

Technical Note

Not peer-reviewed version

Brand Hate Detector: An R Shiny Application for Automated Detection and Multilevel Classification of Brand Hate in Consumer Reviews

[Mohamed Assoud](#) *

Posted Date: 18 August 2025

doi: 10.20944/preprints202508.1213.v1

Keywords: brand hate; natural language processing; shiny app; open source



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Technical Note

Brand Hate Detector: An R Shiny Application for Automated Detection and Multilevel Classification of Brand Hate in Consumer Reviews

Mohamed Assoud

Hassan First University of Settat, Morocco; m.assoud@uhp.ac.ma

Abstract

The Brand Hate Detector is an open-source R Shiny application designed to identify, classify, and visualize brand hate in online consumer reviews. Unlike traditional sentiment analysis tools that reduce feedback to simple positive–negative polarity, the Brand Hate Detector employs a hybrid lexicon-based approach that integrates sentiment polarity, emotion profiling using the NRC Emotion Lexicon, and rule-based classification to capture multiple hate intensity levels. The workflow includes automated scraping of brand-specific reviews from ConsumerAffairs.com, negative sentiment filtering, stopword removal, emotion analysis, and a two-stage classification process that assigns reviews to Mild, Moderate, Strong, or Hybrid hate categories. Results are presented through interactive visualizations, including category distributions, reason-specific breakdowns, emotion word clouds, and sentiment–intensity bubble plots. Designed for transparency, reproducibility, and accessibility, the tool bridges academic theory and managerial practice by providing actionable insights into the drivers and emotional dynamics of consumer hostility toward brands. Its modular architecture also allows future extensions to other review platforms, additional languages, and machine learning–based classification methods.

Keywords: brand hate; natural language processing; shiny app; open source

Introduction

In the current digital landscape, consumer voice has evolved from mere opinion into a decisive market force (Kucuk, 2020). Online reviews influence the majority of purchase decisions, shaping brand reputations and directly affecting sales (Parveen, 2025). Negative reviews, amplified by the psychological negativity bias, often travel faster, resonate longer, and prove more persuasive than positive feedback (Zhang et al., 2025). Digital platforms such as Amazon, TripAdvisor, and ConsumerAffairs now enable consumers to challenge, criticize, or even mobilize against brands at scale, fundamentally reshaping brand–consumer relations (Lis & Fischer, 2020).

Within this context, brand hate has emerged as a critical construct for understanding how dissatisfaction escalates into active hostility (Assoud & Berbou, 2025a). It extends beyond disappointment, reflecting deep emotional rejection that may be moral, symbolic, or personal in nature (Zhang & Laroche, 2020). Prior research (Fetscherin, 2019; Hegner et al., 2017; Kucuk, 2019; Zarantonello et al., 2016) has mapped its antecedents—negative past experiences, symbolic incongruity, and ideological incompatibility—and identified its emotional complexity, often involving anger, sadness, fear, disgust, and surprise (Sternberg & Sternberg, 2008).

Despite its conceptual richness, brand hate remains methodologically underserved (Assoud & Berbou, 2023; Mushtaq et al., 2024; Yadav & Chakrabarti, 2022). Most studies rely on conventional methods (surveys, interviews) ill-suited to capturing how hostility unfolds in real time, in the raw language of online discourse (Assoud & Berbou, 2025b). While Natural Language Processing (NLP) is now widely applied in marketing (Albladi et al., 2025; Kang et al., 2020; Yu & Chauhan, 2025) for sentiment analysis and topic detection, few tools exist that can detect brand hate specifically, rather

than generic negativity; classify its intensity and emotional composition; and provide interpretable, actionable results for managers. Furthermore, Existing solutions such as the one developed by Mednini et al. (2024), often act as black boxes, lack emotional granularity, and demand significant computational or financial resources, limiting their adoption, especially by SMEs.

To address these gaps, we developed Brand Hate Detector, an open-source R Shiny application that integrates lexicon-based sentiment analysis, psychological models of brand hate (Kucuk, 2019; Sternberg & Sternberg, 2008; Zhang & Laroche, 2020), and interactive visualization. The tool processes consumer reviews in real time, classifies brand hate into four categories, and profiles its emotional drivers using the NRC Emotion Lexicon. It is designed to be transparent, low-cost, and accessible—bridging academic theory with managerial utility.

Implementation

The Brand Hate Detector was developed using R Shiny (version X.X.X) to provide an interactive, modular workflow for the automated detection, classification, and visualization of brand hate in consumer reviews. The app integrates data acquisition, text preprocessing, lexicon-based emotion analysis, rule-based classification, and multiple visualization options. All computations are executed server-side, and all outputs are downloadable for reproducibility. The application is structured into five functional modules (Figure 1).

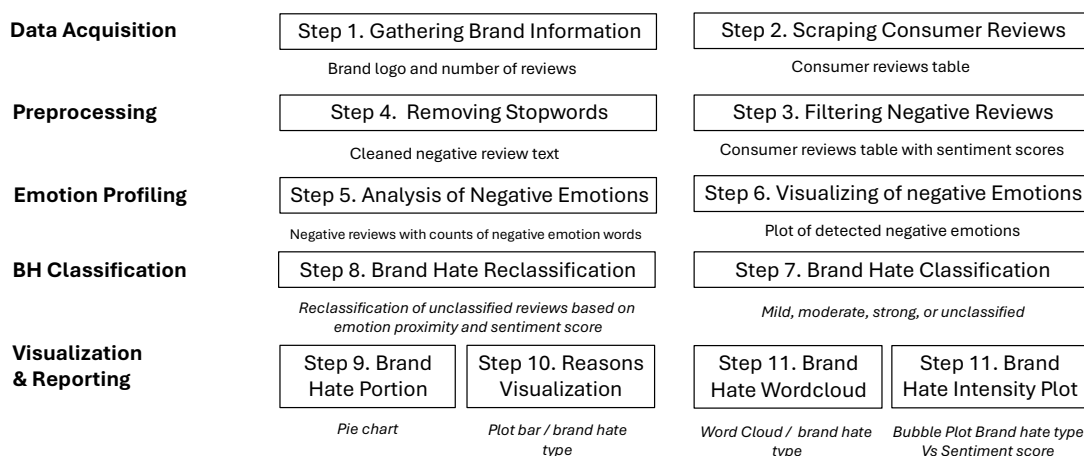


Figure 1. Analytical pipeline of the NLP-based Brand Hate Detetor.

1. Data Acquisition Module

The workflow begins with Brand Information. Here, the user provides a ConsumerAffairs brand URL, which the application uses to retrieve and display key metadata—such as the brand’s name, logo, and total number of available reviews. This initial step ensures the target context is correct before proceeding. Once confirmed, the process transitions seamlessly into Scraping Reviews. Leveraging the *rvest* package (Wickham & Bryan, 2023), the application automatically navigates paginated review pages, extracting the full review text, optional “reasons” metadata, and basic sentiment cues. The collected data are immediately presented in an interactive *DT table*, allowing users to inspect, search, and filter entries. For transparency and reproducibility, the raw review set can be exported at this stage as a CSV file for independent analysis.

2. Preprocessing Module

Following data acquisition, the workflow advances to Negative Review Filtering. In this step, each review is scored for sentiment polarity using the *sentimentr* package (Rinker, 2017). Only those with a negative polarity score are retained, narrowing the dataset to content most likely to contain brand hate signals in line with theoretical assumptions. With the filtered set-in place, the process

transitions to Stopword Removal. Using the *tidytext* and *stopwords* packages (Wickham & Bryan, 2023), the retained reviews are tokenized and stripped of common stopwords, punctuation, and numerical characters. Both the cleaned text and its original counterpart are preserved, ensuring transparency and allowing for full traceability in subsequent analysis.

3. Emotion Profiling Module

Once the textual data are cleaned, the workflow moves into NRC Emotion Analysis. At this stage, tokens are matched to the NRC Emotion Lexicon (Mohammad & Turney, 2013), assigning each word to one or more of eight primary emotions—anger, fear, anticipation, trust, surprise, sadness, joy, and disgust—along with positive or negative polarity. For the purposes of brand hate detection, emphasis is placed on negative emotions most strongly associated with hostility, including anger, sadness, fear, disgust, and surprise (Kucuk, 2019; Sternberg & Sternberg, 2008; Zhang & Laroche, 2020). Counts are aggregated at the review level, producing a structured emotional profile for each entry. Building on these results, the process transitions into Emotion Visualization. Using *ggplot2*, the application generates interactive bar charts (Wickham & Bryan, 2023) that summarize the overall distribution of emotions within the dataset. These visualizations provide a high-level emotional “fingerprint” of the brand’s negative reviews, setting the stage for classification into distinct hate intensity levels.

4. Brand Hate Classification Module

With the emotional profiles established, the workflow advances to the Initial Rule-Based Classification stage. Here, the application combines aggregated sentiment statistics—mean and standard deviation of polarity scores—with the presence of specific negative emotion patterns to assign each review to one of three categories: Mild Hate, Moderate Hate, or Strong Hate. Reviews that do not meet the criteria for any of these categories are temporarily labeled as Unclassified. This step ensures that only reviews with clear emotional or sentiment-based signals are assigned an intensity level at this stage.

From this point, the process transitions to a Reclassification Step, where unclassified reviews are re-evaluated using a purely emotion-based approach. This secondary pass identifies complex cases in which multiple high-intensity emotions co-occur, even if sentiment thresholds were not met. Such cases are labeled as Hybrid Hate, reflecting the combined presence of diverse emotional triggers. This two-tiered approach ensures a comprehensive categorization of all reviews, capturing both direct hostility and more nuanced emotional blends.

5. Visualization & Reporting Module

With classification complete, the workflow transitions to the Visualization & Reporting stage, where results are synthesized into interactive, exportable outputs. The application provides multiple perspectives on the data to support both academic analysis and managerial decision-making.

The Sentiment Distribution plot shows the proportion of negative reviews relative to the total dataset, providing a quick gauge of overall consumer dissatisfaction. The Hate Category Distribution chart displays the relative prevalence of mild, moderate, strong, and hybrid hate, enabling easy comparison of intensity levels.

To explore underlying drivers, the Reasons Analysis module aggregates the “reasons” metadata—such as customer service, product quality, or pricing—by hate category. This view highlights which aspects of the brand are most associated with different levels of hostility.

Lexical patterns are further illustrated in the Emotion Word Clouds, where words are color-coded by emotion category. These clouds provide a visually intuitive summary of the most salient terms within each hate level.

Finally, the Hate Intensity Bubble Plot maps sentiment scores against hate categories, revealing the relationship between emotional diversity and perceived hostility. Outliers and clusters are easily spotted, aiding in the identification of unusual or particularly impactful reviews.

All visualizations can be downloaded in publication-quality PNG or TIFF formats, and all processed data tables can be exported as CSV files for further offline analysis or integration into other research workflows.

Conclusions

The Brand Hate Detector provides a transparent, accessible, and academically grounded solution for the real-time detection and classification of brand hate in consumer reviews. By integrating lexicon-based emotion analysis, a two-stage rule-based classification framework, and interactive visualization, the application bridges the gap between theoretical constructs in brand hate research and practical tools for brand monitoring.

Crucially, the tool moves beyond conventional polarity classifications of positive, negative, or neutral. Instead, it distinguishes between multiple hate intensity levels—mild, moderate, strong, and hybrid—and profiles the emotional composition underlying each. This granularity allows marketing researchers to capture the complexity of negative brand relationships and enables brand managers to identify not only whether hostility exists, but how intensely it is expressed and what emotions drive it.

The tool's design prioritizes interpretability and reproducibility. All classification rules, sentiment statistics, and emotion mappings are visible within the interface, and all outputs are exportable in standard formats. These features make it well-suited for academic replication studies, teaching applications, and practitioner adoption, including by organizations with limited technical or financial resources.

Looking ahead, the Brand Hate Detector could be extended in several ways. Performance on larger datasets may be optimized through parallel processing; additional languages and lexicons could expand its global applicability; and integration with supervised machine learning could provide adaptive, domain-specific classification models. Furthermore, the data acquisition module—currently optimized for ConsumerAffairs.com—could be expanded to support multiple review platforms, e-commerce sites, and social media channels, thereby broadening its utility across diverse consumer contexts. Nonetheless, even in its current form, the application represents a significant methodological contribution to both brand management research and applied consumer analytics.

Availability and requirements

- Project name: Brand Hate Detector
- Project home page: <https://github.com/MEDASSOUD/brandhatedetector.git>
- Operating system(s): Platform independent
- Programming language: R
- Other requirements: Dependent on R packages.
- License: MIT
- Any restrictions to use by non-academics: Refer to ConsumerAffairs.com terms of use.

Acknowledgements: The development of the Brand Hate Detector benefited from discussions with colleagues in marketing analytics and computational linguistics, whose feedback guided both methodological choices and interface design. The author acknowledges the ConsumerAffairs.com platform as the primary source of publicly available consumer reviews for this study.

Authors' contributions: M.A. developed the Shiny application, designed the analytical framework, and wrote the manuscript text. M.A. also curated the datasets, implemented the classification algorithms, and revised the final manuscript.

Funding: No external funding was received for this work. The project was self-funded by the author.

Data availability: No proprietary datasets were generated during this study. All data were obtained from publicly accessible reviews on ConsumerAffairs.com. The Shiny application source code and demonstration datasets are available upon request.

Ethics approval and consent to participate: Not applicable.

Consent for publication: Not applicable.

Competing interests: The author declares no competing interests.

References

- Albladi, A., Islam, M., & Seals, C. (2025). Sentiment Analysis of Twitter Data Using NLP Models: A Comprehensive Review. *IEEE Access*, 13, 30444-30468, <https://doi.org/10.1109/access.2025.3541494>
- Assoud, M., & Berbou, L. (2023). Brand Hate: One Decade of Research A Systematic Review. *International Journal of Latest Research in Humanities and Social Science*, 06(11), 146-165, <https://doi.org/10.2139/ssrn.5037055>
- Assoud, M., & Berbou, L. (2025a). Conceptualizing Brand Hate Escalation: A Dual SEM-PLS / Necessary Condition Analysis Approach. *African Journal of Human Resources, Marketing and Organisational Studies*, 2(1), 55-81, https://hdl.handle.net/10520/ejc-aa_ajhrmos_v2_n1_a4
- Assoud, M., & Berbou, L. (2025b). Mapping the Evolution of Brand Hate: A Comprehensive Bibliometric Analysis. *Journal of Consumer Satisfaction, Dissatisfaction and Complaining Behavior*, 38(1), 21-58, <https://www.jcsdcb.com/index.php/JCSDCB/article/view/1086>
- Fetscherin, M. (2019). The five types of brand hate: How they affect consumer behavior. *Journal of Business Research*, 101, 116-127, <https://doi.org/10.1016/j.jbusres.2019.04.017>
- Hegner, S. M., Fetscherin, M., & van Delzen, M. (2017). Determinants and outcomes of brand hate. *Journal of Product and Brand Management*, 26(1), 13-25, <https://doi.org/10.1108/JPBm-01-2016-1070>
- Kang, Y., Cai, Z., Tan, C.-W., Huang, Q., & Liu, H. (2020). Natural language processing (NLP) in management research: A literature review. *Journal of Management Analytics*, 7(2), 139-172, <https://doi.org/10.1080/23270012.2020.1756939>
- Kucuk, S. U. (2019). Consumer Brand Hate: Steam rolling whatever I see. *Psychology and Marketing*, 36(5), 431-443, <https://doi.org/10.1002/mar.21175>
- Kucuk, S. U. (2020). *Consumer Voice: The Democratization of Consumption Markets in the Digital Age*. Springer International Publishing. <https://doi.org/10.1007/978-3-030-53983-2>
- Lis, B., & Fischer, M. (2020). Analyzing different types of negative online consumer reviews. *Journal of product & brand management*, 29(5), 637-653, <https://doi.org/10.1108/jpbm-05-2018-1876>
- Mednini, L., Noubigh, Z., & Turki, M. D. (2024). Natural Language Processing for Detecting Brand Hate Speech. *Journal of Telecommunications and the Digital Economy*, 12(1), 486-509, <https://doi.org/10.18080/jtde.v12n1.859>
- Mohammad, S. M., & Turney, P. D. (2013). CROWDSOURCING A WORD-EMOTION ASSOCIATION LEXICON. *Computational Intelligence*, 29(3), 436-465, <https://doi.org/10.1111/j.1467-8640.2012.00460.x>
- Mushtaq, F. M., Ghazali, E. M., & Hamzah, Z. L. (2024). Brand hate: a systematic literature review and future perspectives. *Management Review Quarterly*, 1-34, <https://doi.org/10.1007/s11301-023-00402-z>
- Parveen, U. (2025). An Analysis of Linguistics Patterns in Online Product or Service Reviews and their Influence on Customer Behavior. *Qlantic Journal of Social Sciences and Humanities*, 6(2), 111-117, <https://doi.org/10.55737/qjssh.vi-ii.25358>
- Rinker, T. (2017). Package 'sentimentr'. Retrieved, 8, 31,
- Sternberg, R. J., & Sternberg, K. (2008). *The Nature of Hate*. Cambridge University Press.
- Wickham, H., & Bryan, J. (2023). *R packages*. "O'Reilly Media, Inc."
- Yadav, A., & Chakrabarti, S. (2022). Brand hate: A systematic literature review and future research agenda. *International Journal of Consumer Studies*, 46(5), 1992-2019, <https://doi.org/10.1111/ijcs.12772>
- Yu, J. H., & Chauhan, D. (2025). Trends in NLP for personalized learning: LDA and sentiment analysis insights. *Education and Information Technologies*, 30(4), 4307-4348, <https://doi.org/10.1007/s10639-024-12988-2>
- Zarantonello, L., Romani, S., Grappi, S., & Bagozzi, R. P. (2016). Brand hate. *Journal of Product & Brand Management*, 25(1), 11-25, <https://doi.org/10.1108/jpbm-01-2015-0799>

Zhang, C., & Laroche, M. (2020). Brand hate: a multidimensional construct. *Journal of Product and Brand Management*, 30(3), 392-414, <https://doi.org/10.1108/JPBM-11-2018-2103>

Zhang, Q., Su, L., Zhou, L., & Dai, Y. (2025). Distance Brings About Beauty: When Does the Influence of Positive Travel Online Reviews Grow Stronger Relative to Negative Reviews? *Journal of Travel Research*, 64(1), 172-188, <https://doi.org/10.1177/00472875231209426>

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.