in Finnish. Our paper utilizes many tried and tested NLP and computational literary studies methods, but to our knowledge it is the first to combine emotion detection / sentiment analysis with traditional literary affect studies exploring tone and mood.

Since the choice of words is of the utmost importance when creating both of these sometimes entangled concepts in literary works, mood, and to some extent tone, detection are perfect subjects for analyzing the use of emotion-associated words in text. It is not easy to determine where specifically mood is created in a text. Does mood reside

¹On the concepts of tone and mood in literary studies, see Richards (1929); Ngai (2005); Flatley (2008).

everywhere, or is it most strongly present in the first chapter of books, or perhaps the first few paragraphs of each chapter? In a larger sense *mood* is not reducible to a single aspect, but generated by a set of textual elements.

We suggest that the computational study of the valence of the lexicon can be valuable in providing an accurate picture on the distribution of the positive and negative valence in a text-continuum, and thus help us to better understand the relationship between the linguistic qualities of a text and its perceived emotional effects, particularly the mood of a text.

2 Previous Work

Literature can be considered a domain where the affective functions of language are of principal importance (Hogan, 2011). Research on literature and emotions is an active field in traditional literary studies, particularly after the "affective turn"; a shift in attitude towards how affect is perceived in literary theory that took place about two decades ago (see e.g. Smith, 2011 and Armstrong, 2014). Research topics range from the study of literature and empathy (Keen, 2007) to the study of literature and cognition (Hogan, 2011), negative affects and tone in texts (Ngai, 2005) to empirical perspectives (Sklar, 2013; Van Lissa et al., 2018), and even emotions specific to Finnish literature (Rossi, 2020; Rossi and Lyytikäinen, 2022).

Recently, the question of a text's overall emotional tone or mood has aroused vivid interest (e.g. Ngai (2005), Lyytikäinen (2017), and Rossi (2020)). Yet a systematic theory of how tone and mood are created and triggered is still under construction. We suggest that a study of the emotional valence of the lexicon measured quantitatively provides a new approach that can help with understanding the components of a text's mood (Öhman and Rossi, 2021).

The field of computational literary studies (CLS)

can generally be considered a field within digital humanities that uses NLP tools to analyze and make new discoveries in texts by quantitative means. It is somewhat rare for studies within the field of CLS to have literary experts working on the project, and many such projects rely heavily on the analysis of the quantitative results also conducted by experts of NLP rather than experts of literature as is true for many other interdisciplinary fields (Bartlett et al., 2018). Kim and Klinger (2018) provide a substantive overview of sentiment analysis as it is used in CLS. Very few, if any, of the different types of such studies discussed in their survey deal with topics that are common in traditional literary analysis. We hope to rectify this situation by bringing CLS and affect studies in literature closer together and this paper will hopefully be a small part of that process.

3 Data

We downloaded the first 1000 books from Project Gutenberg², where (1) the language was Finnish, and (2) the text was in utf-8 plain text format. As far as we are aware, there is no way of filtering out texts by their original language, so our dataset includes many translated works. We used the simple gutenberg-cleaner³ to get rid of the preamble and the legal text at the end of the book, then we created a regex to extract key information such as the title, the name of the author, the year of publication, and whether the book was originally written in Finnish. The translation status of the book was extracted based on whether the terms *suomentaja*, *suomennettu*, *suomentanut*, or any version of *kääntäjä/käännös/käännetty* etc. were present within the first ten lines of text after the preamble was removed.

Due to encoding issues, our final corpus consists of 975 books of which roughly half were originally written in Finnish. A vast majority (95+%) were written or translated between the years 1850 and 1925 and over 90% after 1880, with only a few instances of older texts meaning that the language used in the texts can be considered Modern Finnish (Forsman Svensson, 2011). The final data consists of 2,938,032 sentences and 41,417,116 tokens.

²<https://www.gutenberg.org/>

³<https://libraries.io/pypi/gutenberg-cleaner>

3.1 The Emotion Intensity Lexicon

We used the Finnish Emotion Intensity Lexicon (FEIL) (Öhman, 2022) as a base for our emotion lexicon. FEIL is based on the NRC emotion lexicons (Mohammad and Turney, 2013) adapted for Finnish and lists words alongside the emotions they are associated with as well as the intensity of the associated emotion as a number between 0 and 1. It is roughly based on Plutchik's wheel of emotions (Plutchik, 1980) and contains the emotions *anger*, *anticipation*, *disgust*, *fear*, *joy*, *sadness*, and *trust*.

4 Method

Finnish is a great language to work with in terms of NLP. Numerous resources not only exist but are also well-curated. There are several researchers and research groups who actively develop new tools and improve upon old ones, and most of these tools are open source (Hämäläinen and Alnajjar, 2021). Thus we had the opportunity to test several different lemmatizers and tokenizers. We further add to this list of tools by having created a Finnish version of the chapterize⁴ package.

After the Project Gutenberg added information was removed, the texts were lemmatized, split into paragraphs, and tokenized. First, we tried fine-tuning Finnish BERT (Virtanen et al., 2019) to work with our texts (as per Gururangan et al., 2022), but the results were not promising and require further work (particularly the vocabulary was not improved sufficiently). We subsequently tried multiple different lemmatization tools, including the Turku Neural Parser (Kanerva et al., 2019), murre (Partanen et al., 2019; Hämäläinen et al., 2021), and both the *experimental* and *news* Finnish spaCy models. In the end we settled for the Turku Neural Parser as the results were the most accurate (see table 1 for an example) and all words were parsed, and parsed correctly in context as well (in the example it was the only one able to correctly parse the nonstandard form *kahvians* – standard form: *kahviansa* – partitive case of 3rd. pers. sing./plur. coffee).

Incidentally, in the dissertation of Airio (2009) *kahviansa* is discussed as an example of "parasite words" since it can be mistakenly split into *kahvi* (coffee) and *ansa* (trap), something none of the lemmatizers did. With careful optimism, we take this

⁴The original: <https://pypi.org/project/chapterize/> and the Finnish version: <https://github.com/sohman/chapterize-fi>

as a demonstration of how good lemmatizers for morphologically complex languages have become in the past decade.

After pre-processing the texts, we identified the first three paragraphs of each book. This was trickier than expected as despite removing the preambles/headers, some metadata remains in the text files and this metadata is of various shapes without uniformity or even commonly recurring pattern of where the actual text of the book starts. For this reason, we created a Finnish version of the chapterize package for Python and used the chapter splits to help recognize opening paragraphs together with sentence and paragraph ids provided by the conllu metadata. We used two different text sections as targets for overall mood detection: the first three paragraphs of each book, and the first 200 tokens from each chapter in each book.

From previous studies (Öhman and Rossi, 2021) we know that certain words can quickly overwhelm the results; when analyzing the novel *Rautatie* by Juhani Aho, the term *rautatie* (railroad) in the lexicon was associated with *trust* and because the store takes place on the railroad and is about the railroad the levels of *trust* in the results were not representative of the level of *trust* in the novel itself. As FEIL contains mostly contemporary words and their contemporary emotion associations we needed to make sure that (1) the most common words in our texts that in our opinion have an emotion association are indeed in the lexicon, and (2) that the most common emotion word matches represent the correct emotions at reasonable intensities. These steps are iterative and continuous in that they should be repeated whenever the lexicon or lemmatization is altered.

The removal of non-emotion associated words is straightforward and fairly uncontroversial. However, re-labeling emotion words or adding new words to the lexicon should be done with utmost care, ideally using multiple annotators who are not the authors and cross-checking the results using inter-annotator agreement scores (van Atteveldt et al., 2021). In this vein, we did not want to bias the lexicon with our own interpretations of emotion intensities so instead we created a large word2vec model of our corpus and used it to look up words in the lexicon with high cosine similarity to the words we wanted to introduce to the lexicon (as per e.g. Maas et al., 2011; Yu et al., 2017; Ye et al., 2018). This led to the association for e.g. the

words *kirkas*, *valkoinen*, and *valkea* to be identical. As the words that needed to be added were relatively few, we manually checked that the emotion associations and intensities made sense. For future projects we intend to employ human annotators in addition to this approach, however, this approach alone showed a lot of promise and was very accurate within the small sample size.

After completing steps (1) and (2), we removed 128 entries from the lexicon and added 203 tokens including *rakastaa*, to love. The exclusion of such an important term from the lexicon exemplifies some of the issues with using a lexicon that was originally created for English where the noun and verb forms are often the same unlike in Finnish where the forms are distinct (cf. to love/a love, to run/a run vs. *rakastaa/rakkaus*, *juosta/juoksu*). Some of these issues were fixed in FEIL by adding both the Finnish noun and verb forms of a single English entry, but many such examples remain (Öhman, 2022). We used this domain- and period-specific version of FEIL to tabulate normalized (per token count for inter-text comparability), intensity scores for each target text. Other future projects should include checking that both noun and verb forms are found in the lexicon.

If we are looking at purely the word choices of the author, tone and mood can be difficult to distinguish from each other and can be intertwined to different degrees. However, the tone of a literary text tends to shift much more even within a chapter and therefore by focusing on the first paragraphs of each chapter, or even the opening paragraphs of the first chapter only, we can get a fairly accurate idea of the mood of the text, with less of a risk of it being confused with tone.

Although the tone of the text may vary within one work (due to e.g. changes in the narrative point of view or mover from description to dialogue narration) the analysis of the beginnings of a text may be indicative of the overall tone of a text. Namely from the perspective of the reader, the beginnings tend to shape the experience of reading. Theories of perception (e.g. Perry 1979) argue that the openings play a crucial role in creating a text's overall emotional disposition: as in everyday life, the first impression matters in reading, too. The emotion effects created in the beginning of a text modify and adjust the reader's general emotional orientation by shaping up modes of perception and organization of information. For instance, the melancholic

Translation	The provost sits down in his rocking chair, stands his pipe on the floor against the table leg, and starts drinking his coffee
Original	Rovasti istuutuu keinutuoliinsa, panee piippunsa lattialle pöydän jalkaa vasten pystyyn ja rupeaa juomaan kahvians
Murre (hist)	(öljy)rovasti istua keinutuoli panna piippu latija pöytä jalka vaste pystyä ja ruveta juoma kahvis
spaCy (news_lg)	Rovasti istuutua keinutuoliinsa panee piippu lattia pöytä jalka vasten pystyyn ja rupeata juoda kahvians
spaCy (exp. /w voikko)	rovasti istuutua keinutuoli panna piippu lattia pöytä jalka vasten pystyyn ja ruveta juoda kahvians
UralicNLP	Rovastilrovasti istuutua panna piippu lattia pöytä jalka vasten pystyynpysty ja ruveta juomaljuoda
Turku Neural Parser	rovasti istuutua keinu#tuoli panna piippu lattia pöytä jalka vasten pystyyn ja ruveta juoda kahvi

Table 1: Example of lemmatization using different lemmatizers for Finnish.

mood created in the beginning of Aho’s *Rautatie*, or the strong effects of disgust in the beginning of Sillanpää’s *Hurskas kurjuus*, are likely to influence the reader experiences later reactions and feelings triggered by narrative events.

5 The Mood in Selected Texts

The results are difficult to present in a small space as they list all the emotion scores per text, therefore we are focusing here on qualitatively evaluating a small subset of the data. The selection is pragmatic, and based on the second author’s area of expertise.

5.1 An overview of the selected texts

The first one is Juhani Aho’s breakthrough novel *Rautatie* (tr. as *The Railroad*, 1884). In this text, from the perspective of the implied reader, the novel evokes emotional effects of melancholia and nostalgia, which are characteristic of Aho’s work.

The second one is Minna Canth’s *Kauppa-Lopo* (no translation, the title refers to the protagonist’s nickname, 1889), a tragic story of poverty and illness. The beginning of the novella, set in prison, underlines the anti-hero’s ugly appearance, but the narrative contrasts the physical ugliness with an inner goodness: she is described as good-hearted and compassionate towards other people. Canth’s naturalism was considered "poor art" by her contemporaries and she was accused of being an admirer of disgust, “destroying the laws of beauty, unfolding ugliness in every sense” (Rossi, 2007, 52).

The third one is Frans Emil Sillanpää’s *Hurskas kurjuus* (tr. as *Meek Heritage*, literally “Sacred Misery”, 1919), which begins with a shocking prologue which anticipates the death of the protagonist: it describes the execution of a poor tenant farmer who had ended up as a Red Guard soldier in the Finnish Civil War (1918). Despite the negative emotions and the tragic events the narrator also expresses trust and comfort in the future of the Finnish nation.

The last one is *Putkinotko* (no translation, 1919-

20) by Joel Lehtonen. Like Sillanpää’s *Sacred Misery*, this novel tracks the tensions that escalated in the Finnish Civil War in 1918. The novel’s protagonist, a good-hearted yet self-willed tenant farmer resigns to obey the landlord and instead resorts to illegal distillery to support the family. The novel is emotionally ambivalent: the idyllic descriptions of Finnish summer nature and the comic elements are likely to arouse positive emotions, while the unembellished description of poverty intends to evoke moral anger and sadness for the social inequality.

5.2 Results

The results for the texts are presented in figures 1 and 2.




Figure 1: Emotion word distribution in first three paragraphs per 1000 words




Figure 2: Emotion word distribution in the first 200 tokens of each chapter per 1000 words

From figure 1 it is possible to see some patterns emerge. In particular, the laconism of *Rautatie* is evident when compared to some of the other au-

thors and the strong emotional impact of *Hurskas Kurjuus* becomes very apparent, with *fear* and *anger*, but also *sadness* being particularly notable. *Fear* and to some extent *anger* are also very present in *Kauppa-lopo*, likely due to the described prison environment and appearance of the protagonist. *Trust* and *joy* are the most notable emotions in *Putkinotko*, perhaps due to the detailed descriptions of the idyllic landscape that dominate the opening chapter.

In general, when comparing the two approaches, first paragraph-only vs. opening paragraphs of each chapter, in the latter the positive emotions are much more prevalent. This indicates that the intended opening mood for these novels is built on negative emotions intended to evoke strong feelings in the reader. In the former, the openings of each chapter are more varied.

6 Discussion

We find that the preliminary results correspond well with established interpretations of mood in select texts when comparing to the emotion word distribution of the first three paragraphs of a literary text. Comparing the valency and intensity of emotions in the opening paragraphs of the book and the opening paragraphs of all chapters, we can see that when looking at all chapters, the emotions approach the distribution of emotion in the lexicon, i.e. they become muddled even though there are some idiosyncrasies that can be assumed to be because of the author's style. Furthermore, the differences between the different texts are also evened out in the all-chapters approach. This could also be in part because the focus becomes more varied and therefore the results average out and start to converge on the distribution of emotions in the lexicon. We recommend that the quest for mood should begin with the opening paragraphs of a text.

It is important to note that evoking emotional effects in literature is not restricted to emotion words or to direct descriptions of the character's emotions. All aspects of the narrative, from description of objects to narrative point of view and style, including tropes and even the rhythm of the text are important aspects in triggering emotional effects in the reader. For instance, the melancholic tone of Juhani Aho's text is not generated by themes of separation and loss alone but also by Aho's style, which favors fragmentation and loosening of syntax, with a recurring mannerism of three points "...", as a

sign of hesitation and withdrawal, even evoking a depressive loss of contact.

The qualitative analysis demonstrates that the selected texts depict and trigger negative emotions in particular: feelings of deception, fear, anxiety, disgust and hatred, anger, moral indignation and melancholia. On one hand, this can be explained by genre-specific emotional effects: a critical naturalist novel tends to shock and challenge its reader by representing and inciting strong negative emotions, which confirm the effect of reality of a text and direct the reader's attention to the social defects described. For instance, the emotion of disgust, which is a genre-specific emotion of the naturalist novel, is a named emotion and salient in Sillanpää's and Canth's novels in particular (Rossi, 2007, 2017, 2020).

The salience of negative emotions can be explained by the importance of negative emotions in literature and art in general. As discussed by Menninghaus et al. (2017) negative emotions are an important resource for the arts, since negative emotions have been shown to be particularly powerful in securing attention, intense emotional involvement, and high memorability, and hence is precisely what artworks strive for.

This dataset will be used for more robust detection of tone and mood in Finnish literature. Our preliminary studies show that the "big data" results support qualitative analyses and further justifies the use of purely lexicon-based methods when dealing with larger collections of text where word choice is an important factor of creating affective states in the reader. Specifically, we can see that the choice of emotion associated words in the first three paragraphs correlates highly with established analyses of mood in the selected texts. We hope to add established emotion categories from literary affect studies (see e.g. Hogan 2011) to the lexicon as a measure to further improve the usability of the FEIL lexicon for the literary domain (Öhman, 2020). Additionally, we would like to expand on the methodologies used in this exploratory study and hopefully create more and more robust approaches to tone and mood detection in literature.

Acknowledgements

This work was supported by JSPS KAKENHI Grant Number 22K18154 and the Waseda University Tokutei-kadai Grant 2022C-544.

References

- Eija Airio. 2009. *Morphological Problems in IR and CLIR. Applying linguistic methods and approximate string matching tools*. Tampere University Press.
- Nancy Armstrong. 2014. The affective turn in contemporary fiction. *Contemporary Literature*, 55(3):441–465.
- Andrew Bartlett, Jamie Lewis, Luis Reyes-Galindo, and Neil Stephens. 2018. The locus of legitimate interpretation in Big Data sciences: Lessons for computational social science from-omic biology and high-energy physics. *Big Data & Society*, 5(1):2053951718768831.
- Jonathan Flatley. 2008. *Affective mapping: melancholia and the politics of modernism*. Harvard University Press Cambridge, MA.
- Pirkko Forsman Svensson. 2011. Virtuaalinen vanha kirjasuomi. http://www.vvks.info/aanne-_ja_muotooppi/heittyminen/.
- Suchin Gururangan, Mike Lewis, Ari Holtzman, Noah Smith, and Luke Zettlemoyer. 2022. **DEMIX layers: Disentangling domains for modular language modeling**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5557–5576, Seattle, United States. Association for Computational Linguistics.
- Mika Härmäläinen and Khalid Alnajjar. 2021. **The current state of Finnish NLP**. In *Proceedings of the Seventh International Workshop on Computational Linguistics of Uralic Languages*, pages 65–72, Syktyvkar, Russia (Online). Association for Computational Linguistics.
- Mika Härmäläinen, Niko Partanen, and Khalid Alnajjar. 2021. **Lemmatization of historical old literary Finnish texts in modern orthography**. In *Actes de la 28e Conférence sur le Traitement Automatique des Langues Naturelles. Volume 1 : conférence principale*, pages 189–198, Lille, France. ATALA.
- Patrick Colm Hogan. 2011. *What literature teaches us about emotion*. Cambridge University Press.
- Jenna Kanerva, Filip Ginter, and Tapio Salakoski. 2019. Universal Lemmatizer: A Sequence to Sequence Model for Lemmatizing Universal Dependencies Treebanks. *arXiv preprint arXiv:1902.00972*.
- Suzanne Keen. 2007. *Empathy and the Novel*. Oxford University Press on Demand.
- Evgeny Kim and Roman Klinger. 2018. A survey on sentiment and emotion analysis for computational literary studies. *arXiv preprint arXiv:1808.03137*.
- Pirjo Lyytikäinen. 2017. How to Study Emotion Effects in Literature: Written Emotions in Edgar Allan Poe’s “The Fall of the House of Usher”. In *Writing Emotions: Theoretical Concepts and Selected Case Studies in Literature*, pages 247–64. Bielefeld: Transcript Verlag.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. **Learning word vectors for sentiment analysis**. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA. Association for Computational Linguistics.
- Winfried Menninghaus, Valentin Wagner, Julian Hanich, Eugen Wassiliwizky, Thomas Jacobsen, and Stefan Koelsch. 2017. The distancing-embracing model of the enjoyment of negative emotions in art reception. *Behavioral and Brain Sciences*, 40.
- Saif M. Mohammad and Peter D. Turney. 2013. Crowdsourcing a word-emotion association lexicon. *Computational Intelligence*, 29(3):436–465.
- Sianne Ngai. 2005. *Ugly feelings*, volume 6. Harvard University Press Cambridge, MA.
- Emily Öhman. 2020. Emotion annotation: Rethinking emotion categorization. In *DHN Post-Proceedings*, pages 134–144.
- Emily Öhman. 2022. **SELF & FEIL: Emotion Lexicons for Finnish**. In *Proceedings of the 6th Digital Humanities in the Nordic and Baltic Countries conference*. CEUR Workshop Proceedings.
- Emily Öhman and Riikka Rossi. 2021. Affect and Emotions in Finnish Literature: Combining Qualitative and Quantitative Approaches. In *The Language of Emotions: Building and Applying Computational Methods for Emotion Detection for English and Beyond*.
- Niko Partanen, Mika Härmäläinen, and Khalid Alnajjar. 2019. **Dialect text normalization to normative standard Finnish**. In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 141–146, Hong Kong, China. Association for Computational Linguistics.
- Menakhem Perry. 1979. Literary dynamics: how the order of a text creates its meanings [with an analysis of Faulkner’s “A Rose for Emily”]. *Poetics today*, 1(1/2):35–361.
- Robert Plutchik. 1980. A general psychoevolutionary theory of emotion. *Theories of emotion*, 1:3–31.
- I.A. Richards. 1929. *Practical Criticism*, volume 1964. Routledge: London.
- Riikka Rossi. 2007. *Le naturalisme finlandais. Une conception entropique du quotidien*. Suomalaisen Kirjallisuuden Seura.
- Riikka Rossi. 2017. Writing disgust, writing realities. *Sabine Schönfellner, Gudrun Tockner (eds.) Writing Emotions*, page 277.

- Riikka Rossi. 2020. *Alkukantaisuus ja tunteet — Primitivismi 1900-luvun alun suomalaisessa kirjallisuudessa*. Number 1456 in Suomalaisen Kirjallisuuden Seuran toimituksia. Suomalaisen Kirjallisuuden Seura.
- Riikka Rossi and Pirjo Lyytikäinen. 2022. Pohjoisia tunteita. *AVAIN-Kirjallisuudentutkimuksen aikakauslehti*, 19(1):3–9.
- Howard Sklar. 2013. *The Art of Sympathy in Fiction : Forms of Ethical and Emotional Persuasion*. Linguistic Approaches to Literature. John Benjamins Publishing Company.
- Rachel Greenwald Smith. 2011. Postmodernism and the affective turn. *Twentieth Century Literature*, 57(3/4):423–446.
- Lewis Turco. 2020. *The book of literary terms: the genres of fiction, drama, nonfiction, literary criticism, and scholarship*. University of New Mexico Press.
- Wouter van Atteveldt, Mariken ACG van der Velden, and Mark Boukes. 2021. The validity of sentiment analysis: Comparing manual annotation, crowd-coding, dictionary approaches, and machine learning algorithms. *Communication Methods and Measures*, 15(2):121–140.
- Caspar J Van Lissa, Marco Caracciolo, Thom van Duuren, and Bram van Leuven. 2018. Difficult Empathy-The Effect of Narrative Perspective on Readers’ Engagement with a First-Person Narrator.
- Antti Virtanen, Jenna Kanerva, Rami Ilo, Jouni Luoma, Juhani Luotolahti, Tapio Salakoski, Filip Ginter, and Sampo Pyysalo. 2019. Multilingual is not enough: BERT for Finnish. *arXiv preprint arXiv:1912.07076*.
- Zhe Ye, Fang Li, and Timothy Baldwin. 2018. [Encoding sentiment information into word vectors for sentiment analysis](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 997–1007, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Liang-Chih Yu, Jin Wang, K Robert Lai, and Xuejie Zhang. 2017. Refining word embeddings using intensity scores for sentiment analysis. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 26(3):671–681.