International Insurance Law Review
ISSN: 0968-2090
Vol 34 S1 (2026): Special Issue
https://lumarpub.com/iilr

Digital Marketing Automation and Sentiment Analysis Using NLP in Business Communication

Dr. Satish Gopichand Athawale[1], Dr. Vikesh Kashyap[2], Dr. Amrita Singh[3], Prof. Munianjinappa.K[4], AKANSH GARG[5], Aparna Vajpayee[6]

[1]Assistant Professor of Marketing Management, Faculty of Management Studies (FMS), Parul Institute of Management and Research, Parul University, Vadodara, Gujarat,satishathawale@gmail.com
[2]Associate Professor, School of Hospitality and Tourism, Bahra University, Solan, Waknaghat ,Himachal Pradesh vikeshkashyap@yahoo.com
[3]Associate Professor, Department of Management, Shivalik College of Engineering, Dehradun Uttarakhand , amritanov1311@gmail.com
[4]Assistant Professor and Research Scholar, Department of English, International Institute of Business Studies, Bengaluru Urban, Bengaluru, Karnataka
[5]Director Array Research Pvt Ltd, 7505264391akg@gmail.com
[6]Professor, School of Liberal Arts and Management, P P Savani University, Dhamdod, Kosamba, Surat, 394125, Gujarat, India, ORCID: https://orcid.org/0000-0003-4616-8194

Abstract

The confluence of digital marketing automation platforms and natural language processing technologies has fundamentally transformed how organizations analyze, generate, and respond to business communications at scale, creating both unprecedented opportunities for personalization and efficiency and novel challenges related to sentiment accuracy, context preservation, and ethical deployment of algorithmic decision systems in customer-facing interactions. This paper presents a comprehensive investigation of NLP-driven sentiment analysis integration with digital marketing automation workflows across a multi-industry dataset comprising 2.3 million customer communication records from e-commerce, financial services, hospitality, and healthcare sectors. We develop and evaluate a hybrid transformer-based sentiment classification architecture that combines domain-adaptive fine-tuning of BERT with lexicon-augmented rule systems to address the domain-specificity and context-dependency limitations of general-purpose sentiment models in professional business communication contexts. The proposed architecture achieves sentiment classification accuracy of 91.7 percent on a curated multi-domain benchmark, representing an improvement of 8.4 percentage points over vanilla BERT and 14.2 percentage points over traditional lexicon-based approaches. Integration of the sentiment analysis pipeline with marketing automation triggers is evaluated through a randomized controlled trial across 180,000 customer interaction records, demonstrating statistically significant improvements in email open rates, conversion rates, and customer satisfaction scores when sentiment-adaptive automation logic replaces uniform rule-based triggers. The analysis further examines the ethical dimensions of automated sentiment analysis in business communication, including bias amplification risks, transparency obligations, and the appropriate scope of algorithmic autonomy in customer relationship management.

Citation: Athawale S. et.al. (2026). Digital Marketing Automation and Sentiment Analysis Using NLP in Business Communication. *International Insurance Law Review*, 34 (S1), 491-500.

Keywords: *Natural Language Processing, Sentiment Analysis, Digital Marketing Automation, BERT, Transformer Models, Customer Communication, Marketing Technology, Text Classification, Business Intelligence, Customer Relationship Management*

I. INTRODUCTION

Digital marketing has undergone a fundamental technological transformation over the past decade, driven by the proliferation of customer touchpoints across email, social media, chatbot interfaces, review platforms, and messaging applications that generate volumes of unstructured text data far exceeding the processing capacity of human marketing and customer service teams [1]. The automation of marketing workflows, including email campaign sequencing, social media posting schedules, lead nurturing sequences, and customer service response routing, has become a strategic imperative for organizations competing across digital channels, with the global marketing automation software market exceeding $6.4 billion in 2023 and projected to grow at compound annual rates exceeding 12 percent through 2028 [2].

Natural language processing, and specifically sentiment analysis as its most directly actionable application for business communication contexts, offers a technically mature pathway for enriching marketing automation systems with semantic understanding of customer intent, emotional valence, and satisfaction levels that rule-based automation systems fundamentally cannot capture [3]. The ability to classify customer communications as expressing positive, negative, or nuanced mixed sentiment, and to further identify specific emotional dimensions including frustration, enthusiasm, confusion, or urgency, enables marketing automation systems to dynamically adjust their response strategies in ways that align with customer emotional states rather than mechanically executing predetermined communication scripts regardless of contextual appropriateness [4].

The technical landscape of NLP-based sentiment analysis has been transformed by the development of large pretrained transformer language models, particularly BERT and its domain-specific and task-specific variants, which have established new performance benchmarks across text classification tasks including sentiment analysis by learning rich contextual representations of language from massive text corpora [5]. However, the translation of general-purpose transformer model performance into reliable business communication sentiment analysis presents challenges related to domain specificity, as the linguistic patterns, terminology, and sentiment expression conventions of professional business communication contexts differ substantially from the social media and review texts that dominate standard sentiment analysis training corpora [6]. Financial services communications, medical patient communications, hospitality service communications, and retail customer service communications each exhibit distinctive linguistic characteristics that require domain adaptation beyond what general pretraining provides [7].

The integration of sentiment analysis outputs into marketing automation decision logic introduces additional challenges beyond classification accuracy, including the appropriate mapping of probabilistic sentiment scores to deterministic automation trigger conditions, the handling of edge cases and classification uncertainty, the latency requirements of real-time response generation, and the governance questions surrounding algorithmic autonomy in customer-facing communication decisions [8]. These integration challenges have received less systematic attention in the academic NLP literature than the upstream sentiment classification problem,

Citation: Athawale S. et.al. (2026). Digital Marketing Automation and Sentiment Analysis Using NLP in Business Communication. *International Insurance Law Review*, 34 (S1), 491-500.

creating a gap between the technical capabilities demonstrated in NLP research settings and the operational deployment challenges faced by marketing technology practitioners [9].

This paper bridges the gap between NLP research and marketing automation practice through a systems-level investigation that addresses the full pipeline from raw customer communication text to automated marketing response, encompassing the sentiment classification architecture, the automation integration logic, the operational performance evaluation methodology, and the ethical governance framework appropriate for enterprise deployment [10]. The contributions include a novel domain-adaptive BERT variant optimized for business communication sentiment analysis, a documented integration architecture for connecting sentiment outputs to automation triggers, empirical evidence on the business performance impact of sentiment-aware automation from a large-scale randomized controlled trial, and a practical ethical framework for responsible deployment of sentiment-driven automation in customer relationship management.

II. OBJECTIVES

- Develop transformer-based sentiment analysis model for multiple business communication domains.
- Design scalable system integrating sentiment analysis with marketing automation platforms.
- Evaluate impact of sentiment-based marketing on engagement, conversions, satisfaction, and churn.
- Analyze bias in sentiment models and develop effective mitigation strategies.
- Establish ethical framework ensuring transparency, consent, accountability in sentiment-driven marketing.

III. RELATED WORKS

The computational analysis of sentiment in text has a history extending to early rule-based lexicon approaches developed by Turney [11] and Pang, Lee, and Vaithyanathan [12], whose foundational work demonstrated that machine learning classifiers applied to bag-of-words text representations could achieve competitive sentiment classification accuracy on movie review corpora. The subsequent decade produced extensive development of supervised machine learning approaches to sentiment analysis including support vector machines with carefully engineered features, neural network classifiers, and ensemble methods, with progressively improving performance on benchmark datasets while remaining constrained by the limitations of shallow text representations that fail to capture contextual nuances, irony, negation scope, and domain-specific sentiment conventions [13].

The introduction of recurrent neural network architectures for sentiment analysis, particularly LSTM and bidirectional LSTM models, addressed some limitations of bag-of-words approaches by capturing sequential dependencies in text, enabling better handling of negation and long-range sentiment modifiers [14]. Attention mechanisms and the transformer architecture introduced by Vaswani et al. [15] represented a paradigm shift in NLP that enabled the development of large pretrained language models learning rich contextual representations from massive text corpora through self-supervised objectives. The BERT model of Devlin et al. [5], pretrained through masked language modeling and next sentence prediction on combined BookCorpus and Wikipedia text, established new state-of-the-art performance across diverse NLP benchmarks

Citation: Athawale S. et.al. (2026). Digital Marketing Automation and Sentiment Analysis Using NLP in Business Communication. *International Insurance Law Review*, 34 (S1), 491-500.

including sentiment classification and demonstrated that fine-tuning pretrained representations on task-specific labeled data required substantially fewer labeled examples than training task-specific models from scratch.

Domain adaptation of pretrained language models for specialized sentiment analysis contexts has been addressed through several approaches including domain-adaptive pretraining on domain-specific corpora before task fine-tuning, as investigated by Gururangan et al. [16] for scientific, biomedical, and news corpora; domain-specific BERT variants pretrained from scratch on domain corpora such as BioBERT [17] for biomedical text and FinBERT [18] for financial communications; and multi-task learning approaches that jointly optimize sentiment classification with related auxiliary tasks to improve generalization. The business communication domain has received less systematic attention than the scientific and medical domains, with existing work primarily addressing customer review sentiment without the full range of professional communication types—email correspondence, chat transcripts, complaint letters, support tickets—that characterize enterprise customer relationship management contexts.

Marketing automation has been analyzed in the management information systems literature primarily from adoption, effectiveness, and strategic implementation perspectives [19]. The technical integration of NLP capabilities into marketing automation platforms has been explored in practitioner-oriented literature and vendor case studies but has received limited rigorous academic treatment, particularly regarding the design principles for connecting probabilistic NLP outputs to deterministic automation logic [20]. The question of appropriate automation autonomy in customer communication has been examined from behavioral and ethical perspectives, with Ostrom et al. [21] documenting that customer acceptance of automated responses depends heavily on the perceived accuracy and empathy of the automation and on the availability of human escalation pathways.

Bias in sentiment analysis systems has been documented across multiple dimensions including racial and gender bias in sentiment lexicons that consistently assign more negative sentiment scores to texts mentioning minority groups, demographic bias in training data that reflects historical patterns of sentiment expression that may not generalize across populations, and domain transfer bias where models trained on product review corpora systematically misclassify sentiment in professional communication contexts [22]. Kiritchenko and Mohammad [23] conducted a systematic evaluation of sentiment analysis system bias and documented consistent disparities across gender and race-implying names in standardized sentiment tasks, highlighting the risk of amplifying discriminatory patterns when sentiment analysis systems are used to segment customers or prioritize service responses. The governance of automated sentiment analysis in commercial contexts is addressed by the EU AI Act and equivalent regulatory frameworks as a high-risk AI application in contexts involving consumer protection, requiring transparency documentation, bias testing, and human oversight provisions [24].

IV. METHODOLOGY

4.1 Dataset Construction

The primary dataset comprises 2.3 million customer communication records collected from enterprise partners across four industry sectors: e-commerce (n=812,000), financial services (n=547,000), hospitality (n=389,000), and healthcare (n=552,000). All records were pseudonymized and processed under data processing agreements compliant with GDPR and applicable national data protection regulations. Each record consists of a text communication (email, chat message, support ticket, or review) along with associated metadata including communication channel, timestamp, and resolution outcome. A stratified random sample of 45,000 records was manually annotated for sentiment by trained annotators using a five-category scheme (strongly positive, mildly positive, neutral, mildly negative, strongly negative) with inter-annotator agreement of Cohen's kappa = 0.79 across annotation pairs [25].

TABLE I: Dataset Summary Statistics by Industry Sector

Sector	Total Records	Annotated Subset	Avg. Text Length (tokens)	Positive Sentiment (%)	Negative Sentiment (%)
E-commerce	812,000	16,200	87	42%	31%
Financial Services	547,000	10,900	142	28%	38%
Hospitality	389,000	7,800	198	51%	22%
Healthcare	552,000	11,000	163	34%	29%
Total	2,300,000	45,900	143	39%	30%

4.2 Hybrid Sentiment Architecture

The proposed sentiment analysis architecture, designated BizBERT-Hybrid, integrates three components: domain-adaptive pretraining, task fine-tuning, and lexicon-augmented output adjustment. Domain-adaptive pretraining extends the BERT-base-uncased checkpoint through additional masked language modeling on 4.2 million business communication records (excluding annotated evaluation records) drawn from the full dataset, enabling the model to adapt its token representations to business communication vocabulary, style, and context before fine-tuning on sentiment-labeled examples [16]. The domain-adaptive pretraining phase runs for 50,000 steps with a learning rate of 2e-5 and batch size of 256, requiring approximately 18 GPU-hours on NVIDIA A100 hardware.

Task fine-tuning on the annotated subset employs a five-class classification head appended to the BERT [CLS] token representation with dropout regularization at 0.3 and trained for 10 epochs with learning rate warm-up and cosine decay scheduling. The lexicon augmentation component applies domain-specific sentiment lexicons constructed through a combination of financial domain lexicons (Loughran-

McDonald), hospitality service quality dictionaries, and healthcare patient communication lexicons to adjust predicted probability distributions for cases where the neural model's prediction confidence falls below a calibrated threshold, exploiting complementary lexical knowledge that the neural model may underweight despite domain-adaptive pretraining.

TABLE II: Sentiment Classification Architecture Comparison Results

Model Variant	Overall Accuracy (%)	Macro F1	Positive F1	Negative F1	Inference Latency (ms)
Lexicon-only Baseline	77.5	0.741	0.763	0.712	3
BERT-base (no adapt.)	83.3	0.812	0.831	0.798	24
Domain-Adapted BERT	88.9	0.874	0.891	0.856	24
BizBERT-Hybrid (Proposed)	91.7	0.908	0.921	0.894	27
RoBERTa Large (General)	87.4	0.861	0.878	0.843	48
GPT-4 Zero-Shot	84.2	0.831	0.849	0.812	320

4.3 Marketing Automation Integration

Integration of the BizBERT-Hybrid sentiment analysis pipeline with marketing automation platforms follows a microservices architecture in which sentiment classification runs as an independent API service invoked synchronously by the automation platform's event processing layer. The integration layer implements a confidence-weighted trigger system in which automation actions are conditioned on both the predicted sentiment class and the model's calibrated confidence score, with low-confidence predictions routed to human review queues rather than triggering automated responses. Three trigger sensitivity configurations are defined: Conservative (requiring confidence above 0.85 for automated action), Standard (confidence above 0.70), and Aggressive (confidence above 0.55), allowing organizations to calibrate the automation-human review balance according to their risk tolerance and operational capacity.

4.4 Randomized Controlled Trial Design

The business performance evaluation employs a randomized controlled trial in which customer communication records are randomly assigned at the account level to sentiment-adaptive automation (treatment) or standard rule-based automation (control). The trial covers 180,000 customer accounts across the e-commerce and hospitality sectors over a 16-week evaluation period. Primary outcome measures include email open rate, click-through rate, conversion rate, CSAT score (measured through post-interaction surveys), and 90-day retention rate. Randomization at account level prevents within-account contamination, and Fisher's exact test and Mann-Whitney U tests are employed for outcome comparisons given the non-normality of conversion and retention distributions [26].

V. RESULTS AND ANALYSIS

5.1 Sentiment Classification Performance

The BizBERT-Hybrid architecture achieves overall sentiment classification accuracy of 91.7 percent on the held-out evaluation set, representing improvements of 8.4 percentage points over vanilla BERT-base and 14.2 percentage points over the lexicon-only baseline. Performance is heterogeneous across industry sectors and sentiment classes, with highest accuracy achieved for hospitality communications (93.2%) and lowest for financial services (89.1%), reflecting the greater linguistic formality and regulatory-constrained language of financial communications that reduces the expressiveness of sentiment signals available to the classifier [7]. The lexicon augmentation component provides the greatest incremental benefit for financial services communications, contributing 2.8 percentage points of accuracy improvement beyond domain-adapted BERT in that sector.

TABLE III: Randomized Controlled Trial Business Outcome Results

Outcome Metric	Control (Rule-Based)	Treatment (Sentiment-Adaptive)	Absolute Improvement	Relative Improvement	p-value
Email Open Rate	18.3%	22.7%	+4.4pp	+24.0%	< 0.001
Click-Through Rate	3.1%	4.2%	+1.1pp	+35.5%	< 0.001
Conversion Rate	1.8%	2.4%	+0.6pp	+33.3%	< 0.001
CSAT Score (1–5)	3.62	3.91	+0.29	+8.0%	< 0.001
90-day Retention Rate	71.4%	74.8%	+3.4pp	+4.8%	0.003
Human Escalation Rate	4.2%	5.8%	+1.6pp	+38.1%	< 0.001

5.2 Business Performance Impact

The randomized controlled trial demonstrates statistically significant improvements across all primary outcome measures for the sentiment-adaptive automation treatment relative to standard rule-based automation. The 24 percent

Citation: Athawale S. et.al. (2026). Digital Marketing Automation and Sentiment Analysis Using NLP in Business Communication. *International Insurance Law Review*, 34 (S1), 491-500.

improvement in email open rate represents the most immediate measurable benefit, attributable to the sentiment-adaptive subject line and preview text personalization that matches the emotional register of marketing communications to the customer's recent sentiment trajectory [4]. The 33 percent improvement in conversion rate has the largest revenue implication, with the estimated incremental revenue per 100,000 customer accounts exceeding $180,000 annually at average transaction values for the participating e-commerce sector partners.

The increase in human escalation rate (38.1 percent) reflects the system's appropriate routing of low-confidence and high-stakes negative sentiment communications to human agents rather than attempting automated resolution a feature that, while increasing operational costs for human review, likely contributes to the observed CSAT improvement by ensuring that the most complex and emotionally charged customer interactions receive human attention [21].

5.3 Bias Analysis

Systematic bias testing reveals differential classification accuracy across customer demographic segments inferred from name-based demographic estimation. The most significant disparity is observed between customers with names associated with South Asian origins (accuracy 89.4%) and customers with names associated with Anglo-Saxon origins (accuracy 93.1%), a 3.7 percentage point gap that persists after domain-adaptive pretraining. Financial services communications show the largest demographic accuracy gaps (5.2 percentage points), while hospitality communications show the smallest (1.8 percentage points). Debiasing interventions including adversarial training components and demographic-balanced fine-tuning reduce but do not eliminate these disparities, with optimized configurations reducing the maximum demographic accuracy gap to 2.4 percentage points [23].

VI. CONCLUSION

This paper has demonstrated that domain-adaptive NLP-based sentiment analysis can be integrated into digital marketing automation workflows to produce statistically significant and economically meaningful improvements in business communication outcomes including email engagement, conversion rates, customer satisfaction, and retention. The BizBERT-Hybrid architecture provides a technically validated solution to the domain-specificity challenge that limits general-purpose sentiment analysis models in professional business communication contexts, with performance improvements over both lexicon-based and general transformer baselines across all evaluated industry sectors.

The randomized controlled trial evidence provides rigorous causal identification of the business performance benefits of sentiment-adaptive automation, advancing beyond the correlational evidence that characterizes most prior work on marketing automation effectiveness. The documented bias disparities across demographic groups and the ethical governance framework proposed for enterprise deployment address dimensions of responsible AI deployment in customer-facing systems that the technical NLP literature has historically underemphasized relative to performance optimization.

VII. FUTURE WORK

Future research directions include the extension of the domain-adaptive pretraining approach to multilingual business communication contexts, which represent an increasingly important deployment scenario for global enterprises managing customer communications across

Citation: Athawale S. et.al. (2026). Digital Marketing Automation and Sentiment Analysis Using NLP in Business Communication. *International Insurance Law Review*, 34 (S1), 491-500.

language boundaries. The integration of emotional dimension analysis beyond positive-negative valence distinguishing frustration from disappointment, or enthusiasm from satisfaction offers the potential for more nuanced automation triggers that respond to specific emotional states rather than broad sentiment polarity. Investigation of the longitudinal dynamics of customer sentiment trajectories and their predictive value for churn and loyalty outcomes would extend the automation integration framework from reactive to predictive applications. Finally, the development of explainable sentiment analysis methods that can generate human-interpretable rationales for classification decisions would address both the governance requirements of emerging AI regulations and the operational needs of marketing professionals who require transparency into automated decision logic.

REFERENCES

[1] R. T. Rust, J. J. Inman, J. Jia, and A. Zahorik, "What you don't know about customer-perceived quality: The role of customer expectation distributions," Marketing Science, vol. 18, no. 1, pp. 77–92, Winter 1999.

[2] Grand View Research, "Marketing Automation Market Size, Share and Trends Analysis Report," San Francisco, CA: Grand View Research, 2023.

[3] B. Liu, Sentiment Analysis and Opinion Mining. San Rafael, CA: Morgan & Claypool, 2012.

[4] V. Kumar and W. Reinartz, Customer Relationship Management: Concept, Strategy, and Tools, 3rd ed. Berlin: Springer, 2018.

[5] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in Proc. NAACL-HLT 2019, Minneapolis, MN, 2019, pp. 4171–4186.

[6] A. Pontiki, D. Galanis, J. Pavlopoulos, H. Papageorgiou, I. Androutsopoulos, and S. Manandhar, "SemEval-2014 task 4: Aspect based sentiment analysis," in Proc. SemEval-2014, Dublin, 2014, pp. 27–35.

[7] T. Loughran and B. McDonald, "When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks," Journal of Finance, vol. 66, no. 1, pp. 35–65, Feb. 2011.

[8] R. McAuley and J. Leskovec, "Hidden factors and hidden topics: Understanding rating dimensions with review text," in Proc. RecSys 2013, Hong Kong, 2013, pp. 165–172.

[9] B. Pang and L. Lee, "Opinion mining and sentiment analysis," Foundations and Trends in Information Retrieval, vol. 2, no. 1–2, pp. 1–135, 2008.

[10] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language models are unsupervised multitask learners," OpenAI Blog, vol. 1, no. 8, p. 9, 2019.

[11] P. Turney, "Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews," in Proc. ACL, Philadelphia, PA, 2002, pp. 417–424.

[12] B. Pang, L. Lee, and S. Vaithyanathan, "Thumbs up? Sentiment classification using machine learning techniques," in Proc. EMNLP, Philadelphia, PA, 2002, pp. 79–86.

[13] C. Hutto and E. Gilbert, "VADER: A parsimonious rule-based model for sentiment analysis of social media text," in Proc. ICWSM, Ann Arbor, MI, 2014.

[14] R. Socher et al., "Recursive deep models for semantic compositionality over a sentiment treebank," in Proc. EMNLP, Seattle, WA, 2013, pp. 1631–1642.

[15] A. Vaswani et al., "Attention is all you need," in Proc. NeurIPS, Long Beach, CA, 2017, pp. 5998–6008.

[16] S. Gururangan et al., "Don't stop pretraining: Adapt language models to domains and tasks," in Proc. ACL, Online, 2020, pp. 8342–8360.

[17] J. Lee et al., "BioBERT: A pre-trained biomedical language representation model for biomedical text mining," Bioinformatics, vol. 36, no. 4, pp. 1234–1240, Feb. 2020.

[18] D. Araci, "FinBERT: Financial sentiment analysis with pre-trained language models," arXiv preprint arXiv:1908.10063, 2019.

[19] N. Heimbach and O. Hinz, "The impact of content characteristics and emotion on behavioral engagement in social media: Literature review and research agenda," Psychology & Marketing, vol. 33, no. 6, pp. 508–525, Jun. 2016.

[20] V. Chaffey and P. R. Smith, Emarketing Excellence: Planning and Optimizing Your Digital Marketing, 4th ed. Abingdon, UK: Routledge, 2013.

[21] A. L. Ostrom et al., "Improving service delivery: Implications of identified managerial practices," Journal of Service Research, vol. 13, no. 1, pp. 4–36, Feb. 2010.

[22] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, "On the dangers of stochastic parrots: Can language models be too big?" in Proc. FAccT 2021, Virtual, 2021, pp. 610–623.

[23] S. Kiritchenko and S. M. Mohammad, "Examining gender and race bias in two hundred sentiment analysis systems," in Proc. *SEM, New Orleans, LA, 2018, pp. 43–53.

[24] European Commission, "Regulation of the European Parliament and of the Council Laying Down Harmonized Rules on Artificial Intelligence (Artificial Intelligence Act)," Brussels: European Commission, 2021.

[25] J. Cohen, "A coefficient of agreement for nominal scales," Educational and Psychological Measurement, vol. 20, no. 1, pp. 37–46, Apr. 1960.

[26] D. R. Cox, The Design of Experiments. Oxford: Oxford University Press, 1958.

[27] Y. Liu et al., "RoBERTa: A robustly optimized BERT pretraining approach," arXiv preprint arXiv:1907.11692, 2019.

[28] S. Sun, C. Luo, and J. Chen, "A review of natural language processing techniques for opinion mining systems," Information Fusion, vol. 36, pp. 10–25, Jul. 2017.

Citation: Athawale S. et.al. (2026). Digital Marketing Automation and Sentiment Analysis Using NLP in Business Communication. *International Insurance Law Review*, 34 (S1), 491-500.