Check for updates

ACM DIGITAL LIBRARY | Association for Computing Machinery | acm open

Latest updates: https://dl.acm.org/doi/10.1145/3749893.3749959

RESEARCH-ARTICLE

SENTIMAP: Spatiotemporal Mapping of Emotions in Historical Newspapers using LLMs

TOPRAK FIRAT, Kadir Has University, Istanbul, Turkey

BESTE ILERI, Koç University, Istanbul, Turkey

Open Access Support provided by:

Kadir Has University

Koç University

PDF Download
3749893.3749959.pdf
30 March 2026
Total Citations: 0
Total Downloads: 341

Published: 05 November 2025

Citation in BibTeX format

Expanded 2025: Conference on Animation and Interactive Art
September 3 - 7, 2025
Linz, Austria

https://doi.org/10.1145/3749893.3749959
ISBN: 9798400715327

SENTIMAP: Spatiotemporal Mapping of Emotions in Historical Newspapers using LLMs

Toprak Firat
Kadir Has University
Istanbul, Turkiye
atoprakfirat@gmail.com

Beste Ileri
Koç University
Istanbul, Turkiye
ileribeste@gmail.coom

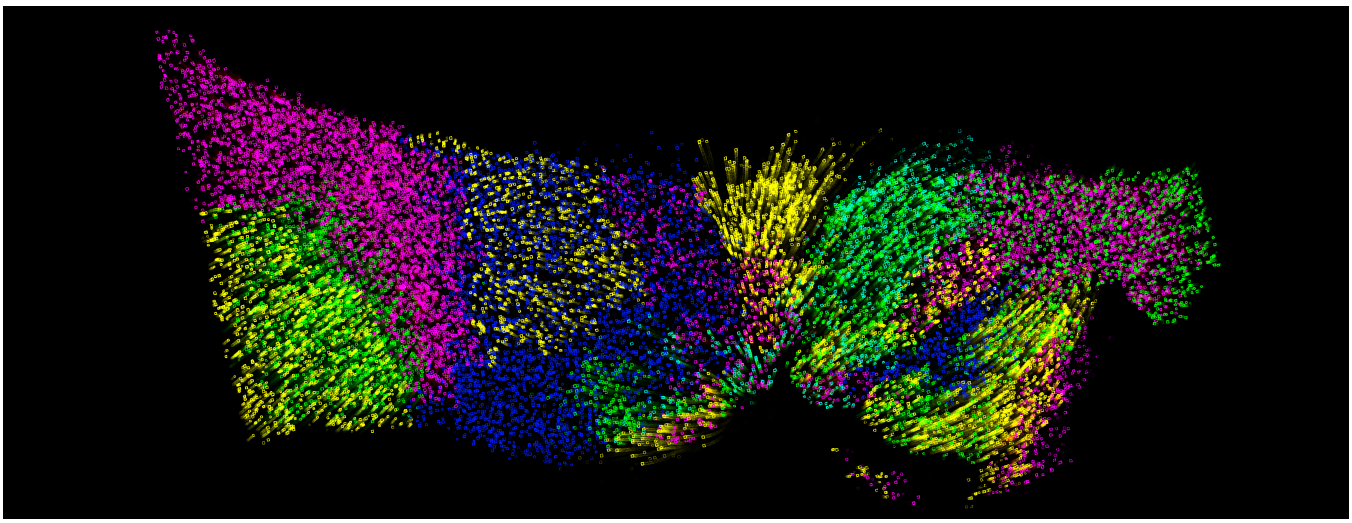


Figure 1: Still frame from the temporal flow map animation (Channel 1) showing district-level dominant emotions in the city of Istanbul. The choropleth visualization updates monthly from 1970 to 2023, assigning each district the color of its most frequently occurring emotion. Color encoding is as follows: yellow for anxious, magenta for happy, cyan for sad, green for hopeful, blue for neutral, and orange for angry. This continuous mapping allows viewers to trace shifts in affective geography over time.

Abstract

SENTIMAP is a data-driven visualization project that maps historical emotional trends across Turkish urban districts from 1970 to 2024, generating district-level maps where colors reflect emotional shifts over time. It simultaneously visualizes periodic emotional summits, emotional apexes, and aggregate emotional volume, offering a rich spatiotemporal understanding of how public feeling accumulates, intensifies, and transforms alongside socio-political developments. This enables a more nuanced reading of the emotional dimensions that unfold in tandem with historical change.

It adopts a more robust and flexible approach by leveraging large language models (LLMs), which excel at capturing context, tone, and nuance within complex linguistic structures. Unlike conventional techniques that rely on custom rule sets, labeled training data, or rigid pipelines, LLMs generalize emotional understanding across varied historical and stylistic texts with minimal preprocessing. This results in a more scalable and accurate emotional extraction process—especially valuable when working with decades of unstructured, archival media.

This approach is particularly significant for Turkish, a language that poses unique challenges for natural language processing (NLP) due to its agglutinative structure, extensive morphology, and complex grammar. These linguistic features often undermine the effectiveness of traditional emotional analysis methods such as lexicon-based scoring or statistical classifiers—challenges that LLMs are uniquely positioned to overcome.

This work is licensed under a Creative Commons Attribution 4.0 International License.
Expanded 2025, Linz, Austria
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1532-7/25/09
https://doi.org/10.1145/3749893.3749959

CCS Concepts

• **Human-centered computing** → **Visualization application domains**; *Empirical studies in HCI*; • **Computing methodologies** → **Natural language processing**; • **Information systems** → *Geographic information systems*; • **Applied computing** → *Digital libraries and archives.*

Keywords

Spatiotemporal Visualization, Emotion Detection, Emotion Mapping, Historical Newspapers, Large Language Models, Turkish NLP,

Zero-Shot Classification, Cultural Analytics, Affective Computing, Digital Humanities, Choropleth Visualization, Media Archaeology

ACM Reference Format:
Toprak Firat and Beste Ileri. 2025. SENTIMAP: Spatiotemporal Mapping of Emotions in Historical Newspapers using LLMs. In *Conference on Animation and Interactive Art (Expanded 2025), September 03–07, 2025, Linz, Austria.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3749893.3749959

1 Introduction

Sentiment analysis has proven highly effective in quantifying public opinion at scale [12]. Financial analysts employ it to anticipate market movements, political scientists track electoral dynamics, and businesses monitor brand perception across millions of reviews [9, 13, 14, 23]. These applications thrive where sentimental valence is explicit—such as tweets signaling approval or product reviews reflecting satisfaction. While binary sentiment classification suffices for marketing or brand monitoring, domains like mental health analysis, political discourse, crisis response, and cultural research require granular emotion detection [4, 15]. As the number of target emotional categories increases, the layered and overlapping nature of emotions introduces significant challenges, including the need for large annotated datasets and increased ambiguity in textual emotion expression.

This challenge becomes especially relevant when analyzing newspapers, which offer vast, structured records of public discourse. Unlike tweets or reviews where emotions are often stated directly, newspapers encode emotion subtly through lexical choices, narrative framing, and emphasis [18]. A headline like "Everything Is Under Control, Officials Assure Amid Rising Concerns" sounds calm and reassuring with words like everything, under control, and assure. Yet, the phrase “amid rising concerns” hints at growing anxiety and mistrust. Traditional sentiment analysis may interpret this as positive or neutral because of the reassuring language, missing the underlying feelings of skepticism and unease conveyed through subtle context.

Unlocking emotional dimensions of news texts requires sophisticated methods capable of interpreting indirect signals across time and cultural context. The challenges are magnified in morphologically complex languages like Turkish [21]. Turkish’s agglutinative nature, where multiple semantic elements combine into single words, breaks standard NLP approaches designed for English. Moreover, Turkish’s flexible word order and rich morphology make lexicon-based sentiment analysis particularly difficult, while the scarcity of labeled Turkish emotional datasets compounds these difficulties [16].

Large language models offer a fundamentally different approach addressing these challenges simultaneously. Unlike traditional methods, LLMs use attention mechanisms to understand entire textual sequences bidirectionally, capturing context, recognizing negation, handling morphological complexity, and detecting subtle emotional signals in seemingly neutral text [7, 20]. Most importantly, LLMs perform few-shot and zero-shot learning for emotional classification without requiring extensive labeled training data, making them uniquely suited for analyzing historical Turkish newspaper corpora [5].

What remains underexamined is the role of public emotion as a long-form narrative unfolding across time and space. While newspapers document events, emotion constitutes their affective subtext—a shifting collective mood responding to crises, elections, and reforms. This perspective is particularly significant in Türkiye, where state-controlled narratives and media censorship shape how events are emotionally coded and circulated.

SENTIMAP is a spatiotemporal mapping system applying zero-shot emotion classification using large language models to Istanbul-based newspaper articles from the 1970s to 2023. It processes large-scale archival corpus, identifying district-level references through place-name detection and geolocating articles accordingly. This enables construction of an affective cartography reflecting the city’s evolving emotional landscape across space and time. The project repurposes tools often associated with mass influence into instruments of counter-analysis, presenting results as a two-channel video installation animating Istanbul’s affective history over 54 years.

2 Gathering and Processing the Data

The study uses *Cumhuriyet Gazetesi*, one of Turkey’s oldest national newspapers, as its primary corpus, covering daily issues from **1970 to 2024** [6]. The research initially sought to incorporate multiple newspapers—including *Milliyet*, *Hürriyet*, and *Agos*—to enhance dataset breadth and enable cross-ideological comparison. However, practical constraints such as newspaper access restrictions due to state repression concerns and inadequate archival infrastructures necessitated exclusive use of *Cumhuriyet*. This newspaper was considered methodologically viable owing to its extensive digital archive and sustained commitment to independent journalism under political pressure.

The archive was provided as scanned page images. We developed a processing pipeline to segment individual articles and extract content via OCR. Each page was binarized, article regions detected using contour-based methods, and fragmented layouts reconstructed based on spatial proximity. All article blocks were passed through Tesseract OCR with Turkish language settings [10], then cleaned and stored. The process was parallelized to handle the archive’s scale. In total, **220,348 pages** were processed, yielding **1,157,974 article segments**. Each article was timestamped at the day/month/year level, with geographic attribution inferred by detecting district names in headlines and lead paragraphs. District-referenced news articles provide key methodological advantages for emotion analysis. Local news often reflects broader national dynamics through a local lens, and in contexts with media censorship, district-specific reporting offers less filtered affective expressions. This bottom-up approach captures emotional margins of public discourse while ensuring longitudinal consistency through uniform geographic filtering. To ensure accuracy, we accounted for historical changes in Istanbul’s administrative boundaries, including district renamings, mergers, and subdivisions, integrating these shifts into temporal mapping.

The corpus forms the foundation for spatiotemporal emotion analysis and SENTIMAP’s affective cartography of Turkish media discourse.



Figure 2: The processing pipeline for extracting and segmenting articles from scanned pages of *Cumhuriyet Gazetesi*. Each page is binarized, and text is first extracted via OCR using Tesseract. Subsequently, the extracted text blocks are grouped into individual articles by reconstructing fragmented layouts based on spatial proximity and contour-based region detection. The resulting structured text is stored for further analysis.

3 Methodology

At the heart of *SENTIMAP* lies a foundational assumption: that historical data of public emotion can be reconstructed from the events that shaped collective life and how these events were mediated through the press. Newspapers rarely describe emotional states directly—especially in politicized or censored contexts—but they chronicle what happened, what was emphasized, and what was silenced. A military coup, a factory strike, the opening of a circus—these are not just news items but affective indicators of the time. By systematically analyzing the flow of such events over five decades, *SENTIMAP* seeks to uncover the emotional undercurrents that animated Istanbul's public sphere.

The selection of emotion labels in SENTIMAP—happy, anxious, neutral, sad, hopeful, and angry—is informed by established psychological frameworks, including the models proposed by Ekman, Plutchik, and the OCC (Ortony, Clore, and Collins) model [8, 17, 19]. These categories were chosen to represent a theoretically grounded and contextually relevant spectrum of emotional states commonly observed in sociopolitical discourse. Basic emotions such as happiness, sadness, anger, and anxiety correspond to widely recognized affective responses, while the inclusion of hopeful reflects the significance of forward-looking sentiments in times of political uncertainty. The neutral category serves as both a baseline for emotional intensity and a potential indicator of media desensitization or suppression. Collectively, these labels enable a nuanced analysis of emotional dynamics within archival newspaper corpora and support a richer understanding of affective public discourse.

Turkish's agglutinative grammar presents significant challenges for standard NLP. Words often embed multiple grammatical layers, complicating surface-level analysis. For example, while *umut* (hope) is positive, its morphologically dense form *umutsuzlaştırabildiklerimizdenmişsiniz* ("you were apparently among those we could make hopeless") conveys complex negation, agency, and evidentiality. Lexicon-based models cannot resolve such shifts, and preprocessing tools like lemmatizers or stemmers fall short in restoring emotional polarity.

To address this challenge, we employed instruction-tuned large language models (LLMs) for zero-shot emotion classification in Turkish. Given the absence of openly available Turkish-specific emotion classification models, we systematically evaluated several multilingual LLMs using a testing dataset of 150 manually annotated Turkish newspaper excerpts balanced across six emotion categories. Our evaluation focused on decoder-only architectures that naturally support prompt-based classification, as encoder-based models present significant limitations for zero-shot tasks. While **GPT-4** demonstrated strong Turkish language comprehension [11], its computational cost proved prohibitive for our corpus scale. We subsequently evaluated **Llama2-7B**, **Zephyr-7B**, **OpenChat-3.5**, and **Mistral-7B** using standard classification metrics including accuracy, precision, recall, and F1-score [1–3, 22]. Each model was deployed using a standardized prompt template designed to ensure consistent zero-shot classification:

> "You are a Turkish emotion analysis bot. You will be given a message, and you will reply with one of the categories: `Happy, Hopeful, Anxious, Sad, Angry, Neutral`, or `X` if you cannot decide on anything. Reply only with the category name, nothing else. If it is an advertisement, reply with `X` as well."

The evaluation results revealed significant performance differences across models. **OpenChat-3.5** achieved the highest performance with 72% accuracy and 0.62 macro F1-score, followed by **Zephyr-7B** (70% accuracy) and **Mistral-7B** (66% accuracy, highest precision at 0.77). **Llama2-7B** significantly underperformed with only 52% accuracy. Based on these results, we selected **OpenChat-3.5** for large-scale analysis of five decades of news coverage, with each article assigned its most prominent emotion label.

4 Design Framework



Figure 3: Installation view of SENTIMAP as exhibited in From Sketch to Pixel, organized by the Ministry of Culture and Tourism of the Republic of Türkiye, at Galataport Istanbul, 2024. Photograph courtesy of Ayyıldız Media Agency, 2024.

SENTIMAP visualizes the emotional trajectory of Turkish public discourse by projecting sentiment data onto a district-level map, updated month by month from January 1970 to August 2023. Each article was assigned a single emotion label, and articles were aggregated by district and month to calculate emotion frequencies across time and geography.

The visualization consists of a choropleth map, where each district is colored based on the most frequently occurring emotion in that region for a given month. Color encoding is as follows: yellow for anxious, magenta for happy, cyan for sad, green for hopeful, blue for neutral, and orange for angry. The artwork, presented as a two-channel video, invites viewers to observe the overall process on the first screen through a flow video (38'53") where the map advances to the subsequent month every 4 seconds. Meanwhile, the second screen features a shorter video (1'40") presenting analyses via infographics, allowing for deeper engagement with the evolving emotional landscape. In addition to the yearly maps, SENTIMAP includes visualizations of several macro-level features.

As the visual encoding method in the second video, the following infographics were employed: The **Periodic Summits** infographic indicated the five years during which each emotion peaked, marked with an "X" on a grid, while a linear animation beneath it depicted the frequency of the emotion over the years. The **Feelings Apex** infographic mapped, across 54 years, the five districts where each emotion reached its highest recorded values, displayed over a city map. While offering a peak view of local emotional intensity, it depicts cross-regional spikes of shared emotional tone. The **Emotion Volume** visualization was constructed by layering spatial zones dominated by each emotion across the 54-year period, generating a volumetric representation of affective intensity over time. Semi-transparent annual layers were stacked to form a cumulative structure, where color saturation encoded regional density, and vertical stratification corresponded to temporal progression. A linear time marker, positioned laterally, anchored the chronology—extending from 1970 at the base to 2023 at the top. The resulting volume offers a longitudinal perspective on emotional distribution across the corpus, while the rotating animation reveals the structure from multiple angles enhancing spatial legibility and allowing viewers to intuitively grasp the temporal and geographic dynamics embedded within the emotional landscape.

Additionally, during the examination of each emotion, the Periodic Summits and Feelings Apex datasets were also presented

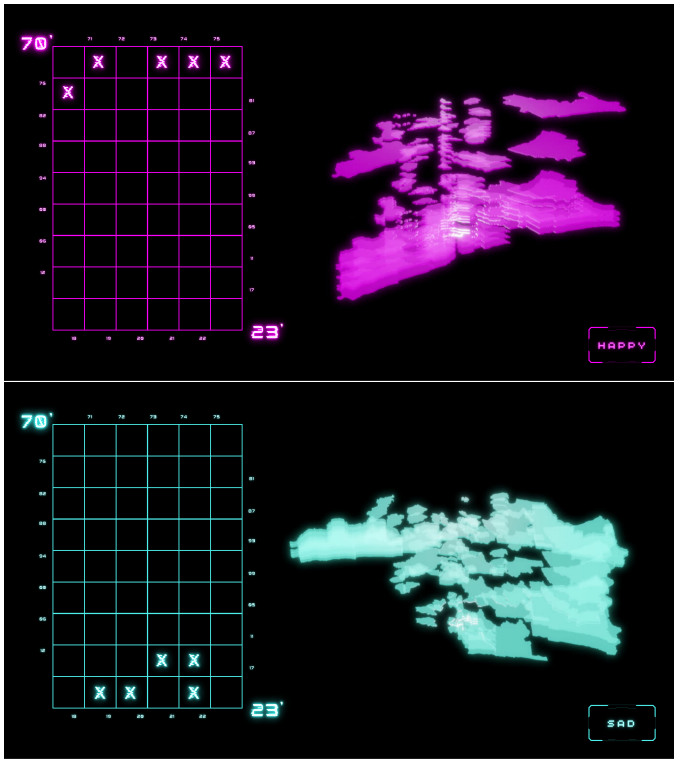


Figure 4: Periodic Summits and Emotion Volume infographics for the emotions happy(left) and sad(right). The summits chart highlights the five peak years in frequency for each emotion (marked with "X"), while the volume visualization shows the spatial accumulation of emotional intensity over 54 years. Volumetric layers correspond to temporal progression from 1970(base) to 2023(top), with color density indicating spatial concentration. Notably, peak years for happiness cluster exclusively in the 1970s, whereas sadness peaks emerge only after 2015. The volumetric forms reflect this shift: happiness contracts in recent decades, while sadness expands—demonstrating alignment between frequency peaks and sustained spatial density.

in text format, facilitating simultaneous comparison with other emotions.

SENTIMAP adopts a cognitively navigable interface that encourages attentional interaction rather than direct manipulation. Through its dual-channel structure, the work balances narrative continuity with comparative analysis: the temporal flow map offers an immersive, sequential traversal of emotional change, while the infographic loop enables cross-emotional, spatial, and temporal readings. Viewers are invited to engage selectively—pausing at familiar districts, aligning emotionally charged periods, or tracing thematic arcs. This interaction unfolds through visual correlation and affective resonance, allowing users to navigate the emotional history of Istanbul according to personal memory, political orientation, or interpretive curiosity. In this way, it functions as a semi-guided system of affective discovery—one that evokes rather

than instructs, and transforms observation into a mode of critical, situated engagement. This visual strategy was designed to balance analytic clarity with emotional immediacy, enabling both researchers and general viewers to intuitively explore emotional geography through time. Emphasis was placed on narrative flow – using colors and transitions that evoke the qualities of the emotions they represent – enabling users to "read" through history not just as a timeline, but as a mood map. Although grounded in technical accuracy, SENTIMAP also operates as a conceptual artwork. It reflects on how emotions are not just individual states but political signals; spatially distributed, historically produced, and collectively felt.

5 Discussion

The dataset collected through district-level localization was initially expected to primarily represent local news content; however, it demonstrated a significantly higher degree of political representativeness than anticipated.

The system revealed several distinct emotional patterns that aligned with major political and societal events in Turkey. By mapping sentiment over both space and time, the system uncovered how emotions clustered regionally, intensified during moments of national upheaval, and sometimes lingered long after the events that triggered them.

Analysis of **periodic summits** revealed that four out of the five peak years for the emotion of **hope** occurred during the 1970s, highlighting a critical decline in collective well-being and future expectations in Turkey. Similarly, five peak years for **happiness** were exclusively concentrated in the 1970s. On the other hand, peak years for **sadness** emerged solely after 2015, while **anger** peaked only after 2013, and **anxiety** after 2016.

The fact that levels of **anger**, **sadness**, and **anxiety** are currently higher than during historically traumatic events such as the 1980 military coup, widely considered the most nationally traumatic period, provides empirical evidence of the profound societal distress associated with the current government's political climate.

Even the highest values for **neutrality** were observed exclusively during the 1970s, suggesting not only a deterioration in overall sentiment but also a structural shift toward a more emotionally charged and provocative news discourse over time.

From **feelings apex**, it was observed that districts exhibiting the highest concentrations of **anxiety** and **sadness** correspond to socioeconomically disadvantaged areas, often associated with urban ghettos. Moreover, districts showing peak values for different **negative emotions** (anger, sadness, and anxiety) largely overlap.

Notably, *Sultangazi*, *Sultanbeyli*, and *Sancaktepe* consistently appear among the top five districts with the highest emotional intensity for each of these three negative emotions across the entire 54-year period. *Ataşehir* appears at the apex for two negative emotions — anger and anxiety — while *Esenyurt* appears for anger and sadness.

Conversely, *Beşiktaş*, *Fatih*, and *Başakşehir* emerge as districts consistently associated with peak values in **positive emotions** (hope and happiness), demonstrating the coexistence and spatial consistency of positive affective patterns within the urban landscape.

5.1 Emotional volume

The **emotion volume** analysis in Istanbul over the past 54 years reveals that anxiety has been the dominant emotion, occupying the largest space. While the volumes of hope and happiness exhibit a marked decline in the upper strata—corresponding to recent years—the volume of sadness displays a notable increase over the same period.

The Emotion Volume analysis demonstrated strong alignment with the Periodic Summits infographic, particularly for emotions of hope, happiness, and sadness. This convergence is significant, as Periodic Summits highlights the years in which each emotion reached its numerical peak, but does not capture whether those peaks reflect sustained emotional intensity over time. In some cases, peak years may occur as isolated spikes and thus risk misrepresenting the broader emotional trajectory. However, the overlap between peak frequency years and the volumetric spatial density of these emotions suggests a mutually reinforcing pattern, strengthening the validity of these indicators.

Notably, fear was found to be spatially non-dominant across nearly the entire corpus, surfacing only once in a geographically significant way in 2022—again corresponding with its identified peak year. While we would have expected anger to be the emotion with the lowest value, in line with the neutral tone of news reporting, the recent emergence of anger aligns with the changing and increasingly provocative nature of news discourse in Turkey. Anxiety, by contrast, displayed a relatively consistent spatial spread across most of the 54-year period, indicating its longstanding presence in public discourse.

Although its peak years cluster toward the recent past, this pattern may reflect not only persistent societal unease but also the tendency for anxiety-inducing topics to dominate news coverage. The case of neutral sentiment presents a particularly thought-provoking discrepancy: while its top five peak years are concentrated in the 1970s, the volumetric analysis shows its spatial dominance emerging in more recent years. This divergence invites further inquiry into the nature of "neutral" reporting—whether interpreted as a marker of journalistic objectivity or, alternatively, as a symptom of media suppression. If the latter, the historical pattern may point to abrupt, censorial responses to traumatic events in the past, whereas the contemporary expansion of neutrality could signal a more systematic and enduring form of editorial control.

5.2 A Temporal Analysis of Key Historical Events

Figure 5 presents the evolution of emotional deviations across Istanbul between 1970 and 2023, as detected through newspaper sentiment analysis. The heatmap visualizes relative emotional deviation, where values are mean-centered and range-scaled for each sentiment category. Positive deviations represent higher-than-average emotional intensities, while negative deviations reflect suppressed emotional expressions. Notable emotional disturbances correspond closely to key political, economic, and natural events throughout the period.

Here's an analysis of how emotions fluctuate around several major historical events, based on changes in the month before, during, and after each event:

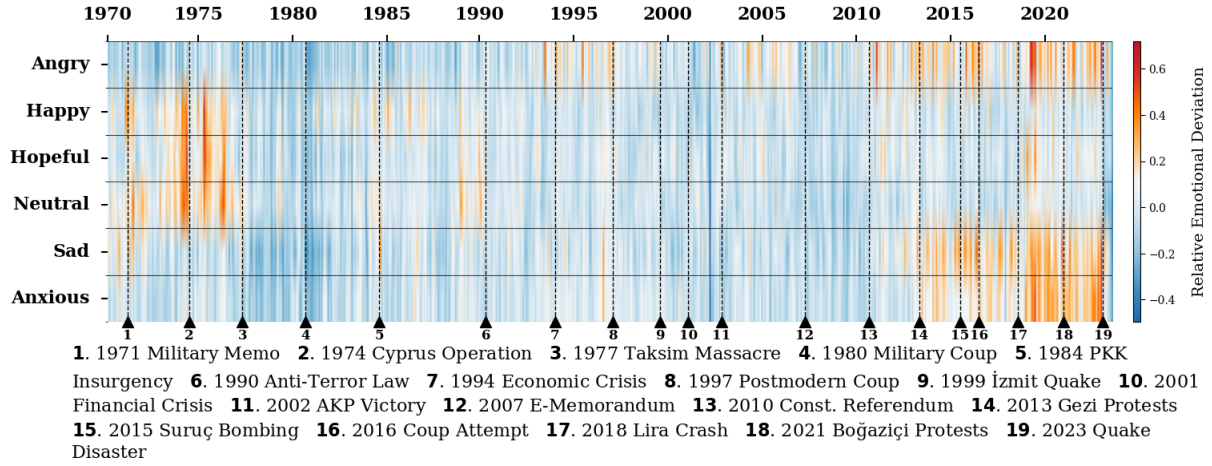


Figure 5: Temporal analysis of emotional fluctuations in Istanbul (1970–2023). The heatmap displays mean-centered, range-scaled deviations in emotional intensity derived from monthly newspaper sentiment data. Positive values indicate above-average emotional expression; negative values reflect relative suppression. Fluctuations closely align with major historical events, including political, economic, and natural crises, highlighting the affective imprint of collective experience over time.

Military Coups and Political Coups. The military coups (1971, 1980, 1997, 2016) consistently triggered sharp increases in sadness and anxiety. Hope levels dropped significantly, while anger remained relatively muted, suggesting a predominance of fear and instability over rage during these periods.

Natural Disasters (Earthquakes). Major earthquakes (1999 İzmit, 2020 İzmir, 2023 Kahramanmaraş) caused massive spikes in sadness and anxiety. These events correspond to the highest sadness values in the dataset, with hope and happiness collapsing immediately afterward.

Economic Crises. The economic crises of 1994, 2001, 2018, and 2022 resulted in consistent drops in hope and happiness, alongside sharp rises in anxiety. Sadness also increased moderately. Anger remained relatively subdued compared to political protest periods.

Protests and Civil Movements. Civil protests, particularly the 2013 Gezi Park protests and the 2021 Boğaziçi University protests, were associated with clear spikes in anger. Hope declined during these periods, while anxiety rose moderately. Gezi Park protests showed the strongest anger surge across the entire study period.

Terror Attacks. The Suruç bombing and the Ankara Peace Rally bombing in 2015 led to immediate and intense rises in sadness and anxiety, with neutral sentiment collapsing. Anger increased but was secondary to the dominant sadness response.

Elections and Political Transformations. Election periods and political openings (2002, 2003, 2005, 2023, 2024) generally produced rises in hope and happiness, reflecting optimism during democratic transitions. However, the 2017 Constitutional Referendum was marked by increased anger, suggesting polarization rather than unified optimism.

Other Major Political Shifts. Events such as the 1978 Ecevit coalition, the 1979 martial law extension, the 1982 Constitution referendum, the 1984 start of PKK armed insurgency, and the 1985 EU application showed moderate emotional shifts. Martial law and insurgency triggered anxiety increases, while EU integration steps slightly boosted hope. The 1999 capture of Öcalan generated a mixed emotional response, combining rises in hope with persistent anger.

6 Conclusion

SENTIMAP introduces a novel framework for tracing the emotional history of a society through the long arc of newspaper archives. By combining large language models, spatial visualization, and historical media analysis, it transforms decades of unstructured Turkish newspaper text into a spatiotemporal map of collective sentiment—revealing how public emotions move, accumulate, and rupture alongside political upheavals, cultural transformations, and regional shifts.

Unlike traditional sentiment analysis pipelines, which flatten affect into binary categories and require language-specific resources, SENTIMAP employs zero-shot LLMs to extract nuanced emotional states directly from noisy, long-form texts. This approach enables scalable emotion mapping across underrepresented languages such as Turkish, without reliance on rigid lexicons, labeled datasets, or brittle rule-based heuristics.

Beyond its technical contributions, SENTIMAP demonstrates how computational methods can enrich the fields of digital humanities, affective media archaeology, and data-driven art. It highlights how ephemeral media not only records facts, but also encodes the affective atmospheres of its time—atmospheres that can now be excavated, visualized, and critically explored.

Future work will expand SENTIMAP along multiple dimensions: incorporating temporal emotional intensity, allowing articles to express blended emotional states; extending the framework to multilingual corpora for comparative cross-cultural analysis; and developing dynamic emotional networks that capture the co-occurrence and diffusion of feelings over space and time using fine-tuned models that specialize in the subject.

Acknowledgments

We sincerely thank Istanbul Museum of Modern Art and curators Ümit Mesci and Nilay Dursun for their kind invitation and thoughtful curatorial guidance throughout the development of SENTIMAP as part of the exhibition Timeless Curiosities.

References

[1] [n. d.]. Mistral 7B. https://www.mistral.ai. Accessed: 2025-06-22.
[2] [n. d.]. OpenChat-3.5. https://huggingface.co/OpenChat/OpenChat-3.5. Accessed: 2025-06-22.
[3] [n. d.]. Zephyr-7B. https://github.com/OpenRLHF/Zephyr. Accessed: 2025-06-22.
[4] Francisca Adoma Acheampong, Chen Wenyu, and Henry Nunoo-Mensah. 2020. Text-based emotion detection: Advances, challenges, and opportunities. *Engineering Reports* 2, 7 (2020), e12189.
[5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
[6] Cumhuriyet Gazetesi. [n. d.]. Cumhuriyet Gazetesi Dijital Arşiv.
[7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. arXiv:1810.04805 [cs.CL] https://arxiv.org/abs/1810.04805
[8] Paul Ekman. 1992. An argument for basic emotions. *Cognition & emotion* 6, 3-4 (1992), 169–200.
[9] Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*. 168–177.
[10] Kurtuluğ Karasu and Muhammet Bağtan. 2015. Turkish OCR on mobile and scanned document images. In *2015 23nd Signal Processing and Communications Applications Conference (SIU)*. 2074–2077. doi:10.1109/SIU.2015.7130278
[11] Mustafa Eray Kılıç. 2023. AI in medical education: a comparative analysis of GPT-4 and GPT-3.5 on Turkish medical specialization exam performance. *medRxiv* (2023), 2023–07.
[12] Bing Liu. 2022. *Sentiment analysis and opinion mining*. Springer Nature.
[13] Tim Loughran and Bill McDonald. 2011. When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of finance* 66, 1 (2011), 35–65.
[14] Kostadin Mishev, Ana Gjorgjevikj, Irena Vodenska, Lubomir T Chitkushev, and Dimitar Trajanov. 2020. Evaluation of sentiment analysis in finance: from lexicons to transformers. *IEEE access* 8 (2020), 131662–131682.
[15] Pansy Nandwani and Rupali Verma. 2021. A review on sentiment analysis and emotion detection from text. *Social network analysis and mining* 11, 1 (2021), 81.
[16] Kemal Oflazer. 1994. Two-level description of Turkish morphology. *Literary and linguistic computing* 9, 2 (1994), 137–148.
[17] Andrew Ortony, Gerald L Clore, and Allan Collins. 2022. *The cognitive structure of emotions*. Cambridge university press.
[18] Flor Miriam Plaza-del Arco, Alba Curry, Amanda Cercas Curry, and Dirk Hovy. 2024. Emotion analysis in NLP: Trends, gaps and roadmap for future directions. *arXiv preprint arXiv:2403.01222* (2024).
[19] Robert Plutchik. 1991. *The emotions*. University Press of America.
[20] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training. (2018).
[21] Kadir Tohma and Yakup Kutlu. 2020. Challenges Encountered in Turkish Natural Language Processing Studies. *Natural and Engineering Sciences* 5, 3 (2020), 204–211. doi:10.28978/nesciences.833188
[22] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023).
[23] Andranik Tumasjan, Timm Sprenger, Philipp Sandner, and Isabell Welpe. 2010. Predicting elections with twitter: What 140 characters reveal about political sentiment. In *Proceedings of the international AAAI conference on web and social media*, Vol. 4. 178–185.