# Bias or Not? Exploring US Press Representations of Law Enforcement in Lynching Coverage, 1789–1963

Mohamed Salama
PhD Student
Merrill College of Journalism
/University of Maryland
College Park, MD, USA
Msalama@umd.edu

## ABSTRACT

This research is inspired by existing scholarly literature [1,3,4,5,8,9,10], which argues for the explicit or implicit involvement of the press and law enforcement in the historical practices of lynching in the US. These arguments are typically based on piecemeal evidence and isolated reports, which may lack comprehensiveness. Leveraging access to 60,000 newspapers spanning from 1789 to 1963, this study rigorously investigates the norm of objectivity by analyzing 1,767 stories specifically mentioning law enforcement to understand how the press portrayed these authorities. Given that the concept of journalistic objectivity was novel during this era, we will scrutinize tonality bias—defined as the positive or negative slant in reporting—based on the framework by Eberl et al [2]. Our methodology employs a carefully designed model that integrates dependency parsing and sentiment analysis techniques. This approach allows for precise detection of objectivity and bias in journalistic reporting, adding a new layer to the investigation that enhances the depth of analysis and offers insights for future research.

## CCS CONCEPTS

• Information systems → Data mining; • Computing methodologies → Natural language processing → Information extraction.

• Computing methodologies → Natural language processing → Parsing; • Information retrieval → Sentiment analysis.

## KEYWORDS

Media Bias, Law Enforcement, Lynching, Natural Language Processing, Historical Analysis

## 1 INTRODUCTION

The investigation into how the press framed law enforcement in the historical context of lynching in the United States is significant for understanding the evolution of journalistic objectivity and media bias. Lynching, a brutal and extrajudicial form of racial violence primarily targeting African Americans, was prevalent from the late 19th century through the early 20th century, leading to the deaths of over 4,000 individuals. This period coincides with significant changes in journalistic practices, including the gradual emergence of the norm of objectivity [6,7].

During the era under study, lynching was often met with varying degrees of response from law enforcement, ranging from passive complicity to active involvement. The role of the press in documenting these events was complex, as newspapers served both as chroniclers of fact and as shapers of public opinion.

Historical accounts indicate that newspapers often reflected and reinforced the prevailing racial prejudices of the time. The coverage of lynching incidents by the press varied widely, influenced by factors such as the political orientation of the newspaper, regional attitudes, and the social and economic status of the African American community. By examining these narratives, we can gain insights into how the press either challenged or perpetuated racial violence and injustice.

This study will utilize a content analysis approach to investigate newspaper articles to identify differences in reporting sentiment towards law enforcement. The goal is to uncover biases and understand how these biases influenced public perception of law enforcement's role in lynching. Such an analysis is essential for comprehending the role of the media in shaping historical and contemporary understandings of justice and accountability.

Building on scholarly work on racial bias in journalism and the history of lynching in America [1,3,4,5,8,9,10] this research aims to contribute a nuanced understanding of the interplay between journalistic practices and social justice. The findings will offer valuable insights into the long-term impacts of media narratives on public perception and the importance of maintaining objectivity in journalism, particularly when reporting on issues of racial violence and law enforcement.

## 2   PREVIOUS WORK

The overarching goal of the conducted research is to examine how the press framed law enforcement regarding the issue of lynching in the US. Since the journalistic norm of objectivity was not common during the study period and began to crystallize only from the early to mid-20th century [6,7], we are interested in a study of bias. More precisely, we aim to investigate whether the press framed law enforcement negatively, positively, or neutrally. This is in part inspired by existing scholarly literature [1,3,4,5,8,9,10], which tends to argue for the explicit or implicit involvement of the press and law enforcement in the practice of lynching. The press has been accused of failure to report lynching events, normalizing the murder of African Americans, and even of indirectly inciting this kind of violence. Arguably, law enforcement has received even more drastic accusations, from complicity in lynching to its failure to prevent this gruesome act. These arguments are developed based on piecemeal evidence and individual accounts and cannot be generalized. To fully examine the allegations, there is a need for an analysis of an extensive sample of US news, which this paper intends to do with a large-scale computational analysis of the same, spanning the period from 1789 to 1963. This will enable us to assess whether bias is present and, if so, the extent of it in the corpus of press materials respecting the phenomenon of lynching. The research question proposed for this study is: **How did the US press portray law enforcement in terms of bias or neutrality in its historical coverage of lynching?** Investigation of bias in general is not a novel task, and there is a wide range of studies that have investigated the issue in different contexts and with the use of different methods. In our case, we shall adopt the understanding of bias as a predisposition to or disproportionate favoring or disfavoring of a person, entity, or group. It must be pointed out that not all bias is necessarily negative—it can be a useful heuristic, or an instance of avoiding harm. Bias is measured with the purpose of detecting systematic departures from the observed norm; bias is, in this sense, inferred from the residuals of the fit. Eberl et al. [2] have argued that bias can manifest in two principal ways: a visibility bias and a tonality bias. The visibility bias is the frequency of reporting on an issue. The tonality bias, in turn, is the systematic positive or negative slant in reporting. This study is interested in tonality bias and how the press framed law enforcement in the historical context of lynching. Drawing from such a large corpus of historical news articles, this study can help provide an understanding of the extent of press bias in lynching reporting, which adds to discussions around press bias and its consequences in historical and modern-day journalism.

## 3   METHODOLOGY

### 3.1   Data Collection and Source

The dataset used in this study is derived from a broader, more encompassing research project, of which the main researcher of the study is a part. It consists of 60,000 newspaper pages that have already been digitized and processed with OCR. These documents, retrieved from the archives of the Library of Congress, encapsulate a broad spectrum of news content spanning from the late 18th century in 1789 to the mid-20th century in 1963.

The selection of these pages was developed with a stratified sampling method designed purposively for strategic capturing of articles relevant to the study of lynching. The selection of pages was thus stringent and filtered with the key words 'lynching' and other terms that are pertinent so that finally selected pages are directly relevant to the thematic focus of the study.

Over the past two years, the research team has conducted an exhaustive review of these newspapers' pages, identifying articles that document incidents of lynching. For the sake of precision, a set of restrictive rules are used that include only articles that document a case of lynching and not those that do not. In this process, on an average, every page has between one or two stories about lynching. Only a few pages have up to 20 stories of lynching, which denote very intense periods of reportage about the act. On identification of these, the textual material of these scanned pages is extracted by the team for subsequent analysis. So far, this systematic procedure allowed for the retrieval and text analysis of about 7,000 stories, which has in turn enabled the current research to make a detailed inspection of historical trends regarding lynching coverage.

In this study, the coverage of the press on lynching was specifically narrowed to focus on how law enforcement was portrayed in incidents of lynching. To accomplish this, we rummaged through to pick out relevant news stories mentioning law enforcement by using a list of keywords generated from a comprehensive review of literature and observations. Terms included "Sheriff," "Officer," "Police," and others. From an initial set of 7,000 stories, this refined approach yielded a subset of 1,767 stories that specifically addressed interactions with law enforcement in the context of lynching coverage.

### 3.2   Data Quality

As the textual data were extracted in the form of CSV files, a preliminary inspection of the data showed that pervasive inconsistencies were present in the data due to the process of extraction. These inconsistencies manifested as broken rows: for example, in one row was the headline of a story, in another a partial sentence, and in the third, the rest of the sentence, rather than in a unified text as would be expected. Such fragmentation posed a challenge, making it necessary to first consolidate such spread-out rows before further analysis. By reassembling the scattered elements in a single coherent row for each story, the researcher kept the integrity of both the narrative and its associated metadata.

Once consistency was achieved in the data set, attention was given to those stories related to law enforcement only, which resulted in a total of 1,767 entries. The stories were cleaned up even further by dividing the text up into complete sentences, in each case placing the text on separate lines. Again, this cleared the presentation of the data.

The final stage of information processing entailed selective extraction of a sentence, or sentences, containing law enforcement-related keywords, as outlined in the list above. For example, if a report included five sentences, and only one of them contained the relevant keyword, that sentence was extracted, and the others were

not. Due to this careful selection process, the resultant focused dataset yielded 3,140 news lines, allowing in-depth analysis on media representation of the law enforcement.

## 3.3  Unit of Analysis

This study endeavors to examine the portrayal of law enforcement in the press, specifically investigating whether such media representations are imbued with bias. To analyze the tonality of the coverage, this research will focus on the syntactic relationships between adjectives and pivotal nouns—specifically those related to law enforcement entities. This investigation will employ dependency parsing techniques, a detailed discussion of which is reserved for subsequent sections. Further, sentiment analysis will be conducted on the adjectives identified in relation to these key nouns, providing an understanding of the emotional undertones present in the text. This aspect will also be elaborated on later.

Therefore, the basic unit of analysis in this study is adjective-noun pairs within sentences. Each pair is analyzed to establish the kind of relationship between the key noun and its modifying adjective. This evaluation is pivotal in determining the extent of objectivity or subjectivity in the text's portrayal of law enforcement. By meticulously analyzing these pairs, we gain insight into the underlying sentiment and tonality of the press coverage, thereby assessing the potential bias in media portrayals of law enforcement.

This approach to linguistic analysis offers multiple layers of insight:

*Adjectives* are used as indicators of sentiment and tonality, which are used to show the coloring of emotion connected to the noun that they modify.

*Nouns* signify critical entities or subjects within the discourse, serving as focal points for the analysis.

*Pairs of adjectives and nouns* reveal how attributes are ascribed to these entities, a crucial factor in discerning the degrees of subjectivity or objectivity embedded within the text.

This methodological approach will ensure that the research undergirds a more holistic assessment of media bias in the portrayal of law enforcement, one that should serve as a clear lens through which to evaluate the press's neutrality or lack thereof.

## 3.4  Dependency Parsing and Building the Model

In this study, we utilized dependency parsing to reveal the syntactic relationships between key nouns associated with law enforcement and their modifying adjectives within sentences. The aim was to explore the connections between these nouns and the adjectives that describe them, focusing on both direct and mediated syntactic links.

Deviation from standard text preprocessing techniques marked a distinctive aspect of the methodology. Specifically, we refrained from common preprocessing practices such as converting text to lowercase, removing stopwords, and lemmatizing. This decision was driven by the understanding that such modifications could disrupt the natural grammatical structure essential for precise dependency parsing. Dependency parsing relies critically on the original syntactic configuration of sentences; any modifications

could potentially introduce errors or inconsistencies in the analysis. Therefore, to maintain the grammatical integrity crucial for accurate parsing, the original structure of the sentences was preserved throughout the study.

In this study, we focused specifically on two types of relationships: direct and indirect.

**A. Direct Relationships**:
In these relationships, the noun acts as the parent node, with the adjective functioning as the child or direct modifier.
Example: "We congratulate and honor our brave policemen."
Here, "policemen" is the central noun, and "brave" is the adjective directly modifying the noun, exemplifying a direct syntactic relationship.

**B. Indirect Relationships**:
The analysis also extended to various forms of indirect relationships, which include:

**Adjectives connected via conjunctions to the noun or a shared parent**: Example: "permitted by a weak and incompetent police department."
In this instance, "weak" and "incompetent" are adjectives connected by the conjunction "and," both modifying the noun "police," thus illustrating an interconnected descriptive framework.

**Adjectives linked to the noun through an auxiliary verb**:
Example: "The officer was considered courageous by his peers."
Here, "courageous" is an adjective associated with the noun "officer" via the auxiliary verb "was considered," highlighting an auxiliary-mediated relationship.

**A specific syntactic pattern involving a determiner, superlative, adjective, and noun**: Example: "The most courageous officer received the highest honor."
This sentence structure includes a determiner ("the"), a superlative modifier ("most"), an adjective ("courageous"), and a noun ("officer"), which together form a complex syntactic pattern delineating a hierarchically structured relationship.

To analyze syntactic relationships, the researcher used Python packages, utilizing precisely the en_core_web_lg model from SpaCy, chosen for its comprehensive capabilities in processing complex linguistic features. This large English model offers extensive word vectors, part-of-speech tagging, named entity recognition, and dependency parsing, enabling nuanced language understanding and high accuracy, especially in parsing sentence structures.

## 3.5  Process and Verification of Model Implementation

Prior to deploying the model, the researcher initiated a preliminary analysis by selecting a random sample comprising 10 percent of approximately 331 news lines. Within this subset, the researcher conducted a manual review to discern and annotate the nouns and adjectives, specifically identifying and recording the presence or absence of syntactic relationships between them. Relationships confirmed during this review were labeled as positive, while their absence was denoted as negative.

Subsequently, the model was applied to this chosen sample by the researcher. It was tasked with identifying nouns and adjectives and assessing the syntactic relationships between them. Similar to the manual process, the model classified these relationships as either positive or negative based on their presence or absence.

The final phase involved a comparative evaluation between the outcomes of manual identification and those generated by the model, conducted by the researcher. This comparison facilitated a rigorous assessment of the model's performance, enabling the calculation of key metrics such as accuracy, precision, recall, and the F1-score.

The model's performance was evaluated, yielding an accuracy of 97.89%, precision of 82.76%, recall of 92.31%, and F1-score of 87.27%.

Additionally, we conducted intercoder reliability by asking another researcher to take a sample of 5 percent to manually check if the identified relations between the adjectives and key nouns were correctly annotated and to compare the results with the initial manual evaluation. We used Cohen's kappa statistic to measure the agreement between the two manual evaluations, achieving a result of 0.82, which indicates substantial agreement.

## 3.6  Sentiment Analysis

After preparing the text and just sorting out the key noun-adjective pair instances, we proceeded to detect and investigate tonality bias, resulting in the detection of 176 relations. We chose a straightforward approach, focusing on word-based sentiment, specifically targeting adjectives or modifiers associated with key nouns within the context of law enforcement. For this purpose, we utilized the VADER lexicon, a dictionary-based model that performs sentiment analysis at the word level. VADER assigns each word a sentiment score reflecting its positive, negative, or neutral sentiment. These individual scores are then aggregated to provide an overall sentiment score for the text.

We selected word-level sentiment analysis for its precision and granularity, which allows for the detailed identification of sentiment-laden words. This level of detail is crucial in contexts like law enforcement, where specific terms and their modifiers can significantly influence the meaning and perceived bias of the text. Furthermore, dictionary-based models like VADER are versatile across various contexts, as they do not rely on domain-specific training data. Unlike pre-trained models that may perform well only within their training domains, VADER can handle diverse text types without significant accuracy loss.

While pre-trained models can also be used for sentiment analysis, we followed the methodology of Laurenz A. Cornelissen et al. (2019), who demonstrated that dictionary-based approaches often yield better results than pre-trained models. This advantage arises because pre-trained models are typically trained on specific domains and may not generalize well outside those areas.

Additionally, we recognize that sentiment analysis can be enhanced by examining sentence-level features such as valence shifters—negators (e.g., never, without), amplifiers (e.g., incredibly, highly), de-amplifiers (e.g., slightly, somewhat), and adversative conjunctions (e.g., although, yet). These elements can significantly alter word polarity; for example, "although" in "The policy was effective, although it was not popular," changes the sentiment from positive to negative, and "never" in "never easy" shifts "easy" to a negative context.

However, we did not use sentence-level analysis for several reasons. Our primary focus is on the pair relationships between key nouns and their adjective modifiers. This focus simplifies the analysis, making it easier to understand and interpret results while still providing valuable insights into the sentiment dynamics within the text. The precision and granularity offered by word-level analysis are particularly important in the context of law enforcement, where specific terms and their modifiers can significantly influence perceived bias.

Furthermore, a sentence-based approach may still be imperfect due to unexpected artifacts or irregularities within sentence structures, which can significantly impact the overall sentiment score. Despite pre-processing, historical data may contain such irregularities that are difficult to account for, potentially skewing the results. Therefore, we believe that evaluating sentiment at the word level, focusing on adjectives and their relationships with key nouns, is more robust and precise within the scope of our study. This method aligns better with the nature of our data and ensures more reliable analysis outcomes.

## 6  RESULTS AND DISCUSSION

The aim of this analysis is to investigate whether the press was biased in its portrayal of law enforcement by measuring tonality. The sentiment analysis results provide a perspective on the spread of sentiment scores for pairs of key nouns related to law enforcement and their modifying adjectives, assessed based on the dependency parsing results, which identified all pair relations (see Figure 1). One significant observation is the clustering of the 0-sentiment score for most of these neutral pairs. Unexpectedly, despite the norm of objectivity not being firmly established at the time, the neutrality of the data suggests that the necessary normalcy of objective journalism during the period of study was mostly adhered to in reports on the topics of law enforcement. In this set of pairs, the average neutrality fell within a range of 0 ± 0.1, highlighting the predominance of neutral coverage. For instance, a neutral report stated, "Extra police were called the prisoner was placed in the county and jail" (St. Paul daily globe, Saint Paul, Minn., 1885). Conversely, a notable fraction of the pairs has positive sentiment scores, varying from 0.2 to 0.6, indicating some favorable reporting towards law enforcement. An example of this positive sentiment is, "We congratulate and honor our brave policemen and officials in their determined and successful effort, at the risk of their own lives, to save the accused negro from the wrath of the mob" (The Birmingham age-herald, Birmingham, Ala., 1920). Additionally, another fraction shows negative sentiment scores, ranging from -0.6 to -0.2, reflecting criticism towards law enforcement actions. An illustration of negative sentiment is, "The plot was conceived in a lawless resort, carried out by hoodlums and

permitted by a weak and incompetent police department" (Phoenix tribune, Phoenix, Ariz., 1920).

Despite these findings, though, there are possible limitations. The algorithm that carries out sentiment analysis is likely to be biased in some way by nature, and this may not model the result accurately, as language is relatively complex and nuanced, with some sensitivity to context. Additionally, sentiment scores would not be able to weigh the context in which terms are used, for example, sarcastic or ironic tones, because these are weighed about other words that an algorithm may misconstrue. For instance: "Sheriff Stahl of Kenosha county, Wis., is an enterprising and hustling public officer, and undoubtedly inspired with great zeal for his calling and love of justice, but the manner in which he manifests his concern that the heavy hand of the law should fall upon the guilty is not altogether in tune with the sentiment of enlightenment that supposedly prevails in civilized communities" (Evening star, Washington, D.C., 1912). The irony in these lines, which praise the sheriff superficially while criticizing his methods, might not be accurately detected by sentiment analysis algorithms. However, the coexistence of positive and negative favors highlights particular moments where the press either condemned or favored the actions of law enforcement. While most of the reporting was neutral in describing the events unfolding and reflects general adherence to the then-budding journalistic norms of objectivity, the coexistence of the positive and negative favors paints a different picture. These favors, at the same time, do not need to be indicative of biased reporting; it may instead be actual condemnation or veneration of law enforcement actions.
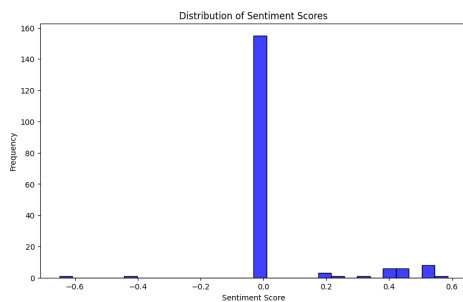


**Figure 1**

## 6   LIMITATIONS AND SUGGESTIONS

Some of the significant inherent limitations of the study pertain to the bias of the sentiment analysis algorithm. This often makes it recognize subtle language incorrectly and not consider the complexity of human communication, especially in picking sarcastic or ironic tones. As a result, this tool lacks accuracy in evaluating the weight of the context in the use of terms, leading to inaccurate sentiment scoring and sentiment misclassification.

Our approach to sentiment analysis, therefore, is at the level of words and not at the level of context or sentence. This seemed like a good idea at first. Still, it has since become clear that this approach will require a significant amount of algorithm development and performance testing until we can reach a point where we get robust and accurate detection of tone.

Of a more positivistic note would be the fact that the dependency parsing model created was highly effective in identifying relationships between key nouns and adjectives with potential as indicators of tone. However, a more robust sentiment analysis is required to complete the detection of bias fully.

In conclusion, a robust sentiment analysis model should be developed. This will help future researchers investigate the coverage of racial violence and navigate through journalistic norms such as objectivity.

## REFERENCES

[1]   Brundage, W. F. 2021. The Press and Lynching. In K. R. Forde and S. Bedingfield Eds. Journalism and Jim Crow: White Supremacy and the Black Struggle for a New America. University of Illinois Press, 72–95. Available at http://public.eblib.com/choice/PublicFullRecord.aspx?p=6810075.

[2]   Eberl, J. M., H. G. Boomgaarden, and M. Wagner. 2015. One Bias Fits All? Three Types of Media Bias and Their Effects on Party Preferences. Communication Research, 44, 8 (2015), 1125–1148. DOI:https://doi.org/10.1177/0093650215614364.

[3]   Forde, K. R., and S. Bedingfield. 2021. Journalism and Jim Crow: White Supremacy and the Black Struggle for a New America. University of Illinois Press. Available at https://search.ebscohost.com/login.aspx?direct=true&AuthType=ip,uid&db=nlebk&AN=3025221&site=ehost-live.

[4]   Perloff, R. M. 2000. The Press and Lynchings of African Americans. Journal of Black Studies, 30 (2000), 315–330.

[5]   Pfeifer, M. J. 2011. The roots of rough justice: Origins of American lynching. University of Illinois Press.

[6]   Schudson, M. 1978. Discovering the news: A social history of American newspapers. Basic Books.

[7]   Schudson, M. 2001. The objectivity norm in American journalism. Journalism, 2, 2 (2001), 149–170. DOI:https://doi.org/10.1177/146488490100200201.

[8]   Seguin, C., and D. Rigby. 2019. National Crimes: A New National Data Set of Lynchings in the United States, 1883 to 1941. Socius: Sociological Research for a Dynamic World, 5 (2019), 1–9.

[9]   Tolnay, S. E., and E. M. Beck. 1995. A festival of violence: An analysis of Southern lynchings, 1882-1930. University of Illinois Press.

[10]  Wood, A. L. 2009. Lynching and Spectacle; Witnessing Racial Violence in America, 1890-1940. University of North Carolina Press.