

Emotionally Aware Moderation: The Potential of Emotion Monitoring in Shaping Healthier Social Media Conversations

XIAOTIAN SU, ETH Zurich, Switzerland

NAIM ZIERAU, University of St.Gallen, Switzerland

SOOMIN KIM, Seoul National University, South Korea

APRIL YI WANG, ETH Zurich, Switzerland

THIEMO WAMBSGANSS, Bern University of Applied Sciences, Switzerland

Social media platforms increasingly employ proactive moderation techniques, such as detecting and curbing toxic and uncivil comments, to prevent the spread of harmful content. Despite these efforts, such approaches are often criticized for creating a climate of censorship and failing to address the underlying causes of uncivil behavior. Our work makes both theoretical and practical contributions by proposing and evaluating two types of emotion monitoring dashboards to enhance users' emotional awareness and mitigate hate speech. In a study involving 211 participants, we evaluate the effects of the two mechanisms on user commenting behavior and emotional experiences. The results reveal that these interventions effectively increase users' awareness of their emotional states and reduce hate speech. However, our findings also indicate potential unintended effects, including increased expression of negative emotions (Angry, Fear, and Sad) when discussing sensitive issues. These insights provide a basis for further research on integrating proactive emotion regulation tools into social media platforms to foster healthier digital interactions.

CCS Concepts: • **Computing methodologies** → **Natural language processing**; • **Human-centered computing** → **Social media**; **User studies**; *Empirical studies in HCI*.

Additional Key Words and Phrases: social media, proactive intervention, emotion monitoring

ACM Reference Format:

Xiaotian Su, Naim Zierau, Soomin Kim, April Yi Wang, and Thiemo Wambsganss. 2025. Emotionally Aware Moderation: The Potential of Emotion Monitoring in Shaping Healthier Social Media Conversations. *Proc. ACM Hum.-Comput. Interact.* 9, 7, Article CSCW291 (November 2025), 28 pages. <https://doi.org/10.1145/3757472>

1 Introduction

Social media provides unprecedented freedom for communication and self-expression. However, this openness is frequently exploited to disseminate violent messages, derogatory comments, and hateful speech [15]. This phenomenon, known as online hate speech, is defined as a “direct and serious attack on any protected category of people based on their race, ethnicity, national origin, religion, sex, gender, sexual orientation, disability or disease” [18, 49].

Hate speech is increasing at an alarming rate across social media platforms [15, 51, 68]. This increase not only leads to a deteriorated environment on the internet but can also translate into violence outside the virtual world [74]. For instance, a study reported that between 15% and 35% of young people have been victims of cyberbullying [27]. Following these incidents, two-thirds of

Authors' Contact Information: [Xiaotian Su](mailto:xiaotian.su@inf.ethz.ch), xiaotian.su@inf.ethz.ch, ETH Zurich, Zurich, Switzerland; [Naim Zierau](mailto:naim.zierau@unisg.ch), naim.zierau@unisg.ch, University of St.Gallen, St. Gallen, Switzerland; [Soomin Kim](mailto:smsoominkim@gmail.com), smsoominkim@gmail.com, Seoul National University, Seoul, South Korea; [April Yi Wang](mailto:april.wang@inf.ethz.ch), april.wang@inf.ethz.ch, ETH Zurich, Zurich, Switzerland; [Thiemo Wambsganss](mailto:thiemo.wambsganss@bfh.ch), thiemo.wambsganss@bfh.ch, Bern University of Applied Sciences, Bern, Switzerland.



This work is licensed under a [Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-nc-sa/4.0/).

© 2025 Copyright held by the owner/author(s).

ACM 2573-0142/2025/11-ARTCSCW291

<https://doi.org/10.1145/3757472>

teens felt a decline in self-esteem, nearly a third reported affected friendships, and 13% indicated a negative impact on their physical health [62]. These numbers illustrate the devastating effect of harassing online messages on many people's well-being and the social fabric overall, highlighting the urgency to foster more inclusive and constructive dialogue on social media platforms.

Most existing interventions for addressing harmful content on social media adopt a reactive approach, focusing on removing offensive content after it has been posted. This is typically achieved through either automated systems [8, 14] or manual moderation [9, 11]. However, reactive methods are criticized for both their potential to suppress free expression [25] and for their delayed intervention, addressing harmful content only after it appears online [32]. Furthermore, such methods often lack the nuance required to address the underlying emotional dynamics of online interactions. Uncivil behavior often originates from regular users caught in highly charged or tense situations, exacerbated by the social and technological affordances of social media platforms [10].

In contrast, interventions like those designed by Peng et al. [55] aim to proactively regulate group emotions, following theories of emotion regulation. However, these approaches still followed a reactive paradigm, intervening only after the content was posted. There are some proactive interventions designed to promote constructive dialogue on social media, encouraging users to reflect on their messages before posting (e.g., [59, 63, 75]). While promising, these methods are generally static and applied uniformly, failing to account for the diverse communication styles and emotional responses of individual users [30]. Recent studies by Wright et al. [85] and Chang et al. [10] have explored personalized, real-time support to reduce toxic communication. However, these approaches rely solely on machine learning to analyze text, overlooking users' emotional states.

In contrast to these existing approaches, our emotion regulation intervention follows a proactive, personalized approach by designing two emotion intervention mechanisms and developing a specialized dashboard. Specifically, we aim to provide insight into the emotional tone of social media posts and assist users with emotion regulation before they post. This dashboard distinguishes between two types of emotion monitoring interventions: self monitoring and peer monitoring. Self monitoring enhances self-awareness and emotional regulation by reflecting the user's own emotions, while peer monitoring enhances peer-awareness and social understanding by reflecting the emotions of others. This study evaluates the benefits and potential pitfalls of each monitoring type, focusing on their impact on users' emotional awareness, social interactions, and the reduction of hate speech. Specifically, we aim to answer the following research questions:

- (1) How does real-time emotion *peer monitoring* influence users' messages in online discussions compared to environments without such monitoring?
- (2) How does real-time emotion *self monitoring* influence users' messages in online discussions compared to environments without such monitoring?

Within the context of our main study, we examined whether our two key interventions—providing feedback on (1) the peer's and (2) users' emotional sentiments—can impact user commenting behaviors and emotional experiences using both quantitative and qualitative data. While the effects noted in our main study were modest and limited by its scale, they, combined with our qualitative insights, indicate enhancing users' emotional awareness could influence social media interactions in various ways. This approach has the potential to contribute positively to existing strategies for promoting healthier, more constructive interactions, but it may also have unforeseen negative impacts on users' emotional experiences as our study shows. This research lays a foundation for future explorations, offering specific avenues for integrating natural language processing-based user-centered tools into social media platforms to make social media discussions more positive and productive. In this research, we present three contributions:

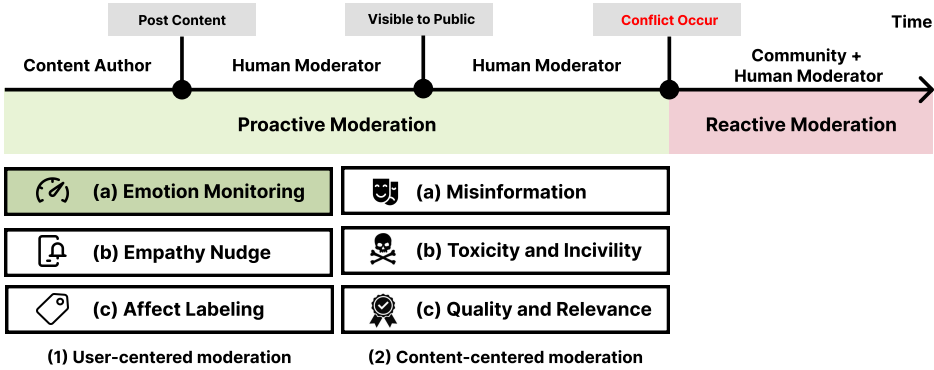


Fig. 1. Our work (Emotion Monitoring) builds on existing text-based moderation practices on social media along the proactive-reactive spectrum. Our emotion monitoring dashboard aims to enhance users' emotional awareness and help them construct more neutral messages.

- We introduce a mechanism designed to enhance users' emotional awareness and reduce toxic speech in social media conversations.
- We develop a pair of user interventions, manifested as an emotion monitoring dashboard, which operationalizes this mechanism within social media platforms.
- We provide evidence from a comprehensive user study with a total of 211 participants to investigate the effect of our intervention. The studies offer preliminary findings on the multi-faceted impact of our intervention, demonstrating their influence on emotional awareness and users' commenting behavior.

2 Related Work

Our work builds on the extensive research on online discussion moderation. In this section, we first survey the types of existing content moderation practices along the reactive-proactive spectrum, characterizing their effectiveness and costs.

2.1 Content Moderation Strategies on Social Media

The escalating spread of harmful material on digital platforms presents a significant societal challenge, encompassing various forms like hate speech, offensive language, bullying and harassment, misinformation, spam, violence, graphic content, sexual abuse, self-harm, and many others. Researchers in the Human-Computer Interaction (HCI) and Computer-Supported Collaborative Work (CSCW) fields have a growing interest in content moderation to foster more positive and constructive online discourse [43, 73, 79]. An emerging line of work proposes distinguishing different moderation practices based on the timing of the intervention relative to the objectionable action.

As shown in Figure 1, proactive moderation is designed to prevent such behavior from occurring in the first place [10, 61], either remind the content author to reflect on their texts before posting or detect and filter out the inappropriate content by human moderators after users posting. Another approach is reactive moderation [65], where people respond to conflicts or negative signals after they occur. This approach utilizes a combination of community involvement and human moderation to address problematic content. The community plays a role in identifying and flagging problematic content through collaborative efforts, such as users on Quora downvoting an unfriendly comment. At the same time, human moderators can intervene and take appropriate action against both the

sender and recipients of the content. For example, a moderator may remove the unfriendly comment and hide it from the user who received it. Various tools have been developed to support moderators' tasks across different platforms [9, 11].

Reactive moderation has inherent weaknesses while widely used and studied [9, 11, 29]. Firstly, it involves taking action against uncivil content that has already been posted, allowing the offending content to be seen and potentially spread before moderators can intervene [32]. This can harm platforms and their users and can also adversely affect moderators who must regularly view and take action on disturbing content, threatening their emotional well-being [38, 47] and even their personal safety [84]. These harms underscore the need to explore alternative moderation approaches that can mitigate such human suffering.

Proactive approaches to moderation aim to discourage undesirable actions or to encourage prosocial behavior and productive conversations [64]. Some studies explored indirect methods, like prompting users with tasks before commenting to prime positive engagement (e.g., completing a captcha [63]). For more direct interventions, Wright et al. [85] reports the toxicity score of users' textual input, providing immediate feedback on potentially harmful content. This is argued as a more effective way to protect online communities from harm because it reduces the amount of uncivil content that gets created in the first place [32].

2.2 Algorithmic Approaches for Proactive Moderation

Many algorithmic approaches have been developed for proactive content moderation. These methods aim to validate the authenticity and quality of public content, addressing critical issues such as misinformation and toxicity (Figure 1.2). For instance, Kim et al. [33] developed a scalable online algorithm to reduce the spread of fake news and misinformation. Risch and Krestel [58] presented various deep learning approaches for detecting toxic comments while Wani et al. [82] proposed a methodology to identify toxic counterfeit news articles across social media platforms. Complementing these efforts, Bian et al. [6] developed a semi-supervised coupled mutual reinforcement framework that simultaneously evaluates content quality and user reputation.

Despite the effectiveness of proactive approaches in detecting toxicity [85] and incivility [10] in certain contexts, these methods lack personalized guidance for user commenting behavior [30]. And they can sometimes exacerbate tensions or make users feel unfairly judged or censored [25]. Such approaches often lack the nuance required to address the underlying emotional dynamics of online interactions, resulting in a one-size-fits-all strategy that overlooks individual differences in communication styles and emotional responses [30]. To address these limitations, we propose integrating emotional awareness, a critical factor in fostering prosocial behavior [37, 45], into content moderation mechanisms. Emotional awareness encompasses both the ability to recognize one's own emotions and the capacity to understand them in context [78]. Extensive research underscores the role of emotional capacities, such as emotion regulation, in enabling individuals to empathize with others and respond to suffering without becoming overwhelmed [23, 26]. Building on this understanding, we advocate for a user-centered moderation approach (Figure 1.1a) that prioritizes users' emotional dynamics over traditional content-centered moderation strategies (Figure 1.2).

2.3 Text-based Emotion Regulation on Social Media

The ability to regulate emotions is essential across nearly every domain of life [24]. There is a growing interest in HCI to envision, design, and evaluate technology-enabled interventions that support users' emotion regulation [70, 71]. Recognizing and managing emotions is particularly challenging in online text-based communication [35]. The lack of nonverbal cues in online communication deteriorates the ability to control emotions and empathize with others [80]. To address these

challenges, researchers have investigated various user-centered approaches (Fig 1.1) to encourage emotion regulation and foster emotional awareness in online discussions. For example, [59] explored designs to support emotion regulation by encouraging reflective behaviors, while [75] explored how different types of empathy nudges can promote prosocial behaviors (Figure 1.1b). Saldías and Picard [60] reported that users produced a significantly higher percentage of neutral content when had access to emotional statistics about their posts. L. de Carvalho et al. [39] explored how labeling emotions in user interfaces for online news articles can assist in emotion regulation (Figure 1.1c). Collectively, these findings highlight the promise of integrating emotional awareness into the design of online platforms as a means of fostering healthier, more supportive digital communities. However, current approaches are predominantly static and do not adapt to the dynamic nature of online interactions or the unique needs of individual users. To effectively enhance interpersonal communication and emotional well-being on social media, it is crucial to implement adaptive emotion regulation strategies that can swiftly respond to the diverse emotional expressions and communication styles of users [70]. To address this challenge, we propose the development of an emotion monitoring dashboard—a novel tool designed to boost emotional awareness and minimize offensive comments by providing users with real-time feedback on the emotional tone of their own and others' messages.

Although emotion regulation is a well-established concept in HCI, there is limited research on emotion regulation interventions in online discussions, with the GremoBot project being a notable exception [55]. We build on the framework proposed by Slovak et al. [70] to understand this work's unique contributions regarding (1) the target of emotion regulation initiatives, (2) their delivery mechanisms, and (3) their design components. First, GremoBot aimed to regulate group behavior in work discussion settings, whereas our study focuses on providing emotion feedback to individual users in social media discussions. Second, while both studies employ a didactic approach with a recommender system to facilitate emotional awareness, GremoBot provides feedback reactively—after a prolonged pause in discussion or when a certain level of negativity is detected. In contrast, our study adopts a proactive approach, offering feedback before a comment is posted. Third, GremoBot delivers analyses and recommendations at the group level, with all participants receiving the same suggestions. In our study, we provide personalized analyses and recommendations tailored to individual users. Taken together, our work marks a significant shift from reactive group-level correction to proactive, preemptive individual-level support. By enabling users to regulate their emotions before posting potentially toxic content, we offer personalized feedback and actionable recommendations. By addressing the emotional and social complexities of online interactions, our study aims to go beyond merely moderating behavior. Instead, we seek to foster genuine understanding and connection among users, paving the way for healthier and more empathetic digital environments. This approach allows us to investigate the timing and effectiveness of these strategies in real-time, offering practical design guidance for emotion regulation tools. Moreover, while the GremoBot project featured an early-stage exploration of an innovative design space with 27 participants, our study advances the field by conducting a large-scale experimental evaluation involving 211 participants. This responds to recent calls for large-scale experimental studies and randomized controlled trials to provide causal evidence on the effectiveness of technology-facilitated emotion regulation interventions [70].

3 Overview of The Methodology

As shown in Table 1, our research encompassed two preliminary studies (Section 4) and one comprehensive main study (Section 5) to rigorously test our hypotheses. In the first pre-study, we explored various topics to select an emotionally engaging subject as the foundation. The second pre-study then aimed to confirm our premise that users' level of emotional awareness can be raised

Study Group	N	Gender			Age		Highest Education		
		Male	Female	Other	Mean	Std.	< BSc	>= BSc	
Pre 1	-	41	46.3%	48.8%	4.9%	36.9	12.6	34.1%	65.9%
Pre 2	-	67	52.4%	40.2%	7.3%	28.5	11.5	44.5%	55.5%
Main	CG	39	20.5%	76.9%	2.6%	28.3	11.3	51.3%	48.7%
	TG1	A1~A28	28.6%	71.4%	0.0%	27.9	7.5	50.0%	50.0%
	TG2	B1~B36	36.1%	61.1%	2.8%	26.1	6.1	47.2%	52.8%

Table 1. Demographics data of participants in two pre-studies (Pre 1: topic involvement in Section 4.1), Pre 2: emotional effect in Section 4.2) and the main study in Section 5. Participants are denoted as A1 to A28 in TG1 and denoted as B1 to B36 in TG2.

in emotionally charged social media interactions. Culminating in the main study, we applied these interventions, specifically emotion peer and self monitoring tools, in a simulated online discussion on abortion. By uncovering qualitative and quantitative insights into participant engagement with and reactions to the emotion monitoring feature, we provide a comprehensive evaluation of our hypotheses in a controlled, yet realistic, social media setting.

4 Pre-studies

Before our main study, we conducted two pre-studies focusing on the topic involvement and emotional awareness. The topic involvement study aims to identify a topic that not only resonates emotionally with social media users but also frequently encompasses hate speech or toxic comments. The emotional awareness study aims to provide initial insights into the potential of increased emotional awareness as a strategy to mitigate toxic communication. The details on the demographics of participants can be found in Table 1.

4.1 Topic involvement

Measurement	Topic	Abortion	Television	Vegetarian
Self-Assessment: Negative Emotion		2.6	2.0	1.9
Self-Assessment: Positive Emotion		1.6	1.7	2.0
Peer-Assessment: Negative Emotion		3.3	3.1	2.0
Peer-Assessment: Positive Emotion		1.5	1.5	2.2
Outcome-Relevant Involvement		5.0	4.2	4.4
Income-Relevant Involvement		4.7	4.2	4.6

Table 2. Comparison of the mean value for six measurements regarding three different topics (abortion, television, and vegetarian). The first four items on emotion assessments were rated on a 1 - 5 Likert Scale (1: not at all, 5: extremely). The last two items on topic involvement were rated on a 1 - 7 Likert Scale (1: low, 4: neutral, 7: high).

This study aims to identify a topic that not only resonates emotionally with social media users but also frequently encompasses hate speech or toxic comments. Research has examined controversial

topics on social media around policy change and activism, such as abortion [88], climate change [66], and racial inequality [13]. In collaboration with a senior researcher in Cognitive Science, we curated three topics known to provoke toxic communication on social media: abortion, the role of television, and vegetarianism. To determine the most emotionally engaging topic, we enlisted 41 participants from Prolific ($M(\text{age}) = 36.9$, $SD(\text{age}) = 12.6$; 46.3% male, 48.8% female, 4.9% other). Each participant engaged in separate social media discussions on these three topics, presented in randomized order. Each discussion scenario consisted of a newspaper article and a user comment, which they were asked to write a response to. After each scenario, they rated the emotions they perceived in both their responses and those of their peers. Emotional perceptions were assessed across four emotions (two negative: anger, anxiety; and two positive: happiness, hopefulness) using a 5-point Likert scale [41]. Additionally, we measured topic involvement with four items using an established scale [40]. As outlined in Table 2, the findings indicate that the topic of abortion resonated most strongly with our participants, eliciting the highest level of engagement. Consequently, we selected abortion as the central theme for our main study. Our rationale was that a topic eliciting such strong engagement would provoke more profound emotional reactions in our controlled laboratory setting. While we acknowledge that using a different topic may influence the specific emotional responses elicited, choosing abortion as the focal topic provided a robust context for investigating interaction effects due to its known ability to generate intense engagement. Moreover, abortion is not only a contentious issue in the US [2] but also globally [52], making it relevant and generalizable for studying social media interactions. This, we anticipated, would create a more authentic scenario, enhancing the external validity of our data analysis.

4.2 Emotional Effect

4.2.1 Procedure. We recruited 67 participants ($M(\text{age}) = 28.3$, $SD(\text{age}) = 11.5$; 52.4% male, 40.2% female, 7.3% other) from Prolific. Adhering to the result from the topic involvement experiment, we narrowed our focus to abortion as the sole topic of discussion. Following their posts, participants were prompted to evaluate the emotions in both their own and their peers' comments. This emotion rating task was intended as a stimulus to increase participants' emotional awareness and encourage more positive commenting behaviors on social media. Finally, we asked participants to revise their initial comments to a format they deemed appropriate for a social discussion.

4.2.2 Results. Only messages exceeding 30 words and those participants that took more than 5 minutes to finish the survey were selected for the analysis. We used the `text2emotion`¹ package and a fine-tuned Google T5² model to investigate the emotions. The `text2emotion` is a Python package that detects emotion-related words and categorizes them into five distinct emotions: Angry, Fear, Sad, Surprise, and Happy. The package outputs a dictionary for each message, with keys as emotion types and values as emotion scores in percentages, indicating the dominant emotion. Detailed description of the Google T5 model see Appendix B.

As we intended, our analysis revealed reductions in scores of Fear and Sad, and an increase in Surprise




Fig. 2. Comparison of the mean value of emotion scores between original comments and rewritten texts detected by the `text2emotion` package.

¹<https://github.com/aman2656/text2emotion-library>

²<https://huggingface.co/mrm8488/t5-base-finetuned-emotion>

and Happy (Figure 2). The results of a Shapiro-Wilk test [67] showed that our data was not normally distributed. Therefore, we performed the non-parametric Mann-Whitney U (MWU) test between the emotion scores on the original and rewritten comments. There was a significant decrease in Sad according to the MWU test ($p = .028$, $U = 2726.5$). Appendix B shows more details about the Google T5 model. The figure shows that after rewriting, the messages classified as Angry, Fear, and Sad each decreased by about five, while those conveying Happy increased from just over 20 to 35. Combining both results, we could conclude that participants tend to rephrase their comments with balanced and positive language after being prompted to rate their own and peer’s emotions. This supports the idea that fostering emotional awareness leads to more positive, constructive dialogue, especially on sensitive topics like abortion.

5 Main study

The two pre-studies help us identify (1) topics that can be controversial and emotional, and (2) the potential of increased emotional awareness as a strategy to mitigate toxic communication. Based on the results from the pre-studies, we conducted a controlled experiment with 103 participants posting comments after viewing a screenshot of a social media discussion on a controversial topic. Participants were randomly assigned to one of three conditions (as shown in Figure 3): the control group (CG), which did not receive any emotion monitoring; treatment group 1 (TG1), which received emotional analysis of a peer’s comment; and treatment group 2 (TG2), which received emotional analysis of their own written comment.

Our goal was to gain empirical insights into how the two types of emotional monitoring influence users’ communication patterns in social media discourse. Consequently, our analysis focused on comparing each treatment condition with the control condition (CG vs. TG1, CG vs. TG2) rather than comparing all three conditions against each other. Notably, these two mechanisms of emotional monitoring may coexist and complement each other when integrated into social media platforms.

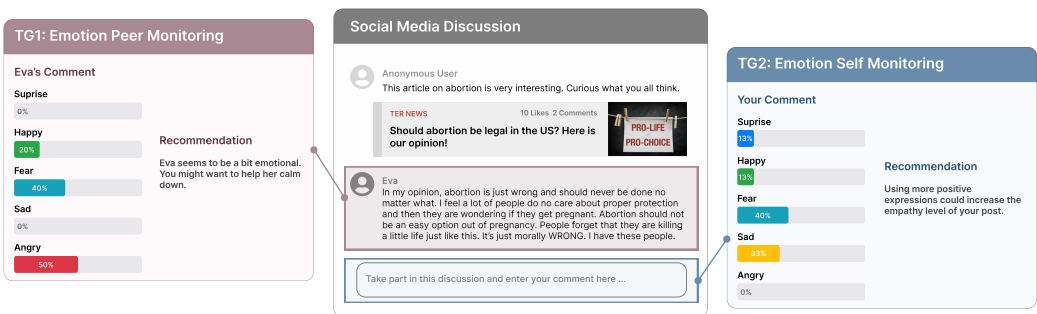


Fig. 3. Overview of experimental setup in the main study: participants of the control group (CG) did not receive any analysis. Participants in treatment group 1 (TG1) received emotional analysis of a given peer’s comment and participants in treatment group 2 (TG2) received emotional analysis of their own written comment. Both treatment groups receive personalized recommendations based on the emotional analysis. The interface of the emotion monitoring system is provided in Appendix C.

5.1 Participants and Task

A total of 161 participants were recruited from the Prolific³ platform (M (age) = 27.4, SD (age) = 8.7; 28.2% male; 69.9% female; 1.9% other), as shown in Table 1. Our objective was to obtain a representative subset of social media users, with a focus on younger individuals, while keeping the study open to participants of all ages. Our study involves the following three phases:

Phase I. We informed participants that they would engage in an online study regarding social media discussion and that their task was to write a comment. Then, we asked participants to fill out the Positive and Negative Affect Schedule (PANAS) scale [83], openness to new experiences, and moral foundations. The PANAS scale is a questionnaire assessing positive and negative affect, with scores measuring the level of agreement with each of the listed positive (active, attentive, determined, inspired) and negative (afraid, alert, hostile, nervous, upset) words on a 7 point Likert scale (1: totally disagree to 7: totally agree). This supported us to check people's initial emotional affect as a control mechanism.

Phase II. Mirroring the experimental paradigm of pre-studies, users were presented with a social media discussion centered on the topic of abortion. In this discussion, an initial comment is made by a user named "Eva", who states that "In my opinion, abortion is just wrong and should never be done no matter what. I feel a lot of people do not care about proper protection and then they are wondering if they get pregnant. Abortion should not be an easy option out of pregnancy. People forget that they are killing a little life just like this. It's just morally WRONG. I have these people." Then, participants are prompted to engage in the dialogue by responding to Eva's perspective. Upon drafting their response, they can press the button either "Emotion Peer Monitoring" (TG1) or "Emotion Self Monitoring" (TG2) to receive emotional feedback on the message they wrote. The subsequent dashboard provides an emotion analysis, showcasing a breakdown of five emotions in percentages on the left. For TG1, it demonstrates an emotional analysis of Eva's comment and TG2 would see an emotional analysis of their own comment. On the right, TG1 is encouraged to address Eva's message to make her less emotional and TG2 is offered personalized recommendations, guiding them in crafting more balanced and neutral comments. Appendix C displays the complete user interface of our developed tool.

Phase III. After posting comments, participants were asked to evaluate the peer's emotions according to the post and their own feelings about the previous discussion with several items. The answers were given on a 5-point Likert scale (1: not at all, 5: extremely) about five emotions (Angry, Fear, Sad, Surprise, and Happy). Next, we controlled for a battery of psychological mechanisms that could potentially interfere with our results comprising of perceptions of topic involvement, communication efficiency, communication fairness, communication effectiveness, and elaboration. Finally, to receive fine-grained information on user experience with the emotion monitoring dashboard, we included three open-ended survey questions about their likes, improvement areas, and additional ideas in social media interactions.

5.2 The Emotion Monitoring Dashboard

We designed an emotion monitoring dashboard for presenting the peer and self monitoring feedback. To inform the design of the emotion dashboard, we drew on an extensive review of existing emotion monitoring tools [17, 19, 20] and literature on emotion visualization for user interpretation [31, 77, 89]. In addition, we organized a design workshop, guided by the principles of research-through-design [90], aiming at crafting a user interface that effectively supports emotion monitoring in social media interactions.

³<https://www.prolific.com/>

5.2.1 Design Workshop for the Emotional Dashboard. We recruited 15 participants from diverse disciplines using mailing lists and leveraging internal university networks. There were seven females and eight males ($M(\text{age}) = 28.7$, $SD(\text{age}) = 2.4$), mainly HCI and UX researchers and practitioners (more demographic details see Appendix A). The workshop spanned approximately 90 minutes and unfolded in three phases: (1) *Identifying requirements*: We established the dashboard's primary purpose, focusing on the need for effective emotion monitoring in social media discussions. (2) *Detailing requirements*: We refined our initial requirements and deliberated on specific components of the user interface. (3) *Implementing requirements*: Our choice of bar graphs to represent sentiment data, with five distinct numerical values, was driven by clarity and immediate user comprehension. Additionally, we incorporated explanatory messages in personalized recommendations to boost user acceptance and understanding, promoting empathetic and reflective content engagement [30].

5.2.2 The Design of the Dashboard. The outcome of the workshop was a user-centric emotion monitoring dashboard, designed to be intuitive and insightful for users engaging in social media discussions. The bar graph format for representing sentiment data was chosen for its clarity and immediate comprehension. The inclusion of explanatory messages was a strategic decision aimed at fostering a more empathetic and reflective interaction with the content. This comprehensive approach ensured that the dashboard not only met the identified user requirements but also aligned with the theoretical and practical aspects of emotion monitoring in digital communications.

5.3 Data Analysis

To examine the impact of emotion peer monitoring and emotion self monitoring, we conducted both quantitative and qualitative analyses. Our measurements include the following.

5.3.1 Data Quality Check. To ensure high data quality, we only kept participants who passed the attention and manipulation checks. For the manipulation check, we asked participants to recall whether they received additional information when contributing to the discussion and what kind of information they received (e.g. on the emotionality of Eva's comment or their own comment). Those who failed to choose the right answer were excluded from the final analysis ($N = 21$ for CG, $N = 22$ for TG1, $N = 15$ for TG2). Furthermore, we again only considered participants when their messages exceeded 30 words and when they took more than 5 minutes to finish the survey. This resulted in a final valid sample of 103 participants ($N = 39$ for CG, $N = 28$ for TG1, $N = 36$ for TG2), which we used to report our main findings. Finally, we verified the PANAS scores for the chosen users, ensuring they had comparable emotional levels at the beginning of the study and that there were no excessively emotional participants.

5.3.2 Measured Emotions and Hate Speech in Users' Comments. To analyze emotions in the written messages, we applied the `text2emotion` package and Linguistic Inquiry and Word Count (LIWC) [56] with the newest dictionary⁴ to examine the emotional states in the written message. LIWC is a tool used for text analysis, particularly in psychological and linguistic research. It analyzes written text to identify emotional, cognitive, and structural components, categorizing words into various psychological and linguistic categories. The results of measured emotions in the user comments are presented in Table 3, Figure 5 and 8. Furthermore, we measured hate speech using the finetuned Twitter-roBERTa-base model [5]⁵. This is a roBERTa-base model trained on around 58M tweets and finetuned for hate speech detection with the TweetEval benchmark. This model is specialized to detect hate speech against women and immigrants. Given a message, it classifies the content by determining the percentage of hate speech. The result is presented in Figure 4.

⁴<https://www.liwc.app/>

⁵<https://huggingface.co/cardiffnlp/twitter-roberta-base-hate>

5.3.3 Linguistic Statistics in Users' Comments. To examine the linguistic profile of happy and sad expressers, the written messages were analyzed using LIWC. It analyzes messages on a word-by-word basis and compares words against a dictionary of words divided into different linguistic dimensions, including pronouns, affect terms, cognition terms, social and communicative processes. Its psychometric properties and external validity have been established in many studies, and have been used to examine the relationship between language and emotion. For the current study, only the variables relevant to our investigation were included in the analysis. The results of the linguistic profile are presented in Table 4, and 5.

5.3.4 Self-Reported Emotions. Additionally, we analyzed user-reported emotions regarding both their peers and themselves, focusing on how users perceived their peers' emotions as well as their own emotional experiences. The results are shown in Figure 6, 7, 9, and 10.

5.3.5 Qualitative Analysis on the Perceptions of Emotional Monitoring. To probe into the perceptions of emotional monitoring, we conducted a qualitative analysis for the final open-ended questions. Based on qualitative analysis guidelines [50], we prioritized the discovery of common themes over strict inter-rater reliability, acknowledging that different coders might interpret meanings in varied ways [50]. The first author and another researcher conducted reflexive thematic analysis [7] to code and develop themes. This process started with descriptive coding of user evaluations of emotion monitoring, followed by categorizing these codes into themes. Both researchers independently coded the entire dataset and then cross-checked each other's work. A third researcher validated and finalized these themes.

6 Results

We organized the findings for the two research questions based on four main themes: (1) participants' perceptions of the monitoring dashboard, (2) measured emotions and hate speech in user comments, (3) the reported emotions of their peers and their own, and (4) the difference in the linguistic usage of treatment groups compared to the control group.






Group						Emotion		Tone	
	Angry	Fear	Sad	Surprise	Happy	Negative	Positive	Negative	Positive
CG	6.6%	8.0%	14.1%	25.8%	45.7%	0.18%	0.64%	4.92%	1.81%
TG1	7.4%	8.6%	22.1%	21.4%	40.4%	0.21%	0.69%	4.81%	1.38%
TG2	10.9%	7.9%	20.2%	21.4%	39.3%	0.54%	0.43%	4.96%	2.30%

Table 3. Statistics summary of the percentages of five measured emotions using the text2emotion package, and the overall tone and emotion measurement using LIWC-22. The numbers in the last two columns represent the percentage of words in the comments that fall into the corresponding category. The positive and negative emotions are restricted to words that strongly imply corresponding emotions. The positive and negative tone dictionaries include words related to positive and negative emotions (e.g., happy, sad, angry) and also words related to those emotions (e.g., birthday, beautiful, kill, funeral) [53].

6.1 RQ1: Effects of Emotion Peer Monitoring

6.1.1 Participants thought the analysis was accurate and tried to adjust their wording to empathize with others. In TG1, we presented participants with the emotion analysis of their peer Eva's message. 71.4% of the participants expressed their appreciation of the breakdown of emotions and real-time analysis (e.g. A17: "I liked the breakdown of the emotions which demonstrates the complexity in




Fig. 4. Hate speech comparison for all three groups (CG, TG1, and TG2), using Twitter-roBERTa-base.

people’s responses”, A9: “The system is very intriguing. I’ve never seen anything like that before. I think it will serve people well in learning about themselves and their reactions to posts on social media.”). Three participants (A3, A8, A27) explicitly expressed that they thought the analysis was accurate and confirmed their feelings about the peer’s comment, for example, A8 wrote: “I felt that it accurately portrayed the emotions of the other commenter. I firmly disagree with their viewpoint. The feedback system validated my suspicion that the commenter was scared and angry.”. Only two participants wrote that the system was inaccurate (A31, A55). Moreover, the system’s ability to facilitate empathetic communication was highlighted by several participants. For instance, A18 mentioned “Because I knew how people felt about Eva’s comment, I tried to put more thought into my comment and tried to be mindful of the views of both sides of the discussion.”. Interestingly, participants also noted that the emotion analysis contributed to their self-reflection. A19 stated “It reminded me to take a step back and frame my response in a way that *maybe* Eva could actually hear me. It felt like I was given a little insight into how to build a better bridge, rather than burn the whole thing down.” These insights underscore the system’s potential to foster more constructive and empathetic online interactions by making users aware of the emotional undertones of communications.

6.1.2 *Heightened negative emotions, reduced positive emotion and reduced hate speech were detected in the message from female participants, while the reversed trend was detected in the message from male participants.* In their final comments, as shown in Table 3, the message written by TG1 exhibited a higher inclination towards all negative emotions (Angry, Fear, and Sad), coupled with a reduction in neutral (Surprise) and positive (Happy) emotions. However, none of these differences were statistically significant. Since over two-thirds of the participants in the main study were female, the overall trends were predominantly influenced by female responses. Analyzing the data by gender, as illustrated in Figure 5, reveals divergent trends among different gender groups. Due to the small number of participants identifying as “Other”, we only consider binary gender in this analysis.

The figure indicates a general increase in negative emotions (Angry, Fear, Sad) and a decrease in Happy among female participants, whereas male participants exhibited opposite trends: a reduction in Angry and Fear, and an increase in Happy. Notably, we observed a significant reduction in Happy for females (Two-sample T-test, $p = .0264$, $T = 2.29$) and a significant increase in Sad (MWU test, $p = .0193$, $U = 182$). Despite the intensified negative emotions among female participants, their comments contained less hate speech, whereas hate speech increased among male participants, as shown in Figure 4.




Fig. 5. Comparison of measured emotions using the *text2emotion* package between CG and TG1.




Fig. 6. Comparison of perceived peer emotions between CG and TG1. Participants were asked to rate to what extent do they feel certain emotions from a 1-5 Likert scale.




Fig. 7. Comparison of perceived self emotions between CG and TG1. Participants were asked to rate to what extent do they feel certain emotions from a 1-5 Likert scale.

It is also worth noting that compared to the control group, TG1 has a reduced hate speech, and this is statistical significance according to the MWU test ($p = .0257$, $U = 722$). Examining the reduction of hate speech from a gender perspective gives us another insight. From Figure 4 we can see that females reduced hate speech TG1 while males slightly increased hate speech. And the reduction of hate speech of females in TG1 is significant ($p = .00647$, $U = 438$).

6.1.3 Participants used more pronouns and fewer articles. To further understand the difference, we tried to examine their use of languages using LIWC. From linguistic dimensions, we found that compared to CG, TG1 participants used pronouns more frequently and articles less frequently. When we examined further, we discovered that this statistical significance actually came from female participants as shown in Table 4. Even though there were similar trends in male participants, the effects were not as significant. The use of personal pronouns is an important aspect of language that reveals information about people’s mental states [46]. Pronouns indicate whether an individual’s focus is on the self (first-person singular pronouns such as “I”, “me”, “my”), on others (second-person pronouns such as “s/he”, “you”, “they”, “others”). Pronouns like “I” and “you” are often used to establish connections between interlocutors, with their increased usage reflecting a more engaged and personalized discourse. Specifically, the frequent use of first-person singular pronouns signals a higher degree of self-reference, often associated with greater emphasis on personal experiences, emotions, or perspectives. Research suggests that this inward focus, as evidenced by the elevated use of first-person singular pronouns, can be indicative of mental distress [16, 91]. This aligns with the heightened negative emotions observed in the messages from female participants, as discussed in Section 6.1.2. Alternatively, this change of wording may reflect a shift towards more direct and concise language. Articles like “the” and “a” are often used to specify or generalize nouns, and their reduced frequency could indicate a preference for brevity or a more streamlined communication style. This might suggest that individuals in TG1 are more focused on conveying their emotions or ideas directly, without excessive elaboration. The observed changes in pronoun and article usage could be attributed, at least in part, to the emotional states experienced by individuals in the treatment group. Negative emotions, such as Sad or Angry, might lead to a greater emphasis on personal experiences and perspectives, as well as a more direct expression of thoughts or feelings.

Group	pronoun	ppron	ipron	i	you	article	auxverb	conj
CG	12.43	6.56	5.88	1.81	1.01	7.51	12.05	7.77
TG1	15.97	8.17	7.79	2.31	2.13	5.45	12.88	7.38
<i>p-value</i>	.005**	.023*	.023*	.038*	.018*	.011*	.407	.516
CG (female)	11.77	6.59	5.18	1.73	1.11	7.31	11.67	7.73
TG1 (female)	15.91	8.34	7.57	2.44	2.55	5.45	12.81	7.82
<i>p-value</i>	.003**	.082	0.02**	.023	.036*	.042*	.335	.889
CG (male)	14.37	6.13	8.24	2.31	0.21	8.99	13.28	7.96
TG1 (male)	16.12	7.75	8.36	1.98	1.08	5.47	13.06	6.29
<i>p-value</i>	.557	.375	.952	.688	.054	.087	.920	.200

Table 4. Comparison of linguistic statistics between CG and TG1. We used the Student’s t-test for normally distributed groups, Welch’s t-test for normally distributed but not equal variance groups, and used the MWU test for not normally distributed groups (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$). **pronoun** indicates total pronouns, including personal pronouns (**ppron**: “I”, “you”, “my”, “me”) and impersonal pronouns (**ipron**: “that”, “it”, “this”, “what”). **i** represents 1st person singular, including “I”, “me”, “my”, “myself”. **you** represents 2nd person, including “you”, “your”, “u”, “yourself”. **article** represents articles, containing “a”, “an”, “the”, “alot”. **auxverb** represents auxiliary verbs, containing “is”, “was”, “be”, “have”. **conj** represents conjunctions, containing “and”, “but”, “so”, “as”. Check the official documentation for more details ⁶.

⁶<https://www.liwc.app/static/documents/LIWC-22%20Manual%20-%20Development%20and%20Psychometrics.pdf>

6.1.4 *Participants did not view the peer’s emotions differently compared to CG and they viewed their own emotions less negatively.* Even though TG1 participants were presented additionally with the emotion analysis of Eva’s comment compared to CG, there was not much difference in how they viewed the emotions in Eva’s comment, no matter whether overall or by gender as shown in Figure 6. However, in their self-reported emotions, both female and male participants reported they experienced less negative emotions, as shown in Figure 7.

6.1.5 *Summary of Findings for RQ1.* The findings suggest that the exposure of participants in TG1 to emotion analysis of peer messages influenced their perception and expression of emotions. This influence is reflected not only in the evaluation of peer comments but also in the linguistic features of their communication. The reduced hate speech and qualitative feedback indicating the accuracy of the analysis further underscore the role of emotion monitoring in shaping language use and interpersonal communication.

6.2 RQ2: Effects of Emotion Self Monitoring

6.2.1 *Participants thought the analysis was inaccurate and reported the discrepancy between emotional feedback and their true feelings.* In TG2, we presented participants with the emotion analysis of their own written messages. In contrast to the positive feedback from TG1, there are more polarized opinions on the emotion self monitoring system in TG2. While over half of the participants (58.3%) in TG2 expressed appreciation for the system and the concept of emotion monitoring, concerns were raised about its accuracy and the alignment of the analysis with their personal feelings. Seven participants complained that they did not think the system was accurate (B4 , B8 , B12 , B15 , B19 , B29 , B33). Additionally, four participants (B15 , B21 , B26 , B34) reported discrepancies between the emotions portrayed in the dashboard and their actual emotional experiences (e.g. B34 : “I’m not really sure how accurate it was. Yes, sadness was a big part of my response, but I felt anger was more dominant in my response.”), potentially leading them to disregard the dashboard’s feedback. Furthermore, two participants (B1 , B29) found it challenging to reduce the negative aspects of their language. B1 noted “*It was difficult to tone down my emotional statuses. I did not think I was being negative, but that is what the emotion dashboard determined.*”

This discrepancy may be due to the following factors. Firstly, the algorithm’s inherent limitations in capturing nuanced contextual cues led to reduced accuracy and made people disregard the feedback. The text2emotion package works by detecting emotion-related words and categorizing them into five predefined emotions. When participants discuss negative topics, the algorithm reports negative emotions even if the participants do not intend to express such emotions. For instance, B17 stated “*In my original comment, I talked about rape situations where abortion is necessary. When I took out the word rape, my fear score went down. Yes, rape is bad, but I’m not actively afraid, I’m just talking about the action as a simple fact.*”. Therefore, there might be a disconnect between the emotions users intend to convey and those that are effectively communicated through their messages. This misalignment highlights the complex interplay between users’ emotional intentions and the expressive nuances of their language.

6.2.2 *Heightened negative emotions and a reduced positive emotion were detected in the message from both females and males, but only female participants reduced hate speech.* Figures 8 and Table 3 present that more negative emotions and less positive emotions were detected in messages from both male and female participants. However, in Figure 4, we can see that only female participants had a reduced hate speech.

6.2.3 *Participants used more auxiliary verbs and fewer conjunctions.* Regarding language usage, TG2, similar to TG1, demonstrated a higher frequency of pronoun usage compared to the control




Fig. 8. Comparison of measured emotions using the *text2emotion* package between CG and TG2.




Fig. 9. Comparison of perceived peer emotions between CG and TG2. Participants were asked to rate to what extent do they feel certain emotions from a 1-5 Likert scale.




Fig. 10. Comparison of perceived self emotions between CG and TG2. Participants were asked to rate to what extent do they feel certain emotions from a 1-5 Likert scale.

group, although this difference was not statistically significant (as shown in Table 5). Conversely, TG2 displayed lower usage of articles, but again, this difference did not reach statistical significance. However, according to the MWU test, TG2 did exhibit a significant increase in the usage of auxiliary verbs ($p = .0161, U = 475$) and a decrease in conjunctions ($p = .0175, U = 927$) compared to the control group. These results suggest a potential shift towards more direct and action-oriented language within TG2.

The increase in auxiliary verb usage in TG2 may indicate a tendency towards expressing actions, states, or conditions more explicitly. Auxiliary verbs often serve to clarify tense, aspect, modality,

and voice, which can make language more precise and direct. For instance, participants in TG2 might be more focused on describing their current emotional states or anticipated actions, reflecting the impact of the emotion self monitoring system on their language patterns. Alternatively, one hypothesis could be that the emotion self monitoring system encouraged participants to narrate their experiences in a step-by-step or instructional manner, which often relies on auxiliary verbs to describe processes and sequences. The decrease in conjunction usage in TG2 suggests a reduction in the complexity of sentence structures. Conjunctions are used to link clauses, phrases, and words, creating more complex and compound sentences. A reduction in conjunction usage may indicate a shift towards simpler, more straightforward communication. This could be a result of participants attempting to convey their emotions more clearly and directly, possibly in response to the feedback from the emotion self monitoring system. Alternatively, this simplification could result from cognitive load: participants might subconsciously reduce sentence complexity when focusing on emotional regulation tasks, prioritizing clarity over nuance.

6.2.4 Participants view the peer's emotions more negatively compared to CG and they viewed their own emotions differently for different genders. From Figure 9 we can see that compared to CG, TG2 participants evaluate their peer's message had more negative emotions for both males and females. For self-reported emotions, there was a gender difference. Female participants reported experiencing more negative emotions while male participants reported experiencing less negative emotions as shown in Figure 10. This trend is reflected in their messages, where more negative emotions were evident, and fewer positive emotions were expressed, as detailed in Table 3. However, statistical analysis revealed that these emotional differences did not reach statistical significance when compared between the groups.

Group	pronoun	ppron	ipron	i	you	article	auxverb	conj
CG	12.43	6.56	5.88	1.81	1.01	7.51	12.05	7.77
TG2	14.44	7.36	7.08	2.24	1.11	7.18	13.70	6.63
<i>p-value</i>	.146	.754	.151	.310	.961	.669	.016*	.035*
CG (female)	11.77	6.59	5.18	1.73	1.11	7.31	11.67	7.73
TG2 (female)	13.84	7.04	6.80	2.02	1.03	7.38	13.46	6.21
<i>p-value</i>	.173	.712	.364	.426	.727	.941	.038*	.013*
CG (male)	14.37	6.13	8.24	2.31	0.21	8.99	13.28	7.96
TG2 (male)	15.99	8.03	7.96	2.79	1/33	7.11	14.03	6.85
<i>p-value</i>	.622	.375	.952	.689	.054	.09	.385	.358

Table 5. Comparison of linguistic statistics between CG and TG2. We used the Student's t-test for normally distributed groups, Welch's t-test for normally distributed but not equal variance groups, and used the MWU test for not normally distributed groups (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$). **pronoun** indicates total pronouns, including personal pronouns (**ppron**: "I", "you", "my", "me") and impersonal pronouns (**ipron**: "that", "it", "this", "what"). **i** represents 1st person singular, including "I", "me", "my", "myself". **you** represents 2nd person, including "you", "your", "u", "yourself". **article** represents articles, containing "a", "an", "the", "alot". **auxverb** represents auxiliary verbs, containing "is", "was", "be", "have". **conj** represents conjunctions, containing "and", "but", "so", "as". Check the official documentation for more details ⁷.

⁷<https://www.liwc.app/static/documents/LIWC-22%20Manual%20-%20Development%20and%20Psychometrics.pdf>

6.2.5 Summary of Findings for RQ2. The findings suggest that the exposure of participants in TG2 to emotion analysis of their own messages influenced their perception and expression of emotions. We observed a reversed trend in how female and male participants evaluate their own emotions. Furthermore, we found the reduction in hate speech and significant changes in linguistic features both came from female participants.

6.3 Transparency and censorship

In addition to the findings from isolated factors, several common concerns emerged across both participant groups. Despite efforts to enhance user comprehension and acceptance through personalized recommendations and explanatory messages, many participants expressed a lack of transparency in explaining how emotional scores are calculated. Participants from both groups emphasized the need for clearer communication regarding the methodology. Over one-third participants in TG1 (A1, A2, A4, A7, A10, A12, A15, A16, A17, A25) and four participants in TG2 (B7, B12, B17, B30) expressed that they struggled to interpret the results and requested insight into the algorithm's workings. For instance, from TG1 A12 mentioned *"It could give more information on how it worked, I was still confused after I looked at the explanation."* and A16 suggested *"Make the methodology more transparent. Even just a blurb explaining how it works."* Similarly, B7 in TG2 noted *"Highlight or underline which words or sections of a comment are triggering certain positive or negative responses from the feedback tool."* This lack of transparency may not only impede users' ability to act upon emotional feedback but also contribute to skepticism and disengagement with the system. As A10 noted *"I did not understand what I was supposed to do. It was confusing and made me feel frustrated."*

Interestingly, participants in TG2 who received emotion self monitoring features reported feeling censored—a perception not observed among TG1 participants. For example, TG2 participant B13 remarked *"Rather than feeling like I'm making sure my comments are civil, it really just feels like I'm being censored."* Similarly, B32 expressed skepticism about the system's utility, stating *"It's extremely hectoring and judgmental. I don't need my sentiment to be analyzed since obviously I know how I feel (and better than an algorithm)."*

7 Discussion

Through a series of studies, we investigated the effects of two types of emotion monitoring on emotional experience and linguistic behavior in online communication of a sensitive topic. Similar to [10], our findings indicate that dynamic, personalized interventions may substantially impact user's commenting behavior. Specifically, our findings show that emotional analytics may enhance emotional awareness in social media discussions.

7.1 Theoretical Implications

Table 6 summarizes the findings by topic through two research questions. The findings indicate that exposure to emotion analysis of peer messages significantly impacts users' perception and expression of emotions. Participants of TG1 exhibited a heightened sensitivity towards negative emotions in peer comments. The qualitative feedback further corroborates these findings, emphasizing the role of emotion monitoring in fostering empathy and guiding users toward adjusting the emotional tone of their messages. This highlights the potential of emotion peer-monitoring to not only influence individual behavior but also facilitate more empathetic and emotionally attuned interactions within online discussions.

Interestingly, female participants in both monitoring groups expressed more negative emotions in their comments. We controlled for participants' emotional states using the PANAS scores before

Topic	RQ1: Effects of Emotion Peer-Monitoring	RQ2: Effects of Emotion Self-Monitoring
Participants reactions	They thought the analysis was accurate and tried to adjust their wordings to empathy with others.	They thought the analysis was inaccurate and reported the discrepancy between emotional feedback and their true feelings.
Inferred emotions in the message	Heightened negative emotions, reduced positive emotion and reduced hate speech from female participants, while the reversed trend from male participants.	Heightened negative emotions and a reduced positive emotion from both genders, but only female participants reduced hate speech.
Self-reported emotions by participants	Participants did not view the peer's emotions differently compared to CG and they viewed their own emotions less negatively.	Participants viewed the peer's emotions more negatively compared to CG and they viewed their own emotions differently for different genders.
Linguistic usage	Participants used more pronouns and fewer articles, statistical significance was detected in female participants.	Participants used more auxiliary verbs and fewer conjunctions, statistical significance was detected in female participants.

Table 6. Summary of findings for RQ1 and RQ2.

they began the conversation. This ensured that all participants had comparable emotional baselines at the start, with no individuals exhibiting excessively high emotional levels. While the heightened negative emotions among female participants could, in part, be attributed to the topic —abortion. It is important to note that existing research suggests mixed findings regarding gender-based attitudes toward abortion. Specifically, while abortion is a cornerstone issue in the women's rights movement, attitudes toward it are not consistently aligned with gender [44, 54]. Besides different genders could have different attitudes towards abortion.

Another explanation for the heightened negative emotions could lie in the differences in how males and females process emotional stimuli. Research indicates that females may engage different neurocircuitry than males during perceptual emotion processing. In some cases, this can lead to more accurate or faster emotional assessments, while in others, it may result in heightened reactivity, potentially increasing vulnerability to affective disorders [22].

This result also highlights a gendered dimension in emotional sensitivity and expression that warrants further exploration. The heightened levels of negative emotions, particularly Angry and Sad, underscore the potential impact of emotion monitoring interventions on emotional expression. These results align with previous research indicating that heightened self-awareness of emotions can sometimes lead to increased negative affect, especially when individuals become overly focused on monitoring and regulating their emotional experiences [69]. As heightened awareness reduces the inclination to suppress genuine feelings, leading to more pronounced negative emotions in responses [86]. Despite the increase in negative emotional expression, we also observed a decrease in the use of hate speech among these participants, indicating a complex interaction between emotional awareness and communicative behavior.

The linguistic analysis of the texts revealed that several verbal dimensions differed significantly across the two monitoring conditions. Both treatment groups (TG1, TG2) have an increase in the use of pronouns and auxiliary verbs, suggesting that emotion monitoring systems might encourage users to articulate their emotions and actions more clearly. This could be beneficial for improving emotional self-awareness and communication.

Additionally, the accuracy of the system's analysis was a key factor in its acceptance and effectiveness. While the majority of participants found the analysis accurate, those who did not raise concerns about the potential misinterpretation of emotions (A31, A55). The discrepancy between users' self-reported emotional states and the emotional content of their messages highlights the complexity of emotional expression in online communication. Users' attempts to adjust their messages to diminish perceived negativity suggest a conscious effort to align their online persona with their desired emotional presentation. However, the challenges they encountered in accurately conveying their emotional states suggest limitations in current emotion detection algorithms and the need for greater transparency in how emotional scores are computed. This feedback also highlights the importance of continually refining the emotion analysis algorithms to ensure reliability and to address the nuances in human emotions effectively.

7.2 Practical Implications

This study holds practical implications for designers and researchers. First, it underscores the significant influence of platform design on user behavior [30]. Specifically, designers should carefully identify features that support users in empathizing with their peers on these platforms [72, 75] and build the self-reflection capacities of users [28, 59]. Second, as algorithmic outputs grow more complex, users require guidance on interpreting and acting on this feedback while the design of explainable systems [1]. Trust in these algorithms hinges on transparent explanations of their workings, making such clarity essential for their acceptance [34, 85]. As Andalibi and Buss [3] highlight, users often have limited understanding of how their emotions are detected and interpreted, which can lead to misinformed interactions with these tools. Designers must prioritize clear and accessible explanations of how these systems work, including the data sources, algorithms, and potential biases involved. Enhancing transparency in algorithmic interface [29, 36, 85] could bridge this trust gap, helping users understand and accept these tools more readily. Finally, these findings suggest that future investigations into emotion regulation must take a holistic, context-sensitive approach. Ultimately, by embedding these contextual considerations into both design and research, future developments in emotion regulation technologies can better support users' well-being.

In conclusion, the preliminary findings suggest that the emotion analysis system is not only well-received but also has a positive impact on users' online interactions. By providing insights into the emotional content of messages, the system encourages empathy, self-reflection, and more mindful communication. These results pave the way for further investigations into the long-term benefits of such systems in promoting healthier and more productive online discourse. Future work will involve larger-scale studies to validate these findings and explore additional features that could enhance the system's utility and user experience.

7.3 Limitations and Future Work

This research comes with some limitations which suggests fruitful avenues for future research. First, our measurement of emotions relies on detecting keywords categorized into specific emotional categories, without considering the context in which these keywords are used. Future research should aim to incorporate contextual nuances into emotion detection systems while ensuring transparency in their operations [21, 76]. Second, while this study investigated the effects of emotion self-monitoring and emotion peer-monitoring independently, future work could explore

the combined impact of providing individuals with emotional analyses of both their own and their peers' written comments. Such an approach may reveal new insights into collaborative dynamics and emotional awareness. Third, as participants in both treatment groups complained about the transparency of the emotion analysis, future research could provide more transparency in the algorithmic result on the emotion or toxicity score [85]. Future research could focus on enhancing the interpretability of these systems to foster trust and usability [85]. Additionally, we encourage the HCI community to prioritize research on trust calibration with AI systems [4] and the development of explainable AI technologies [1, 12] to promote responsible and effective AI usage [81]. Finally, future work could expand beyond linguistic markers to enhance emotion detection and communication. Exploring multimodal features such as speech [42], facial expressions [87?], and tactile cues like keystroke strength [48] could significantly advance the field.

8 Conclusion

Hate speech and toxic communication on social media platforms are on the rise threatening the fabric of society. In addition to reactive moderation, proactive moderation that provides users with real-time, personalized support to engage more constructively in dialogue may be key to further reducing problematic content. Based on the assumptions that discussions on social media platforms are susceptible to lower levels of emotional awareness, we designed and tested an emotion monitoring dashboard to help users perceive the emotions in their own and their peers posts.

Our study illuminates the intricate interplay between emotion monitoring, linguistic expression, and users' emotional experiences in online environments. The observed increase in negative emotional expression, particularly anger and sadness, suggests that emotion monitoring can heighten self-awareness and potentially lead to more genuine emotional expressions. This aligns with existing research indicating that increased self-awareness of emotions can sometimes exacerbate negative affect, especially when individuals focus intensely on monitoring and regulating their emotional states. The discrepancy between users' self-reported emotions and the emotional content of their messages points to the complexities of emotional expression in digital communication. Users' attempts to moderate their message tone to align with their desired emotional presentation indicate a conscious effort to manage their online persona. However, the challenges in accurately conveying these emotional states underscore the limitations of current emotion detection technologies. Future research should focus on enhancing the accuracy and granularity of emotion detection algorithms to better support users in expressing their true emotions and improving the overall quality of online interactions. By advancing our understanding of how emotion monitoring can shape online conversations, we can develop tools that foster healthier and more empathetic digital communities.

Ethical Statement

The experiment was done according to the ethical guidelines of our university. An ethical assessment was done and reviewed according to the internal criteria. An external approval were not required for the study on human participants in accordance with the local legislation and institutional requirements of a Western European university.

References

- [1] Ashraf Abdul, Jo Vermeulen, Danding Wang, Brian Y. Lim, and Mohan Kankanhalli. 2018. Trends and Trajectories for Explainable, Accountable and Intelligible Systems: An HCI Research Agenda. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3173574.3174156
- [2] Anna L. Altshuler, Helen L. Gerns Storey, and Sarah W. Prager. 2015. Exploring abortion attitudes of US adolescents and young adults using social media. *Contraception* 91, 3 (2015), 226–233. doi:10.1016/j.contraception.2014.11.009

- [3] Nazanin Andalibi and Justin Buss. 2020. The Human in Emotion Recognition on Social Media: Attitudes, Outcomes, Risks. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3313831.3376680
- [4] Tita Alissa Bach, Amna Khan, Harry Hallock, Gabriela Beltrão, and Sonia Sousa. 2022. A systematic literature review of user trust in AI-enabled systems: An HCI perspective. *International Journal of Human-Computer Interaction* (2022), 1–16.
- [5] Francesco Barbieri, Jose Camacho-Collados, Luis Espinosa Anke, and Leonardo Neves. 2020. TweetEval: Unified Benchmark and Comparative Evaluation for Tweet Classification. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 1644–1650. doi:10.18653/v1/2020.findings-emnlp.148
- [6] Jiang Bian, Yandong Liu, Ding Zhou, Eugene Agichtein, and Hongyuan Zha. 2009. Learning to recognize reliable users and content in social media with coupled mutual reinforcement. In *Proceedings of the 18th International Conference on World Wide Web* (Madrid, Spain) (WWW '09). Association for Computing Machinery, New York, NY, USA, 51–60. doi:10.1145/1526709.1526717
- [7] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative research in sport, exercise and health* 11, 4 (2019), 589–597.
- [8] Tommaso Caselli, Valerio Basile, Jelena Mitrović, and Michael Granitzer. 2021. HateBERT: Retraining BERT for Abusive Language Detection in English. In *Proceedings of the 5th Workshop on Online Abuse and Harms (WOAH 2021)*, Aida Mostafazadeh Davani, Douwe Kiela, Mathias Lambert, Bertie Vidgen, Vinodkumar Prabhakaran, and Zeerak Waseem (Eds.). Association for Computational Linguistics, Online, 17–25. doi:10.18653/v1/2021.woah-1.3
- [9] Eshwar Chandrasekharan, Chaitrali Gandhi, Matthew Wortley Mustelier, and Eric Gilbert. 2019. Crossmod: A Cross-Community Learning-Based System to Assist Reddit Moderators. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 174 (nov 2019), 30 pages. doi:10.1145/3359276
- [10] Jonathan P. Chang, Charlotte Schluger, and Cristian Danescu-Niculescu-Mizil. 2022. Thread With Caution: Proactively Helping Users Assess and Deescalate Tension in Their Online Discussions. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 545 (nov 2022), 37 pages. doi:10.1145/3555603
- [11] Frederick Choi, Tanvi Bajpai, Sowmya Pratipati, and Eshwar Chandrasekharan. 2023. ConvEx: A Visual Conversation Exploration System for Discord Moderators. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW2, Article 262 (oct 2023), 30 pages. doi:10.1145/3610053
- [12] Michael Chromik and Andreas Butz. 2021. Human-XAI Interaction: A Review and Design Principles for Explanation User Interfaces. In *Human-Computer Interaction – INTERACT 2021: 18th IFIP TC 13 International Conference, Bari, Italy, August 30 – September 3, 2021, Proceedings, Part II* (Bari, Italy). Springer-Verlag, Berlin, Heidelberg, 619–640. doi:10.1007/978-3-030-85616-8_36
- [13] Munmun De Choudhury, Shagun Jhaver, Benjamin Sugar, and Ingmar Weber. 2016. Social media participation in an activist movement for racial equality. In *Proceedings of the international aaai conference on web and social media*, Vol. 10. 92–101.
- [14] Gloria del Valle-Cano, Lara Quijano-Sánchez, Federico Liberatore, and Jesús Gómez. 2023. SocialHaterBERT: A dichotomous approach for automatically detecting hate speech on Twitter through textual analysis and user profiles. *Expert Systems with Applications* 216 (2023), 119446. doi:10.1016/j.eswa.2022.119446
- [15] Nicholas Diakopoulos and Mor Naaman. 2011. Towards quality discourse in online news comments. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work* (Hangzhou, China) (CSCW '11). Association for Computing Machinery, New York, NY, USA, 133–142. doi:10.1145/1958824.1958844
- [16] To'Meisha Edwards and Nicholas S. Holtzman. 2017. A meta-analysis of correlations between depression and first person singular pronoun use. *Journal of Research in Personality* 68 (2017), 63–68. doi:10.1016/j.jrp.2017.02.005
- [17] Mohammed-Amr Abd El-Migid, Damon Cai, Thomas Niven, Jeffrey Vo, Kashumi Madampe, John Grundy, and Rashina Hoda. 2022. Emotimonitor: A Trello power-up to capture and monitor emotions of Agile teams. *Journal of Systems and Software* 186 (2022), 111206. doi:10.1016/j.jss.2021.111206
- [18] Mai ElSherief, Vivek Kulkarni, Dana Nguyen, William Yang Wang, and Elizabeth Belding. 2018. Hate Lingo: A Target-Based Linguistic Analysis of Hate Speech in Social Media. *Proceedings of the International AAAI Conference on Web and Social Media* 12, 1 (Jun. 2018). doi:10.1609/icwsm.v12i1.15041
- [19] Mohamed Ez-zaouia and Elise Lavoué. 2017. EMODA: A Tutor Oriented Multimodal and Contextual Emotional Dashboard. In *Proceedings of the Seventh International Learning Analytics & Knowledge Conference* (Vancouver, British Columbia, Canada) (LAK '17). Association for Computing Machinery, New York, NY, USA, 429–438. doi:10.1145/3027385.3027434
- [20] Mohamed Ez-zaouia, Aurélien Tabard, and Elise Lavoué. 2020. Emodash: A dashboard supporting retrospective awareness of emotions in online learning. *International Journal of Human-Computer Studies* 139 (2020), 102411. doi:10.1016/j.ijhcs.2020.102411

- [21] Jose Maria Garcia-Garcia, Victor M. R. Penichet, and Maria D. Lozano. 2017. Emotion detection: a technology review. In *Proceedings of the XVIII International Conference on Human Computer Interaction* (Cancun, Mexico) (*Interacción '17*). Association for Computing Machinery, New York, NY, USA, Article 8, 8 pages. doi:10.1145/3123818.3123852
- [22] Richard Gater, Michele Tansella, Ailsa Korten, Bea G. Tiemens, Venos G. Mavreas, and Michael O. Olatawura. 1998. Sex Differences in the Prevalence and Detection of Depressive and Anxiety Disorders in General Health Care Settings: Report From the World Health Organization Collaborative Study on Psychological Problems in General Health Care. *Archives of General Psychiatry* 55, 5 (05 1998), 405–413. arXiv:https://jamanetwork.com/journals/jamapsychiatry/articlepdf/203838/yoa6182.pdf doi:10.1001/archpsyc.55.5.405
- [23] Elena Geangu, Oana Benga, Daniel Stahl, and Tricia Striano. 2011. Individual differences in infants' emotional resonance to a peer in distress: Self–other awareness and emotion regulation. *Social Development* 20, 3 (2011), 450–470.
- [24] James J Gross. 1998. The emerging field of emotion regulation: An integrative review. *Review of general psychology* 2, 3 (1998), 271–299.
- [25] Oliver L. Haimson, Daniel Delmonaco, Peipei Nie, and Andrea Wegner. 2021. Disproportionate Removals and Differing Content Moderation Experiences for Conservative, Transgender, and Black Social Media Users: Marginalization and Moderation Gray Areas. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 466 (oct 2021), 35 pages. doi:10.1145/3479610
- [26] Sébastien Henschel, Jean-Louis Nandrino, and Karyn Doba. 2020. Emotion regulation and empathic abilities in young adults: The role of attachment styles. *Personality and Individual Differences* 156 (2020), 109763.
- [27] Sameer Hinduja and Justin W. Patchin. 2010. Bullying, Cyberbullying, and Suicide. *Archives of Suicide Research* 14, 3 (2010), 206–221. arXiv:https://doi.org/10.1080/13811118.2010.494133 doi:10.1080/13811118.2010.494133 PMID: 20658375.
- [28] Krithika Jagannath, Katie Salen, and Petr Slovák. 2020. "(We) Can Talk It Out...": Designing for Promoting Conflict-Resolution Skills in Youth on a Moderated Minecraft Server. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW1, Article 49 (may 2020), 26 pages. doi:10.1145/3392855
- [29] Shagun Jhaver, Darren Scott Appling, Eric Gilbert, and Amy Bruckman. 2019. "Did You Suspect the Post Would be Removed?": Understanding User Reactions to Content Removals on Reddit. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 192 (nov 2019), 33 pages. doi:10.1145/3359294
- [30] Przemysław Kazienko, Julita Bielaniewicz, Marcin Gruza, Kamil Kanclerz, Konrad Karanowski, Piotr Miłkowski, and Jan Kocoń. 2023. Human-centered neural reasoning for subjective content processing: Hate speech, emotions, and humor. *Information Fusion* 94 (2023), 43–65. doi:10.1016/j.inffus.2023.01.010
- [31] Renato Kempter, Valentina Sintsova, Claudiu Musat, and Pearl Pu. 2014. EmotionWatch: Visualizing fine-grained emotions in event-related tweets. In *Proceedings of the international AAI conference on web and social media*, Vol. 8. 236–245.
- [32] Sara Kiesler, Robert Kraut, Paul Resnick, and Aniket Kittur. 2012. Regulating behavior in online communities. *Building successful online communities: Evidence-based social design* 1 (2012), 4–2.
- [33] Jooyeon Kim, Behzad Tabibian, Alice Oh, Bernhard Schölkopf, and Manuel Gomez-Rodriguez. 2018. Leveraging the Crowd to Detect and Reduce the Spread of Fake News and Misinformation. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining* (Marina Del Rey, CA, USA) (*WSDM '18*). Association for Computing Machinery, New York, NY, USA, 324–332. doi:10.1145/3159652.3159734
- [34] Sunnie S. Y. Kim, Elizabeth Anne Watkins, Olga Russakovsky, Ruth Fong, and Andrés Monroy-Hernández. 2023. "Help Me Help the AI": Understanding How Explainability Can Support Human-AI Interaction. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (*CHI '23*). Association for Computing Machinery, New York, NY, USA, Article 250, 17 pages. doi:10.1145/3544548.3581001
- [35] Joel Kiskola, Thomas Olsson, Heli Väättäjä, Alekski H. Syrjämäki, Anna Rantasila, Poika Isokoski, Mirja Ilves, and Veikko Surakka. 2021. Applying Critical Voice in Design of User Interfaces for Supporting Self-Reflection and Emotion Regulation in Online News Commenting. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (*CHI '21*). Association for Computing Machinery, New York, NY, USA, Article 88, 13 pages. doi:10.1145/3411764.3445783
- [36] René F. Kizilcec. 2016. How Much Information? Effects of Transparency on Trust in an Algorithmic Interface. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (*CHI '16*). Association for Computing Machinery, New York, NY, USA, 2390–2395. doi:10.1145/2858036.2858402
- [37] Sara H. Konrath, Edward H. O'Brien, and Courtney Hsing. 2011. Changes in Dispositional Empathy in American College Students Over Time: A Meta-Analysis. *Personality and Social Psychology Review* 15, 2 (2011), 180–198. arXiv:https://doi.org/10.1177/1088868310377395 doi:10.1177/1088868310377395 PMID: 20688954.
- [38] Yubo Kou. 2021. Punishment and Its Discontents: An Analysis of Permanent Ban in an Online Game Community. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 334 (oct 2021), 21 pages. doi:10.1145/3476075

- [39] Mariana L. de Carvalho, Thomas Olsson, and Joel Kiskola. 2021. Exploration of user interface mechanisms with affect labeling to enhance online discussion. In *Proceedings of the 24th International Academic Mindtrek Conference (Tampere/Virtual, Finland) (Academic Mindtrek '21)*. Association for Computing Machinery, New York, NY, USA, 210–220. doi:10.1145/3464327.3464332
- [40] Q. Vera Liao and Wai-Tat Fu. 2013. Beyond the Filter Bubble: Interactive Effects of Perceived Threat and Topic Involvement on Selective Exposure to Information. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Paris, France) (CHI '13)*. Association for Computing Machinery, New York, NY, USA, 2359–2368. doi:10.1145/2470654.2481326
- [41] Rensis Likert. 1932. A technique for the measurement of attitudes. *Archives of psychology* (1932).
- [42] Wei-Cheng Lin and Dimitra Emmanouilidou. 2022. Toxic Speech and Speech Emotions: Investigations of Audio-based Modeling and Intercorrelations. In *2022 30th European Signal Processing Conference (EUSIPCO)*. 115–119. doi:10.23919/EUSIPCO55093.2022.9909856
- [43] Baoxi Liu, Peng Zhang, Yubo Shu, Zhengqing Guan, Tun Lu, Hansu Gu, and Ning Gu. 2022. Building a Personalized Model for Social Media Textual Content Censorship. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 499 (nov 2022), 31 pages. doi:10.1145/3555657
- [44] Mary-Kate Lizotte. 2015. The Abortion Attitudes Paradox: Model Specification and Gender Differences. *Journal of Women, Politics & Policy* 36, 1 (2015), 22–42. arXiv:https://doi.org/10.1080/1554477X.2015.985151 doi:10.1080/1554477X.2015.985151
- [45] Patricia L Lockwood, Ana Seara-Cardoso, and Essi Viding. 2014. Emotion regulation moderates the association between empathy and prosocial behavior. *PLoS one* 9, 5 (2014), e96555.
- [46] Minna Lyons, Nazli Deniz Aksayli, and Gayle Brewer. 2018. Mental distress and language use: Linguistic analysis of discussion forum posts. *Computers in Human Behavior* 87 (2018), 207–211. doi:10.1016/j.chb.2018.05.035
- [47] Renkai Ma, Yue You, Xinning Gui, and Yubo Kou. 2023. How Do Users Experience Moderation?: A Systematic Literature Review. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW2, Article 278 (oct 2023), 30 pages. doi:10.1145/3610069
- [48] Aicha Maalej and Ilhem Kallel. 2020. Does Keystroke Dynamics tell us about Emotions? A Systematic Literature Review and Dataset Construction. In *2020 16th International Conference on Intelligent Environments (IE)*. 60–67. doi:10.1109/IE49459.2020.9155004
- [49] Binny Mathew, Anurag Illendula, Punyajoy Saha, Soumya Sarkar, Pawan Goyal, and Animesh Mukherjee. 2020. Hate begets Hate: A Temporal Study of Hate Speech. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW2, Article 92 (Oct. 2020), 24 pages. doi:10.1145/3415163
- [50] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and Inter-Rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 72 (nov 2019), 23 pages. doi:10.1145/3359174
- [51] Miljana Mladenović, Vera Ošmjanski, and Staša Vujičić Stanković. 2021. Cyber-Aggression, Cyberbullying, and Cyber-Grooming: A Survey and Research Challenges. *ACM Comput. Surv.* 54, 1, Article 1 (jan 2021), 42 pages. doi:10.1145/3424246
- [52] Katarzyna Molek-Kozakowska and Michał Wanke. 2019. Reproductive rights or duties? The rhetoric of division in social media debates on abortion law in Poland. *Social Movement Studies* 18, 5 (2019), 566–585. arXiv:https://doi.org/10.1080/14742837.2019.1629279 doi:10.1080/14742837.2019.1629279
- [53] Dario Monzani, Laura Vergani, Silvia Francesca Maria Pizzoli, Giulia Marton, and Gabriella Pravettoni. 2021. Emotional tone, analytical thinking, and somatosensory processes of a sample of Italian tweets during the first phases of the COVID-19 pandemic: Observational study. *Journal of Medical Internet Research* 23, 10 (2021), e29820.
- [54] Danny Osborne, Yanshu Huang, Nickola C Overall, Robbie M Sutton, Aino Petterson, Karen M Douglas, Paul G Davies, and Chris G Sibley. 2022. Abortion attitudes: An overview of demographic and ideological differences. *Political Psychology* 43 (2022), 29–76.
- [55] Zhenhui Peng, Taewook Kim, and Xiaojuan Ma. 2019. GremoBot: Exploring Emotion Regulation in Group Chat. In *Companion Publication of the 2019 Conference on Computer Supported Cooperative Work and Social Computing (Austin, TX, USA) (CSCW '19 Companion)*. Association for Computing Machinery, New York, NY, USA, 335–340. doi:10.1145/3311957.3359472
- [56] James W Pennebaker, Martha E Francis, and Roger J Booth. 2001. Linguistic inquiry and word count: LIWC 2001. *Mahway: Lawrence Erlbaum Associates* 71, 2001 (2001), 2001.
- [57] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *J. Mach. Learn. Res.* 21, 1, Article 140 (jan 2020), 67 pages.
- [58] Julian Risch and Ralf Krestel. 2020. Toxic comment detection in online discussions. *Deep learning-based approaches for sentiment analysis* (2020), 85–109.

- [59] Kathleen Van Royen, Karolien Poels, Heidi Vandebosch, and Bieke Zaman. 2022. Think twice to be nice? A user experience study on a reflective interface to reduce cyber harassment on social networking sites. *International journal of bullying prevention* 4, 1 (2022), 23–34.
- [60] F. Belèn Saldías and Rosalind W. Picard. 2019. Tweet Moodifier: Towards giving emotional awareness to Twitter users. In *2019 8th International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 1–7. doi:10.1109/ACII.2019.8925533
- [61] Charlotte Schluger, Jonathan P. Chang, Cristian Danescu-Niculescu-Mizil, and Karen Levy. 2022. Proactive Moderation of Online Discussions: Existing Practices and the Potential for Algorithmic Support. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 370 (nov 2022), 27 pages. doi:10.1145/3555095
- [62] Security.org. 2024. Cyberbullying Facts & Statistics. <https://www.security.org/resources/cyberbullying-facts-statistics/>. Accessed: 2024-10-30.
- [63] Joseph Seering, Tianmi Fang, Luca Damasco, Mianhong 'Cherie' Chen, Likang Sun, and Geoff Kaufman. 2019. Designing User Interface Elements to Improve the Quality and Civility of Discourse in Online Commenting Behaviors. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3290605.3300836
- [64] Joseph Seering, Robert Kraut, and Laura Dabbish. 2017. Shaping Pro and Anti-Social Behavior on Twitch Through Moderation and Example-Setting. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (Portland, Oregon, USA) (CSCW '17). Association for Computing Machinery, New York, NY, USA, 111–125. doi:10.1145/2998181.2998277
- [65] Joseph Seering, Tony Wang, Jina Yoon, and Geoff Kaufman. 2019. Moderator engagement and community development in the age of algorithms. *New Media & Society* 21, 7 (2019), 1417–1443.
- [66] Alexandra Segerberg and W Lance Bennett. 2011. Social media and the organization of collective action: Using Twitter to explore the ecologies of two climate change protests. *The Communication Review* 14, 3 (2011), 197–215.
- [67] Samuel Sanford Shapiro and Martin B Wilk. 1965. An analysis of variance test for normality (complete samples). *Biometrika* 52, 3/4 (1965), 591–611.
- [68] Amit Sheth, Valerie L. Shalin, and Ugur Kursuncu. 2022. Defining and detecting toxicity on social media: context and knowledge are key. *Neurocomputing* 490 (2022), 312–318. doi:10.1016/j.neucom.2021.11.095
- [69] Paul J Silvia. 2002. Self-awareness and emotional intensity. *Cognition & Emotion* 16, 2 (2002), 195–216.
- [70] Petr Slovak, Alissa Antle, Nikki Theofanopoulou, Claudia Daudén Roquet, James Gross, and Katherine Isbister. 2023. Designing for Emotion Regulation Interventions: An Agenda for HCI Theory and Research. *ACM Trans. Comput.-Hum. Interact.* 30, 1, Article 13 (March 2023), 51 pages. doi:10.1145/3569898
- [71] Wally Smith, Greg Wadley, Sarah Webber, Benjamin Tag, Vassilis Kostakos, Peter Koval, and James J. Gross. 2022. Digital Emotion Regulation in Everyday Life. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 444, 15 pages. doi:10.1145/3491102.3517573
- [72] Wiktor Soral, Katarzyna Malinowska, and Michał Bilewicz. 2022. The role of empathy in reducing hate speech proliferation. Two contact-based interventions in online and off-line settings. *Peace and Conflict: Journal of Peace Psychology* (2022).
- [73] Lee Sproull and Sara Kiesler. 1986. Reducing social context cues: Electronic mail in organizational communication. *Management science* 32, 11 (1986), 1492–1512.
- [74] Daniel M. Stewart and Eric J. Fritsch. 2011. School and Law Enforcement Efforts to Combat Cyberbullying. *Preventing School Failure: Alternative Education for Children and Youth* 55, 2 (2011), 79–87. arXiv:<https://doi.org/10.1080/1045988X.2011.539440> doi:10.1080/1045988X.2011.539440
- [75] Samuel Hardman Taylor, Dominic DiFranzo, Yoon Hyung Choi, Shruti Sannon, and Natalya N. Bazarova. 2019. Accountability and Empathy by Design: Encouraging Bystander Intervention to Cyberbullying on Social Media. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 118 (nov 2019), 26 pages. doi:10.1145/3359220
- [76] Daniela Teodorescu, Tiffany Cheng, Alona Fyshe, and Saif M Mohammad. 2023. Language and Mental Health: Measures of Emotion Dynamics from Text as Linguistic Biosocial Markers. *arXiv preprint arXiv:2310.17369* (2023).
- [77] Megan K Torkildson, Kate Starbird, and Cecilia Aragon. 2014. Analysis and visualization of sentiment and emotion on crisis tweets. In *Cooperative Design, Visualization, and Engineering: 11th International Conference, CDVE 2014, Seattle, WA, USA, September 14-17, 2014. Proceedings 11*. Springer, 64–67.
- [78] Esther Van Leeuwen, Daan Van Knippenberg, and Naomi Ellemers. 2003. Continuing and changing group identities: The effects of merging on social identification and ingroup bias. *Personality and social psychology bulletin* 29, 6 (2003), 679–690.
- [79] JOSEPH B. WALTHER. 1992. Interpersonal Effects in Computer-Mediated Interaction: A Relational Perspective. *Communication Research* 19, 1 (1992), 52–90. arXiv:<https://doi.org/10.1177/009365092019001003> doi:10.1177/009365092019001003

- [80] Joseph B. Walther. 1993. Impression development in computer-mediated interaction. *Western Journal of Communication* 57, 4 (1993), 381–398. arXiv:<https://doi.org/10.1080/10570319309374463> doi:10.1080/10570319309374463
- [81] Qiaosi Wang, Michael Madaio, Shaun Kane, Shivani Kapania, Michael Terry, and Lauren Wilcox. 2023. Designing Responsible AI: Adaptations of UX Practice to Meet Responsible AI Challenges. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 249, 16 pages. doi:10.1145/3544548.3581278
- [82] Mudasir Ahmad Wani, Mohammad ELAffendi, Kashish Ara Shakil, Ibrahim Mohammed Abuhaimed, Anand Nayyar, Amir Hussain, and Ahmed A. Abd El-Latif. 2023. Toxic Fake News Detection and Classification for Combating COVID-19 Misinformation. *IEEE Transactions on Computational Social Systems* (2023), 1–0. doi:10.1109/TCSS.2023.3276764
- [83] David Watson, Lee Anna Clark, and Auke Tellegen. 1988. Development and validation of brief measures of positive and negative affect: the PANAS scales. *Journal of personality and social psychology* 54, 6 (1988), 1063.
- [84] Donghee Yvette Wohn. 2019. Volunteer Moderators in Twitch Micro Communities: How They Get Involved, the Roles They Play, and the Emotional Labor They Experience. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3290605.3300390
- [85] Austin P. Wright, Omar Shaikh, Haekyu Park, Will Epperson, Muhammed Ahmed, Stephane Pinel, Duen Horng (Polo) Chau, and Diyi Yang. 2021. RECAST: Enabling User Recourse and Interpretability of Toxicity Detection Models with Interactive Visualization. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 181 (apr 2021), 26 pages. doi:10.1145/3449280
- [86] Dana Yagil and Hana Medler-Liraz. 2014. Feel Free, Be Yourself: Authentic Leadership, Emotional Expression, and Employee Authenticity. *Journal of Leadership & Organizational Studies* 21, 1 (2014), 59–70. arXiv:<https://doi.org/10.1177/1548051813483833> doi:10.1177/1548051813483833
- [87] Zihan Yan, Yufei Wu, Yang Zhang, and Xiang 'Anthony' Chen. 2022. EmoGlass: An End-to-End AI-Enabled Wearable Platform for Enhancing Self-Awareness of Emotional Health. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 13, 19 pages. doi:10.1145/3491102.3501925
- [88] Amy X Zhang and Scott Counts. 2016. Gender and ideology in the spread of anti-abortion policy. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 3378–3389.
- [89] Jian Zhao, Liang Gou, Fei Wang, and Michelle Zhou. 2014. PEARL: An interactive visual analytic tool for understanding personal emotion style derived from social media. In *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)*. 203–212. doi:10.1109/VAST.2014.7042496
- [90] John Zimmerman, Jodi Forlizzi, and Shelley Evenson. 2007. Research through design as a method for interaction design research in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '07). Association for Computing Machinery, New York, NY, USA, 493–502. doi:10.1145/1240624.1240704
- [91] Johannes Zimmermann, Timo Brockmeyer, Matthias Hunn, Henning Schauenburg, and Markus Wolf. 2017. First-person pronoun use in spoken language as a predictor of future depressive symptoms: Preliminary evidence from a clinical sample of depressed patients. *Clinical psychology & psychotherapy* 24, 2 (2017), 384–391.

A Demographics data of participants to the design workshop

Participant	Age	Gender	Occupation
P1	34	Female	UX Research
P2	32	Female	UX Designer
P3	28	Male	Product Manager
P4	28	Male	Product Designer
P5	26	Female	Graphic Designer
P6	28	Female	Graphic Designer
P7	27	Male	UX Researcher
P8	28	Female	UX Researcher
P9	31	Male	Product Designer
P10	28	Female	UX Researcher
P11	32	Female	Digital Marketer
P12	28	Male	Engineer
P13	27	Male	Engineer
P14	26	Male	UX Researcher
P15	28	Male	Graphic Designer

Table 7. Demographics data of fifteen participants to the design workshop for emotion monitoring dashboard.

B Emotional effects: Google T5 analysis

To corroborate our results in the emotional effects study, we also used Google's T5-base model [57] fine-tuned for emotion recognition⁸ to evaluate the original comments and reframed comments. Instead of detecting the percentage of emotions in each message as *text2emotion* did, this model classifies each message into only one emotion class. However, it is worth noticing that the model contains only four emotions: Angry, Fear, Sad, and Happy. Our findings, illustrated in Figure 11, show an increase in the number of messages that express Happy and a decrease in the number of messages that present negative emotions post-reframing.




Fig. 11. Comparison of the number of messages in four emotion classes (Angry, Fear, Sad, Happy) between original and rewritten comments detected by the fine-tuned T5 model from HuggingFace.

⁸<https://huggingface.co/mrm8488/t5-base-finetuned-emotion>

C Interface of Emotion Monitoring System.

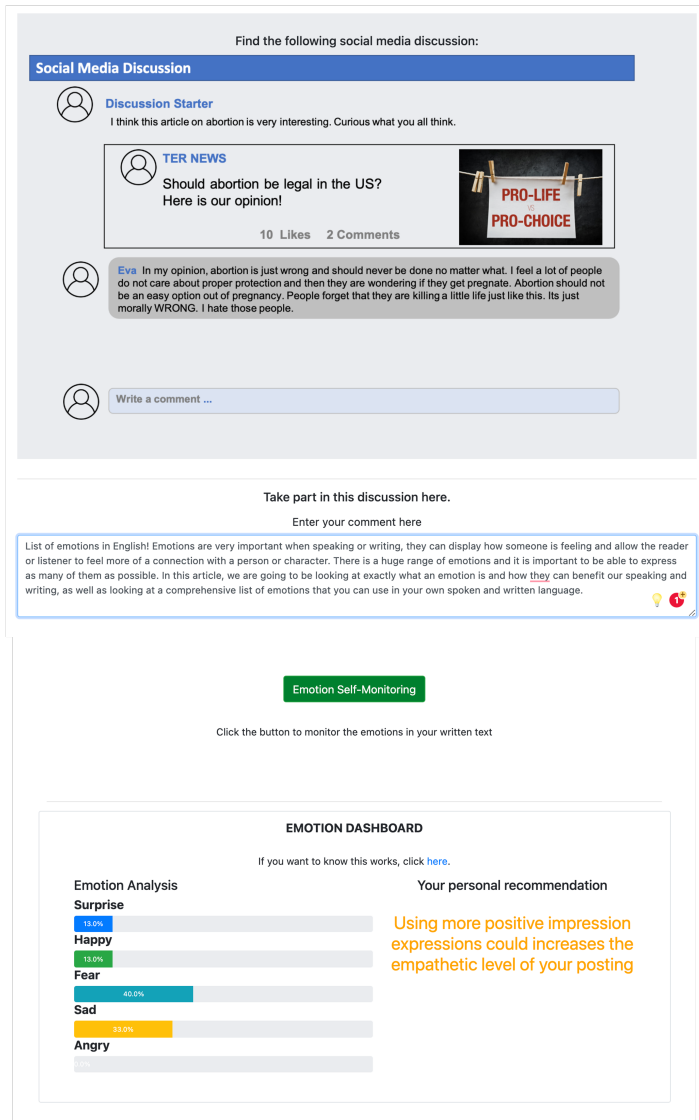


Fig. 12. Interface of Treatment Group 2 where users receive emotion self-monitoring. Participants are presented with a social media discussion containing the news and another user’s message, they are asked to write a comment to address the previous user’s view. Then they can click the "Emotion Self-Monitoring" button and receive emotional feedback on their own messages. The emotion dashboard presents the final result, it shows emotion analysis on five categories with a personal recommendation.

Received July 2024; revised December 2024; accepted March 2025.