Annotation and Detection of Emotion Polarity in *I Promessi Sposi*: Dataset and Experiments

Rachele Sprugnoli[1,*], Arianna Redaelli[1]

[1]*Università di Parma, Via D'Azeglio, 85, 43125 Parma, Italy*

Abstract

Emotions play a crucial role in literature and are studied by various disciplines, e.g. literary criticism, psychology, anthropology and, more recently, also with computational methods in NLP. However, studies in the Italian context are still limited. This work therefore aims to advance the state of the art in the field of emotion analysis applied to historical texts by proposing a new dataset and describing the results of a set of emotion polarity detection experiments. The text analyzed is "I Promessi Sposi" in its final edition (published in 1840), one of the most important novels in the Italian literary and linguistic canon.

Keywords

emotion analysis, annotation, fine-tuning, Italian, literary texts

1. Introduction

Emotions play a key role in literature, representing a bridge between the author's purposes, the text, and the reader's personal background: literature collects experiences and contains the emotions that accompany them, in turn generating new experiences and new emotions. Therefore, studying emotions in literary texts implies the possibility of providing valuable insights into the deeper meanings and intentions behind a work, the form it may take, and the readers' engagement with it. This field of study has recently experienced a flourishing national and international development involving different disciplines, from literary criticism to philosophy, from anthropology to psychology. For example, in the Italian context, Ginzburg et al. [1] analyzed how Matte Blanco's psychoanalytic theories on emotions are applied to literary criticism, taking into account authors like Tozzi, Pirandello, and Svevo, while Guaragnella [2] explored the complex interaction between humor and sadness in 20th-century Italian literature from a both philosophical and literary point of view.

However, some literary works remained underexplored. One such work is Alessandro Manzoni's "I Promessi Sposi". Despite its emotional richness, the novel has often been regarded as monolithic and static, both because of the narrated events, strongly influenced by the author's religious spirit and social and political polemic, and because it quickly became a model of the Italian language, stably included in school curricula as mandatory study material. This has led to a certain degree of reluctance and lack of enthusiasm among the readers.

As a consequence, a study of emotions in "I Promessi Sposi" can be beneficial from both an academic and educational standpoint. Academically, it can provide new insights into a classic text, encouraging new interpretations and scholarly discussions. For didactic purposes, analyzing the emotions in "I Promessi Sposi" can make the novel more relatable and appealing for students, revealing the depth and complexity of the characters' experiences in the context in which they live, and encouraging a closer connection with them and with Manzoni's social issues.

Given this context, computational methods, already widely applied especially on user-generated contents (such as reviews and social media posts), can be profitably tested on the fictional text after developing specific datasets for training and evaluating new models. The present work takes as a basis a preliminary annotation of the Manzoni's novel, expanding the number of manually labeled sentences and proposing the development of some models of varying complexity.

More specifically, two are the main contributions of our work: i) we release[1] a new dataset made of more than 3.000 sentences taken from "I Promessi Sposi" manually annotated with four emotion polarity classes (i.e. POSITIVE, NEGATIVE, NEUTRAL, MIXED); ii) we test various approaches for emotion polarity detection using the new dataset as-is but also augmenting it with other annotated Italian resources.

CLiC-it 2024: Tenth Italian Conference on Computational Linguistics, Dec 04 — 06, 2024, Pisa, Italy

*Corresponding author.

†This paper is the result of the collaboration between the two authors. For the specific concerns of the Italian academic attribution system: Rachele Sprugnoli is responsible for Sections 2, 3.2, 3.3, 4; Arianna Redaelli is responsible for Sections 1, 3, 3.1, 5.

rachele.sprugnoli@unipr.it (R. Sprugnoli); arianna.redaelli@unipr.it (A. Redaelli)

0000-0001-6861-5595 (R. Sprugnoli); 0000-0001-6374-9033 (A. Redaelli)

[1]https://github.com/RacheleSprugnoli/Emotion_Analysis_Manzoni

2. Related Work

Emotion analysis, that is the automatic recognition of emotions conveyed in a text, is a Natural Language Processing (NLP) task applied to various types of texts. In fact, although most datasets and systems are developed to process social media posts and reviews, there are also applications on news [3], songs [4] and personal narratives [5]. After the so-called affective-turn in literary studies [6], the attention towards this task has significantly increased also in the humanities with studies on both historical and ancient languages and on various textual genres.[2] Among these we mention, as examples, drama plays [9, 10], fairy tales [11], poems [12] and children's literature [13]. As for novels, Mohammad [14] compared fairy tales and novels from the point of view of emotions identified with the NRC Emotion Lexicon [15]; Zehe et al. [16] used emotion analysis for discriminating between German novels with and without happy endings; Stankovic et al. [17] presented various experiments on Serbian novels; Kim [18] tested the dictionary-based tool Syuzhet[3] on a set of 19th-century British novels. As regards the literary domain, however, the works on Italian are few: for example, Rebora [19] analyzed the annotation of a short story by Pirandello as performed by a group of students; Pavan [20] applied a lexicon-based software to 16 novels and poems written in the twentieth century; and Zhang et al. [21] released a dataset of opera verses with which they performed various emotion recognition experiments.

The present work wants to advance the state of the art in the field of emotion analysis applied to historical novels; specifically, a previous preliminary annotation of "I Promessi Sposi" is taken up [22], expanding the number of manually labelled sentences (from 338 to 3,095) and proposing new experiments for the automatic identification of emotion polarity. Although the novel in question is considered one of the most important in the history of Italian literature and language, as far as we know, this study is the first to address the topic of emotions in Manzoni's work through computational methods, developing specific resources and models.

3. Dataset Creation

The dataset is composed of 3,095 manually split[4] sentences from 12 chapters (about 30% of the total chapters of the novel) chosen to cover various phases of the plot, different characters and types of content. Specifically, we used Chapter III, in which Renzo (one of the protagonists) goes to the lawyer Azzeccagarbugli in an attempt to resolve the legal obstacle preventing him from marrying his beloved Lucia. However, this results in a misunderstanding and the ultimate failure of his endeavor. Chapters IV and V describe the conversion of Fra Cristoforo, a religious figure and friend of the betrothed couple, and his heated discussion with Don Rodrigo, the lord who is preventing Renzo and Lucia's marriage, which also ends in failure. Chapters IX and X introduce the ambivalent story of the Nun of Monza, chosen as the protector of Lucia who is fleeing from Don Rodrigo. Chapters XIV and XV depict Renzo's involvement in the bread riot in Milan, after which he gets drunk at the Full Moon Tavern, is arrested, and eventually manages to escape. Chapters XX and XXI describe Lucia's arrival at the house of the Unnamed, the worst baron of that time, who, at Don Rodrigo's request, kidnaps her – only to later repent in a tormenting process of conversion to the Christian faith. Chapter XXVIII contain an historical digression on Milan, devastated by famine, the invasion of the Lansquenets, and the threat of the plague. Chapter XXXIII portrays Don Rodrigo on his deathbed, suffering from the plague, and a flashback to Renzo, who, having recovered from the disease, sets out to find Lucia. Finally, the last chapter, Chapter XXXVIII, depicts the conclusion of the story, with the serene reunion of the couple, now ready to embark on their married life. As can be seen, our choice provided very lively parts, others more introspective, and others that contain descriptions and historical digressions.

The annotation was carried out by the two authors of this paper independently, following the guidelines reported below and using a spreadsheet having a sentence per row. While annotating, the annotators did not have access to each other annotator's score. Each chapter consists of between approximately 170 and 330 sentences and the average annotation time per chapter was about 1.5 hours (18 hours in total). Subsequently, the results of the independently conducted annotations were placed in parallel columns to allow each annotator to revise any obvious errors or oversights. This preliminary phase was followed by a direct discussion between the two annotators to address the most problematic cases and achieve the gold annotation (see Sections 3.2 and 3.3).

3.1. Guidelines

The annotation was carried out at sentence level and was based on both the lexicon used and the images evoked by the author, for example through the use of rhetorical figures. The annotation followed the flow of the text, so the annotator can take into account the previous sentences

[2] For a complete overview of sentiment and emotion analysis in the field of literary studies please refer to the survey papers by Kim and Klinger [7] and Rebora [8].

[3] https://github.com/mjockers/syuzhet

[4] Sentence splitting was done manually because automatic segmentation presented significant challenges for the models currently available for Italian, largely due to the novel's intricate punctuation. Details are described in [23].

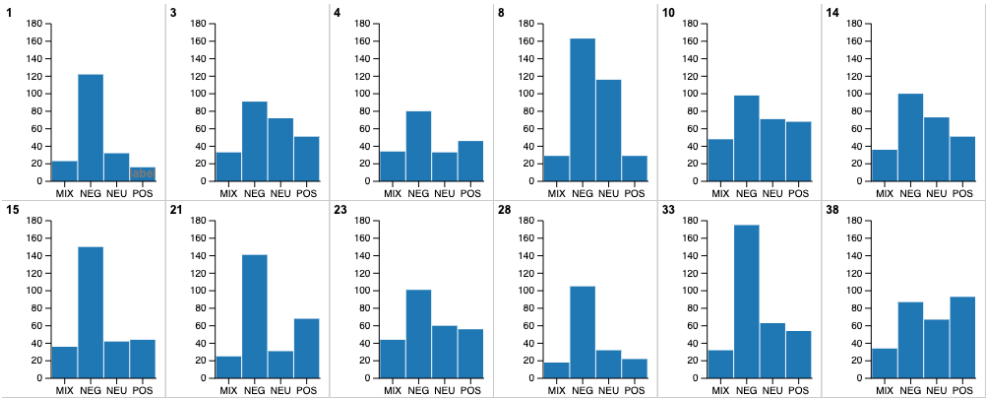

Figure 1: Bar charts displaying the distribution of the four classes in the twelve annotated chapters.

but not the following ones. The polarity to be annotated was the one expressed by the author, either through the narrator or through the characters who take part in the events told, and not the one felt by the annotator while reading the sentence. The polarity could also concern emotions related to a different time from that of the main story. To assign the correct label, the annotator had to answer the question *how are the emotions evoked by the author in the sentence being analyzed?* with one of the following options:

- predominantly or solely positive (label POSITIVE), such as caring, joy, relief, amusement;
- predominantly or solely negative (label NEGATIVE), such as confusion, nervousness, annoyance, resignation, disapproval, fear, disappointment, embarrassment, sadness, pain, anger and remorse;
- of the opposite type, thus it is not possible to find a clearly prevalent emotion (label MIXED);
- absent (label NEUTRAL).

This distinction was inspired by previous annotation efforts, such as the one underlying the SENTIPOLC shared task in which the four labels were applied to tweets [24, 25]. The guidelines have been revised and enriched after the analysis of the disagreements, as will be described in the next subsection.

3.2. Agreement

The Cohen's kappa calculated for each chapter recorded a minimum value of 0.51 (on Chapter III) and a maximum value of 0.71 (on Chapter XXIII). On average, therefore, a moderate agreement of 0.62 was obtained. Specifically, the most difficult class to annotate was MIXED (k = 0.50), while for the other labels the differences were less marked: 0.63 for NEUTRAL, 0.65 for POSITIVE and 0.70 for NEGATIVE.

From the analysis of the disagreements, it emerged that some uncertainties were related to the presence of irony. It was therefore decided to annotate these cases as MIXED, since in such sentences two polarities coexist, i.e. the one expressed by the literal meaning and the one due to the presence of irony. It is important to note that, in our annotation, irony was considered as a sentiment shifter that changes the polarity of the literal meaning of a sentence. This interpretation of irony is much more narrowed compared to that of Manzoni's literary criticism. In fact, in "I Promessi Sposi" irony is a complex rhetorical device that can subtly influence the reader's perception and understanding on multiple levels of the novel [26]. Consequently, the term *irony* refers not only to irony in its strict sense but also to humor, sarcasm, innuendo, and other related concepts, which the author uses to suggest more in-depth information into characters, situations, linguistic uses, and social problems [27]. However, for our purposes, it was not practical to apply this broader concept of irony because it often requires a deep understanding of the author's intentions that goes far beyond the sequential interpretation of individual sentences. Another aspect revised and better detailed in the guidelines was the annotation of approval expressions (such as "Sì, signore.", EN: *Yes, sir*), that it was decided to annotate as NEUTRAL and not as POSITIVE, unless they were accompanied by other elements expressing positive emotions. Descriptive sentences also had to be annotated as NEUTRAL if they did not contain words that evoked specific emotions. For example, "Era un guazzabuglio di steli, che facevano a soverchiarsi l'uno con l'altro nell'aria" (EN: *It was a jumble of stems, which tried to overwhelm each other in the air.*) should have been annotated as NEGATIVE for the presence of words that evoke confusion and oppression; on the contrary, "Per un

buon pezzo, la costa sale con un pendìo lento e continuo" (EN: *For a good while, the coast rises with a slow and continuous slope*) should have been annotated as NEUTRAL. Lastly, courtesy titles (such as "reverendissimo", EN: *most reverend*) also had to be assigned the NEUTRAL label because they represent a formal requirement and not a true positive emotional involvement. Annotating dialogue turns proved to be particularly difficult, especially when dealing with very short sentences, composed of 1 to 3 words. In these cases, the preceding context but also the presence of punctuation and interjections were essential for assigning the polarity label.

3.3. Final Dataset

The dataset resulting from the consolidation of disagreements is made up of 1,413 sentences annotated as NEGATIVE (corresponding to 46% of the total sentences), 692 NEUTRAL sentences (22%), 598 POSITIVE sentences (19%) and 392 MIXED ones (13%). The distribution of the four classes in the various chapters is shown by the bar graphs in Figure 1. The fact that most of the sentences have a negative polarity is in line with the topics covered in the novel: kidnappings, misunderstandings, plague. The only chapter in which the POSITIVE label prevails is the last one (XXXVIII) which tells the happy ending of the novel, that is, the marriage and the new happy life of the two protagonists. It is interesting to note that compared to the first tests of annotating emotion polarity [22], the NEUTRAL class is no longer the most frequent in the data. Since then, the guidelines had been enriched with details regarding the specific emotions to be considered as positive and negative: this allowed the annotators to be more precise in identifying the prevalent type of emotion even in the case of minimal nuances.

4. Experiments

The annotated dataset described in the previous Section was used to train and evaluate various approaches of different complexity, namely:

- a Linear Support Vector classifier (SVC) developed using the `scikit-learn` library with default parameters and to be considered as a baseline;
- a fine-tuned model of `bert-base-italian-xxl-cased`[5] using the AdamW optimizer (learning rate: 2e-5, epsilon: 1e-8) and 2 epochs[6];
- a fine-tuned model of multilingual XLM-RoBERTa [28] using an Hugging Face PyTorch implementation[7] and the following hyperparameters: 32 for batch size, 2e-5 for learning rate, 6 epochs, AdamW optimizer;
- a lexicon-based script employing both a polarity lexicon created for contemporary Italian (i.e., W-MAL, Weighted-Morphologically-inflected Affective Lexicon) [29] and one derived from 19th-century Italian narrative texts[8]. A score is computed for each sentence by summing the polarity values of the tokens. If the score is greater than 0, the label is POSITIVE; if it is less than 0, the label is NEGATIVE; if it is equal to 0 because all tokens have this value or are not present in the lexicon, the label is NEUTRAL; if it is equal to 0 because the sum of tokens with positive and negative polarities is balanced, the label is MIXED.

The experiments were performed using the dataset consisting only of the novel's chapters (divided into training, development and test sets according to the proportions 80/10/10) but also adding data from other Italian linguistic resources annotated with emotions in order to have more training examples. In particular, the resources used to augment the original dataset are the following:

- MultiEmotions-it: a multi-labelled emotion dataset made of comments posted on Facebook and YouTube annotated following Plutchik's basic emotions (`anger`, `disgust`, `fear`, `joy`, `sadness`, `surprise`, `trust`, `anticipation`) and dyads (such as `love` and `disappointment`) [30];
- FEEL-IT: a benchmark corpus of tweets annotated with four emotions, that is `fear`, `joy`, `sadness`, `anger` [31];
- EMit: a dataset of multi-labelled tweets annotated with Plutchik's basic emotions plus `love` and `neutral` [32];
- XED: a multilingual emotion dataset in which the annotation performed on Finnish and English sentences are projected on the corresponding items in 30 languages, including Italian, using parallel corpora [33]. The eight Plutchik's basic emotions are adopted for the annotation;

[5] Provided by the MDZ Digital Library team of the Bavarian State Library through the Hugging Face framework: https://huggingface.co/dbmdz/bert-base-italian-xxl-cased. We tested all the BERT models available in the MDZ Digital Library repository, that is `bert-base-italian-cased`, `bert-base-italian-uncased`, `bert-base-italian-xxl-cased` and `bert-base-italian-xxl-cased`; the latter performed better than the others, so we will focus on it in the remainder of the paper.

[6] We adapted the notebook https://www.kaggle.com/code/neerajmohan/fine-tuning-bert-for-text-classification.

[7] We adapted the following implementation: https://gist.github.com/sayakmisra/b0cd67f406b4e4d5972f339eb20e64a5.

[8] https://github.com/RacheleSprugnoli/Emotion_Analysis_Manzoni

Table 1
Results in terms of F1 for the tested supervised approaches. In bold the best F1 achieved for each class and the best macro average score. The last column displays how many instances are in each class in the test set.

	SVC			Fine-tuned BERT			Fine-tuned XLM-RoBERTA			Support
	Manzoni	M-M-E	All	Manzoni	M-M-E	All	Manzoni	M-M-E	All	
POSITIVE	0.29	0.30	0.09	0.53	0.49	0.50	**0.60**	0.55	0.59	105
NEGATIVE	0.49	0.47	0.46	0.58	0.58	0.54	**0.59**	**0.59**	0.55	102
NEUTRAL	0.32	0.28	0.22	0.52	0.57	0.46	0.56	**0.65**	0.55	78
MIXED	0.08	0.04	0.09	0.13	0.23	0.18	0.09	**0.27**	0.32	39
Macro Avg.	0.29	0.27	0.28	0.44	0.47	0.42	0.46	**0.53**	0.50	

Table 2
F1 score obtained with the lexicon-based approach.

F1	Lexicon-Based Approach	
	W-MAL	**XIX cent.**
POSITIVE	0.45	0.44
NEGATIVE	0.35	0.31
NEUTRAL	0.15	0.48
MIXED	0.00	0.19
Macro Avg.	0.24	0.35

- TwIT: a corpus of tweets annotated with six different emotions (i.e., `happiness`, `trust`, `sadness`, `anger`, `fear` and `disgust`) [34];
- AriEmozione 2: a dataset of verses of opera arias written in 18th-century Italian annotated with one out of six emotions (i.e., `love`, `joy`, `admiration`, `anger`, `sadness`, `fear`) [21].

The original emotion labels of the aforementioned resources were mapped onto our four classes on the basis of their polarity. Data labelled with ambiguous emotions (such as `surprise` and `anticipation`) were left out. Please note that only MultiEmotions-it and EMit contain the class NEUTRAL and that their multi-label structure allowed us to convert the original annotation to the MIXED class when the emotions assigned to the same sentence were of opposite polarity.

Based on the characteristics of the aforementioned datasets, three training sets were prepared: one with only sentences taken from "I Promessi Sposi" (`Manzoni`, 2,771 instances), one adding MultiEmotions-it and EMit to the sentences taken from Manzoni's novel (`Manz-Multi-EMit`, 10,755 instances), and one joining all the available datasets (`All`, 21,923 instances).

4.1. Results

Tables 1 and 2 show the results of the experiments carried out reporting the F1 score for each class and the macro average. As for the supervised approaches (Table 1) scores are given considering each one of the set used for training or fine-tuning the models.

The lexicon-based approach outperforms the baseline (i.e., the Support Vector Classifier); the latter does not benefit from increasing the size of the training set and performs very poorly in recognizing sentences annotated as MIXED (F1 < 0.1). Using an in-domain lexicon specially created starting from nineteenth-century texts yields better results with respect to using the W-MAL lexicon. This improvement is noted both in terms of macro average F1 (+ 0.11) and in the recognition of NEUTRAL and MIXED instances, +0.33 and +0.19 respectively. The fine-tuned XLM-RoBERTa model achieves the best F1 both overall (0.53) and for all classes even if using different training sets. Interestingly, in the case of fine-tuned models (both using BERT and XLM-RoBERTa) the `All` training set, although significantly larger than the others, does not provide the greatest benefits. Indeed, the most beneficial training set is `Manz-Multi-EMit` which combines the most similar datasets from the annotation point of view, as both MultiEmotions-it and EMit contain NEUTRAL and MIXED sentences.

Figure 2 shows the confusion matrix for the best model. We can notice an over-prediction of the NEGATIVE label even if this is not the most frequent class of the dataset, covering 35.8% of the total (while the POSITIVE class represents 38.1% of the total). Examples of sentences incorrectly classified as NEGATIVE are:

- "Per i nostri fu una nuova cuccagna." EN: *For our people it was a new bonanza.* Gold label = POSITIVE
- "Già principiava a farsi buio." EN: *It was already starting to get dark.* Gold label = NEUTRAL
- "Io ho perdonato tutto: non ne parliam più: ma me n'avete fatti dei tiri." EN: *I've forgiven everything: we don't talk about it anymore: but you played tricks on me.* Gold label = MIXED

5. Conclusions

This paper presents a new manually annotated dataset and a set of experiments for the automatic detection of emotion polarity. More specifically, the dataset contains

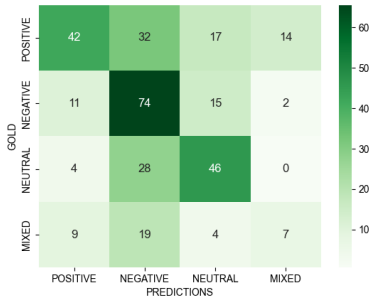

Figure 2: Confusion matrix for the XLM-RoBERTA model fine-tuned with the `Manz-Multi-EMit` training set.

3,095 sentences taken from "I Promessi Sposi" and the experiments cover different approaches, namely lexicon-based, SVC and the fine-tuning of an Italian BERT model and of the multilingual XLM-RoBERTa model. The impact of the training set size is also evaluated by increasing the in-domain dataset by combining other annotated Italian resources.

We are aware that for the emotion analysis task, as for all NLP tasks, Large Language Models are now widely used [35] but these require computational powers currently not available to the authors of the paper. In the future, our work will focus on this aspect in order to be in line with the current state of the art. Another future work will concern the annotation of emotions with more granular labels, extending an activity already started on Chapter VIII only, on which the label scheme proposed for the GoEmotions dataset [36] was applied [22]. Additionally, we plan to pay greater attention to the annotation of irony, a crucial aspect of the novel. This could be incorporated into the dataset using a binary 0/1 value to indicate its presence or absence, as we have already begun to implement [9]. Finally, we would like to explore the applications of our work in the school context. Concerning the study of emotions in Manzoni's novel, computational methods and tools could provide inputs and data useful for didactic practical activities, such as visual representations of affective scenes, role-playing exercises, or even crowd-sourced annotation that allows students to express their personal interpretations of the characters' emotions in different chapters and situations. Activities like these can make the whole learning experience more dynamic and captivating, promoting a deeper connection between the students and the novel and, meanwhile, improving their critical thinking and empathy.

[9] https://github.com/RacheleSprugnoli/Emotion_Analysis_Manzoni

Acknowledgments

The authors thank Giovanni Moretti for the help given with the fine-tuning scripts.
Questa pubblicazione è stata realizzata da ricercatrice con contratto di ricerca cofinanziato dall'Unione europea - PON Ricerca e Innovazione 2014-2020 ai sensi dell'art. 24, comma 3, lett. a, della Legge 30 dicembre 2010, n. 240 e s.m.i. e del D.M. 10 agosto 2021 n. 1062.

References

[1] A. Ginzburg, R. Luperini, V. Baldi (Eds.), Emozioni e letteratura. La teoria di Matte Blanco e la critica letteraria contemporanea, Fabrizio Serra, Firenze, 2009.

[2] P. Guaragnella, I volti delle emozioni. Riso, sorriso e malinconia nel Novecento letterario italiano, Società Editrice Fiorentina, Firenze, Italy, 2015.

[3] A. M. Patronella, Covering Climate Change: A Sentiment Analysis of Major Newspaper Articles from 2010-2020, Inquiries Journal 13 (2021).

[4] D. Edmonds, J. Sedoc, Multi-emotion classification for song lyrics, in: Proceedings of the Eleventh Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis, 2021, pp. 221–235.

[5] A. Tammewar, A. Cervone, E.-M. Messner, G. Riccardi, Annotation of emotion carriers in personal narratives, in: Proceedings of the Twelfth Language Resources and Evaluation Conference, 2020, pp. 1517–1525.

[6] S. Keen, Introduction: Narrative and the Emotions, Poetics Today 32 (2011) 1–53. URL: https://doi.org/10.1215/03335372-1188176. doi:10.1215/03335372-1188176.

[7] E. Kim, R. Klinger, A survey on sentiment and emotion analysis for computational literary studies, Wolfenbüttel, 2019.

[8] S. Rebora, Sentiment analysis in literary studies. a critical survey, DHQ: Digital Humanities Quarterly 17 (2023).

[9] T. Schmidt, K. Dennerlein, C. Wolff, Emotion classification in German plays with transformer-based language models pretrained on historical and contemporary language, in: S. Degaetano-Ortlieb, A. Kazantseva, N. Reiter, S. Szpakowicz (Eds.), Proceedings of the 5th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature, Association for Computational Linguistics, Punta Cana, Dominican Republic (online), 2021, pp. 67–79. URL: https://aclanthology.org/2021.latechclfl-1.8. doi:10.18653/v1/2021.latechclfl-1.8.

[10] F. Debaene, K. van der Haven, V. Hoste, Early Modern Dutch comedies and farces in the spotlight: Introducing EmDComF and its emotion framework, in: R. Sprugnoli, M. Passarotti (Eds.), Proceedings of the Third Workshop on Language Technologies for Historical and Ancient Languages (LT4HALA) @ LREC-COLING-2024, ELRA and ICCL, Torino, Italia, 2024, pp. 144–155. URL: https://aclanthology.org/2024.lt4hala-1.17.

[11] E. P. Volkova, B. Mohler, D. Meurers, D. Gerdemann, H. H. Bülthoff, Emotional perception of fairy tales: achieving agreement in emotion annotation of text, in: Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text, 2010, pp. 98–106.

[12] R. Sprugnoli, F. Mambrini, M. Passarotti, G. Moretti, The Sentiment of Latin Poetry. Annotation and Automatic Analysis of the Odes of Horace, IJCoL. Italian Journal of Computational Linguistics 9 (2023) 53–71.

[13] S. Rebora, M. Lehmann, A. Heumann, W. Ding, G. Lauer, Comparing ChatGPT to Human Raters and Sentiment Analysis Tools for German Children's Literature, in: CHR 2023: Computational Humanities Research Conference, Paris, France, 2023, pp. 333–343.

[14] S. Mohammad, From once upon a time to happily ever after: Tracking emotions in novels and fairy tales, in: K. Zervanou, P. Lendvai (Eds.), Proceedings of the 5th ACL-HLT Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities, Association for Computational Linguistics, Portland, OR, USA, 2011, pp. 105–114. URL: https://aclanthology.org/W11-1514.

[15] S. M. Mohammad, P. D. Turney, Crowdsourcing a word-emotion association lexicon, Computational Intelligence 29 (2013) 436–465.

[16] A. Zehe, M. Becker, L. Hettinger, A. Hotho, I. Reger, F. Jannidis, Prediction of happy endings in German novels based on sentiment information, in: 3rd Workshop on Interactions between Data Mining and Natural Language Processing, Riva del Garda, Italy, volume 5, 2016.

[17] R. Stankovic, M. Kosprdic, M. Ikonic-Nesic, T. Radovic, Sentiment Analysis of Sentences from Serbian ELTeC corpus, in: Proceedings of the SALLD-2 Workshop at Language Resources and Evaluation Conference (LREC), Marseille, France, 2022, pp. 31–38.

[18] H. Kim, Sentiment analysis: Limits and progress of the syuzhet package and its lexicons., DHQ: Digital Humanities Quarterly 16 (2022).

[19] S. Rebora, et al., Shared Emotions in Reading Pirandello. An Experiment with Sentiment Analysis, Marras, C., Passarotti, M., Franzini, G., and Litta, E.(eds), Atti del IX Convegno Annuale AIUCD. La svolta inevitabile: sfide e prospettive'per l'Informatica Umanistica. Università Cattolica del Sacro Cuore, Milano (2020) (2020) 216–221.

[20] L. Pavan, A survey of some Italian literature works using sentiment analysis, International Journal of Linguistics, Literature and Translation (2022) 117–121.

[21] S. Zhang, F. Fernicola, F. Garcea, P. Bonora, A. Barrón-Cedeño, AriEmozione 2.0: Identifying Emotions in Opera Verses and Arias, IJCoL. Italian Journal of Computational Linguistics 8 (2022) 7–26.

[22] R. Sprugnoli, A. Redaelli, How to Annotate Emotions in Historical Italian Novels: A Case Study on I Promessi Sposi, in: Proceedings of the Third Workshop on Language Technologies for Historical and Ancient Languages (LT4HALA)@ LREC-COLING-2024, 2024, pp. 105–115.

[23] A. Redaelli, R. Sprugnoli, Is Sentence Splitting a Solved Task? Experiments to the Intersection Between NLP and Italian Linguistics, in: Proceedings of the 10th Italian Conference on Computational Linguistics (CLiC-it 2024), Pisa, Italy, 2024.

[24] V. Basile, A. Bolioli, V. Patti, P. Rosso, M. Nissim, Overview of the evalita 2014 sentiment polarity classification task, in: Proceedings of the First Italian Conference on Computational Linguistics CLiC-it 2014 & and of the Fourth International Workshop EVALITA 2014: 9-11 December 2014, Pisa, Pisa University Press, 2014, pp. 50–57.

[25] F. Barbieri, V. Basile, D. Croce, M. Nissim, N. Novielli, V. Patti, et al., Overview of the evalita 2016 sentiment polarity classification task, in: EVALITA Evaluation of NLP and Speech Tools for Italian. Proceedings of the Final Workshop, AILC, 2016, pp. 146–155.

[26] E. Raimondi, L'ironia polifonica in manzoni, in: La dissimulazione romanzesca: antropologia manzoniana, Il mulino, Bologna, 1990, pp. 45–80.

[27] L. Manfreda, Figure dell'ironia nei Promessi sposi : il ruolo doppio a rovescio dei personaggi, Metauro, Pesaro, 2006.

[28] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, V. Stoyanov, Unsupervised cross-lingual representation learning at scale, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, online, 2020, pp. 8440–8451. URL: https://aclanthology.org/2020.acl-main.747. doi:10.18653/v1/2020.acl-main.747.

[29] M. Vassallo, G. Gabrieli, V. Basile, C. Bosco, et al., Polarity imbalance in lexicon-based sentiment analysis, in: Proceedings of the Seventh Italian Conference on Computational Linguistics, CEUR, 2020,

pp. 1–7.
[30] R. Sprugnoli, MultiEmotions-it: A new dataset for opinion polarity and emotion analysis for Italian, in: Proceedings of the Seventh Italian Conference on Computational Linguistics (CLiC-it 2020), Accademia University Press, online, 2020, pp. 402–408.
[31] F. Bianchi, D. Nozza, D. Hovy, FEEL-IT: emotion and sentiment classification for the Italian language, in: The 16th Conference of the European Chapter of the Association for Computational Linguistics, Association for Computational Linguistics, online, 2021, pp. 76–83.
[32] O. Araque, S. Frenda, R. Sprugnoli, D. Nozza, V. Patti, et al., EMit at EVALITA 2023: overview of the categorical emotion detection in Italian social media task, in: EVALITA Proceedings of the Eighth Evaluation Campaign of Natural Language Processing and Speech Tools for Italian Final Workshop, AILC - Associazione Italiana di Linguistica Computazionale, 2023, pp. 37–44.
[33] E. Öhman, M. Pàmies, K. Kajava, J. Tiedemann, XED: A multilingual dataset for sentiment analysis and emotion detection, in: D. Scott, N. Bel, C. Zong (Eds.), Proceedings of the 28th International Conference on Computational Linguistics, International Committee on Computational Linguistics, Barcelona, Spain (Online), 2020, pp. 6542–6552. URL: https://aclanthology.org/2020.coling-main.575. doi:10.18653/v1/2020.coling-main.575.
[34] A. Chiorrini, C. Diamantini, A. Mircoli, D. Potena, E. Storti, EmotionAlBERTo: Emotion Recognition of Italian Social Media Texts Through BERT, in: 2022 26th International Conference on Pattern Recognition (ICPR), IEEE, 2022, pp. 1706–1711.
[35] C. Diamantini, A. Mircoli, D. Potena, S. Vagnoni, et al., An experimental comparison of large language models for emotion recognition in italian tweets., in: ITADATA 2023. Italian Conference on Big Data and Data Science 2023, CEUR Workshop Proceedings, Napoli, Italy, 2023.
[36] D. Demszky, D. Movshovitz-Attias, J. Ko, A. Cowen, G. Nemade, S. Ravi, GoEmotions: A dataset of fine-grained emotions, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Online, 2020, pp. 4040–4054. URL: https://aclanthology.org/2020.acl-main.372. doi:10.18653/v1/2020.acl-main.372.