

Written Goodbyes: How Genre and Sociolinguistic Factors Influence the Content and Style of Suicide Notes

Lucia Busso^{1,*}, Claudia Roberta Combei^{2,*}

¹Aston Institute for Forensic Linguistics, Aston University, Birmingham (UK)

²Dipartimento di Studi Umanistici, Università di Pavia (Italy)

Abstract

The study analyses a novel corpus of 76 freely available English authentic suicide notes (SNs) (letters and social media posts), spanning from 1902 to 2023. By using NLP and corpus linguistics tool, this research aims at decoding patterns of content and style in SNs. In particular, we explore variation in linguistic features in SNs across sociolinguistic factors (age, gender, addressee, time period) and between text type – referred to as genre – (letters vs. online posts). To this end, we use topic models, subjectivity analysis, and sentiment and emotion analysis. Results highlight how both discourse and emotion expression, show differences depending on genre, gender, age group and time period. We suggest a more nuanced approach to personalized prevention and intervention strategies based on insights from computer-assisted linguistic analysis.

Keywords

suicide notes, topic modelling, sentiment and emotion analysis, subjectivity analysis

1. Introduction

This paper investigates the language of suicide notes, with the goal of uncovering patterns of discourse, topics, and emotional expression across various sociolinguistic factors and relationship dynamics, spanning over 100 years. A suicide note (SN) has been defined in the literature as "any available text by a suicide which was authored shortly before death" ([1]: 26).

The importance of a detailed analysis of suicide notes has been acknowledged in the scholarly debate ([2]). In fact, SNs have been widely studied in linguistics, sociology, and psychology starting with the publication in 1959 of Osgood and Walker's seminal work ([3]). Since then, the language of SNs has been investigated mainly through Genre Analysis ([4]), with some scholars working with corpus methods ([1, 5]). Lately, big corpora of SNs have been collected through the Web and used for computational analyses (*inter alia* [6, 7, 8]).

Research on SNs is naturally practical, being focused on suicide prevention ([9]), identification ([10]), and authenticity ([11]). For instance, the study by [6] uses classification algorithms to help mental health professionals

distinguish between genuine and elicited suicide notes. This – the authors claim – can in turn help developing a prediction strategy of repeated suicide attempts, as suicide notes offer valuable insights into specific personality states and mindsets. Similarly, [7] suggests that analysing SNs may contribute to assessing the risk of repeated suicide attempts.

Despite the area being well-researched, especially in forensic linguistics, current analyses of SNs present several shortcomings. Given the difficulty of accessing data, scholars have either used dubious source material (such as the letters published on the blog "The Holy Dark"), or have reused and reanalysed SNs written by famous people (such as Virginia Woolf and Kurt Cobain, e.g., [12, 13]). Moreover, there is no study to date – to the best of the authors' knowledge – that analyses of SNs using text type, which we refer to as genre, or sociolinguistic factors (such as gender, age, addressee, or time period) as covariates.

In the present paper, we set out to perform corpus and computational analyses on a novel dataset of authentic suicide notes. Specifically, we aim to explore whether and to what extent SNs style and content vary according to genre (letter vs. online post) and sociolinguistic factors (the victim's gender and age, as well as the addressee and time period of the SN). To this end, we employ Structural Topic Modelling ([14]) and keyword analysis, subjectivity analysis ([15]), and sentiment and emotion analysis ([16, 17]).

2. Data

Despite the presence in the literature of various datasets of suicide letters, none – to the authors' best knowledge –

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

*Corresponding authors

[†]Although both authors are equally responsible for the conceptualization and the contents of this article, for the purposes of Italian academia, we specify what follows: L. Busso authored Section 1, Section 2, and Section 3; C. R. Combei authored Section 4, Section 5, and Section 6.

✉ l.busso@aston.ac.uk (L. Busso); claudiaroberta.combei@unipv.it (C. R. Combei)

ORCID: <https://orcid.org/0000-0002-5665-771X> (L. Busso);

<https://orcid.org/0000-0003-1884-8205> (C. R. Combei)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



are freely available to other researchers. Furthermore, existing corpora are usually either very small, and hence not suitable for quantitative analysis, or too big, and hence not controlled for the parameters we are interested in analysing. Therefore, we decided to collect a new dataset of genuine suicide notes to fill this gap, and to make it available to researchers interested in the topic. Given the sensitivity of the topic, corpus files are available upon requests to the authors. Using the semi-automated software *Bootcat* ([18], we collected a corpus consisting of 76 suicide letters and social media posts¹. The SNs have the following characteristics:

- freely available on the open internet (i.e., not behind paywalls or log-in platforms)
- taken from reputable news websites to ensure authenticity (i.e. not taken from blogs or other non-official sources)
- only notes that were reproduced in full (i.e. not from extracts or quotes in other texts)

The resulting corpus contains 26,214 tokens, and includes texts from 1902 to 2023. Unavoidably, the distribution is skewed towards more recent texts (only 5 texts are from before 1950, and only 14 are from before 1990. The majority of the corpus (75%) includes SNs from 1990 to the present day). However, the corpus is balanced for textual type (genre), with 43 letters (51% of the tokens) and 33 social media posts (49% of the tokens). The SNs also cover a wide range of addressees, including messages directed to family, life-partners (including ex-partners), friends, the internet, or cases where the addressee is unspecified.

3. Topic Modelling and keywords analysis

3.1. Structural Topic Models

Topic Models (TM) are a family of unsupervised learning algorithms that cluster co-occurring words across documents into thematic nodes, or "topics" ([19]). These algorithms require a substantial human input, as the topics retrieved should be interpretable by the researcher assigning meaning to the patterns discovered ([20, 21]).

In this study we use Structural Topic Modelling ([14]), a type of TM that allows to model topics distribution as a function of document-level covariates in regression-like schemes. The STM analyses are performed in R[22]. We select a number of topics $K=3$ based on mathematical fit

¹We based our data retrieval on the sources provided by [7], and expanded on them through targeted Google searches. For privacy reasons, online posts were only collected if reported by newspaper articles, and were not retrieved on social media platforms themselves.

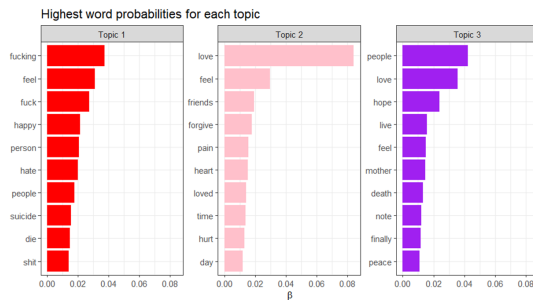


Figure 1: Top 10 word probabilities for each of the 3 topics

Table 1 regression-like analysis on differences in topical prevalence between genres

	estimate	t value	p value
Explanations	-0.04	-0.6	.6
Anguish	-0.27	-3.5	<.001
Connectedness	0.31	4.2	<.001

and ease of interpretation, and we model the effect of the "genre" covariate on topic content (i.e. lexical content used within topics) and prevalence (i.e. the frequency with which a topic is discussed).

Figure 1 shows the top 10 word probabilities for the 3 topics in the corpus. Following extensive concordance analysis to explore the keywords in context, the three topics have been labelled:

1. Topic 1: *Explanations*. This topic clusters words related to reasons, motives, and emotions associated with the act of suicide.
2. Topic 2: *Anguish*. This topic clusters words related to the intimate feelings of pain and hurt that accompany suicidal ideation.
3. Topic 3: *Connectedness*. This topic clusters words that refer to close connections to other people in the victim's life.

As mentioned above, we model the effect of genre (letter vs. online post) for topical content and prevalence. While we find no statistical differences ($p > .05$) for topical content, some interesting differences arise in topical prevalence, as can be seen in Table 1 and in Figure 2. Specifically, we observe that online posts discuss significantly less private feelings of anguish and pain (Topic 2) and significantly more interpersonal relationships (Topic 3).

3.2. Keyword Analysis

To explore the corpus further, beside the "black box" of the STM algorithm, we performed a keyword analysis.




Figure 2: Difference in topic prevalence between letters and online posts



Figure 3: Top 50 idiosyncratic keywords for Letters and Online Posts

Using SketchEngine ([23]), we extract keywords for both letters and social media posts using EnTenTen21 as reference corpus. To ensure that we only consider words that are used throughout the corpus, we discarded instances with a low ARF (average reduced frequency) score ([24]). Not surprisingly, many keywords are shared across the two subcorpora, reflecting "universal" themes of suicidal ideation such as apologies, goodbyes, and explanations. However, idiosyncratic keywords paint an interesting picture (see Figure 3), as online posts seem to display a lower prevalence of intimate feelings, and more polarized emotion words and swearwords.

4. Subjectivity analysis

Subjectivity analysis investigates what is generally labelled as a "private state", namely opinions, feelings, be-

liefs, speculations ([15]: 674), typically classifying a text on a scale ranging from high objectivity to high subjectivity.

Our paper uses this analysis because we see subjectivity as a relevant stylistic and content-related element, useful for understanding suicidal ideation. Although this is a preliminary study, we believe that findings from subjectivity, sentiment, and emotion analysis, supported by the exploration of psychosocial factors (not the object of this paper), could be useful for evaluating the risk of (repeated) suicide attempts. In particular, we expect that highly subjective texts may signal intense personal turmoil, which has, in fact, been reported as a potential risk factor for suicide ([25]).

This research uses the TextBlob library for Python that provides tools for various textual analyses, including subjectivity, as part of its sentiment analysis function². The tool uses a pattern analyzer and a pre-defined dictionary of word polarity and subjectivity. It also incorporates intensity, accounting for the impact of modifiers, which can increase or reduce the measured subjectivity score. Each SN is processed to extract its overall subjectivity score that ranges from 0 (i.e., highly objective) through 1 (i.e., highly subjective). To discuss the effect of genre and sociolinguistic factors on the subjectivity score, we present the results of statistical analyses conducted in R[22].

First of all, the mean subjectivity score at the corpus level ($M = 0.56$, $SD = 0.12$) indicates that SNs are characterized by a level of subjectivity that falls above the midpoint of the scale (0.50); there is, thus, a tendency toward greater subjectivity than objectivity. Interestingly, however, the mean subjectivity scores and their distributions are nearly identical between letters ($M = 0.56$, $SD = 0.13$) and social media posts ($M = 0.57$, $SD = 0.12$).

Next, based on Figure 4, SNs written from 1950-1969 seem to have the highest subjectivity score ($M = 0.72$, $SD = 0.15$). In contrast, the lowest subjectivity is found for SNs written from 1990-1999 ($M = 0.49$, $SD = 0.12$), followed by those from 1970-1989 ($M = 0.52$, $SD = 0.15$). SNs written before 1950 ($M = 0.56$, $SD = 0.11$), from 2000-2019 ($M = 0.56$, $SD = 0.11$), and from 2020-now ($M = 0.56$, $SD = 0.13$) have identical subjectivity scores.

The results displayed in Figure 5 indicate that subjectivity scores of SNs addressed to life-partners ($M = 0.61$, $SD = 0.06$) are the highest, followed by those addressed to family ($M = 0.60$, $SD = 0.09$). This suggests that SNs addressed to people with whom the victim has a close relationship are characterized by a deeper personal engagement and a more vivid linguistic expression than those addressed to the internet ($M = 0.56$, $SD = 0.12$), to friends ($M = 0.55$, $SD = 0.08$), and to other addressees (M

²The sentiment analysis score itself obtained from the TextBlob tool is not used in this study, as more advanced methods for investigating sentiment are preferred (see Section 5)




Figure 4: Subjectivity as a function of the year group

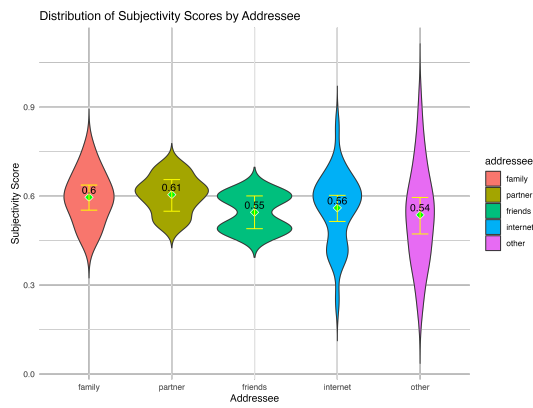


Figure 5: Subjectivity as a function of addressee

$= 0.54, SD = 0.16$). The standard deviations for most addressees (i.e., partner, family, friends) are relatively small, suggesting limited variation within these groups.

As regards the victim’s gender, the average subjectivity score for females ($M = 0.58, SD = 0.11$) is slightly higher than the score for males ($M = 0.54, SD = 0.14$), but the standard deviations point out that the ranges overlap to a large extent. Finally, no consistent tendency emerges from the distribution of subjectivity scores with respect to the victim’s age. In fact, there is substantial variation within each age group, meaning that the degree of subjectivity in SNs is influenced by other factors.

5. Sentiment and emotion analysis

In order to obtain a more fine-grained image of the emotional dimension of the SNs, and to complement the previously discussed findings on the topics and subjectivity of these texts, we also present and discuss the results of




Figure 6: Distribution of probabilities for positive, neutral, and negative sentiment

sentiment and emotion analysis. Sentiment analysis is defined as “the task of finding the opinions of authors about specific entities” ([26]: 82). Emotion analysis (also emotion classification), on the other hand, is often seen as a more refined version of sentiment analysis, since it deals with the identification of primary emotions in a text ([27]).

For this research we employ the latest version available (at the time of writing) of Twitter-roBERTa-base for sentiment analysis, a model trained on over 124 million tweets that is fine-tuned for this task with the TweetEval benchmark ([28, 29, 30]). For emotion classification, we use the Emotion English DistilRoBERTa-base model ([31]) to extract Ekman’s six basic emotions ([32]): anger, disgust, fear, joy, sadness, and surprise, along with a neutral class. The model is a fine-tuned version of DistilRoBERTa-base, trained on six balanced datasets, each containing 2,811 observations per emotion, for a total of almost 20,000 observations.

Our analysis reveals that the average probability of negative sentiment ($M = 0.61, SD = 0.31$) is roughly three times higher than the average probability of neutral ($M = 0.22, SD = 0.15$) and positive sentiment ($M = 0.17, SD = 0.28$). Then, the dominant sentiment in each SN is determined by identifying the highest probability among the three sentiment classes. We find that 73% of the SNs have negative sentiment as the highest probability, 17.1% positive sentiment, and 9.2% neutral sentiment. This trend is also supported by Figure 6 that shows the distribution of sentiment probabilities, confirming that most SNs have a higher likelihood of expressing negative sentiment. We interpret these results as a reflection of the emotional distress tied to both writing the suicide notes and the thoughts surrounding the act of suicide itself.

Some interesting tendencies are observed from the analysis of sentiment distribution across sociolinguistic factors and genre. First, Figure 7 illustrates a consistent difference between the two genres: online posts have a higher prevalence of negative sentiment (90.9%) compared to letters (60.5%).

Next, all SNs from 1970-1989 show negative sentiment as being dominant (100%). A high presence of negative




Figure 7: Sentiment as a function of genre




Figure 8: Sentiment as a function of gender

sentiment (88.5%) is also present in SNs written from 2020-now. Interestingly, SNs from 1990-1999 display a balanced sentiment distribution (50% negative and 50% positive), marking the only period in our corpus with such a high presence of positive sentiment. This situation could be due to the fact that the authors of these (very long) SNs are well-known celebrities (e.g., Kurt Cobain and OJ Simpson). Even if the letters were not intended for the general public, the idea these texts might eventually become public could have influenced the victims to transmit more positive messages.

Some patterns of sentiment distribution are traceable when considering the addressee of the SN. Positive sentiment is more common when the addressee is the victim's partner (40%) or family (35.7%). Contrarily, a very high percentage of negative sentiment is observed in SNs addressed to the general public on the internet (93.1%). Figure 8 shows that the negative sentiment is slightly more frequent in SNs written by female victims (72.7%) compared to male victims (68.2%). As for the victim's age, a distinct pattern is difficult to identify, but, negative sentiment is the most frequent (over 65%) in SNs written by teenagers (10s) and people in their 20s, 30s, 40s, and 60s.

Moving on to emotion analysis, the average probability of SNs conveying sadness ($M = 0.48$, $SD = 0.37$) is four times higher than the average probability of conveying anger ($M = 0.12$, $SD = 0.22$), fear ($M = 0.12$, $SD = 0.21$), and neutrality ($M = 0.12$, $SD = 0.18$). Sadness (53.9%) is, indeed, the dominant emotion in the corpus, followed by neutrality (13.2%), anger (11.8%), and fear (7.9%). This is determined by identifying the highest probability among the seven emotion classes for each individual SN.

We can pinpoint some interesting outcomes from the analysis of emotions across genres and sociolinguistic factors. As concerns genre, Figure 9 depicts an obvious difference between letters and online posts. On the one hand, sadness is more frequent in online posts (59.2%)




Figure 9: Emotions as a function of genre

compared to letters (40%). On the other hand, neutrality and joy, the only two non-negative emotions, are more frequent in letters (14.6% and 8.8%, respectively) than in online posts (9.5% and 3.2%, respectively).

The analysis reveals that sadness is the most prevalent emotion across all time periods. In particular, the presence of sadness exceeds 50% in SNs from 1970-1989 and from 2020-now. Then, the SNs written from 1970-1989 are also characterized by a definite presence of disgust (22.2%). In line with the sentiment analysis results, SNs from 1990-1999 contain the lowest presence of sadness (40.8%) and generally the lowest presence of negative emotions overall, compared to other periods. SNs written before 1950 display the highest presence of fear (17.3%) in the corpus, although sadness still remains the most prevalent emotion in this period.

From Figure 10, we can identify a clear disparity between the emotions transmitted by female and male victims. Sadness appears more frequently in SNs written by females (53.1%) compared to males (41.5%). Additionally, anger is more prevalent in SNs written by males (17.1%), ranking as their second most common emotion (after sadness).




Figure 10: Emotions as a function of gender




Figure 11: Emotions as a function of age group

Although Figure 11 illustrates a complex distribution of emotions across the age groups of the victims, some patterns still emerge. Sadness is the most common emotion in the SNs of all age groups except for those written by people in their 30s, where neutrality prevails (36.2%). Interestingly, teenagers express the lowest neutrality (3.4%) and the highest sadness (60.1%). Additionally, fear is prominent among SNs written by people over 70 years old (31.8%), making it the second most frequent emotion for this age group. Fear is also the second most common emotion for SNs written by teenagers (14.6%).

6. Conclusions

This mixed-methods study analysed the content and style of 76 SNs written over the course of a century, using genre, several sociolinguistic factors, and relationship dynamics as covariates. First of all, three main topics emerged from our corpus, that we labelled as *Explanations*, *Anguish*, and *Connectedness*. Looking at the differences in topical prevalence between the two text types, we observed that online posts displayed less private feel-

ings (e.g., anguish and pain) and greater polarized emotion words and swearwords.

Subjectivity analysis revealed that SNs tended to be more subjective than objective, irrespective of the genre. Some differences based on addressees were identified in the corpus; for example, SNs directed toward close relationships (i.e., life-partners and family) showed higher subjectivity scores, suggesting a more profound and personal style, compared to those directed toward the broader (internet) public.

As far as sentiment analysis is concerned, negative sentiment was dominant in the corpus (i.e. three times more frequent than neutral or positive sentiment), especially in online posts. Then, the analysis of emotions revealed that sadness was the main emotion in the corpus. This evident presence of sadness and negative sentiment reflects the complex emotional challenges and inner struggles that victims experienced at the time they wrote their SNs. Although sadness was the most common emotion in both letters and online posts, it occurred more frequently in the latter text type. Also, letters tended to convey more positive emotions (e.g., joy) more frequently than online posts. Finally, the analysis revealed that sadness was more common in the SNs written by female victims and by teenagers.

All in all, our results reveal that the content, discourse, and emotional expression in SNs vary as a function of genre, sociolinguistic factors, and relationship dynamics. These differences uncover the need of taking into account specific social, demographic, and cultural variables when designing and implementing suicide prevention and intervention strategies. In this sense, we believe that corpus-based and NLP research on SNs can contribute to the improvement of these personalized strategies.

Acknowledgments

The research presented in this paper was conducted while C. R. Combei benefited from support provided by the project "PON Ricerca e Innovazione 2014–2020 - Linea Innovazione (D.M. 1062/2021)".

References

- [1] J. J. Shapero, The language of suicide notes, Ph.D. thesis, University of Birmingham, 2011.
- [2] G. Tellari, C. Zanchi, Il suicidio di universitari nei media italiani: Uno studio basato su corpus, in: S. Matiola, M. Milicevic Petrovic (Eds.), CLUB – Working Papers in Linguistics, volume 8, AMS Acta AlmaDL, Bologna, 2024, pp. 1–20.
- [3] C. E. Osgood, E. G. Walker, Motivation and language behavior: a content analysis of suicide notes.,

- The Journal of Abnormal and Social Psychology 59 (1959) 58–67.
- [4] B. Samraj, J. M. Gawron, The suicide note as a genre: Implications for genre theory, *Journal of English for Academic Purposes* 19 (2015) 88–101.
- [5] A. E. Jaafar, H. A.-S. Jasim, A corpus-based stylistic analysis of online suicide notes retrieved from reddit, *Cogent Arts & Humanities* 9 (2022) 2047434.
- [6] J. Pestian, H. Nasrallah, P. Matykiewicz, A. Bennett, A. Leenaars, Suicide note classification using natural language processing: A content analysis, *Biomedical informatics insights* 3 (2010) BII–S4706.
- [7] S. Ghosh, A. Ekbal, P. Bhattacharyya, Cease, a corpus of emotion annotated suicide notes in english, in: *Proceedings of the twelfth international conference on language resources and evaluation*, 2020, pp. 1618–1626.
- [8] A. M. Schoene, A. Turner, G. R. D. Mel, N. Dethlefs, Hierarchical multiscale recurrent neural networks for detecting suicide notes, *IEEE Transactions on Affective Computing* 14 (2021) 153–164.
- [9] M. Chatterjee, P. Kumar, P. Samanta, D. Sarkar, Suicide ideation detection from online social media: A multi-modal feature based technique, *International Journal of Information Management Data Insights* 2 (2022) 100103.
- [10] T. Zhang, A. M. Schoene, S. Ananiadou, Automatic identification of suicide notes with a transformer-based deep learning model, *Internet interventions* 25 (2021) 100422.
- [11] M. Ioannou, A. Debowska, Genuine and simulated suicide notes: An analysis of content, *Forensic science international* 245 (2014) 151–160.
- [12] E. T. Sudjana, N. Fitri, Kurt cobain’s suicide note case: Forensic linguistic profiling analysis, *International Journal of Criminology and Sociological Theory* 6 (2013) 217–227.
- [13] N. Malini, V. Tan, Forensic linguistics analysis of virginia woolf’s suicide notes, *International Journal of Education* 9 (2016) 53–58.
- [14] M. E. Roberts, B. M. Stewart, D. Tingley, Stm: An R package for structural topic models, *Journal of statistical software* 91 (2019) 1–40.
- [15] A. Montoyo, P. Martínez-Barco, A. Balahur, Subjectivity and sentiment analysis: An overview of the current state of the area and envisaged developments, *Decision Support Systems* 53 (2012) 675–679.
- [16] L. Bing, *Sentiment Analysis: Mining Opinions, Sentiments, and Emotions*, Cambridge University Press, Cambridge, 2015.
- [17] C. Strapparava, R. Mihalcea, Learning to identify emotions in text, in: *Proceedings of the 2008 ACM Symposium on Applied Computing, SAC ’08*, Association for Computing Machinery, New York, NY, USA, 2008, pp. 1556–1560.
- [18] M. Baroni, S. Bernardini, Bootcat: Bootstrapping corpora and terms from the web., in: *Proceedings of the fourth international conference on language resources and evaluation*, Lisbon, Portugal, 26–28 May 2004, 2004, pp. 1313–1316.
- [19] D. M. Blei, Probabilistic topic models, *Communications of the ACM* 55 (2012) 77–84.
- [20] J. Chang, S. Gerrish, C. Wang, J. Boyd-Graber, D. Blei, Reading tea leaves: How humans interpret topic models, *Advances in neural information processing systems* 22 (2009) 288–296.
- [21] M. E. Roberts, B. M. Stewart, D. Tingley, C. Lucas, J. Leder-Luis, S. K. Gadarian, B. Albertson, D. G. Rand, Structural topic models for open-ended survey responses, *American journal of political science* 58 (2014) 1064–1082.
- [22] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2023. URL: <https://www.R-project.org/>.
- [23] A. Kilgarriff, V. Baisa, J. Bušta, M. Jakubíček, V. Kovář, J. Michelfeit, P. Rychlý, V. Suchomel, The sketch engine: ten years on, *Lexicography* 1 (2014) 7–36.
- [24] J. Hlaváčová, P. Rychlý, Dispersion of words in a language corpus, in: *Text, Speech and Dialogue: Second International Workshop, TSD’99 Plzen, Czech Republic, September 13–17, 1999 Proceedings 2*, Springer, 1999, pp. 321–324.
- [25] A. Schuck, R. Calati, S. Barzilay, S. Bloch-Elkouby, I. Galynker, Suicide crisis syndrome: A review of supporting evidence for a new suicide-specific diagnosis, *Behavioral Sciences and the Law* 37 (2019) 223–239.
- [26] R. Feldman, Techniques and applications for sentiment analysis, *Communications of the ACM* 56 (2013) 82–89.
- [27] C. R. Combei, A. Luporini, Sentiment and emotion analysis meet appraisal: A corpus study of tweets related to the COVID-19 pandemic, *Rassegna Italiana di Linguistica Applicata* 53 (2021) 115–136.
- [28] J. Camacho-Collados, K. Rezaee, T. Riahi, A. Ushio, D. Loureiro, D. Antypas, J. Boisson, L. Espinosa Anke, F. Liu, E. Martínez Cámara, TweetNLP: Cutting-edge natural language processing for social media, in: *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, Association for Computational Linguistics, Abu Dhabi, UAE, 2022, pp. 38–49.
- [29] D. Loureiro, F. Barbieri, L. Neves, L. Espinosa Anke, J. Camacho-collados, TimeLMs: Diachronic language models from Twitter, in: *Proceedings of the 60th Annual Meeting of the Association for Com-*

- putational Linguistics: System Demonstrations, Association for Computational Linguistics, Dublin, Ireland, 2022, pp. 251–260.
- [30] F. Barbieri, J. Camacho-Collados, L. Espinosa Anke, L. Neves, TweetEval: Unified benchmark and comparative evaluation for tweet classification, in: T. Cohn, Y. He, Y. Liu (Eds.), Findings of the Association for Computational Linguistics: EMNLP 2020, Association for Computational Linguistics, Online, 2020, pp. 1644–1650.
- [31] J. Hartmann, Emotion English DistilRoBERTa-base, <https://huggingface.co/j-hartmann/emotion-english-distilroberta-base/>, 2022.
- [32] P. Ekman, Basic emotions, in: T. Dalgleish, T. Power (Eds.), *The Handbook of Cognition and Emotion*, John Wiley & Sons, Ltd, Sussex, U.K., 1999, pp. 45–60.