



EChat: An Emotion-Aware Adaptive UI for a Messaging App

Radha Kumaran*
rkumaran@ucsb.edu

University of California Santa Barbara
Santa Barbara, California, USA

Viral Niraj Doshi*
viraldoshi789@gmail.com
IIITDM Jabalpur
Jabalpur, Madhya Pradesh, India

Sherry X. Chen
xchen774@ucsb.edu

University of California Santa Barbara
Santa Barbara, California, USA

Avinash Ajit Nargund
anargund@ucsb.edu

University of California Santa Barbara
Santa Barbara, California, USA

Tobias Höllerer
holl@cs.ucsb.edu

University of California Santa Barbara
Santa Barbara, California, USA

Misha Sra
sra@cs.ucsb.edu

University of California Santa Barbara
Santa Barbara, California, USA




Figure 1: The EChat Pipeline

ABSTRACT

While online forums provide a convenient platform for people to interact anonymously with others who share similar interests, they have to deal with large amounts of hate speech and inappropriate content, often posted by users in the heat of the moment. This can have a negative impact on the psychological state of other forum users and moderators, who are tasked to identify and delete such content. We investigate a preventative approach to this problem with the design of EChat, a proof-of-concept augmentation to online forums that helps users attend to their emotional state. The user's current emotional state is detected using facial emotion recognition, and the aesthetics of the UI are adapted to reflect this emotion. In case of an emotion with negative valence such as anger or sadness, the UI aesthetic is gradually transitioned to one that evokes a more positive emotion. Semi-structured interviews with EChat users confirm the potential of emotion-aware design to reduce hateful content, and also highlight important design considerations.

CCS CONCEPTS

• Human-centered computing → User interface design.

KEYWORDS

Adaptive Interfaces, Social Media, Facial Emotion Recognition

ACM Reference Format:

Radha Kumaran, Viral Niraj Doshi, Sherry X. Chen, Avinash Ajit Nargund, Tobias Höllerer, and Misha Sra. 2023. EChat: An Emotion-Aware Adaptive

*Both authors contributed equally to this research.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

UIST '23 Adjunct, October 29–November 01, 2023, San Francisco, CA, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0096-5/23/10.

<https://doi.org/10.1145/3586182.3616698>

UI for a Messaging App. In *The 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23 Adjunct), October 29–November 01, 2023, San Francisco, CA, USA*. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3586182.3616698>

1 INTRODUCTION

Social media has brought people closer and has had a widespread influence on how we communicate and engage not only with friends and family, but also strangers and others with shared interests [7, 25]. While the anonymity of discussions on community-centric platforms such as Reddit [11] offers many benefits [31], it is also correlated with increased amounts of hateful content posted on these forums [12, 17, 23, 29]. Currently, most forums use a combination of reactive and participatory solutions to the problem including: 1) automated identification and filtering out of hateful content, 2) reliance on users to flag inappropriate content; and/or 3) employing moderators to parse posts to remove hate speech, misinformation and inappropriate content [6, 10, 15]. Given the challenges with effective and fair automatic moderation of content on online platforms [4, 10, 30], many forums tend to favor human moderators for performing these time-intensive tasks. Content moderators who censor objectionable content have to go through many disturbing posts on a regular basis, making content moderation an intensely stressful job [5] [8]. Additionally, viewing inappropriate content regularly can have a negative impact on their mental health [9, 27].

In this work, we investigate a preventative approach to the problem, i.e., an attempt to prevent people from posting hateful content in the heat of the moment. We attempt this by regulating a user's emotions in real time through manipulations of the user interface, as they participate in conversations on the forum. Our work builds on prior work that has shown the aesthetics of user interface design to have an impact on the user's emotional response [3, 28, 32]. We present EChat (Figure 1), a proof-of-concept augmentation to online discussion forums that aims to put the reader into a more positive emotional state by altering visual elements of the user interface, demonstrated here in a one-on-one messaging app. In




Figure 2: Angry to Neutral Theme Transition

contrast to work that focuses on conveying the sender's emotions to the recipient of the message [1, 24], EChat enables the reader to attend to their emotions, which could help with emotion regulation [2] and by extension, cause them to pause and breathe before posting content. The idea of keeping emotions in check while using social media can be vital for reducing the amount of hurtful content posted and helping manage moderator stress levels. We did a pilot study through semi-structured interviews with three users. Feedback indicates potential of adaptive interfaces to reduce hateful content online, and also highlights some design considerations.

2 SYSTEM DESIGN AND PILOT STUDY

EChat adaptively changes the aesthetics of a messaging interface in four stages as described in this section.

Initial Emotion Detection. Images captured from the device's front-facing camera at regular intervals are used as input to Deepface [20]. Although Deepface outputs one of six emotion classes, for the purpose of our system we only consider images that belong to any of four classes - angry, happy, sad and neutral. The emotion with the highest frequency among recent images is considered to be the user's current emotional state.

UI Adaptation to Reflect Emotion. Prior work has shown that users' emotions are impacted by colors [19, 33] and shapes [13, 14, 18, 26] of the interface they interact with. We use this in our design of the EChat themes corresponding to each emotion, with more rounded chat boxes for emotions with lower arousal, and colors for each theme chosen based on Plutchik's mapping[21].

User Interface Transition. If the user's detected emotion shows negative valence [22] (e.g., anger or sadness), the aesthetics of the interface are gradually transitioned to the target emotion (neutral) over 60 seconds (Figure 2), during which time user emotion is recorded but not acted upon. For the purposes of this poster, which focuses on negativity reduction, we currently only transition from negative emotion to neutral.

Emotional Feedback. On completion of transition to the target theme, the adaptation is considered successful if the user's emotion matches the target emotion.

In the pilot study, three users provided feedback on our system in an exploratory investigation. They experienced a conversation taken from the DailyDialog [16] dataset in two interfaces - the EChat interface, and a similar interface which did not change based




Figure 3: Average self-reported ratings for anger(left) and sadness(right) over time, by 3 users.

on emotion (i.e., stayed in the neutral theme for the entire conversation). The chosen conversation had one of the conversation partners feeling angry during the conversation, and pilot users were asked to embody this person's emotional state. This approach was used to ensure consistency in the emotions induced, since free-form conversations and their emotion elicitation would have been more difficult to control. After using both interfaces, semi-structured interviews were conducted to obtain feedback on the concept of theme transitions based on emotion, as well as the EChat system.

3 DISCUSSION AND FUTURE WORK

All users indicated that EChat changed their awareness of their emotional state and how they would respond in a conversation, because the interface aesthetic influenced their emotional state ("The EChat interface felt more empathetic because of the color changes, and that is more likely to calm me down faster, so my responses would probably be more friendly too", U3). Self-reports of emotion over time also indicate a similar trend (Figure 3). Two users preferred EChat to the static interface, because "The color of the user interface can represent my emotion", (U1), and "It felt more personal, more human and less robotic", (U3). This suggests that aesthetic changes have the potential to influence emotional state, and that emotion-aware and responsive interface design could help improve user experience and also prevent strongly negative content on social media. Given the recent rise of LLMs potentially trained on hateful speech content from public forums, we believe EChat offers an option to help improve the overall quality of training data for LLMs.

There are, however, important considerations in this design paradigm. First, an emotionally responsive design lowers user agency, which could alter overall user experience ("I didn't like the color changing automatically without control (over the adaptation)", U2). The agency-intelligence trade-off is a concern in many intelligent systems, and a possible solution is to let users decide the level of agency they want EChat to have. Next, with regards to the aesthetic adaptation of interfaces, some people might associate certain colors/aesthetics with specific applications ("I associate specific colors with specific (social media) apps, changed colors (within the same app) is confusing", U2). Making aesthetic choices for different emotional themes that conform to the application's main theme could reduce this dissonance. Future studies will be needed to determine the optimal duration of each theme transition, as well as the emotional impact of changes in the UI when users communicate their own emotional state rather than just embodying emotional states.

REFERENCES

[1] Jussi Ängeslevä, Carson Reynolds, and Sile O'Modhrain. 2004. EmoteMail. In *ACM SIGGRAPH 2004 Posters*. 9.

[2] Lisa Feldman Barrett and James J Gross. 2001. Emotional intelligence: A process model of emotion representation and regulation. (2001).

[3] Upasna Bhandari, Klarissa Chang, and Tillmann Neben. 2019. Understanding the impact of perceived visual aesthetics on user evaluations: An emotional perspective. *Information & management* 56, 1 (2019), 85–93.

[4] Reuben Binns, Michael Veale, Max Van Kleek, and Nigel Shadbolt. 2017. Like trainer, like bot? Inheritance of bias in algorithmic content moderation. In *International conference on social informatics*. Springer, 405–415.

[5] Iris Birman. 2018. Moderation in different communities on Reddit—A qualitative analysis study. (2018).

[6] Svenja Boberg, Tim Schatto-Eckrodt, Lena Frischlich, and Thorsten Quandt. 2018. The moral gatekeeper? Moderation and deletion of user-generated content in a leading news forum. *Media and Communication* 6, 4 (2018), 58–69.

[7] Moira Burke and Robert E Kraut. 2014. Growing closer on Facebook: Changes in tie strength through social network site use. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 4187–4196.

[8] Eshwar Chandrasekharan, Chaitrali Gandhi, Matthew Wortley Mustelier, and Eric Gilbert. 2019. Crossmod: A cross-community learning-based system to assist reddit moderators. *Proceedings of the ACM on human-computer interaction* 3, CSCW (2019), 1–30.

[9] Bryan Dosono and Bryan Semaan. 2019. Moderation practices as emotional labor in sustaining online communities: The case of AAPI identity work on Reddit. In *Proceedings of the 2019 CHI conference on human factors in computing systems*. 1–13.

[10] Robert Gorwa, Reuben Binns, and Christian Katzenbach. 2020. Algorithmic content moderation: Technical and political challenges in the automation of platform governance. *Big Data & Society* 7, 1 (2020), 2053951719897945.

[11] Miriam Greenberg. 2019. A Study of Community Engagement on Discord for Game Marketing—Case TJR Games Oy. (2019).

[12] Michelle CP Hendriks and Ad JJM Vingerhoets. 2006. Social messages of crying faces: Their influence on anticipated person perception, emotions and behavioural responses. *Cognition and Emotion* 20, 6 (2006), 878–886.

[13] Jesús Ibáñez. 2011. Minimalist approach to show emotions via a flock of smileys. *Journal of network and computer applications* 34, 4 (2011), 1283–1291.

[14] Jesús Ibáñez and Carlos Delgado-Mata. 2013. Can the same Visual Modality Express Arousal or Valence Depending on the other Modalities it is Combined with? *Procedia Technology* 7 (2013), 424–435.

[15] Yubo Kou and Xinning Gui. 2021. Flag and Flaggability in Automated Moderation: The Case of Reporting Toxic Behavior in an Online Game Community. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–12.

[16] Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. Dailydialog: A manually labelled multi-turn dialogue dataset. *arXiv preprint arXiv:1710.03957* (2017).

[17] Mainack Mondal, Leandro Araújo Silva, and Fabrício Benevenuto. 2017. A measurement study of hate speech in social media. In *Proceedings of the 28th ACM conference on hypertext and social media*. 85–94.

[18] Erik C Nook, Kristen A Lindquist, and Jamil Zaki. 2015. A new look at emotion perception: Concepts speed and shape facial emotion recognition. *Emotion* 15, 5 (2015), 569.

[19] Leonhard Obersacher and Michael Gallmetzer. 2003. Colour and emotion. In *Proceedings of AIC*. 370–374.

[20] Omkar M Parkhi, Andrea Vedaldi, and Andrew Zisserman. 2015. Deep face recognition. (2015).

[21] Robert Plutchik. 2001. The nature of emotions: Human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice. *American scientist* 89, 4 (2001), 344–350.

[22] James A Russell. 1980. A circumplex model of affect. *Journal of personality and social psychology* 39, 6 (1980), 1161.

[23] Koustuv Saha, Eshwar Chandrasekharan, and Munmun De Choudhury. 2019. Prevalence and psychological effects of hateful speech in online college communities. In *Proceedings of the 10th ACM conference on web science*. 255–264.

[24] J Alfredo Sánchez, Norma P Hernández, Julio C Penagos, and Yulia Ostróvskaya. 2006. Conveying mood and emotion in instant messaging by using a two-dimensional model for affective states. In *Proceedings of VII Brazilian symposium on Human factors in computing systems*. 66–72.

[25] Shabnoor Siddiqui, Tajinder Singh, et al. 2016. Social media its impact with positive and negative aspects. *International journal of computer applications technology and research* 5, 2 (2016), 71–75.

[26] Paul J Silvia and Christopher M Barona. 2009. Do people prefer curved objects? Angularity, expertise, and aesthetic preference. *Empirical studies of the arts* 27, 1 (2009), 25–42.

[27] Miriah Steiger, Timir J Bharucha, Sukrit Venkatagiri, Martin J Riedl, and Matthew Lease. 2021. The psychological well-being of content moderators: the emotional labor of commercial moderation and avenues for improving support. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. 1–14.

[28] Alistair Sutcliffe. 2009. Designing for user engagement: Aesthetic and attractive user interfaces. *Synthesis lectures on human-centered informatics* 2, 1 (2009), 1–55.

[29] Madeline Touibia and Charlene Zietsma. 2017. The message is on the wall? Emotions, social media and the dynamics of institutional complexity. *Academy of Management Journal* 60, 3 (2017), 922–953.

[30] Kristen Vaccaro, Christian Sandvig, and Karrie Karahalios. 2020. "At the End of the Day Facebook Does What It Wants" How Users Experience Contesting Algorithmic Content Moderation. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–22.

[31] Emily Van der Nagel and Jordan Frith. 2015. Anonymity, pseudonymity, and the agency of online identity: Examining the social practices of r/Gonewild. *First Monday* (2015).

[32] Jian Wang and Yen Hsu. 2020. The relationship of symmetry, complexity, and shape in mobile interface aesthetics, from an emotional perspective—A case study of the smartwatch. *Symmetry* 12, 9 (2020), 1403.

[33] Sandy Wolfson and Gill Case. 2000. The effects of sound and colour on responses to a computer game. *Interacting with computers* 13, 2 (2000), 183–192.