EmoArc: Interactive Emotion Graph for Human-AI Collaborative Writing

Rifat Mehreen Amin
LMU Munich
Munich, Germany
rifat.amin@ifi.lmu.de

Seyit Baran Özdemir
LMU Munich
Munich, Germany
s.ozdemir@campus.lmu.de

Sven Mayer
TU Dortmund University
Dortmund, Germany
info@sven-mayer.com

Andreas Butz
LMU Munich
Munich, Germany
butz@ifi.lmu.de

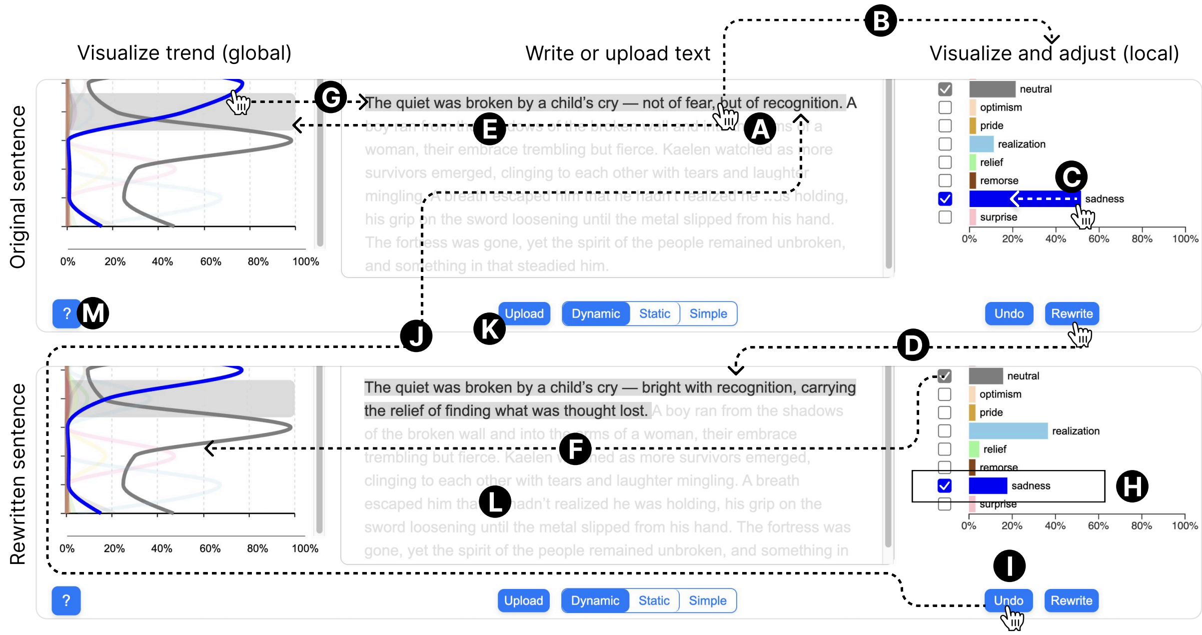

Figure 1: *EmoArc* consists of three components: the main text editor (center, L), a line graph synchronized with the text (left), and a bar graph (right). Writers can either draft text or (K) upload a file. (A) When a sentence is selected, it is highlighted and simultaneously linked to both (E) the global emotional arc and (B) the local emotion bars. (C) Adjusting the bars and (D) pressing Rewrite generates a new sentence with (H) the desired emotion intensities. (F) Toggles in the legend allow selectively focusing on specific emotions, while (G) clicking on the line chart highlights the corresponding sentence in the editor. (I) Pressing Undo (J) restores the original sentence, and the visualizations readapt.

Abstract

With the advent of large language models, creative writers can get support for refinement, co-writing, and text generation. Yet, writers often struggle to understand how model-driven rewrites affect the emotional trajectory of their narratives. We investigate how interactive visualizations can help writers gain awareness, agency, and control over the emotional flow of their texts. First, we designed an interactive visualization-based writing tool that allows writers to adjust emotional pacing, explore different emotional directions, and revise the emotional storyline in real time. Second, we conducted a user study with creative writers ($N = 24$) to evaluate its effectiveness. Our results show a significantly improved enjoyment, exploration, and perceived value. Writers highlighted increased emotional awareness and narrative coherence; simultaneously, they expressed concerns about authorship. Thus, support can enrich the creative process when designed with transparency, user agency, and adaptability in mind, contributing to the understanding of augmenting human writing creativity.

CCS Concepts

• **Human-centered computing → Visualization**; **Interactive systems and tools**; • **Computing methodologies → Artificial intelligence**.

Keywords

Human-AI collaboration, Co-creative systems, Visualization, LLMs, GenAI, emotions, creative writing

AVI '26, Venice, Italy

ACM ISBN 979-8-4007-2342-1/26/06
https://doi.org/10.1145/3811427.3811429

ACM Reference Format:
Rifat Mehreen Amin, Seyit Baran Özdemir, Sven Mayer, and Andreas Butz. 2026. EmoArc: Interactive Emotion Graph for Human-AI Collaborative Writing. In *Proceedings of the 2026 International Conference on Advanced Visual Interfaces (AVI '26), June 08–12, 2026, Venice, Italy*. ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3811427.3811429

1 Introduction

Emotion plays a central role in storytelling, shaping how narratives unfold and how writers intend their work to be experienced. While writers often have an intuitive sense of emotional tone, maintaining awareness of emotional pacing across a text, especially during revision, can be challenging. Subtle tonal imbalances may go unnoticed. Therefore, maintaining a well-balanced emotional flow throughout a story can be challenging [27]. Reflecting on emotional structure often requires rereading and guesswork. Prior work shows that visualizing emotional information, using timelines, color mappings, or circumplex layouts, can make affective patterns in text more accessible [36]. Yet many existing systems analyze emotion post hoc or at coarse granularity, offering limited support for iterative refinement during writing [38]. What remains missing is an approach that combines fine-grained emotion analysis with interactive visualizations and immediate feedback, giving authors agency over their narratives' emotional flow.

To address this gap, we developed and evaluated *EmoArc*, an interactive writing tool that integrates sentence-level emotion detection, visualization, and AI-assisted rewriting. *EmoArc* enables sentence-level exploration of emotional direction without word-level micromanagement, which better accommodates contextual nuance than isolated lexical approaches [35]. Sentence-level emotional shifts play a key role in narrative pacing [24], motivating tools that support fine-grained emotional guidance during writing [25]. *EmoArc* links local edits with global emotional trends and separates emotion analysis from text generation: all sentences are analyzed post hoc by an emotion classification model, while a generative model is used only to propose rewrites conditioned on adjusted emotional targets. In a study with 24 participants, we compared *EmoArc* to a static visualizing-only tool and a baseline without emotional writing support. We asked participants to complete three writing tasks and evaluated the usability and workload. Our research was guided by the following questions:

- **RQ1** How can interactive emotion graphs help authors to adjust the emotional flow of their narratives?
- **RQ2** What impact does integrating emotion graphs and AI-assisted rewriting have on writers' creativity and enjoyment?
- **RQ3** How does real-time emotional feedback shape authors' perceived authorship, control, and confidence while writing?

The results from our study showed that *EmoArc* significantly improved enjoyment, exploration, expressiveness, and perceived value compared to a simple text editor without support, while maintaining low mental load. It also increased enjoyment over the static visualization-only UI, though tensions around authorship and personal involvement emerged. These findings suggest that emotion-centered visualizations can enrich creative exploration while reshaping how writers reflect on agency and ownership.

2 Related Work

We review prior work on (1) emotion visualization and multimodal systems, and (2) emotion in creative writing and narrative studies.

2.1 Emotion Visualization and Multimodal Systems

Building on psychological and computational models of emotion, prior work has explored visualizing and interacting with emotional patterns in text and other media. Many systems adopt dimensional models such as VAD to show emotional variation over time. For example, PEARL [38] visualize emotional dynamics using line- and color-based encodings, while large-scale systems like We Feel [17] and NewsViz [12] focus on aggregate mood patterns in social and narrative data. Related approaches also extend beyond text, visualizing emotion in music or multimodal content using established affective models [9, 26]. However, many of these systems rely on lexicon-based analysis and struggle with context, irony, or cultural variation. More recent work emphasizes interactivity and multimodality. Systems such as StanceVis Prime [16], and EmotionVis [39] provide richer temporal and semantic views of affective information, while multimodal tools like Speak From Heart [18] integrate emotional signals across text, audio, or video.

2.2 Emotion in Creative Writing and Narrative Studies

Early research on emotion in narrative examined emotional arcs and pacing without AI-driven generation. Mohammad showed how emotional expression varies across genres and authors [23, 24], while systems such as Ashida et al. [2] visualized characters' emotional trajectories to support narrative coherence. Later work integrated computational affect into writing workflows; for example, Mori et al. [25] provided emotion-guided story continuation using sentence-level affective targets. Subsequent systems expanded narrative support through visualization and structural interaction. StoryPrint [37] visualizes narrative structure together with character emotion. Recent LLM-based writing systems foreground emotional dynamics in co-creation, including Textoshop's editing-inspired manipulation of affective tone [20]. However, Textoshop operates at the level tone adjustments (e.g., formality, sentiment, complexity) and does not explicitly represent emotional structure across a narrative. Beyond such editing-based approaches, other systems explore higher-level structural interaction with generative text. TaleBrush [6] and PatchView [7] enable high-level structural sketching and visual steering of text generation. Recent work on visual story writing further extends this direction by allowing authors to manipulate visual representations of narrative elements to suggest text revisions [21]. Other LLM-based tools explore narrative and emotional support through vocabulary learning, procedural creativity, and multimodal prompting [29, 31].

Overall, *EmoArc* builds on prior work in emotion visualization and narrative support while integrating analysis and interaction. Unlike systems focused on post hoc analysis or localized tone edits, it allows writers to directly shape the emotional progression of a story through interactive visualizations. This positions *EmoArc* between emotion visualization and generative writing, supporting both reflection and control.

3 EmoArc

We conducted informal design sessions with UX experts and the authors, involving brainstorming, sketching, prototyping, and pilot testing to develop a dynamic, interactive UI for emotion analysis. Through this process, we explored multiple concepts and layouts. The resulting system, *EmoArc*, was guided by four design goals that shaped its architecture and interaction model.

3.1 Design Goals

- **DG1 Support meaningful emotion analysis and manipulation in text.** Emotional expression in writing is nuanced and context-dependent, and generative AI tools often increase metacognitive load through repeated prompt reformulation [32]. In contrast, tools supporting direct analysis and manipulation of emotion can reduce this burden and enable intentional authoring [14].
- **DG2 Ensure real-time responsiveness and interactivity.** Emotion analysis should offer immediate, interpretable feedback to support fluid writing. Static, linear UIs can disrupt creative flow by requiring repeated prompt reformulation [15], whereas fragment-based and responsive interactions enable more continuous engagement with AI tools [4].
- **DG3 Enable bidirectional emotional consistency.** Writing tools should preserve narrative coherence while supporting exploration of alternatives. Prior work shows that authorship can diminish when continuity and control are weakened [15, 22]. Consistently reflecting emotional adjustments across text and visualization helps maintain a sense of authorship.
- **DG4 Provide intuitive emotion visualizations.** Emotion analysis is most effective when presented clearly and is interpretable. Prior work shows that simple visual encodings, such as line and bar graphs, support interpretation and exploration [28, 36], while others emphasize making temporal and structural aspects of writing visible [3, 8].

3.2 User Interface and Interaction

EmoArc turns emotion from a hidden property of text into a visible and editable layer of narrative structure. The UI (see Figure 1) integrates three primary components: text editor, bar graph, and line graph, plus supporting tools (file upload, Rewrite, Undo, tutorial).

3.2.1 *Provide Text (Write or Upload).*

The text editor (Figure 1-L) provides a familiar writing environment. Writers can compose text directly or upload a `.txt` file by clicking on the Upload button, which is automatically segmented into sentences for analysis and editing. Emotion analysis runs automatically during pauses in writing, providing timely feedback while maintaining interface responsiveness (DG2), supporting the aim of making progress visible and concrete enough for immediate feedback [33].

3.2.2 *Select a Sentence.*

Selecting a sentence applies a grey backdrop in the editor (Figure 1-A) and highlights its emotional intensities across both bar graph (Figure 1-B) and line graph (Figure 1-E), coupling text with emotion analysis through visualizations (DG1).

3.2.3 *Inspect Locally (Bar graph).*

Bar graphs provide an intuitive way to represent and compare emotion intensities [36]. In *EmoArc*, they display and allow adjustment of the normalized emotion distribution of a selected sentence, providing clear sentence-level control (Figure 2) and visual clarity (DG4). Hovering reveals exact percentages (Figure 2-F), and dragging a bar updates the target intensity.

3.2.4 *Adjust & Rewrite.*

Users begin with the original sentence and its emotion intensities (Figure 2-A) and adjust the bars directly as sliders, for example, lowering sadness and increasing surprise to set a target emotional distribution (Figure 2-B). Pressing the Rewrite button (Figure 2-C) rewrites a revised sentence aligned with these targets while preserving context (Figure 2-D) (DG3). The new sentence and its emotion intensities are shown immediately.

3.2.5 *Undo.*

If needed, the Undo button restores the original sentence (and visualizations will follow), which supports safe exploration because it makes it easy to recover from mistakes or dead ends (Figure 2-E).

3.2.6 *Inspect Globally (Line graph).*

While the bar graph provides a sentence-level view, the line graph offers a global perspective on emotional flow across the text. Line graphs are well suited for visualizing narrative arcs and emotional variation over a document [36]. In *EmoArc*, each emotion is shown as a line over the document's progression (Figure 3-D), enabling writers to observe global trends. The line graph is scroll-synchronized with the editor (Figure 1-G) so that the visualization reflects the currently visible text.

3.2.7 *Filter & Focus (Legend Toggles).*

To reduce clutter, the line graph includes an interactive legend (Figure 3-B) that allows writers to toggle individual emotions. Selecting an emotion highlights its curve. For example, isolating confusion and disappointment to examine their balance over time (Figure 3-C), while non-selected emotions are grayed out but remain visible (Figure 3-E). This filtering supports focused inspection of emotional arcs without overwhelming the display (DG1).

3.2.8 *Navigate & Synchronize.*

Hovering over a line in the graph dynamically displays the associated emotion label near the cursor (Figure 3-F), while transparent interactive zones along each row allow hover and click interactions. Clicking anywhere in the graph (Figure 3-A) highlights the corresponding sentence with a semi-transparent rectangle and synchronizes the editor view, automatically scrolling to center the selected text using easing and clamping for smooth navigation.

3.2.9 *Learning and Support.*

To lower the entry barrier, a tutorial appears on first use and can be reopened via a help button ? (Figure 1-M). Additional scaffolds, such as tooltips, color coding, and smooth transitions, support quick understanding of the workflow.

The interaction follows a cyclical process: *write* → *visualize* → *adjust* → *rewrite* → *reflect.* **Writers can revise text either by using AI-assisted rewriting or by editing it themselves, with the visualization updating in both cases. They may choose to adjust emotional targets and invoke AI-assisted rewriting, or instead revise the text themselves and see the visualization update accordingly.** By tightly coupling writing, visualization, and rewriting, the system supports deliberate shaping of emotional arcs while preserving writer control (see the supplementary material for an example scenario).

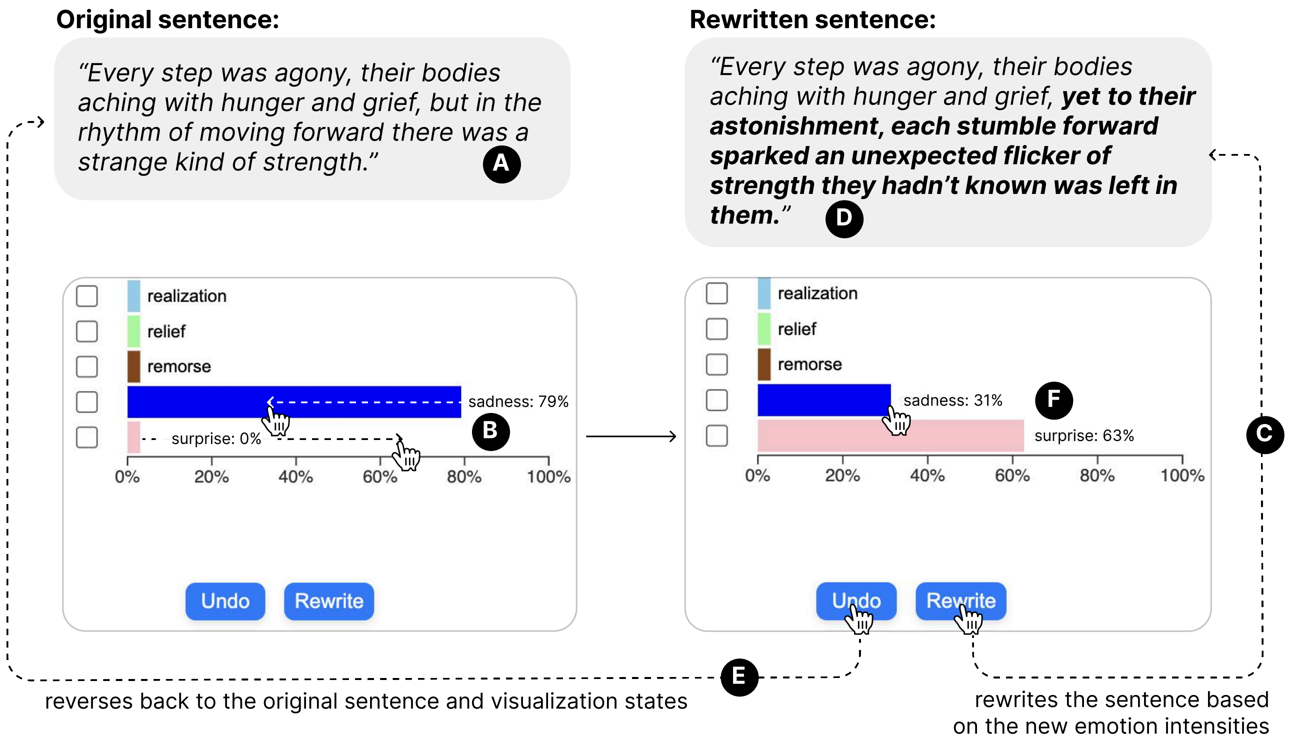

Figure 2: Bar graph for local inspection and adjustment of emotional tone. (A) The original sentence. (B) Users can manipulate the bars directly. (C) The Rewrite button generates a new sentence to align with the new emotional distribution. (D) The generated sentence is displayed, while (E) the Undo button restores the original sentence and visualization state.

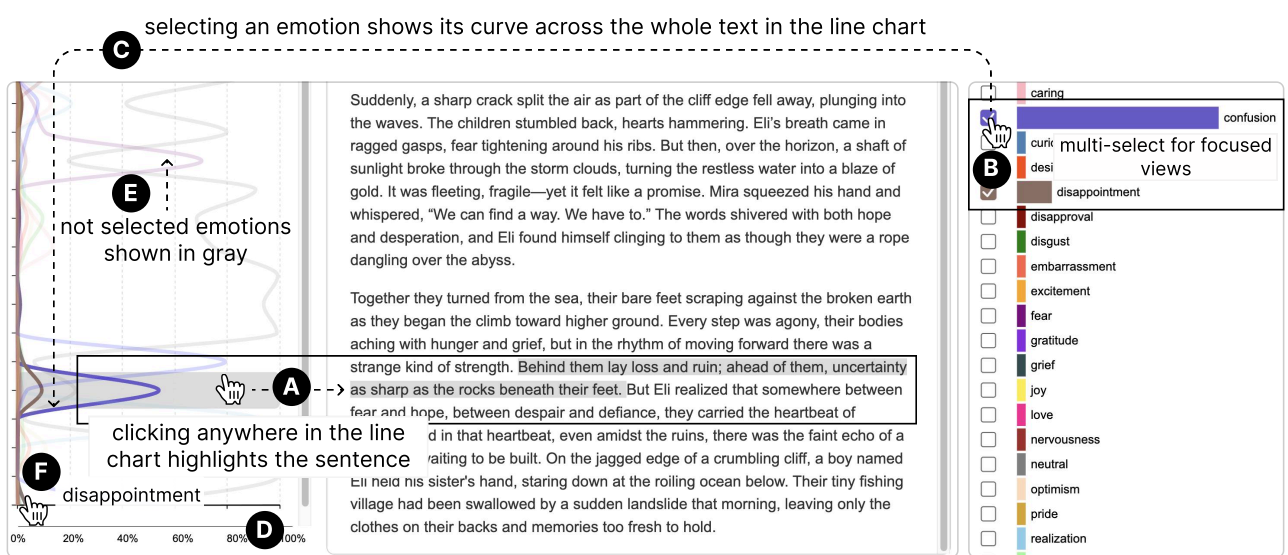

Figure 3: Line graph for global inspection of emotional trajectories. (A) Clicking anywhere in the graph highlights the corresponding sentence in the editor. (B, C) Users can select or multi-select emotions from the legend to focus on specific trajectories, while (E) non-selected emotions are shown in gray. (F) Individual curves represent normalized emotion intensities across the text, and (G) scroll synchronizes the editor with the visualization.

3.3 Technical Implementation

The system is implemented as a modular client–server application, with a web-based frontend (D3.js) and a Python (Flask) backend communicating asynchronously via a DataService module. We base emotion analysis on the GoEmotions dataset, which provides 28 fine-grained emotion categories beyond coarse polarity or Ekman's basic emotions [10]. Emotion classification is performed using a RoBERTa-based model, while rewriting is handled by a GPT-based model conditioned on user-adjusted emotion targets. Importantly, analysis and generation are separated, enabling transparent feedback and controllable rewriting. Further implementation details are provided in the supplementary material.

4 User Study

We conducted a within-subjects lab study with varied creative writing experience and varying familiarity with AI writing tools (see the supplementary material for details).

We evaluated *EmoArc* (dynamic UI) against two conditions: a baseline simple text editor (simple UI, only the center part in Figure 1) and a visualization-only variant without AI-assisted rewriting (static UI). To isolate the effect of visualization, we varied visual feedback only, as additional conditions would increase study complexity and participant burden. The study included a pre-study questionnaire, a standardized interface tutorial, post-task surveys, and a semi-structured interview. Screen and audio were recorded for qualitative analysis and interview coding. Participants completed three short creative writing tasks (details in supplementary materials), each requiring the addition of 2-3 sentences to a given text, with task and interface order counterbalanced across participants.

Measurements. We collected quantitative measures including the Creativity Support Index (CSI), NASA Task Load Index (NASA-TLX), and Likert-scale ratings on usability, authorship, ownership, and desirability of use, and overall tool preferences. Interviews captured perceptions of the emotion graph, creativity and control, authorship and AI collaboration, and usability.

Participants. We recruited 24 participants (ages 18-45, $M = 29.36$, $SD = 6.70$). Participants were compensated 15€ for a 90-minute session. Our institution's ethics board approved the study EK-MIS-2025-0395-FT-d01.

5 Results

In this section, we present the results of our lab study, structured by the research questions they address. For analysis, we performed Shapiro-Wilk tests for normality and used paired t-tests or Wilcoxon signed-rank tests accordingly, with Holm–Bonferroni correction for multiple comparisons. Qualitative data were analyzed through open coding by two authors, with themes refined collaboratively.

5.1 How can interactive emotion graphs help writers to adjust the emotional flow of their narratives? (RQ1)

Participants described the emotion graphs as a useful and often novel aid for reflecting on and guiding emotional flow. Many noted increased awareness of emotional pacing and support for real-time wording or tone adjustments.

5.1.1 Changes to Planning and Story Structure. The degree to which the graph influenced participants' narrative planning varied. A majority indicated that they followed their initial story plans and only occasionally adjusted wording. For instance, P5 explained: *"I followed my initial plan. The graph didn't make me change structure"*. However, others reported concrete changes when the graph highlighted discrepancies between their intended emotion and what was detected. P7 described: *"In the 1st story I was going to write something else. But then, when I looked at the graph, I saw that the sentence was not projecting exactly what I wanted to write. So I had to change one or two words, and then the final outcome was according to what I wanted"*. These suggest that the graph not only supported evaluation but also triggered moments of revision and redirection.

5.1.2 Intuitiveness of Manipulation. Participants perceived the interface for manipulating emotions using the bar graph as straightforward. P1 described it as *"simple and intuitive,"* while P10 emphasized: *"It was very easy to use."* Others highlighted the value of immediate responsiveness, as P21 explained: *"I liked how instant it is. For example, the rewriting process takes only 20 seconds. That is impressive... And the thing is, as I'm writing, the graph keeps changing. That's great."* P19 said, *"It was quite simple. The Bar graphs, when I was increasing and decreasing the emotions, the leveling of those were what I could understand."* Nonetheless, participants identified areas where usability could improve (P14, P17). P17 pointed out that some functions were not self-explanatory: *"The interface was quite intuitive. The only thing not obvious was using 'shift' to select multiple emotions."* Some reported moments of confusion when the system's interpretation diverged from their own, e.g., P2 explained: *"Sometimes it delivers tone-deaf sentences... If the tool says a sentence is 80% admiration, I might feel it's actually 50 or 60%."*. These mismatches introduced moments of doubt and minor frustration.

5.1.3 Cognitive Load and Effort (NASA-TLX and General Perception). We analyzed raw NASA-TLX ratings across the three interfaces to evaluate perceived workload. A Friedman test revealed no significant effects of condition on any NASA-TLX factor, nor on the overall NASA-TLX score (see supplementary materials for details). This indicates that participants experienced all three interfaces as lightweight and manageable.

Qualitative feedback helps explain these patterns. Several participants noted that the tool reduced editing burden by making emotional adjustments more direct. As P17 explained, the system allowed them to "quickly adjust a sentence to portray a certain emotion without overthinking". In contrast, some participants found the static UI more burdensome due to its reduced interactivity. P4 described a sense of "mental pressure" when interaction options were limited, and others noted that static interactions sometimes felt restrictive. These concerns were not raised for the dynamic UI, which participants described as smoother and less obstructive. Overall, while statistical tests did not reveal significant differences, qualitative findings suggest that the dynamic UI offered the best balance of perceived success and manageable effort, whereas static and simple UIs involved trade-offs between effort and engagement.

5.2 What impact does integrating emotion graphs and AI-assisted rewriting have on writers' creativity and enjoyment? (RQ2)

We assessed perceived creativity support using the Creativity Support Index (CSI). Results suggest that combining emotion graphs with AI-assisted rewriting increased enjoyment and perceived creative potential compared to simpler designs.

5.2.1 Creativity Support Index (CSI). CSI scores showed that the dynamic UI consistently outperformed the simple UI across several dimensions (see Figure 4, and supplementary material). A Friedman test revealed significant effects of the UIs on Enjoyment ($\chi^2(2) = 9.18$, $p = .010$), Exploration ($\chi^2(2) = 20.74$, $p < .001$), Expressiveness ($\chi^2(2) = 15.58$, $p < .001$), Results Worth Effort ($\chi^2(2) = 15.93$, $p < .001$), and the Overall CSI score ($\chi^2(2) = 10.81$, $p = .004$). No significant effect was found for Immersion ($\chi^2(2) = 1.28$, $p = .53$).

Enjoyment was significantly higher in the dynamic UI ($p < .001$) compared to the simple UI, and also higher in the static UI ($p = .042$) compared to the simple UI; the dynamic UI was additionally rated more enjoyable than the static UI ($p = .042$). Exploration followed a similar pattern, with the dynamic UI ($p < .001$) rated higher than the simple UI, and the static UI ($p < .001$) also surpassing the simple UI. Expressiveness was likewise higher in the dynamic UI ($p = .004$) than in the simple UI, and in the static UI ($p = .002$) compared to the simple UI. For results worth effort, the dynamic UI ($p = .012$) scored higher than the simple UI, and the static UI ($p < .001$) also outperformed the simple UI.

Overall, CSI ratings reflected the same trend, with the dynamic UI ($p = .004$) significantly higher than the simple UI, and the static UI ($p = .003$) also higher than the simple UI. No significant differences were found for immersion or between the dynamic and static UIs across most dimensions (see supplementary materials for detailed CSI scores). Overall, both the dynamic and static UIs supported creativity more effectively than the simple UI, with the dynamic UI showing the most consistent benefits for enjoyment, exploration, expressiveness, and overall creative support.

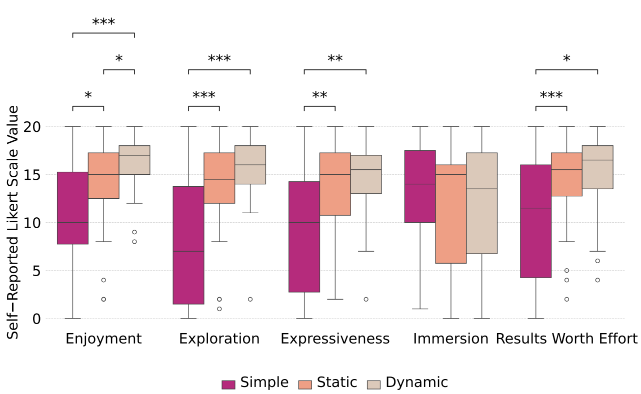

Figure 4: The CSI scores for our three conditions. In line with prior studies [1, 30], we excluded the Collaboration factor in CSI, as our study did not involve collaborative settings.

5.2.2 Enjoyment and Exploring New Creative Directions. Combining emotion graphs with rewriting encouraged participants to explore directions they might not have considered yet. Many reported that sliders and rewrites sparked new ideas and enabled alternative ways of expressing the same sentence. As P1 noted, *"it opened new ways of thinking about my sentences,"* while P8 reflected, *"it showed me that a sentence can be written in different ways to express different emotions... that really made me feel creative."* For some, this led to iterative refinement, with rewrites inspiring further improvements (P14). Participants also described the system as enjoyable and creatively stimulating, often highlighting a sense of playfulness and increased confidence, supported by higher CSI ratings for the dynamic UI. However, a minority reported reduced creativity, particularly when relying heavily on system-generated rewrites or when outputs conflicted with their personal style. As P9 remarked, *"I just let the AI do it for me... so I didn't feel too creative."* Overall, the system's creative benefits were most evident when used as a tool for exploration rather than as a substitute for authorship.

5.2.3 Alignment with Writers' Intentions. While the system was generally successful at modifying emotional tone, participants reported mixed experiences in whether the rewrites aligned with their intentions. Many felt the tool reliably reflected their adjustments. For instance, P6 noted: *"Yes, they mostly reflected what I wanted."* However, others found the results inconsistent or stylistically unsatisfying. P21 emphasized: *"I was not very happy with the rewriting option (in dynamic UI)... It didn't have that human touch."* These responses suggest that while the system adjusted emotional content effectively, maintaining stylistic quality and authorial voice remains a challenge.

5.3 How does real-time emotional feedback shape writers' perceived authorship, control, and confidence while writing? (RQ3)

Participants reflected on how the system influenced their sense of authorship, their control over emotional direction, and how they positioned the tool in relation to their own writing.

5.3.1 Quantitative Results. A Friedman test revealed significant effects of UI on participants' sense of authorship ($\chi^2(2) = 8.54$, $p = .014$) and whether the story felt personal ($\chi^2(2) = 8.96$, $p = .011$). We observed no significant differences for perceived control over the outcome, the extent to which the ending reflected participants' writing style, or overall satisfaction with the final result.

Post-hoc Wilcoxon results revealed differences in how participants evaluated authorship and personal involvement across the three UIs (see Figure 5 and the supplementary material). Participants reported higher authorship in the simple UI compared to the dynamic UI ($p = .040$), and in the static UI compared to the dynamic UI ($p = .024$). This suggests that fewer system interventions led to a stronger sense of authorship. For personal involvement, the dynamic UI was rated higher than the simple UI ($p = .009$), indicating that dynamic interaction did not increase participants' sense of personal engagement. No significant differences were found for control (*"I had control on the final outcome"*), satisfaction (*"I feel satisfied with the final result"*), or style (*"The ending of the story reflects my personal writing style"*).

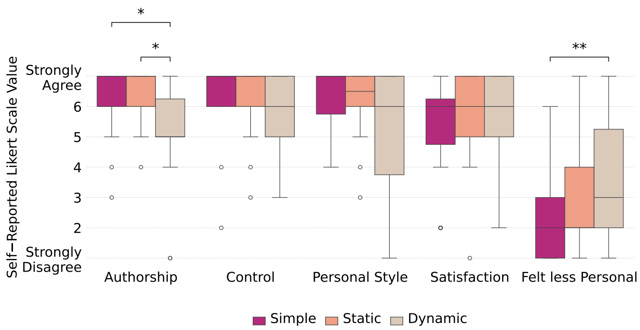

Figure 5: Self-reported perceptions of authorship, control, personal style, satisfaction, and personal involvement across the three UIs.

5.3.2 Authorship, Ownership, and Control. Participants emphasized that authorship remained with them, describing the tool as supportive rather than replacing their creative agency. P7 stressed: *"Even though it is an AI, the original thought is mine. I'm just modifying the way I want my sentence to be perceived by the audience."* Real-time feedback further supported this sense of control; as P4 explained, *"I can control the emotion or direction of the story. Yes, I did feel that."* At the same time, several participants pointed out limitations when the system's output did not align with their intent. P10 noted: *"I would have preferred finer control, because one of the sentences did not reflect the emotion correctly"*. These experiences suggest that authorship and agency remained intact, though occasional imprecision disrupted participants' sense of control.

5.3.3 Collaboration with AI. Most participants described the AI as a helpful assistant that supported their work without taking over creative ownership. P8 explained: *"It was an assistant that helped me to finish my story in a better way."* A smaller group, however, described the tool as taking on a more active role. P14 characterized it as *"someone who gives their own idea to me, someone who shares their insight, someone who's not just going to be a yes-man,"* and P19 went further: *"I think it felt like it's in my head. So it's definitely a co-author."* These differing views suggest that while most writers saw the tool as an assistant, some experienced it as a creative partner, blurring the line between assistance and co-authorship.

5.3.4 Granularity of Control. Participants reflected on the unit of text at which emotional editing should occur. Many were satisfied with sentence-level manipulation, finding it intuitive and sufficient for short passages. P8 noted: *"I think having the opportunity to click on a sentence and know what emotion it's expressing is sufficient, and I can alter the sentence according to the emotion that I'm preferring."* At the same time, a substantial group expressed interest in more flexible granularity, either at the level of phrases, words, or paragraphs. P2 expected selecting specific parts of a sentence: *"The user could select a certain part of a sentence ... and edit those emotions more granularly."* Others preferred paragraph-level control to capture broader context. As P15 explained: *"If it's possible to change the entire paragraph, not only one sentence, I think it will be more convenient, and it will more reflect how people actually write."* These perspectives suggest that while sentence-level control is effective for many use cases, additional options for finer or broader control would better accommodate diverse writing practices.

5.4 Overall feedback

In the final survey, 83.3% of participants reported preferring the dynamic UI, while 16.7% preferred the static UI, and none selected the simple UI. This distribution highlights a strong preference for real-time interaction and control. Participants' explanations in the survey responses further show these preferences and align closely with themes observed in the interview data.

5.4.1 Comparisons Across UIs. The dynamic UI was consistently described as novel and distinct from conventional sentiment tools. As P3 noted, *"The simple UI is similar to a regular sentiment tool like VOYANT… dynamic UI is something new."* Participants also praised the Dynamic UI for supporting creativity, expressiveness, and control. P14 highlighted its breadth: *"It allows me to explore different ideas, structure of sentences, flow of writing – all of these collectively make the whole process of creative writing more fun and fulfilling."* Others emphasized usability, with P22 stating, *"I like the dynamic UI as it adapts in real time, making interactions feel smooth and more human-like. It makes complex tasks feel simple and intuitive."* In contrast, a minority preferred the Static UI for its simplicity and stronger sense of authorship. As P21 explained, *"I was not happy with the rewriting option (dynamic UI)… the static UI gave me a reading of the emotional tone, and I prefer rewriting/editing myself."* These reflections align with interview findings: visualization and emotional manipulation were engaging, but mismatches between system interpretations and user intent caused frustration.

5.4.2 Perceived Usability and Adoption. Participants rated all three UIs as easy to use (Figure 6), with 96% of participants agreeing or strongly agreeing for the static UI, 92% for the dynamic UI, and 79% for the simple UI. Despite its additional interactivity, the dynamic UI was not perceived as harder to use. Willingness to use the tool regularly showed a clearer preference for the dynamic UI (79%), compared to the static (67%) and simple (38%) UIs. Overall, while all interfaces were considered accessible, only the dynamic version combined usability with features that motivated long-term use.

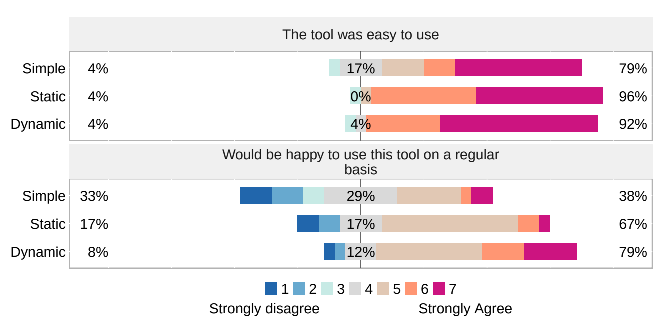

Figure 6: Perceived usability and adoption across our UIs.

6 Discussion & Future Work

In this work, we investigated how interactive emotion graphs support creative writing. We examined (RQ1) how emotion visualizations help writers adjust emotional flow, (RQ2) how they affect creativity and enjoyment, and (RQ3) how real-time feedback shapes authorship and control. Our findings highlight both the potential of emotion-centered interfaces for reflection and exploration and the tensions they introduce around support and authorship.

Limitations. Our study is limited in duration and participant pool, with tasks restricted to short text completions. Although we aimed for ecological validity through task design and participants with a literature background, the findings may not generalize to other domains or longer, real-world workflows; longer-term studies with professional writers could provide additional insights. We relied on user-reported experience without collecting detailed interaction logs, limiting fine-grained analysis. Future work could incorporate logging and external evaluations by expert or independent evaluators to compare perceived and actual output quality.

Visualization as a Creativity Scaffold. Participants described the emotion graph as both a mirror that revealed unintended emotional shifts and a planning aid for shaping emotional trajectories. This aligns with prior work showing that externalizing abstract dimensions of writing can support metacognition [32]. They valued the visualization even without AI rewriting for making relationships visible and surfacing different interpretations, thus the visualization can support reflection and divergent thinking **(RQ1)**.

AI-Assisted Rewriting as a Catalyst for Exploration. Reactions to AI-generated rewrites were mixed. Some participants found outputs incoherent or overly "AI-like," echoing critiques of style loss in LLM co-writing [22]. These tensions suggest that productive surprise requires balancing novelty with coherence and authorial style. At the same time, AI-assisted rewrites encouraged playful, divergent exploration **(RQ2)**, aligning with prior work showing that novelty can catalyze creativity [14], even when used primarily as inspiration.

Preserving Authorship and Control. Participants framed the system as an assistant rather than a co-author, aligning with prior work showing that users accept AI suggestions while maintaining ownership of final decisions [4]. Increased system intervention reduced perceptions of sole authorship, reflecting findings from screenplay co-writing systems [22], suggesting that interactivity can both enhance engagement and complicate authorship **(RQ3)**. Differences in collaboration framing highlight an open design question around system agency. Some participants favored the AI as a critical assistant that offers alternatives, emphasizing transparency and optionality. Opaque prompting has been shown to increase metacognitive burden [1]; making AI reasoning visible and editable may help preserve writer control while supporting collaboration.

Managing Cognitive Demands. Generative AI tools often increase metacognitive load by requiring users to manage complex prompts [32], yet emotion graphs did not substantially increase workload in our study. NASA-TLX ratings were low across UIs, with *EmoArc* associated with higher perceived success and manageable effort, echoing prior findings on dynamic prompting interfaces [1, 34]. Participants described the graphs as reducing editing burden and enabling quick tone adjustments, though some noted distraction when visualizations diverged from expectations, consistent with prior work on UI overload [11]. These results suggest the importance of customizable feedback levels, allowing writers to balance monitoring and intervention to match their creative flow.

Design Opportunities and Future Directions. Our findings point to several directions for extending emotion-centered writing support. First, multilingual emotion analysis is a promising avenue, as emotions can carry culturally specific meanings [13]. Investigating how emotional arcs are preserved across translations could support translators and inform multimedia applications such as video captioning, where current approaches often prioritize factual content while overlooking emotional cues [5]. Second, customization and transparency are important, as participants expressed a need for both fine-grained and high-level control. While the categorical emotion model (GoEmotions) enables such granularity, it may not fully align with how writers perceive or articulate their intentions, which are often more continuous and context-dependent. Third, Justifications for changes help maintain AI as a tool for augmentation. Finally, linking emotional arcs to recommendation systems (e.g., Goodreads) could enable new ways of discovering books based on emotional journeys rather than genre or plot, building on prior work in emotion-aware recommender systems based on reviews [19].

7 Conclusion

In this work, we introduced a system, *EmoArc*, that couples emotion visualization with AI-assisted rewriting to support creative writing. Our study with 24 participants showed that the interface enhanced awareness of emotional pacing, encouraged exploration, and supported reflection on tone, while maintaining manageable cognitive demands. Participants valued the system's transparency and responsiveness, though some raised concerns about authorship and stylistic consistency. These findings suggest that emotion-centered interfaces can enrich human-AI collaboration in writing by helping authors craft more intentional emotional trajectories and opening new directions for emotion-aware design in creative writing.

Acknowledgments

Funded by Elitenetzwerk Bayern. This work has been partly supported by the Research Center Trustworthy Data Science and Security (https://rc-trust.ai), one of the Research Alliance centres within the UA Ruhr (https://uaruhr.de).

References

[1] Rifat Mehreen Amin, Oliver Hans Kühle, Daniel Buschek, and Andreas Butz. 2025. Composable Prompting Workspaces for Creative Writing: Exploration and Iteration Using Dynamic Widgets. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 144, 11 pages. doi:10.1145/3706599.3720243

[2] Atsushi Ashida, Masataka Tokumaru, and Tomoko KOJIRI. 2019. Novel Writing Support System by Target Readers' Story Arcs and Characters' Emotional Changes. *International Conference on Computers in Education* (12 2019). doi:10.58459/icce.2019.283

[3] Michael Mose Biskjaer, Jonas Frich, Lindsay MacDonald Vermeulen, Christian Remy, and Peter Dalsgaard. 2019. How Time Constraints in a Creativity Support Tool Affect the Creative Writing Experience. In *Proceedings of the 31st European Conference on Cognitive Ergonomics (ECCE '19)*. Association for Computing Machinery, New York, NY, USA, 100–107. doi:10.1145/3335082.3335084

[4] Daniel Buschek. 2024. Collage is the New Writing: Exploring the Fragmentation of Text and User Interfaces in AI Tools. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference (DIS '24)*. Association for Computing Machinery, New York, NY, USA, 2719–2737. doi:10.1145/3643834.3660681

[5] Weidong Chen, Cheng Ye, Peipei Song, Lei Zhang, Yongdong Zhang, and Zhendong Mao. 2026. Subjective-Objective Emotion-Correlated Generation Network for Subjective Video Captioning. *IEEE Transactions on Image Processing* 35 (2026), 540–555. doi:10.1109/TIP.2025.3649363

[6] John Joon Young Chung, Wooseok Kim, Kang Min Yoo, Hwaran Lee, Eytan Adar, and Minsuk Chang. 2022. TaleBrush: Sketching Stories with Generative Pretrained Language Models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 209, 19 pages. doi:10.1145/3491102.3501819

[7] John Joon Young Chung and Max Kreminski. 2024. Patchview: LLM-powered Worldbuilding with Generative Dust and Magnet Visualization. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24)*. Association for Computing Machinery, New York, NY, USA, Article 77, 19 pages. doi:10.1145/3654777.3676352

[8] Hai Dang, Karim Benharrak, Florian Lehmann, and Daniel Buschek. 2022. Beyond Text Generation: Supporting Writers with Continuous Automatic Text Summaries. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22)*. Association for Computing Machinery, New York, NY, USA, Article 98, 13 pages. doi:10.1145/3526113.3545672

[9] Januka Dharmapriya, Lahiru Dayarathne, Tikiri Diasena, Shiromi Arunathilake, Nihal Kodikara, and Primal Wijesekera. 2021. Music Emotion Visualization through Colour. In *2021 International Conference on Electronics, Information, and Communication (ICEIC)*. 1–6. doi:10.1109/ICEIC51217.2021.9369788

[10] Paul Ekman. 1992. An argument for basic emotions. *Cognition and Emotion* 6, 3-4 (1992), 169–200. doi:10.1080/02699939208411068

[11] Krzysztof Z. Gajos, Mary Czerwinski, Desney S. Tan, and Daniel S. Weld. 2006. Exploring the design space for adaptive graphical user interfaces. In *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI '06)*. Association for Computing Machinery, New York, NY, USA, 201–208. doi:10.1145/1133265.1133306

[12] E Hanser, P McKevitt, TF Lunney, and J Condell. 2010. NewsViz: emotional visualization of news stories. In *Unknown Host Publication*. Association for Computational Linguistics, United States, 125–130. Proc. of the NAACL-HLT Workshop on Computational Approaches to Analysis and Generation of Emotion in Text ; Conference date: 01-06-2010.

[13] Shreya Havaldar, Sunny Rai, Bhumika Singhal, Langchen Liu, Sharath Chandra Guntuku, and Lyle Ungar. 2023. Multilingual Language Models are not Multicultural: A Case Study in Emotion. In *Proceedings of the 13th Workshop on Computational Approaches to Subjectivity, Sentiment, & Social Media Analysis*. Association for Computational Linguistics, Toronto, Canada, 202–214. doi:10.18653/v1/2023.wassa-1.19

[14] Daphne Ippolito, Ann Yuan, Andy Coenen, and Sehmon Burnam. 2022. Creative Writing with an AI-Powered Writing Assistant: Perspectives from Professional Writers. arXiv:2211.05030 [cs.HC] https://arxiv.org/abs/2211.05030

[15] Max Kreminski. 2024. The Dearth of the Author in AI-Supported Writing. In *Proceedings of the Third Workshop on Intelligent and Interactive Writing Assistants (In2Writing '24)*. Association for Computing Machinery, New York, NY, USA, 48–50. doi:10.1145/3690712.3690725

[16] Kostiantyn Kucher, Rafael, M. Martins, Carita Paradis, and Andreas Kerren. 2020. StanceVis Prime: Visual Analysis of Sentiment and Stance in Social Media Texts. *Journal of Visualization* 23, 6 (2020), 1015–1034. doi:10.1007/s12650-020-00684-5

[17] Mark E. Larsen, Tjeerd W. Boonstra, Philip J. Batterham, Bridianne O'Dea, Cecile Paris, and Helen Christensen. 2015. We Feel: Mapping Emotion on Twitter. *IEEE Journal of Biomedical and Health Informatics* 19, 4 (2015), 1246–1252. doi:10.1109/JBHI.2015.2403839

[18] Chenxiao Liu, Zheyong Xie, Sirui Zhao, Jin Zhou, Tong Xu, Minglei Li, and Enhong Chen. 2024. Speak From Heart: An Emotion-Guided LLM-Based Multimodal Method for Emotional Dialogue Generation. In *Proceedings of the 2024 International Conference on Multimedia Retrieval (ICMR '24)*. Association for Computing Machinery, New York, NY, USA, 533–542. doi:10.1145/3652583.3658104

[19] Elena-Ruxandra Luţan and Costin Bădică. 2023. Emotion-Based Literature Books Recommender Systems. In *2023 18th Conference on Computer Science and Intelligence Systems (FedCSIS)*. 275–280. doi:10.15439/2023F8647

[20] Damien Masson, Young-Ho Kim, and Fanny Chevalier. 2025. Textoshop: Interactions Inspired by Drawing Software to Facilitate Text Editing. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1087, 14 pages. doi:10.1145/3706598.3713862

[21] Damien Masson, Zixin Zhao, and Fanny Chevalier. 2025. Visual Story-Writing: Writing by Manipulating Visual Representations of Stories. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology (UIST '25)*. Association for Computing Machinery, New York, NY, USA, Article 70, 15 pages. doi:10.1145/3746059.3747758

[22] Piotr Mirowski, Kory W. Mathewson, Jaylen Pittman, and Richard Evans. 2023. Co-Writing Screenplays and Theatre Scripts with Language Models: Evaluation by Industry Professionals. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 355, 34 pages. doi:10.1145/3544548.3581225

[23] Saif M. Mohammad. 2012. From once upon a time to happily ever after: Tracking emotions in mail and books. *Decision Support Systems* 53, 4 (2012), 730–741. doi:10.1016/j.dss.2012.05.030 1) Computational Approaches to Subjectivity and Sentiment Analysis 2) Service Science in Information Systems Research : Special Issue on PACIS 2010.

[24] Saif M. Mohammad, Tony, and Yang. 2013. Tracking Sentiment in Mail: How Genders Differ on Emotional Axes. arXiv:1309.6347 [cs.CL] https://arxiv.org/abs/1309.6347

[25] Yusuke Mori, Hiroaki Yamane, Ryohei Shimizu, and Tatsuya Harada. 2022. Plug-and-Play Controller for Story Completion: A Pilot Study toward Emotion-aware Story Writing Assistance. In *Proceedings of the First Workshop on Intelligent and Interactive Writing Assistants (In2Writing 2022)*. Association for Computational Linguistics, Dublin, Ireland, 46–57. doi:10.18653/v1/2022.in2writing-1.6

[26] James A. Russell. 1980. A circumplex model of affect. *Journal of Personality and Social Psychology* 39, 6 (12 1980), 1161–1178. doi:10.1037/h0077714

[27] Marie-Luise C. R. Schmidt, Julia R. Winkler, Markus Appel, and Tobias Richter. 2023. Tracking Emotional Shifts During Story Reception: The Relationship Between Narrative Structure and Affective Responses. *Scientific Study of Literature* (Aug 2023). doi:10.61645/ssol.177

[28] Vidya Setlur, Sarah E. Battersby, Melanie Tory, Rich Gossweiler, and Angel X. Chang. 2016. Eviza: A Natural Language Interface for Visual Analysis. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16)*. Association for Computing Machinery, New York, NY, USA, 365–377. doi:10.1145/2984511.2984588

[29] Nikhil Singh, Guillermo Bernal, Daria Savchenko, and Elena L. Glassman. 2023. Where to Hide a Stolen Elephant: Leaps in Creative Writing with Multimodal Machine Intelligence. *ACM Trans. Comput.-Hum. Interact.* 30, 5, Article 68 (Sept. 2023), 57 pages. doi:10.1145/3511599

[30] Sangho Suh, Meng Chen, Bryan Min, Toby Jia-Jun Li, and Haijun Xia. 2024. Luminate: Structured Generation and Exploration of Design Space with Large Language Models for Human-AI Co-Creation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 644, 26 pages. doi:10.1145/3613904.3642400

[31] Sangho Suh, Jian Zhao, and Edith Law. 2022. CodeToon: Story Ideation, Auto Comic Generation, and Structure Mapping for Code-Driven Storytelling. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22)*. Association for Computing Machinery, New York, NY, USA, Article 13, 16 pages. doi:10.1145/3526113.3545617

[32] Lev Tankelevitch, Viktor Kewenig, Auste Simkute, Ava Elizabeth Scott, Advait Sarkar, Abigail Sellen, and Sean Rintel. 2024. The Metacognitive Demands and Opportunities of Generative AI. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 680, 24 pages. doi:10.1145/3613904.3642902

[33] Jason Chew Kit Tham. 2022. *Keywords in design thinking: A lexical primer for technical communicators & designers*. WAC Clearinghouse.

[34] Priyan Vaithilingam, Elena L. Glassman, Jeevana Priya Inala, and Chenglong Wang. 2024. DynaVis: Dynamically Synthesized UI Widgets for Visualization Editing. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 985, 17 pages. doi:10.1145/3613904.3642639

[35] Fuhai Wang. 2024. Comparative Evaluation of Sentiment Analysis Methods: From Traditional Techniques to Advanced Deep Learning Models. *Applied and Computational Engineering* 105 (2024), 23–29. doi:10.54254/2755-2721/105/2024TJ0056

[36] Jiaqi Wang, Tianyi Gui, Mingzhi Cheng, Xuan Wu, Ruolin Ruan, and Meng Du. 2023. A survey on emotional visualization and visual analysis. *Journal of Visualization* 26, 1 (2023), 177–198. doi:10.1007/s12650-022-00872-5

[37] Katie Watson, Samuel S. Sohn, Sasha Schriber, Markus Gross, Carlos Manuel Muniz, and Mubbasir Kapadia. 2019. StoryPrint: an interactive visualization of stories. In *Proceedings of the 24th International Conference on Intelligent User Interfaces (IUI '19)*. Association for Computing Machinery, New York, NY, USA, 303–311. doi:10.1145/3301275.3302302

[38] 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)*. IEEE, 203–212. doi:10.1109/VAST.2014.7042496

[39] Chris Zimmerman, Mari-Klara Stein, Daniel Hardt, Christian Danielsen, and Ravi Vatrapu. 2016. emotionVis: Designing an Emotion Text Inference Tool for Visual Analytics. In *Tackling Society's Grand Challenges with Design Science*. Springer International Publishing, Cham, 238–244.