Distribution Agreement

In presenting this thesis as a partial fulfillment of the requirements for a degree from Emory University, I hereby grant to Emory University and its agents the non-exclusive license to archive, make accessible, and display my thesis in whole or in part in all forms of media, now or hereafter now, including display on the World Wide Web. I understand that I may select some access restrictions as part of the online submission of this thesis. I retain all ownership rights to the copyright of the thesis. I also retain the right to use in future works (such as articles or books) all or part of this thesis.

Hava Collins April 9, 2025

Reason and Rhetoric: Measuring Emotionality in Supreme Court Opinions

by

Hava Collins

Dr. Sandeep Soni
Adviser

Quantitative Theory and Methods

Dr. Sandeep Soni

Adviser

Dr. Emma Davenport

Committee Member

Dr. Kevin Quinn

Committee Member

2025

Reason and Rhetoric: Measuring Emotionality in Supreme Court Opinions

By

Hava Collins

Dr. Sandeep Soni
Adviser

An abstract of
a thesis submitted to the Faculty of Emory College of Arts and Sciences
of Emory University in partial fulfillment
of the requirements of the degree of
Bachelor of Science with Honors

Quantitative Theory and Methods

2025

Abstract

Reason and Rhetoric: Measuring Emotionality in Supreme Court Opinions
By Hava Collins

This thesis investigates the presence and evolution of emotional rhetoric in U.S. Supreme Court opinions, challenging longstanding assumptions about the Court as a purely rational institution. While legal scholars have noted emotionality in individual dissents or landmark rulings, no large-scale, systematic study has quantified how emotion manifests across judicial writing—or how it changes over time. To address this gap, I develop a computational framework for measuring emotionality in Supreme Court opinions using both traditional lexicon-based methods and advanced vector-space modeling techniques.

The analysis draws from a large corpus of tokenized Supreme Court opinions, spanning multiple decades and encompassing majority, dissenting, and concurring opinions. First, a ratio-based scoring method uses domain-specific dictionaries to compute the prevalence of emotional versus rational language. Second, a Doc2Vec-based model captures semantic nuance by projecting each opinion into a vector space and scoring emotionality relative to custom emotion-reason axes. The results reveal a nuanced historical trajectory: emotionality in dissenting opinions has increased significantly since the 1970s, particularly during moments of heightened ideological division and legal controversy, while majority opinions have remained more rhetorically restrained.

Crucially, the vector-based approach uncovers emotionally strategic language that traditional methods miss—suggesting that even "neutral" judicial rhetoric may carry emotional weight through subtler linguistic cues. This dual-method approach demonstrates that the emotionality of judicial opinions is not merely a function of vocabulary but of rhetorical framing and historical context.

The findings carry important interdisciplinary implications. They challenge the notion of the Court's neutrality, highlight the rhetorical role of emotion in legal argumentation, and establish computational emotionality analysis as a viable tool in legal and political scholarship. Ultimately, this thesis offers a replicable methodology for analyzing sentiment in legal texts and contributes to a broader understanding of how emotion shapes legal reasoning and public perception of the judiciary.

Reason and Rhetoric: Measuring Emotionality in Supreme Court Opinions

By

Hava Collins

Dr. Sandeep Soni
Adviser

A thesis submitted to the Faculty of Emory College of Arts and Sciences
of Emory University in partial fulfillment
of the requirements of the degree of
Bachelor of Science with Honors

Quantitative Theory and Methods

2025

Acknowledgements

I would like to thank Dr. Soni, Dr. Davenport, and Dr. Quinn for their continued support over the course of this past year, and to the entire Emory QTM department for their exceptional wisdom and expertise. To Laura Salem, Sarah Lande, Mica Einhorn, and Zachary de Haaff, thank you for your input, advice, and encouragement. Finally, to my family, and especially my parents Shoshana and Pablo Collins, I could not have done any of this without you. Thank you.

Table of Contents

Reason and Rhetoric: Measuring Emotionality in Supreme Court Opinions

Chapter 1: Introduction

1.1 Problem Definition

The language used by the Supreme Court is inherently influential, shaping not only the immediate legal landscape but also broader historical and social understandings of judicial decision-making. Although Supreme Court justices are traditionally perceived as exemplars of rational deliberation and objective reasoning, the Court's written opinions frequently incorporate subtle forms of emotional rhetoric, adding persuasive power and nuance to legal arguments. For example, Justice Scalia's dissent in *Obergefell v. Hodges* vividly showcases emotional rhetoric aimed at reinforcing judicial arguments, while Justice Ginsburg's passionate advocacy for gender equality in cases such as *United States v. Virginia* underscores emotional language's role in emphasizing moral and social imperatives. Despite these notable examples, the systematic study of emotionality within Supreme Court opinions remains significantly underexplored, especially from a quantitative perspective.

Existing scholarship on Supreme Court opinions has predominantly examined ideological framing, linguistic complexity, or rhetorical devices, but it rarely rigorously quantifies emotional content. Traditional general-purpose sentiment analysis methods, commonly employed in other domains, prove insufficient for accurately discerning emotional nuances within the highly specialized legal context. These generic approaches frequently misclassify legal language due to the distinctive syntax, formalized style, and specialized vocabulary inherent to judicial writing, thus creating a significant methodological gap. This thesis addresses this gap by systematically exploring the presence, role, and historical evolution of emotional rhetoric in Supreme Court opinions through advanced computational methods tailored explicitly to judicial language.

1.2 Challenges and Methodological Innovation

Analyzing emotionality in Supreme Court texts presents substantial methodological challenges due to the sophisticated, context-dependent, and formal rhetorical strategies used by justices. Unlike everyday communication or political discourse, judicial opinions often embed emotional content subtly within precise and formal legal arguments rather than overt emotional declarations. Thus, accurately detecting and quantifying emotionality within judicial texts necessitates advanced computational methods designed to understand and interpret the nuanced and context-sensitive nature of legal language.

Prior methodologies primarily relied on dictionary-based or general sentiment analysis methods, both of which are prone to misclassification errors when applied to judicial texts. Recognizing these limitations, this thesis advances beyond traditional approaches by developing and applying innovative computational frameworks, specifically designed for the legal domain. Initially, baseline methods such as dictionary-based emotional-to-rational term ratio scoring and TF-IDF weighting are employed to establish a foundational understanding of emotional rhetoric.

Subsequently, to capture deeper contextual nuances, advanced methods incorporating vector-space modeling—specifically Doc2Vec, a neural embedding approach that jointly captures word-level and document-level semantic information—are introduced. Emotionality scores are computed by projecting each judicial opinion onto an emotional-rational axis derived from carefully constructed lexicons, providing a more precise and contextually sensitive measurement of emotional content.

1.3 Anticipated Findings

Employing these advanced methods, this thesis anticipates uncovering several key patterns and trends in the emotional rhetoric of Supreme Court opinions. First, even baseline methods are expected to reveal a general increase in emotionality over historical periods, reflecting broader shifts in judicial communication and rhetorical norms. Second, dissenting opinions are predicted to consistently exhibit higher emotionality scores than majority opinions, a pattern indicative of strategic rhetorical differences rooted in the persuasive function of dissents. Moreover, significant spikes in emotional rhetoric are expected to emerge prominently around the mid-to-late twentieth century, coinciding with landmark civil rights legislation, social movements, and periods of heightened ideological tension within the Court.

These anticipated results provide critical insights into how historical contexts and judicial roles influence rhetorical strategies, reflecting emotionality as a deliberate judicial communication tactic rather than an incidental or undesirable feature of legal argumentation.

1.4 Implications and Significance

The findings of this research bear important implications across scholarly, institutional, and public domains. First, systematically quantifying emotional rhetoric in Supreme Court opinions challenges traditional assumptions about judicial objectivity and neutrality, opening a more nuanced dialogue about how emotion strategically functions within institutional rhetoric. Second, this research contributes methodologically, demonstrating the efficacy and necessity of computational linguistic tools explicitly designed for legal contexts, thereby bridging significant interdisciplinary gaps between legal scholarship, computational linguistics, and historical analyses.

Furthermore, at a moment when judicial institutions face increasing scrutiny concerning their legitimacy, impartiality, and role in public life, understanding the impact of emotional rhetoric on public perceptions and judicial authority becomes crucial. By exploring how emotionality might reinforce or undermine judicial legitimacy, this thesis also provides empirical grounding for ongoing debates about judicial communication strategies and their broader societal consequences.

1.5 Contributions of this Thesis

In addressing these challenges and pursuing the research outlined above, this thesis offers several significant contributions:

1. Methodological Advancement: By developing a computational approach tailored explicitly to legal texts—moving from baseline ratio-scoring methods to sophisticated vector-space modeling (Doc2Vec)—this work establishes a rigorous and replicable methodological foundation for future studies of judicial emotionality.
2. Empirical Insights: Through comprehensive computational analyses of an extensive corpus of Supreme Court opinions, this research systematically identifies previously unexplored patterns in judicial rhetoric, such as the comparative emotionality of dissents versus majority opinions and historical shifts tied to broader societal and institutional changes.
3. Interdisciplinary Integration: Bridging computational linguistics, legal analysis, and historical inquiry, this research illustrates the benefits of interdisciplinary collaboration, offering novel insights that advance scholarship within and beyond traditional disciplinary boundaries.

In summary, by rigorously quantifying and analyzing emotional rhetoric in Supreme Court opinions, this thesis not only addresses a notable gap in existing scholarship but also provides essential guidelines and methodological frameworks to inspire and facilitate future interdisciplinary research into the complex interplay of emotion, rhetoric, and judicial reasoning.

Chapter 2: Related Works

Understanding how emotion is expressed and managed in legal texts is increasingly recognized as critical within legal studies, computational linguistics, and political science. This chapter reviews existing literature organized into three key subsections: Section 2.1 addresses the impact of emotional language in legal discourse, Section 2.2 explores computational approaches to legal text analysis, and Section 2.3 evaluates existing emotion detection methods, emphasizing their adaptation to the legal domain. Throughout this chapter, existing works are compared explicitly with the approach and innovations presented in this thesis.

2.1 Impact of Sentiment in Legal Discourse

Legal texts—including judicial opinions, briefs, and legal arguments—are traditionally regarded as purely rational, devoid of emotional influence. However, recent scholarship reveals that emotion plays a substantial role in legal argumentation and judicial communication. Krewson (2019), for instance, introduces the concept of “strategic sensationalism,” arguing that justices intentionally incorporate emotive language to influence public perception and judicial outcomes subtly. Similarly, Bryan and Ringsmuth (2016) analyze Supreme Court dissents, highlighting emotive rhetoric's dual function as both persuasive critique and a signal of ideological stakes.

Empirical studies (Escalante et al, 2015) further underscore that emotional rhetoric in legal briefs can negatively impact an attorney's perceived credibility. Rice and Zorn (2021) add methodological depth to these discussions by constructing corpus-based sentiment dictionaries tailored specifically for legal vocabularies, emphasizing the need for specialized lexicons capable of capturing nuanced emotional content that generic sentiment tools overlook.

In contrast to Krewson’s qualitative case-study method and Rice and Zorn’s lexicon-based dictionary approaches, this thesis employs advanced computational methods explicitly tailored to judicial opinions. By moving beyond qualitative descriptions and static lexicons, this project systematically quantifies emotional rhetoric across an extensive corpus, ensuring scalability, precision, and greater generalizability.

2.2 Computational Analysis of Legal Texts

Parallel to qualitative examinations, computational techniques have grown significantly, enabling more extensive and nuanced analysis of emotional rhetoric in legal texts. Recent advances in machine learning and natural language processing (NLP) have facilitated the detection and analysis of emotional patterns within large-scale corpora. Notably, studies such as "Classification of US Supreme Court Cases using BERT-Based Techniques" (2023) demonstrate that sophisticated transformer models can capture subtle linguistic cues indicative of judicial sentiment.

Similarly, Eliot’s "Legal Sentiment Analysis and Opinion Mining" (2020) outlines a comprehensive framework integrating AI techniques into legal reasoning. A rule-based

approach discussed by Sahoo et al. (2020) further illustrates that combining NLP with domain-specific rules provides meaningful insights into partisan sentiment embedded in legal texts. Additionally, multimodal studies such as those analyzing courtroom audio signals (Mujtaba et al., 2023) have expanded methodological boundaries, connecting emotional signals to legal outcomes directly.

Despite their advancements, these approaches often face significant limitations, such as computational complexity, interpretability issues, and domain-specific contextual misclassification. In contrast, this thesis offers a balanced computational approach employing vector-space modeling (including Doc2Vec) that enhances interpretability without compromising scalability, providing transparent and replicable analyses tailored explicitly to judicial opinions.

2.3 Existing Emotion Detection and Analysis Methods in the Legal Domain

Emotion detection in legal texts presents unique challenges due to specialized vocabulary, formality, and the strategic embedding of emotional cues. Researchers have historically employed diverse methods, ranging from early lexicon-based approaches to advanced transformer models, each with notable strengths and weaknesses.

2.3.1 Lexicon-Based Approaches

Lexicon-based sentiment analysis, exemplified by Rice and Zorn's (2021) legal dictionaries, assigns sentiment scores based on predefined word lists. Their methodology offers interpretability but faces significant limitations regarding context-sensitivity and dynamic adaptability over time.

In contrast, this thesis extends beyond static lexicons by integrating both baseline ratio methods and advanced vector-space models, allowing more dynamic, context-aware analysis. Unlike lexicon-based approaches, vector-space modeling can capture evolving rhetorical strategies within judicial texts.

2.3.2 Deep Learning and Transfer Learning Methods

Deep learning, particularly in low-resource contexts, has emerged as a flexible yet computationally intensive solution. Methods employing transfer learning (Sahoo et al., 2020) demonstrate effectiveness in capturing complex emotional expressions but suffer from low interpretability, limiting their application within transparent judicial contexts.

This thesis addresses these limitations explicitly by favoring computational transparency, prioritizing interpretable vector-space models like Doc2Vec over black-box deep learning methods, thus balancing performance with explanatory power essential in legal analysis.

2.3.3 Transformer-Based Techniques

Transformer models such as BERT provide powerful sentiment analysis capabilities through their contextualized embeddings (Costa et al). However, despite their impressive accuracy, transformers are resource-intensive and often opaque, complicating their deployment in legal contexts where interpretability is critical.

Contrary to transformer-based approaches, the current research focuses on computationally efficient and interpretable methods, utilizing structured emotionality scoring based on vector-space representations that maintain analytical transparency without sacrificing scalability.

2.3.4 Hybrid and Qualitative Approaches

Hybrid methodologies combining qualitative discourse analysis with computational sentiment tools offer detailed, context-sensitive insights but are limited by their scalability and reliance on manual analysis. For instance, studies exploring how judicial language evolves strategically (Black et al., 2012) provide invaluable qualitative depth yet fail to scale effectively for large corpus analyses.

This research overcomes these limitations by adopting a fully computational framework capable of systematically analyzing extensive judicial corpora, thereby enabling consistent and reproducible results at scale.

2.3.5 Datasets and Domain Adaptation

Availability of appropriate datasets presents another significant challenge. General-purpose sentiment datasets, such as GoEmotions (Google Research), are insufficiently tailored to the complexities of legal discourse. Applying these datasets to judicial texts frequently results in inaccuracies due to domain mismatch.

Unlike existing methods dependent on generic sentiment datasets, this thesis addresses domain adaptation explicitly, constructing specialized emotionality scoring systems directly aligned with judicial language characteristics. By integrating customized lexicons and computationally sophisticated vector-space methods, this approach ensures precise emotional content detection uniquely suited to judicial texts.

Summary

While existing research in legal sentiment analysis has significantly advanced, notable limitations persist. Lexicon-based methods provide interpretability but lack context sensitivity, deep learning methods excel in flexibility but lack transparency, and transformer models offer high accuracy but limited interpretability and resource-intensiveness. Qualitative and hybrid approaches, though rich in detail, cannot scale effectively. Moreover, reliance on generic datasets exacerbates domain-specific classification challenges.

In contrast, this thesis addresses these gaps through methodological innovation, leveraging scalable, interpretable vector-space modeling explicitly adapted for Supreme Court opinions. By refining sentiment analysis techniques tailored to legal discourse, this research offers new methodological frameworks, empirical insights, and interdisciplinary advances, substantially contributing to the understanding of emotional rhetoric in judicial decision-making.

Chapter 3: Methods

3.1 Data Collection and Preprocessing

This study analyzes Supreme Court opinions spanning multiple decades, sourced from CourtListener, Harvard CaseLaw and other publicly available legal archives. The dataset includes majority, dissenting, and concurring opinions, allowing for a comparative analysis of emotional expression across judicial perspectives. Each opinion is accompanied by metadata, including the case name, decision date, docket number, and opinion type.

The raw dataset consists of structured JSONL files, where each entry contains a tokenized Supreme Court opinion. During preprocessing, a document ID prefix was added to each opinion to ensure that all text could be traced back to its original case for downstream analysis. The preprocessing pipeline included several steps to prepare the dataset for computational analysis. First, tokenization was applied to segment each opinion into individual words using natural language processing tools such as nltk and spaCy. Next, the text was cleaned by removing non-substantive elements, including case citations, footnotes, and procedural statements that did not contribute to the linguistic analysis. Common legal stopwords, such as "court" and "plaintiff," were filtered out to reduce noise, while all text was converted to lowercase to maintain consistency. Finally, lemmatization was applied to reduce words to their base forms, ensuring that variations of the same word were treated as equivalent.

Once preprocessing was complete, the structured data was stored in a high-efficiency format to facilitate large-scale computation. All preprocessing was performed in a Python-based environment and optimized for efficiency using parallel processing on Emory University's QTM computational server. This approach allowed for scalable and reproducible text processing, ensuring that the dataset remained consistent across different analytical methods.

3.2 Dictionary-Based Emotionality Scoring

The first approach to measuring emotionality relied on pre-constructed sentiment dictionaries specifically designed for legal text analysis. While general-purpose sentiment lexicons, such as LIWC and the NRC Emotion Lexicon, are widely used in sentiment analysis, they lack the specificity required to capture the nuances of legal language. As a result, this study employed custom-built emotion and reason dictionaries, sourced from Gennaro & Ash (2021) in *Emotion and Reason in Political Language* as a baseline and supplemented with additional manually curated term lists tailored for Supreme Court opinions. These dictionaries were designed to distinguish between emotional and rational or 'cognitive' terms, thereby providing a structured framework for quantifying emotionality in legal writing.

To compute emotionality scores, a ratio-based approach was applied, which calculates the proportion of emotion-related words relative to reason-related words within each opinion. The

emotionality ratio for an opinion was computed using the formula:

$$E_{\text{ratio}} = \frac{\text{Count of Emotion Words}}{\text{Count of Rational Words} + \epsilon}$$

Figure 1: Ratio Scoring Construction

where ϵ is a small smoothing factor used to prevent division by zero. This method allows for an intuitive interpretation of emotionality: higher ratios indicate a greater presence of emotional language relative to rational language. Scores were computed separately for majority, dissenting, and concurring opinions, enabling direct comparisons across opinion types. This dictionary-based approach, while straightforward and interpretable, has inherent limitations, including potential misclassification of words and sensitivity to lexicon choice, which may introduce variability in results.

3.3 Vector-Based Emotionality Scoring

To address the limitations of dictionary-based scoring and provide a more nuanced, context-sensitive measure of emotionality, this study employed a vector-based approach using word embeddings. Unlike lexicon methods, which rely on predefined lists of words, vector-based models capture the semantic relationships between words and can recognize context-dependent meaning. This method enables a deeper understanding of emotionality by considering the broader linguistic structure of judicial opinions.

The vector-based scoring method was implemented using Doc2Vec, a neural network model trained specifically on the Supreme Court opinion corpus. Unlike traditional Word2Vec, which learns word-level embeddings, Doc2Vec generates embeddings for entire documents, allowing for a holistic representation of each opinion's semantic structure. The training process involved feeding tokenized opinions into the model, which then learned a multi-dimensional representation of each document based on its contextual usage of words. These embeddings captured the latent emotional and rational characteristics of each opinion, independent of predefined lexicons.

To quantify emotionality, opinion embeddings were projected onto an emotion-reason axis. This axis was constructed by aggregating word vectors for emotion-related and reason-related terms, using the pre-constructed sentiment dictionaries as anchor points. The emotionality score of an opinion was then computed as its cosine similarity to the emotion vector, relative to the reason vector. This projection allowed for a continuous measure of emotionality, distinguishing opinions based on their overall rhetorical style rather than the frequency of individual words.

$$\text{Emotionality Score} = \cos(\theta) = \frac{\text{Opinion Vector} \cdot (\text{Emotion Vector} - \text{Reason Vector})}{||\text{Opinion Vector}|| \cdot ||\text{Emotion Vector} - \text{Reason Vector}||}$$

Figure 2: Vector Scoring Construction

3.4 Statistical Analysis and Comparative Evaluation

Once emotionality scores were computed using both dictionary-based ratio scoring and vector-based scoring, statistical analyses were conducted to identify trends over time and compare emotionality across opinion types. The primary objective was to assess how emotional expression in Supreme Court opinions evolved historically and whether distinct patterns emerged between majority, dissenting, and concurring opinions.

To analyze these trends, mean emotionality scores were computed for each year, allowing for a longitudinal assessment of changes in judicial rhetoric. Regression models were used to examine the relationship between emotionality scores and key metadata attributes, including the opinion author, case type, and ideological alignment of the Court.

A critical component of this study was the comparison between ratio-based and vector-based emotionality scoring. By juxtaposing the results from both methods, the analysis revealed key differences in how emotionality is measured. The ratio-based method provided a broad historical overview but exhibited misclassification issues due to the rigid nature of sentiment dictionaries. In contrast, the vector-based approach offered a more context-aware assessment of emotionality, mitigating some of the errors introduced by lexicon-based classification. These findings underscore the importance of choosing appropriate computational methods when analyzing judicial rhetoric.

3.5 Validation and Limitations

Ensuring methodological robustness required multiple validation steps. First, the consistency of dictionary-based scoring was tested using alternative lexicons, revealing that different sentiment dictionaries produced slightly different results, emphasizing the subjectivity inherent in word classification. Second, the stability of the Doc2Vec model was evaluated by training on varied legal text subsets, confirming that embeddings remained relatively stable across different samples. Third, a convergence test was conducted to examine whether ratio-based and vector-based scoring methods produced similar patterns, helping to assess the reliability of each approach.

Several methodological limitations remain. Lexicon-based methods may misclassify words that have different connotations in legal contexts, while vector-based scoring relies on pre-trained embeddings that may not fully capture rhetorical nuance. Additionally, Supreme Court opinions vary widely in length and subject matter, meaning that some differences in emotionality scores

may reflect case content rather than judicial tone. These limitations suggest that future research should explore hybrid approaches that combine machine learning techniques with legal-specific lexicons for more precise sentiment analysis.

3.6 Computational Infrastructure

All data processing and analysis were conducted within a dedicated computational environment hosted on Emory University's QTM server. The implementation relied on Python-based workflows, leveraging libraries such as pandas, scikit-learn, gensim, and Matplotlib. Large-scale computations, including the training of Doc2Vec models and processing of hundreds of thousands of opinions, were executed using SLURM workload management software. This parallelized approach significantly improved efficiency, allowing the analysis of Supreme Court opinions at an unprecedented scale.

By integrating dictionary-based and vector-based methods, longitudinal statistical analysis, and robust validation techniques, this methodology provides a comprehensive framework for measuring emotionality in judicial opinions. The findings contribute to ongoing scholarly discussions on judicial rhetoric, sentiment analysis, and the intersection of computational linguistics and legal studies.

Chapter 4: Results

4.1 Overview of Emotionality Trends in Supreme Court Opinions

This study examines emotionality in Supreme Court opinions using both ratio-based and vector-based scoring methods, with findings revealing differences in the trajectory and magnitude of emotional expression depending on the method used. While both approaches indicate shifts in emotionality over time, their underlying methodologies produce distinct patterns in the data.

The ratio-based scoring method, which measures emotionality by calculating the proportion of emotional words relative to rational words, suggests a long-term decline in emotionality in majority opinions, a steady rise in dissenting opinions, and relative stability in concurring opinions. This method highlights an overall trend toward less overt emotional language in majority opinions, possibly reflecting increasing formalization in judicial writing. However, it also shows a sharp increase in emotionality in dissenting opinions after 1970, coinciding with greater ideological polarization on the Court.

The vector-based scoring method, which projects opinions onto an emotion-reason axis in a high-dimensional space, provides a more nuanced view of emotionality by capturing contextual meaning rather than relying solely on predefined word lists. Unlike the ratio method, vector-based scoring suggests that majority opinions have not necessarily become less emotional, but instead exhibit greater rhetorical complexity and strategic emotional framing. This distinction indicates that while explicit emotional language may have decreased in majority opinions, the underlying sentiment and persuasive strategies may still carry emotional weight. The vector method also reveals that dissents are not uniformly increasing in emotionality, but rather fluctuate depending on the nature of the cases and the historical context in which they were written.

Overall, the results indicate that emotionality in judicial writing has undergone considerable shifts, reflecting historical, institutional, and ideological transformations within the Court. However, the extent and nature of these changes depend on the analytical method used. While ratio-based scoring provides a broad historical trajectory of explicit emotional expression, vector-based scoring captures the underlying rhetorical structures that shape judicial sentiment. The differences between these methods highlight the complexity of analyzing legal sentiment computationally, underscoring the importance of methodological transparency and interdisciplinary approaches in understanding how Supreme Court opinions evolve over time.

4.2 Ratio-Based Emotionality Trends Over Time

The ratio-based approach, which measures the proportion of emotional words relative to rational words, demonstrates long-term trends in the use of emotional language in Supreme Court opinions. As shown in Figure 3, majority opinions have exhibited a steady average decline in emotionality from the mid-19th century onward. The sharp decline in the early 20th century coincides with the standardization of judicial opinion writing, reflecting the increasing formalization of the Court's language and the adoption of a more rigid institutional style.

However, emotionality in dissenting opinions spikes sharply in the 1950s and 1960s, aligning with the issuance of landmark civil rights decisions and subsequent rulings that shaped desegregation and voting rights. These findings support the hypothesis that periods of significant social and political upheaval are accompanied by increased emotional expression in judicial reasoning.

Figure 3: Ratio-Based Emotion Scores Over Time, Differentiated by Opinion Type

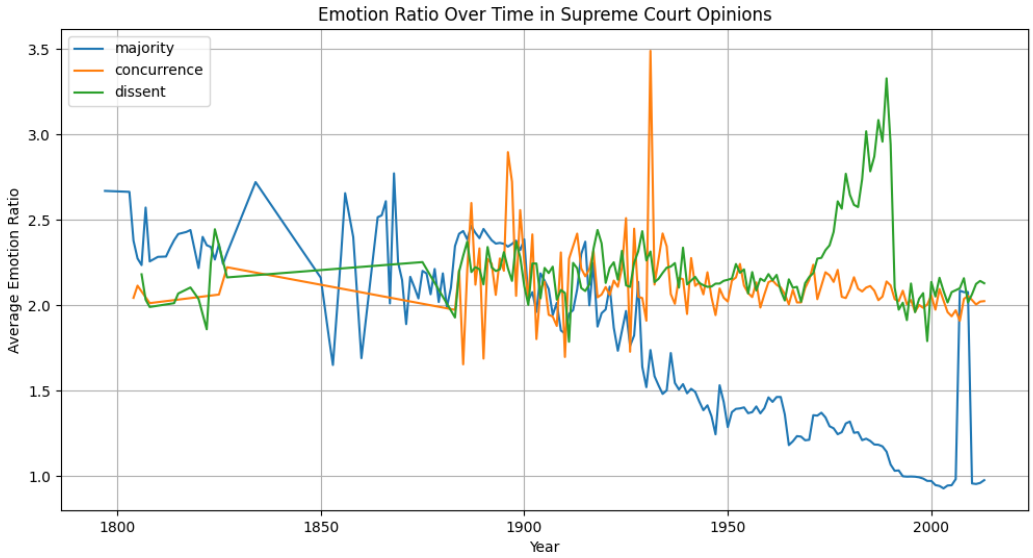

Dissenting opinions, in contrast, remain relatively stable in their emotional expression until the late 1960s, at which point they begin to rise steadily, before a stark decline to former levels around 1990 (Figure 3). The increase in dissenting emotionality during this period coincides with growing ideological divisions within the Court, particularly as contentious issues surrounding civil rights, reproductive rights, and federal power came to the forefront. The results from the ratio-based scoring method show that the hypothesized differences in emotionality, on average, between emotionality in dissents in majority is primarily due to an overall decrease in emotionality in majority opinions rather than an increase in dissents. Concurring opinions, as illustrated in Figure 3, display relatively low and stable emotionality over time, with fluctuations that are smaller in magnitude compared to both majority and dissenting opinions. The relative stability of concurrences supports the interpretation that these opinions, which often serve to clarify rather than contest majority reasoning, are less prone to the emotive rhetoric that characterizes dissents.

A broader analysis of undifferentiated ratio scoring using an unmodified dictionary, shown in Figure 4, reveals similar but slightly differing results. The undifferentiated ratio scores over time reveal an overall upward trend in emotionality in Supreme Court opinions, which is consistent with our initial hypothesis. However, the differentiated model (Figure 5) reveals that this is due

primarily to a significant increase in *majority* opinion scores over time. While patterns in concurring and dissenting opinions almost exactly match those of the modified dictionary method, majority opinion scores are reciprocal to the previous method. These results emphasize the importance of refining and testing various dictionary methods in future work.

Figure 4: Undifferentiated Ratio Scores Over Time Using an Unmodified Dictionary

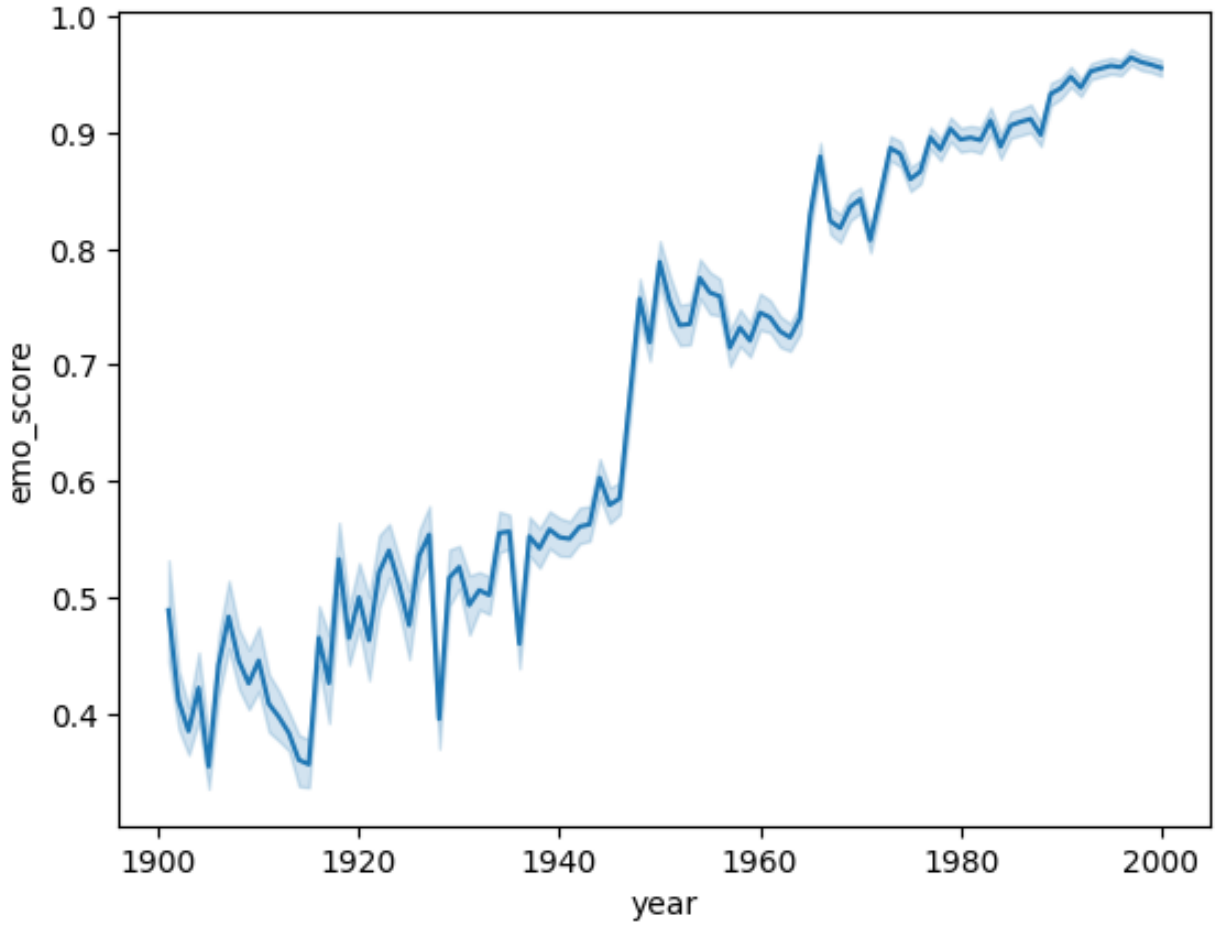

Figure 5: Differentiated Ratio Scores Over Time Using an Unmodified Dictionary

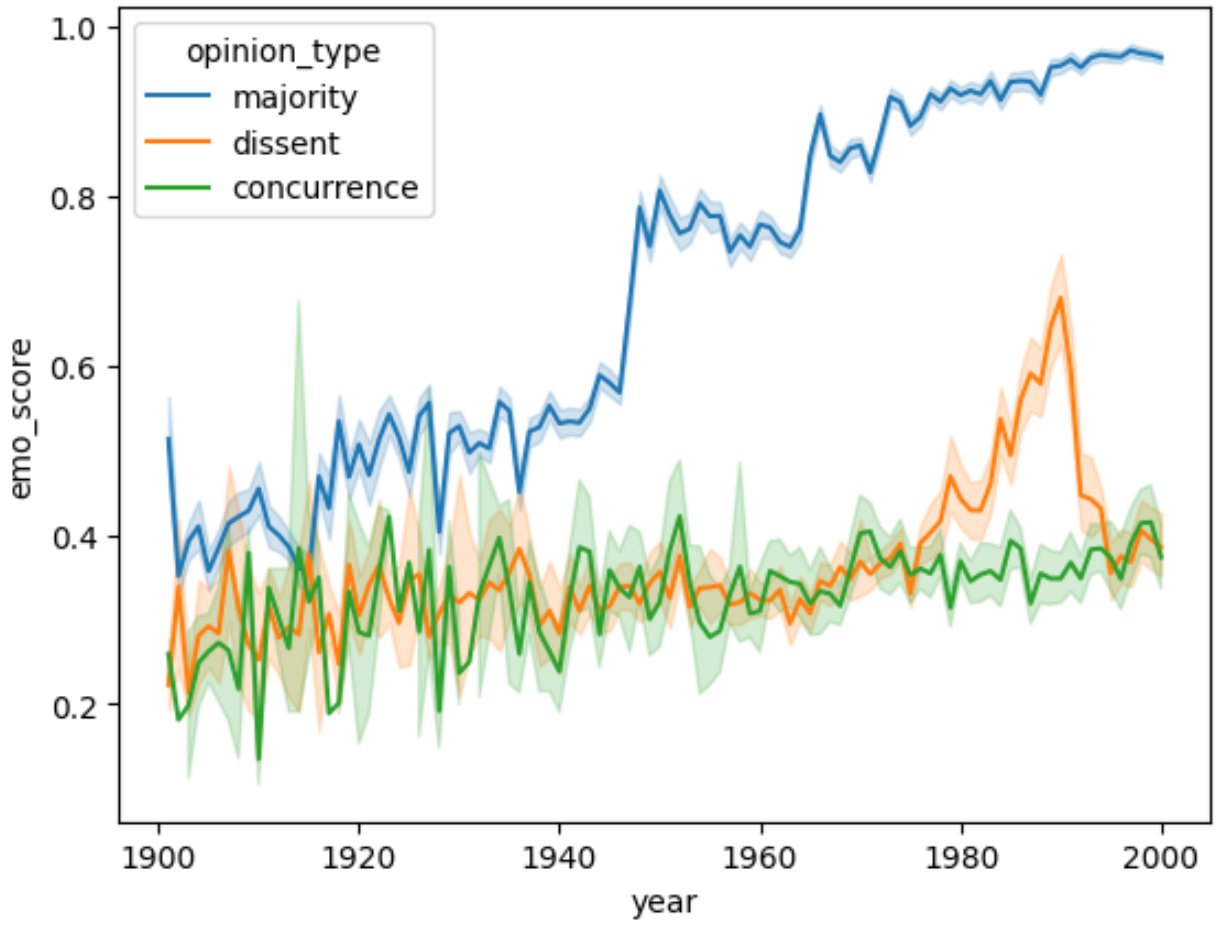

4.3 Vector-Based Emotionality Trends Over Time

While ratio-based scoring provides useful insights into long-term trends, it has notable limitations, particularly in its reliance on predefined word lists. The vector-based approach, which projects opinions onto an emotion-reason axis, refines the analysis by capturing semantic relationships between words rather than relying solely on predefined emotional terms. The results from this method, as shown in Figure 6, largely confirm the ratio-based trends but offer a more nuanced view of emotionality in judicial writing.

Figure 6: Emotionality Vector Scores Over Time, Differentiated by Opinion Type

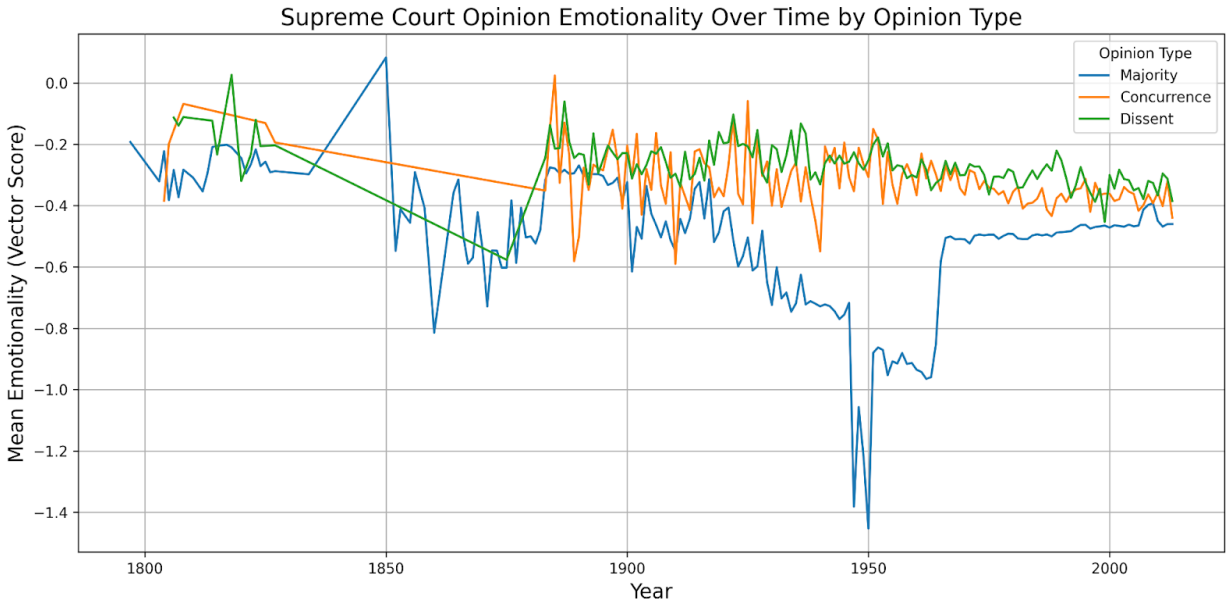

Majority opinions display a general downward trajectory in vector-based emotionality scores, but unlike the ratio method, which suggests a relatively consistent decline, the vector model reveals greater variability in emotional expression. This variability indicates that while overt emotional language may have decreased, majority opinions still retain rhetorical strategies that implicitly convey emotion, such as strategic phrasing and loaded terminology. The vector-based approach also highlights a stark increase in majority emotionality scores around the 1960s, leveling out at that around -0.5. This, in conjunction with the results from the unmodified dictionary methods, suggest that the increase in overall emotionality in court documents is *not* from an increase in dissent emotionality as previously hypothesized, but rather from an increase in *majority* opinion emotion scores.

A particularly interesting finding is that the vector-based method shows greater stability in majority opinions after 1960 than the ratio-based method does, suggesting that emotional expression in majority opinions has become more subtle and rhetorically sophisticated over time. One potential explanation is that majority justices, confident in their rulings, may rely on a detached and formal tone that minimizes overt emotional language while still embedding persuasive elements in their reasoning. By contrast, the increasing emotionality in dissents may reflect frustration or an effort to sway public opinion through heightened rhetorical appeal.

4.4 Case Studies of Highly Emotional Opinions

To further illustrate these trends, several opinions with exceptionally high emotionality scores were examined. The highest-scoring majority opinion, *Geiger v. Jenkins et al.*, contains charged language describing the plaintiff's claims as "confused" and "conclusory," demonstrating how judicial rhetoric can shape the perceived legitimacy of an argument. Similarly, *Drury v. United States* (2005), a murder case, scores highly due to the nature of its subject matter rather than the rhetorical choices of the justices. These examples highlight a key limitation of sentiment

analysis: emotionality scores can sometimes reflect the topic of a case rather than the intent of the author.

Among dissenting opinions, *Marler v. California* (1977) stands out for its use of sarcasm and overt critique of the majority's handling of the case. The opinion contains language suggesting procedural unfairness and judicial overreach, demonstrating how dissents can employ emotional rhetoric to challenge prevailing legal interpretations. Similarly, the concurring opinion with the highest emotionality score, *Guinan v. Delo* (1993), appears to derive much of its emotional weight from the nature of the case itself (a murder conviction) rather than from intentional rhetorical choices, reinforcing the need for careful contextual interpretation of sentiment analysis results.

4.5 Differences Between Ratio and Vector Scoring Methods

A key takeaway from this analysis is that ratio-based and vector-based methods do not always produce identical results. The ratio method, while simple and interpretable, is prone to misclassification due to its reliance on predefined lexicons. Words such as *murder*, *support*, and *credit* may be erroneously classified as emotional, even in neutral legal contexts. The vector method, by contrast, captures semantic meaning more effectively and can distinguish between emotionally charged rhetoric and neutral legal terminology.

However, the vector method also has limitations. It is computationally expensive, requiring the training of word embeddings and document vectors, and its results are highly dependent on the training corpus. Different embeddings can yield slightly different projections, making cross-study comparisons more challenging. The findings underscore the importance of methodological transparency in computational legal analysis and suggest that future studies should employ approaches that involve human annotation to account for subtleties and lexical variations that neither approach can properly account for.

4.6 Conclusion

This study provides a comprehensive analysis of emotionality in Supreme Court opinions, demonstrating how emotional rhetoric varies by opinion type and over time. As shown, different methods yield different results. However, a consistent theme is an increase in dissent emotionality scores around the time of the Civil Rights Movement. The differences between ratio-based and vector-based scoring highlight the challenges and complexities of sentiment analysis in legal texts, emphasizing the need for context-aware methodologies. These findings contribute to a growing understanding of judicial rhetoric and sentiment analysis in legal scholarship, offering a foundation for future research at the intersection of law, language, and computational analysis.

Chapter 5: Discussion & Conclusion

5.1 Discussion

The findings of this study demonstrate a clear increase in emotionality over time across Supreme Court opinions, regardless of the method used for measurement. While both ratio-based scoring and vector-based scoring revealed long-term shifts in judicial language, the overall trend suggests that emotion plays an increasingly prominent role in judicial writing. However, the distribution of this emotionality is neither uniform nor straightforward.

A particularly striking observation is the relative consistency of emotionality in majority opinions, especially after 1960. One possible explanation for this trend is the increasing standardization of majority opinions over time, a result of institutional norms and the expectation that such opinions reflect a neutral, authoritative voice. Another, more nuanced, interpretation is that this stability in majority opinion emotionality is linked to growing partisanship within the Court. If majority justices are increasingly at odds with the dissenting faction, they may feel a greater sense of confidence or detachment, reinforcing a tone of objectivity while allowing dissenters to adopt more rhetorically charged language. This dynamic mirrors psychological patterns observed in interpersonal disputes, where one party maintaining composure can amplify the frustration of the opposing side.

Dissenting opinions, on the other hand, show a clear rise in emotionality post-1970, a trend that corresponds with increasing ideological divisions on the Court. This rise suggests that dissents are becoming more expressive and emotionally charged, perhaps as a rhetorical strategy to influence future legal interpretations or mobilize public opinion. Justices in the dissent may feel a stronger need to appeal to emotion-driven arguments to highlight the perceived gravity of the majority's decisions. These findings reinforce the notion that emotion is not inherently irrational but is often strategically deployed to strengthen legal arguments.

From a methodological perspective, the differences between ratio-based and vector-based scoring reveal important insights into how sentiment in legal texts should be measured. The ratio-based method provided a broad historical overview but struggled with word misclassification, as it treated all words in the lexicon equally without considering their context. This issue was evident when terms such as *argue* were categorized as emotional despite their frequent neutral usage in legal writing. By contrast, the vector-based method mitigated this issue by capturing semantic relationships between words, allowing for a more context-sensitive measurement of emotionality. However, the vector-based approach also had limitations, particularly in its reliance on pre-trained embeddings, which sometimes failed to recognize more complex legal terminology.

The historical analysis further highlights key moments of heightened emotionality in Supreme Court opinions, particularly in the 1950s and 1960s. This period coincides with landmark civil rights decisions, suggesting that cases dealing with fundamental societal changes elicit more emotionally charged language. This spike in emotionality aligns with prior research indicating

that emotion is more pronounced in judicial opinions addressing contentious political and social issues. Additionally, the sharp rise in dissenting emotionality in the 1980s and 1990s corresponds with increasing polarization in high-profile cases involving reproductive rights, affirmative action, and federal power, suggesting that the ideological stakes in these rulings may contribute to greater rhetorical intensity.

Ultimately, these findings challenge the long-standing perception of judicial opinions as purely rational artifacts. Instead, they suggest that emotion is embedded within judicial decision-making and is often a critical tool in shaping legal discourse. The results of this study provide a quantitative foundation for understanding how Supreme Court rhetoric evolves over time, paving the way for more advanced analyses of judicial sentiment.

5.2 Future Research Directions

While this study presents a foundational framework for analyzing emotionality in Supreme Court opinions, several areas of improvement and expansion remain. One major avenue for refinement is the development of a more specialized legal sentiment dictionary. The existing dictionaries, while effective, contained misclassified words such as *argue* that led to potential inaccuracies in ratio-based scoring. A revised lexicon, carefully constructed through legal expert annotation and computational refinement, would allow for more precise sentiment categorization in judicial texts. Additionally, incorporating words that capture legal rhetorical strategies, such as *egregious* or *patently unjust*, into the emotional lexicon would improve the ability to detect nuanced expressions of judicial sentiment.

Another critical improvement involves enhancing vector-based scoring through advanced embedding models. The Doc2Vec model used in this study, while effective in capturing document-level sentiment, failed to recognize some complex legal terms due to limitations in its training data. Future research could explore more state-of-the-art models, such as BERT-based embeddings fine-tuned on legal corpora, to create context-aware representations of emotionality. Transformer models such as CaseLawBERT could significantly improve sentiment analysis by capturing subtle shifts in legal rhetoric and argumentation.

A further methodological enhancement would involve human annotation to train a supervised model for judicial emotionality analysis. While this approach would require significant resources —including time and funding to train annotators—it would provide an invaluable gold-standard dataset for calibrating computational models. A hybrid approach, where machine learning techniques are combined with human-in-the-loop annotation, could improve the overall accuracy of emotion classification in judicial texts.

Beyond computational refinements, this research also has practical applications for legal scholars and political scientists. A refined NLP framework for analyzing legal sentiment would allow researchers to examine large-scale trends in judicial rhetoric, beyond just emotionality. This methodology could be applied to lower court rulings, Congressional records, and legal

briefs, enabling systematic analyses of judicial communication and its influence on policymaking. Additionally, the findings of this study could help political scientists assess the extent to which reason and emotion are intertwined in judicial decision-making, revealing whether judicial rhetoric reflects broader ideological trends in American governance.

5.3 Interdisciplinary Impact

This study sits at the intersection of law, linguistics, and computational social science, contributing to multiple fields by combining natural language processing techniques with legal and political analysis. From a legal studies perspective, this research challenges traditional assumptions about judicial writing, demonstrating that emotion is not incidental but strategically embedded in legal argumentation. By quantifying how emotion is expressed across different opinion types, this study provides empirical evidence for legal scholars examining judicial rhetoric and decision-making.

For computational linguistics, this research highlights the unique challenges of sentiment analysis in legal texts. Unlike typical sentiment analysis tasks applied to social media or product reviews, legal sentiment analysis must account for domain-specific language and formal reasoning structures. The development of more sophisticated legal embeddings and emotion-reason scoring frameworks could significantly enhance NLP applications in the legal field.

In political science, this study provides quantitative insights into the intersection of law and ideology. The findings suggest that as the Court has become more polarized, emotional expression has increased, particularly in dissents. This aligns with broader trends in American political discourse, where rhetorical intensity has risen in response to ideological divisions. By tracking shifts in judicial emotionality over time, this research offers a new lens for understanding how the Supreme Court communicates with the public, influences legal precedent, and reflects broader political transformations.

5.4 Conclusion

This study offers a systematic computational analysis of emotionality in Supreme Court opinions, demonstrating that emotional language has increased over time, particularly in dissenting opinions. While majority opinions have remained relatively stable, their consistency may reflect both increasing standardization and growing ideological divides, where justices in the majority present their rulings with confidence while dissenters adopt more expressive rhetorical strategies.

The comparison between ratio-based and vector-based sentiment analysis underscores the importance of context in measuring judicial emotionality. While ratio scoring provides a clear historical trajectory, vector scoring allows for a more nuanced understanding of semantic relationships and rhetorical intent. The differences between these methods highlight the challenges of computational legal analysis, suggesting that future work should integrate hybrid approaches for greater accuracy.

Ultimately, these findings contribute to a growing body of research at the intersection of law, language, and computational analysis, offering new insights into how the Supreme Court communicates through written opinions. Future research in legal NLP, judicial rhetoric, and political science can build upon this foundation to further explore the role of emotion in judicial decision-making, ensuring that scholars and practitioners alike have the tools to analyze legal texts at scale.

Bibliography

Bryan, A. C., & Ringsmuth, E. M. (2016). Jeremiad or weapon of words?: The power of emotive language in Supreme Court dissents. *Journal of Law and Courts, 4*(1), 159–185. https://www.jstor.org/stable/26558093

Card, D., Chang, S., Becker, C., Mendelsohn, J., Voigt, R., Boustan, L., Abramitzky, R., & Jurafsky, D. (2022). Computational analysis of 140 years of U.S. political speeches reveals more positive but increasingly polarized framing of immigration. *Proceedings of the National Academy of Sciences, 119*(31), e2120510119. https://doi.org/10.1073/pnas.2120510119

Chulvi, B., Molpeceres, M., Rodrigo, M. F., Toselli, A. H., & Rosso, P. (2024). Politicization of immigration and language use in political elites: A study of Spanish parliamentary speeches. *Journal of Language and Social Psychology, 43*(2), 164–194. https://doi.org/10.1177/0261927X231175856

Costa, J. A. F., Dantas, N. C. D., & Silva, E. D. S. A. (2023). Evaluating text classification in the legal domain using BERT embeddings. In P. Quaresma, D. Camacho, H. Yin, T. Gonçalves, V. Julian, & A. J. Tallón-Ballesteros (Eds.), *Intelligent data engineering and automated learning – IDEAL 2023* (Vol. 14404). Springer, Cham. https://doi.org/10.1007/978-3-031-48232-8_6

Effective approach to develop a sentiment annotator for legal domain in a low resource setting. (2020). In G. Ratnayaka, N. de Silva, A. S. Perera, & R. Pathirana (Eds.), *Proceedings of the 34th Pacific Asia Conference on Language, Information and Computation (PACLIC 2020)*. Association for Computational Linguistics. https://aclanthology.org/2020.paclic-1.29.pdf

Google Research. (2021). *GoEmotions: A dataset for fine-grained emotion classification.* http://research.google/blog/goemotions-a-dataset-for-fine-grained-emotion-classification/

Krewson, C. N. (2019). Strategic sensationalism: Why justices use emotional appeals in Supreme Court opinions. *Justice System Journal, 40*(4), 319–336. https://doi.org/10.1080/0098261X.2019.1672597

Mudalige, C. R., Karunarathna, D., Rajapaksha, I., de Silva, N., Ratnayaka, G., Perera, A. S., & Pathirana, R. (2020, November). SigmaLaw-ABSA: Dataset for aspect-based sentiment analysis in legal opinion texts. *Proceedings of the 2020 IEEE 15th International Conference on Industrial and Information Systems (ICIIS)* (pp. 488–493). IEEE.

Plaza-del-Arco, F. M., Curry, A., Curry, A. C., & Hovy, D. (2024). Emotion analysis in NLP: Trends, gaps and roadmap for future directions. *arXiv preprint*, arXiv:2403.01222. https://arxiv.org/abs/2403.01222

Rice, D. R., & Zorn, C. (2021). Corpus-based dictionaries for sentiment analysis of specialized vocabularies. *Political Science Research and Methods, 9*(1), 20–35. https://doi.org/10.1017/psrm.2019.10

SCOTUSblog. (2015). *An empirical analysis of emotional language in legal briefs before the Supreme Court.* https://www.scotusblog.com/2015/12/an-empirical-analysis-of-emotional-language-in-legal-briefs-before-the-supreme-court/