

MIT Open Access Articles

*Affective Typography: The Effect of AI-Driven
Font Design on Empathetic Story Reading*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Shen, Jocelyn, Jin, Kathryn, Zhang, Ann, Breazeal, Cynthia and Park, Hae Won. 2023. "Affective Typography: The Effect of AI-Driven Font Design on Empathetic Story Reading."

As Published: <https://doi.org/10.1145/3544549.3585625>

Publisher: ACM|Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems

Persistent URL: <https://hdl.handle.net/1721.1/150644>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Affective Typography: The Effect of AI-Driven Font Design on Empathetic Story Reading

Jocelyn Shen
joceshen@media.mit.edu
MIT Media Lab
Cambridge, MA, USA

Kathryn Jin*
kjin@mit.edu
MIT CSAIL
Cambridge, MA, USA

Ann Zhang*
annzhang@mit.edu
MIT CSAIL
Cambridge, MA, USA

Cynthia Breazeal
cynthiab@media.mit.edu
MIT Media Lab
Cambridge, MA, USA

Hae Won Park
haewon@media.mit.edu
MIT Media Lab
Cambridge, MA, USA

ABSTRACT

When people listen to stories, the words are colored with emotions conveyed through prosodic features beyond the text alone. Visual font design provides an opportunity to enhance the empathic quality of a story compared to plain text. In this paper, we present the design, implementation, and evaluation of Affective Typography (AFFTYPE), an AI-driven system that extracts prosodic information and sentiment from speech and maps these properties to typographic styles.¹ We conduct a crowdsourced study ($N = 140$) to assess how different font design elements impact readers' empathy with personal stories. While our empathy survey results were not statistically significant, we found that participants had a preference for color to express emotion and saw an increase in average empathy for stories with color-based text alterations. In addition, we offer design insights as to what display features best convey emotional qualities of personal stories for future applications that use affective fonts to create more expressive digital text.

CCS CONCEPTS

• **Human-centered computing** → **User interface design; Empirical studies in HCI.**

KEYWORDS

typography, voice input, speech representation, emotional communication, empathy

ACM Reference Format:

Jocelyn Shen, Kathryn Jin, Ann Zhang, Cynthia Breazeal, and Hae Won Park. 2023. Affective Typography: The Effect of AI-Driven Font Design on Empathetic Story Reading. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*, April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3544549.3585625>

*Both authors contributed equally to this research.

¹<https://github.com/jocelynshen/affective-text-interface>

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).

CHI EA '23, April 23–28, 2023, Hamburg, Germany

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

ACM ISBN 978-1-4503-9422-2/23/04.

<https://doi.org/10.1145/3544549.3585625>

1 INTRODUCTION

Many of the most meaningful human-human connections are fostered through sharing of personal stories and empathizing with emotional experiences [11, 21, 27]. Although people consume pages of digital content everyday, digitized text can reduce the true empathic impact of a story to a stale, desensitized version if design of the text is not carefully considered. In contrast, speech conveys information beyond words alone – there is emotional richness in the way people tell stories. Representing these prosodic features as visual cues could enhance empathy in reading experiences by injecting the human quality back into the text. In addition, such a system could help convey emotion in the voice to people with hearing impairments or trouble with emotion identification.

In this paper, we specifically answer the question: What design elements should affective fonts have in order to visually convey emotion and improve empathy for spoken stories? We present the design and evaluation of Affective Typography (AFFTYPE), an AI-driven system that converts prosodic and sentiment-related information present in speech into typographic properties and emojis. Prior works have explored how to convey speech features through text in order to improve accuracy of the perceived emotion [2, 6]. However, we are specifically interested in how affective typography impacts a reader's ability to empathize with a storyteller and create more expressive text, not just accuracy of the conveyed emotion. Our system implementation leverages deep learning and lexicon-based sentiment models, as well as speech feature extraction pipelines to generate affective text. In addition, our system is easily transferable to digital text, as we use boldness, spacing, color, and emojis, which are available across most front-end applications and web browsers. Through a crowdsourced study ($N = 140$), we evaluate our system's ability to improve empathy on personal stories with different emotional arcs. We do not find statistically significant results across font conditions, but we do identify a preference for color to express emotions. Our preliminary findings motivate more extensive work on using intelligent visual design to convey affective information from modalities that may be less accessible in digital settings and which can enhance the emotions in text.

2 RELATED WORK

Stories have the ability to elicit empathy and help us connect with others [1, 4, 10]. In fact, some research has shown that when telling a story to a second listener, speakers and listeners couple their brain activity, indicating the neurological underpinnings of story sharing [12, 24]. Our work is interested in how AI-driven fonts can be used to convey emotional information present in spoken stories, potentially improving salience and emphasis of a storyteller’s emotions, and ultimately, empathy towards the narrator.

Within the human-computer interaction community, many existing interfaces have explored how to creatively convey emotion information to users. Example interfaces include intelligent mirrors that generate emotionally-relevant poetry [20], personalized animated movies for self reflection [19], and visual user interfaces as prosthetics to improve emotional memory [17]. We focus on using fonts as our visual communication tool, since text is low resource and easy to transfer across a wide variety of applications.

Designers understand the importance of typography in conveying meaning [5, 14], and using computational techniques to render intelligent typography has been explored by prior research. These works used acoustical features of speech such as loudness, pitch, and speed, which were represented by corresponding variations in font, and found that changing the appearance of text over time through kinetic typography enabled users to add emotional content to texts during instant messaging [3, 16, 25]. However, most of these early works were not automatically voice-driven, as the text effects were manually applied, and more recent advances in affective computing allow for advanced computer systems that better capture emotions present in speech and text. Utilizing these methods to automatically detect features related to emotion can be leveraged to create more expressive digital fonts.

Most relevant to our work are prior works that have explored how to bring text to life through speech modulated typography and voice-driven type design, as well as understanding what qualities of font better convey information from speech [7, 26]. For example, more recent works used automatically generated chat balloons for texts shared during instant messaging [2, 6]. These designs show promising results towards using automated systems to support personal expression and emotion understanding through visually conveying information such as arousal and valence. Our work builds on existing work by incorporating both prosodic features and sentiment analysis of the user’s recorded speech, which we automatically translate into variations in typographic design. We combine methods used in existing interfaces and convey affective information through font weight (loudness), font spacing (pace), color (positive/negative sentiment), and emojis (general sentiment). In contrast to prior works, we focus on measuring the effect of font design elements on readers’ empathy, not just accuracy of the perceived emotion.

3 SYSTEM DESIGN

Our system takes in speech, extracts text and prosodic information, and then renders the text with a font conveying the predicted emotion and speech features (Fig. 1). Sentiment is represented by font color and emojis, and prosodic features, such as loudness and pace, are represented by letter boldness and letter spacing. We chose

these mappings based on prior research in voice driven and kinetic typography [7, 26], as well as how people can intuitively interpret louder speech as more bold text, and letter spacing naturally maps to speech pacing. Additionally, prior work demonstrates that people associate certain colors with different types of emotions, such as green, which evokes positive emotions [15]. An example output of a story is shown in Fig. 2.




Figure 1: System design of AFFTYPE. Our system uses prosody and text to generate visual cues including emoji, sentiment, loudness, and pace.

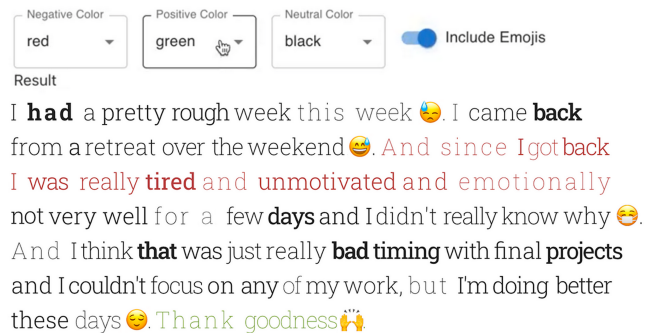


Figure 2: Sample output from the system

The story audio file is passed to the Google Cloud Speech API for automatic speech recognition (ASR) in order to extract the text, as well as timestamps for the approximate start and end times of each individual spoken word. We then extract the prosodic features of the speech using OpenSmile [8]. Specifically, we use the loudness feature, where values are normalized over the entire audio clip using min-max normalization. We then calculate a loudness factor for each word by merging the timestamps from the transcript and the OpenSmile outputs. This factor is used to determine font weight. Next, the text transcript is passed into two sentiment models: DeepMoji [9], which selects the top emoji for the emotion of a given sentence, and VADER [13], which we use to choose text color based on whether the sentence is positive, negative, or neutral. Originally, we used speech affect recognition to map these sentences to colors based on the six basic emotions, but found that this made the text too difficult to read. Font weight and spacing are applied at the word level, and color and emoji are applied at the sentence level in order to better contextualize the emotion. Note that we have redundancy in how sentiment is conveyed, as we wanted our system to offer multiple sources of interpretations of emotion, whether that be through emojis or colors.

We synchronize the text and audio using timestamps outputted by ASR and calculate a proxy for speaker pace by taking the number of letters divided by the length of the time interval. While this approximation is not as accurate as using the number of syllables divided by the length of the time interval, since syllables are more closely related to speaker pace than number of letters, we found that this approximation was good enough for capturing the information needed to display pace in an intelligible way. Based on the model outputs from the server, we render the final font using CSS properties. Our system is fully automated, and we open source our code to aid further research in this area. In addition, we implemented a frontend web interface in React where users can upload audio stories to our Flask server and see their responses in real time.

4 EXPERIMENT

To evaluate the effect of our generated fonts on empathy during story reading, we conducted a crowdsourced study with $N = 140$ participants through Prolific [18]. Our study procedure was approved by our university's Institutional Review Board committee as an exempt study. Participants were predominantly between 18-24 years of age (32, 27 age 25-29, 22 age 55+, 20 age 30-34, 15 age 35-39, 11 age 40-45, 7 age 50-55, and 6 age 45-49), predominantly women (88 women, 46 men, 4 non-binary, 1 transgender woman, and 1 transgender man), and predominantly white (102, 11 Hispanic, 11 Black, 6 Asian, 2 Middle Eastern, 8 other).

At the beginning of the study, we asked for demographic information of the reader, including age range, gender, and ethnicity. We then randomized participants into one of seven conditions (20 participants per condition): (1) control/regular text (abbr. regular) (2) font conveying loudness (abbr. bold) (3) font conveying pace (abbr. spacing) (4) font conveying sentiment with emojis (abbr. emoji) (5) font conveying sentiment with color (abbr. color), (6) font conveying loudness, pace, and sentiment with color (abbr. bold + spacing + color), and (7) font conveying loudness, pace, and sentiment with both emojis and color (abbr. all features). We hypothesized that compared to the regular text condition, participants' empathic responses would increase in all other conditions. Participants in each condition were asked to read two different stories that were run through AFFTYPE and rate the extent to which they empathized with the story as well as answer a free response question about what they empathized with in the story. To assess empathy towards the narrator, we used the 12-item State Empathy Scale [22], which takes into account the affective, cognitive, and associative components of empathy when receiving messages. The stories we ran through the interface were chosen from StoryCorps,² a site containing short recordings of personal stories. Our selected stories are included in the Appendix. We selected two stories with different emotional trajectories to control for variations in empathy based on the overall emotional tone of the story. Note that for each participant, we randomized the order the stories were presented in to control for ordering effects, but the two stories were rendered with the same features. At the end of the task, we asked participants to rate how well the font properties corresponded to speech features, as well as free responses on their likes, dislikes, and improvements for the way the text was presented.

²<https://storycorps.org/>

5 RESULTS

5.1 Effects on Empathy

We first analyze results for the State Empathy Scale [22] across conditions. To calculate p-values, we use a one-sided Mann Whitney U-test, as we identify through a Shapiro-Wilk test across conditions that the data distribution is not normal. We determine statistical significance using a p-value threshold of 0.0083, adjusted using Bonferroni correction for six comparisons (control/regular condition compared to each six experimental condition). Note that reported p-values are relative to the control/regular text condition. The experimental conditions, bold (mean = 2.72, std = 0.80, $p = 0.11$), spacing (mean = 2.72, std = 0.68, $p = 0.071$), color (mean = 2.81, std = 0.82, $p = 0.014$), emoji (mean = 2.73, std = 0.68, $p = 0.043$), and bold + spacing + color (mean = 2.57, std = 0.74, $p = 0.31$), show average increases in empathy, although not significantly to when compared to the regular text condition (mean = 2.47, std = 0.75). Interestingly, the only condition where participants' empathy decreased relative to the regular text condition was the all features condition (mean = 2.18, std = 1.28, $p = 0.71$), and this last condition had the greatest standard deviation in empathy scores. From looking at our qualitative data, we hypothesize that this is because participants found the combination of all features jarring and distracted from the underlying emotional meaning of the story. As shown in Fig. 3, participants in the color condition had the greatest increase in average empathy over the regular text condition, followed by emoji, spacing, bold, bold + spacing + color, and all features.

Looking at psychometric survey data alone offers us one view on how the font elements affected empathy with the stories. In addition, we look at what participants say in their free responses to what they empathized with in the stories. To analyze this, we use dimensions from LIWC (linguistic inquiry and word count) [23]. In particular, we look at the total count of emotional language (emo_pos + emo_neg) used in free responses across conditions (Fig. 4). We find that when compared to the regular text condition (mean = 1.45, std = 2.49), participants in the color condition use the most emotional language on average (mean = 2.73, std = 3.18, $p = 0.017$), followed by all features (mean = 1.8, std = 3.0, $p = 0.37$), emoji (mean = 1.77, std = 3.23, $p = 0.54$), spacing (mean = 1.66, std = 2.82, $p = 0.41$), and bold (mean = 1.5, std = 3.83, $p = 0.74$). Again, although the differences are not statistically significant, participants in the color condition used, on average, the most emotional language relative to the regular text condition. We hypothesize that this could be due to the fact that color draws attention to the emotions present in the text. To validate this, we also asked participants if they felt that the font design helped them perceive emotions in the text. We found that, consistent with the LIWC results, participants in the color condition reported the highest average agreement with this statement (mean = 1.95, std = 0.94) when compared to the regular text condition (mean = 1.5, std = 0.93), although not statistically significantly so ($p = 0.03$).

While we found that the participants in the all features condition were distracted by too many changes in the font, many participants in this condition wrote meaningful responses to the empathy free response survey question. For example, one participant self-disclosed, *"I live in a bordertown and I often think about my grandparents coming to America. They fled the pogroms in Russia. I*

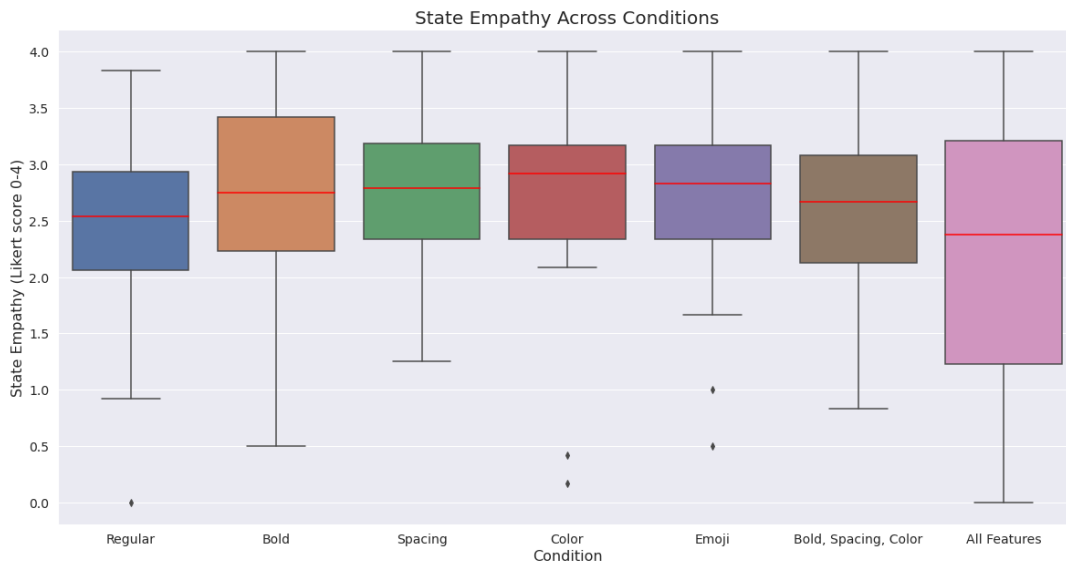


Figure 3: Empathy scores across conditions, where scores range from 0 to 4. Red line indicates the median.

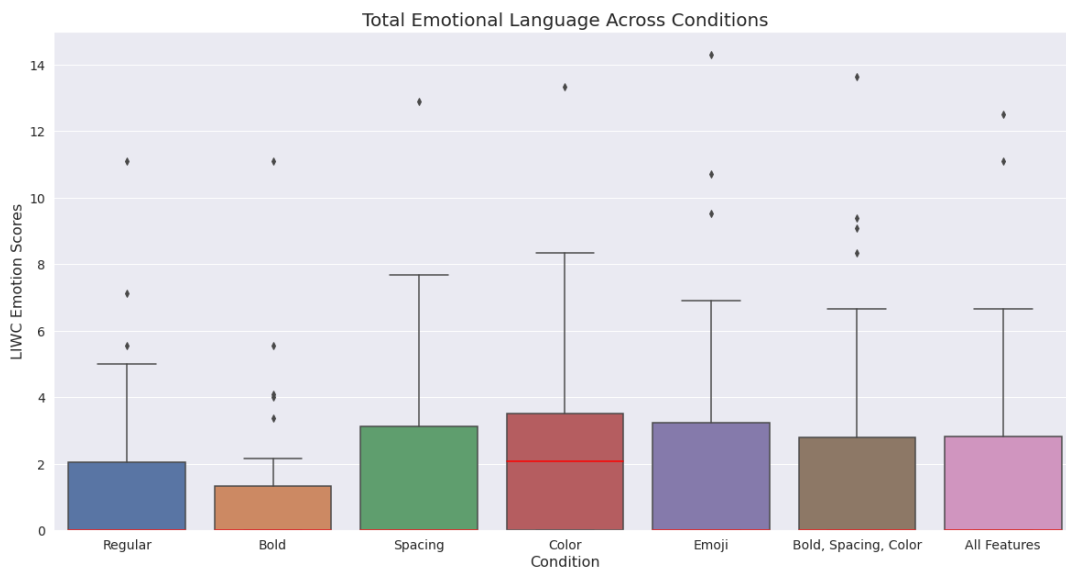


Figure 4: Emotional language participants used in free response to "What did you empathize with in the story?" across conditions. Red indicates the median.

know we have people fleeing their homeland for various humanitarian reasons. I worry about how unwelcoming we have become. I do not know what the solution is and how we can actually help other's be able to stay safely in their homelands." Another wrote, "I somewhat empathized with the feelings of guilt over escaping a bad situation. This story tangentially reminded me of my experience as an LGBTQ+ individual and how although I've experienced oppression and hate, others in the community have experienced it to a much harsher extent." Future work can further explore this relationship between the

font properties and self disclosure when empathizing with another person's story.

5.2 Design Considerations

We hypothesize that some challenges in conveying emotion through AFFTYPE are limitations in the speech to font mappings. At the end of the study, we asked each participant to rate their agreement with how effective each font design change (bold, spacing, color, or emoji) was in capturing the intended speech characteristic (loudness, pace, positive/negative sentiment, or general sentiment). Overall, we

found that participants were, on average, neutral to these mappings, with a slight preference towards boldness (mean = 2.19, std = 1.22), followed by color (mean = 1.94, std = 1.31), emoji (mean = 1.86, std = 1.28), and spacing (mean = 1.82, std = 1.24). Although the alterations we made to the text were motivated by prior works [7, 26], in our application, the effectiveness of these mappings could be improved. In the rest of the section, we provide insight from participants' responses on their likes, dislikes, and suggestions for what font design elements could improve empathy with personal stories.

For each of the following analyses, we used qualitative coding to identify core themes in participants' responses. Three researchers independently coded the survey responses, and commonalities were extracted as major themes. As shown in Table 1, participants liked the more natural and human quality to the text, commenting on how "it felt like someone was speaking to me" and "it looked more personal." Some participants preferred spacing for matching the pace of the story and readily understood that color was associated with emotion. Other participants found boldness helpful in drawing emphasis to specific points. From Table 2, we see that participants disliked the way the font interrupted the flow of the story and commented that there was a lack of correlation between the style and meaning of the story due to too many text alterations. In addition, participants commented that emojis were not effective in promoting empathy, as they affected how the writing style was perceived and made it appear more childish.

Theme	Example
Styles gave the text more personality	"It looked more personal, like handwriting, since it doesn't look like the typical typed text"
Styles conveyed spoken elements to some participants	"I think it felt like someone was speaking to me. Completely telling the story from their view."
Spacing sometimes helped users match the pacing of the story	"I think it is helpful for some to feel the pauses in the text as spaces that were placed."
Color was easily understood to emphasize emotion	"Colored text I read more passionately and felt more emotion from it"
The style helped draw attention to certain parts of the story	"The boldness of text was helpful for emphasizing points."

Table 1: What participants liked about AFFTYPE with respect to empathetic story reading

Finally, participants expressed what they wished was different about the way the text was displayed in order to increase empathy with the story. As shown in Table 3, a major theme was using standard writing to convey emotion in the story with minimal unnatural text alterations. For example, participants suggested including explanations of how something was spoken, such as the tone they used, if they sighed, or what their facial expression was. Others suggested using common text formatting like italics and paragraphs, indicating that seamless integration of the narrators

Theme	Example
Styles visually interrupted the flow of the story	"Not that it was necessarily a hinderance, but some of the spacing was different which made me think about that rather than the story."
Emojis affected how the writing style was perceived	"The emojis made it feel a but silly and not serious, like I was reading a facebook post."
There was sometimes a lack of correlation between style and meaning	"...emphasizing all kinds of words in the story, not just the ones that made sense for the emotional impact."

Table 2: What participants disliked about AFFTYPE with respect to empathetic story reading.

spoken emotions into the text is an important property of the system. Finally, participants suggested using photos instead of emojis to augment the story and preserve the formal quality of the writing as well as using the text to contextualize the narrator's experiences better.

Theme	Example
Explanations of how something was spoken	Adding in explanations of tone, pauses, sighs, and facial expressions.
Use of other text formatting like italics, indentation, and paragraphs	"I would break up the stories into multiple paragraphs, and I would use italics to emphasize important points."
Use photos to augment the story	"I would maybe add a photo/image to help people visualize the person who is telling the story in some way"
Contextualizing the story	"more information about the author, and information about the setting"

Table 3: What participants would change about AFFTYPE with respect to empathetic story reading.

Based on participant feedback and survey results, we summarize the following design insights for AI-driven empathetic fonts: (1) readable – alterations in text should not distract from clarity of the story, (2) natural – speech to font mappings should be intuitive, (3) colorful – colors represent emotions well, (4) appropriate – alterations to text should not affect how the writing style is perceived (eg. emojis make writing more informal), (5) explainable – speech characteristics should be explained directly by the text, and (6) personalized and culturally sensitive – use of features could be interpreted differently across people and cultures.

6 CONCLUSION AND FUTURE WORK

Our work expands on interfaces that can better convey human qualities through computational means, without the reduction that

often occurs when human stories are digitized. While we did not prove our original hypotheses, we found that participants preferred colored text for empathetic story reading and desired more features such as explanations of how something was spoken and improved readability.

There are limitations to our interface, which we identified through our user study. In particular, it is possible that long-term exposure to empathetic text could lead to fatigue or loss of sensitivity to the empathy the text is intended to foster. Our current design does not take this into account. Furthermore, the limitations of our user study are the short duration and scale of the study. If participants had interacted with the system over a longer period of time, perhaps some of the text features would not have been so jarring, or they might have become desensitized to the font changes. In addition, we only asked participants to read two hand-curated stories. Using stories with more diverse emotional trajectories could also have improved the user understanding of how the font features correlate with emotions present in the narrator's voice. Finally, we did not look at how different demographic characteristics affect the usability of the system. For example, younger people might be more receptive to features like emojis.

Based on participant feedback, there is future work that could improve the capabilities of our system in expressing emotions and fostering empathy in spoken stories. In particular, the text can be integrated with spoken emotions in a more natural and seamless manner. For example, one participant commented that the way something is written is important for empathizing with a piece of text. One idea could be to use prompting methods with large language models to generate explanations of the way something is spoken. Therefore, instead of using uncommon visual elements like letter spacing or boldness, the language itself could bring life to the emotions and contextualize the narrator's experiences.

For future applications, it would be interesting to explore how this work could be used in video captioning to generate more empathic captions. Additionally, this system could be used to help individuals reflect on emotions present in stories and help them notice communication patterns. For example, one could easily look at the affective text and see if their words are laced with negativity or if their voice became quieter at the end of sentences due to lack of confidence. Rendering text in this way can be a creative means for users to engage with personal, spoken stories. Further work can explore using the system to convey elements of speech to people with hearing impairments, trouble with emotion identification, and more generally, in storytelling applications to improve empathy and human connection.

ACKNOWLEDGMENTS

We would like to thank our participants and all of our teammates who have contributed constructive comments to our project. This work was supported by an NSF GRFP under Grant No. 2141064.

REFERENCES

- [1] Mary E. Andrews, Bradley D. Mattan, Keana Richards, Samantha L. Moore-Berg, and Emily B. Falk. 2022. Using first-person narratives about healthcare workers and people who are incarcerated to motivate helping behaviors during the COVID-19 pandemic. 299 (2022), 114870. <https://doi.org/10.1016/j.socscimed.2022.114870>
- [2] Toshiki Aoki, Rintaro Chujo, Katsufumi Matsui, Saemi Choi, and Ari Hautasaari. 2022. EmoBalloons - Conveying Emotional Arousal in Text Chats with Speech
- [3] Balloons. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–16. <https://doi.org/10.1145/3491102.3501920>
- [4] Kerry Bodine and Mathilde Pignol. 2003. Kinetic Typography-Based Instant Messaging. In *CHI '03 Extended Abstracts on Human Factors in Computing Systems (CHI EA '03)*. Association for Computing Machinery, New York, NY, USA, 914–915. <https://doi.org/10.1145/765891.766067>
- [5] Guilherme Brockington, Ana Paula Gomes Moreira, Maria Stephani Buso, Sérgio Gomes da Silva, Edgar Altszyler, Ronald Fischer, and Jorge Moll. 2021. Storytelling Increases Oxytocin and Positive Emotions and Decreases Cortisol and Pain in Hospitalized Children. *Proceedings of the National Academy of Sciences* 118, 22 (June 2021), e2018409118. <https://doi.org/10.1073/pnas.2018409118>
- [6] Heloisa Candello, Claudio Pinhanez, and Flavio Figueiredo. 2017. Typefaces and the Perception of Humanness in Natural Language Chatbots. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 3476–3487. <https://doi.org/10.1145/3025453.3025919>
- [7] Qinyue Chen, Yuchun Yan, and Hyeon-Jeong Suk. 2021. Bubble Coloring to Visualize the Speech Emotion. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–6. <https://doi.org/10.1145/3411763.3451698>
- [8] Caluã de Lacerda Pataca and Paula Dornhofer Paro Costa. 2020. Speech Modulated Typography: Towards an Affective Representation Model. In *Proceedings of the 25th International Conference on Intelligent User Interfaces*. ACM, Cagliari Italy, 139–143. <https://doi.org/10.1145/3377325.3377526>
- [9] Florian Eyben, Martin Wöllmer, and Björn Schuller. 2010. Opensmile: The Munich Versatile and Fast Open-Source Audio Feature Extractor. In *Proceedings of the 18th ACM International Conference on Multimedia (MM '10)*. Association for Computing Machinery, New York, NY, USA, 1459–1462. <https://doi.org/10.1145/1873951.1874246>
- [10] Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad Rahwan, and Sune Lehmann. 2017. Using Millions of Emoji Occurrences to Learn Any-Domain Representations for Detecting Sentiment, Emotion and Sarcasm. *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (2017), 1615–1625. <https://doi.org/10.18653/v1/D17-1169> arXiv:1708.00524
- [11] Melanie C. Green and Timothy C. Brock. 2000. The Role of Transportation in the Persuasiveness of Public Narratives. *Journal of Personality and Social Psychology* 79, 5 (Nov. 2000), 701–721. <https://doi.org/10.1037/0022-3514.79.5.701>
- [12] Sara D. Hodges, Kristi J. Kiel, Adam D. I. Kramer, Darya Veach, and B. Renee Villanueva. 2010. Giving Birth to Empathy: The Effects of Similar Experience on Empathic Accuracy, Empathic Concern, and Perceived Empathy. *Personality and Social Psychology Bulletin* 36, 3 (March 2010), 398–409. <https://doi.org/10.1177/0146167209350326>
- [13] Christopher J. Honey, Christopher R. Thompson, Yulia Lerner, and Uri Hasson. 2012. Not Lost in Translation: Neural Responses Shared Across Languages. *The Journal of Neuroscience* 32, 44 (Oct. 2012), 15277–15283. <https://doi.org/10.1523/JNEUROSCI.1800-12.2012>
- [14] C. Hutto and Eric Gilbert. 2014. VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text. *Proceedings of the International AAAI Conference on Web and Social Media* 8, 1 (May 2014), 216–225.
- [15] Sarah Hyndman. 2016. *Why Fonts Matter: a multisensory analysis of typography and its influence from graphic designer and academic Sarah Hyndman* (1st edition ed.). Virgin Books, London.
- [16] Naz Kaya and Helen H. Epps. 2004. Relationship between color and emotion: a study of college students. *College Student Journal* 38, 3 (Sept. 2004), 396–406. <https://go.gale.com/ps/i.do?p=AONE&sw=w&cissn=01463934&v=2.1&it=r&id=GALE%7CA123321897&sid=googleScholar&linkaccess=abs> Publisher: Project Innovation Austin LLC.
- [17] Joonhwan Lee, Soojin Jun, Jodi Forlizzi, and Scott E. Hudson. 2006. Using Kinetic Typography to Convey Emotion in Text-Based Interpersonal Communication. In *Proceedings of the 6th Conference on Designing Interactive Systems (DIS '06)*. Association for Computing Machinery, New York, NY, USA, 41–49. <https://doi.org/10.1145/1142405.1142414>
- [18] Daniel McDuff, Amy Karlson, Ashish Kapoor, Asta Roseway, and Mary Czerwinski. 2012. AffectAura: An Intelligent System for Emotional Memory. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Austin Texas USA, 849–858. <https://doi.org/10.1145/2207676.2208525>
- [19] Stefan Palan and Christian Schitter. 2018. Prolific.ac—A subject pool for online experiments. *Journal of Behavioral and Experimental Finance* 17 (2018), 22–27. <https://doi.org/10.1016/j.jbef.2017.12.004>
- [20] Fengjiao Peng, Veronica Crista LaBelle, Emily Christen Yue, and Rosalind W. Picard. 2018. A Trip to the Moon: Personalized Animated Movies for Self-reflection. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–10.
- [21] Nina Rajcic and Jon McCormack. 2020. Mirror Ritual: An Affective Interface for Emotional Self-Reflection. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–13. <https://doi.org/10.1145/3313831.3376625>

- [21] Mahnaz Roshanaei, Christopher Tran, Sylvia Morelli, Cornelia Caragea, and Elena Zheleva. 2019. Paths to Empathy: Heterogeneous Effects of Reading Personal Stories Online. In *2019 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*. IEEE, Washington, DC, USA, 570–579. <https://doi.org/10.1109/DSAA.2019.00072>
- [22] Lijiang Shen. 2010. On a Scale of State Empathy During Message Processing. *Western Journal of Communication* 74, 5 (Oct. 2010), 504–524. <https://doi.org/10.1080/10570314.2010.512278>
- [23] Yla R. Tausczik and James W. Pennebaker. 2010. The Psychological Meaning of Words: LIWC and Computerized Text Analysis Methods. *Journal of Language and Social Psychology* 29, 1 (March 2010), 24–54. <https://doi.org/10.1177/0261927X09351676>
- [24] Kiran Vodrahalli, Po-Hsuan Chen, Yingyu Liang, Christopher Baldassano, Janice Chen, Esther Yong, Christopher Honey, Uri Hasson, Peter Ramadge, Kenneth A. Norman, and Sanjeev Arora. 2018. Mapping between fMRI Responses to Movies and Their Natural Language Annotations. *NeuroImage* 180 (Oct. 2018), 223–231. <https://doi.org/10.1016/j.neuroimage.2017.06.042>
- [25] Hua Wang, Helmut Prendinger, and Takeo Igarashi. 2004. Communicating emotions in online chat using physiological sensors and animated text. In *CHI '04 Extended Abstracts on Human Factors in Computing Systems*. ACM, Vienna Austria, 1171–1174. <https://doi.org/10.1145/985921.986016>
- [26] Matthias Wölfel, Tim Schlippe, and Angelo Stitz. 2015. Voice Driven Type Design. In *2015 International Conference on Speech Technology and Human-Computer Dialogue (SpED)*. 1–9. <https://doi.org/10.1109/SPED.2015.7343095>
- [27] Kevin Wright. 2002. Motives for Communication within On-line Support Groups and Antecedents for Interpersonal Use. *Communication Research Reports* 19, 1 (Jan. 2002), 89–98. <https://doi.org/10.1080/08824090209384835>