

Deep Learning Applications in Digital Marketing of Luxury Brands : Consumer Sentiment Analysis Using YSL as a Case Study

Jiaxin Lei*

Dulwich International Hengqin High School, Zhuhai, Guangdong, 519031, China

*Corresponding author: Jiaxin Lei, jiaxin_lei@163.com

Copyright: 2026 Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY-NC 4.0), permitting distribution and reproduction in any medium, provided the original author and source are credited, and explicitly prohibiting its use for commercial purposes.

Abstract: The world has gone digital; therefore, luxury fashion brands have had to adapt their brand-building methods to some extent. Apply various deep learning methods, such as Long Short-Term Memory (LSTM) networks, Bidirectional Encoder Representations from Transformers (BERT), and Convolutional Neural Networks (CNN), to conduct consumer sentiment analysis for Yves Saint Laurent (YSL) in the luxury fashion industry in this paper. Based on a corpus of over 120,000 user-generated content (UGC) items collected from Weibo, Instagram and other brand-owned platforms between 2020 and 2024, multiple deep learning models for multi-class sentiment classification and aspect-level analysis have been constructed and evaluated in this study. According to the above experiments, the finely tuned BERT model has achieved a classification accuracy of 91.3% and exceeded the results of the traditional machine-learning baseline by 14.7%. Based on the above results, the main reasons for changes in public opinion about YSL are quality issues, brand history and influencer endorsements. Based on the above analysis, some practical suggestions are put forward for managers of luxury brands, and at the same time, a data-driven digital marketing system for real-time sentiment analysis and campaign optimization is proposed.

Keywords: Deep Learning; Sentiment Analysis; Luxury Brand Marketing; Digital Marketing; BERT; LSTM; YSL; Natural Language Processing; Consumer Behavior

Published: Aug 28, 2026

DOI: <https://doi.org/10.62177/apemr.v3i7.1678>

1. Introduction

The global luxury goods market, valued at approximately USD 352 billion in 2023 and projected to surpass USD 500 billion by 2030, occupies a singular position within the broader consumer economy. Luxury brands derive their competitive advantage not merely from product superiority but from an intricate ecosystem of aspirational identity, cultural heritage, and curated exclusivity. The digital revolution has disrupted this ecosystem profoundly: consumers no longer passively receive brand narratives but actively co-create them through social media commentary, online reviews, and influencer discourse. For luxury houses such as Yves Saint Laurent (YSL)—a brand synonymous with Parisian haute couture since its founding in 1961—navigating this participatory media environment constitutes one of the most complex strategic challenges of the contemporary era ^[1].

Traditional marketing research tools—consumer surveys, focus groups, and retrospective sales analyses—are increasingly

inadequate for capturing the velocity, volume, and nuance of online consumer sentiment. The emergence of deep learning as a dominant paradigm in natural language processing (NLP) offers a transformative alternative. Deep learning models, particularly transformer-based architectures such as BERT^[2], have demonstrated remarkable capacity for semantic understanding, contextual disambiguation, and cross-lingual generalization—capabilities that are indispensable for analyzing the multilingual, slang-laden, and culturally stratified discourse characteristic of luxury brand conversations on platforms such as Weibo and Instagram.

Despite the growing body of literature on sentiment analysis and luxury marketing separately, the systematic application of advanced deep learning architectures to luxury brand digital marketing strategy remains an underexplored domain. Existing studies have predominantly employed lexicon-based or shallow machine learning approaches that fail to capture the subtle irony, aspirational framing, and polarity complexity inherent in luxury consumption discourse. Furthermore, the aspect-level granularity required for actionable brand management insights—distinguishing negative sentiment toward pricing from positive sentiment toward craftsmanship—has rarely been addressed in luxury-specific NLP research.

This paper addresses these gaps through a comprehensive empirical study centered on YSL's digital marketing ecosystem. The research objectives are threefold: (1) to develop and benchmark deep learning models for multi-class and aspect-level sentiment classification of luxury brand UGC; (2) to identify the principal determinants of consumer sentiment toward YSL across multiple digital platforms and temporal periods; and (3) to derive strategic implications for luxury digital marketing practice from these computational insights. By integrating technical rigor with marketing theory, this study aspires to advance both scholarly understanding and practical management of brand sentiment in the luxury sector.

2. Literature Review

2.1 Luxury Brand Marketing in the Digital Age

The democratization of digital media has forced luxury brands to reconcile the inherent tension between exclusivity and accessibility. Kapferer and Bastien articulated the 'anti-laws' of luxury marketing, emphasizing that traditional luxury strategy deliberately eschews mass-market principles. Subsequent scholarship has recognized that digital platforms create participatory spaces that simultaneously threaten and amplify brand desirability. The concept of the 'digital luxury paradox' captures this duality: social media engagement broadens awareness and fosters emotional brand attachment, yet risks diluting the aura of unattainability that underpins luxury premiums.

Empirical research on luxury social media marketing has identified several key drivers of brand equity and purchase intention, including content quality, perceived authenticity, parasocial interaction with brand ambassadors, and community belonging. Godey et al.^[3] conducted a large-scale survey of 845 luxury brand consumers across four countries, developing a structural equation model that demonstrates robust links between social media marketing activities and brand preference, price premium tolerance, and consumer loyalty. Their findings confirm that luxury brands operating on social media must carefully manage both informational and hedonic content dimensions to sustain brand equity.

In the Chinese market—where YSL generates a substantial proportion of global revenues—Weibo, WeChat, and Xiaohongshu (Little Red Book) have emerged as pivotal battlegrounds for luxury brand positioning. The distinctive local consumer psychology emphasizing social signaling and Key Opinion Leader (KOL) influence adds further complexity to sentiment dynamics that any analytical framework must address.

2.2 Sentiment Analysis: Methods and Evolution

Sentiment analysis—the computational identification and categorization of opinions expressed in textual data—has evolved through three principal methodological generations. The first generation relied on manually curated sentiment lexicons that assign polarity scores to individual words or phrases; while computationally efficient, these approaches are brittle in the face of negation, sarcasm, domain-specific vocabulary, and contextual polarity reversal. The second generation employed statistical machine learning algorithms—Naive Bayes, Maximum Entropy, Support Vector Machines—trained on labeled corpora, achieving moderate performance improvements but requiring substantial feature engineering effort.

The deep learning revolution inaugurated the third generation, characterized by end-to-end representation learning that circumvents manual feature design. Recurrent Neural Networks (RNNs) and their gated variants—Long Short-Term Memory

(LSTM) ^[4] and Gated Recurrent Units—captured sequential dependencies in text, demonstrating substantial performance gains on benchmark sentiment datasets. Convolutional Neural Networks adapted for text classification ^[5] proved particularly effective at extracting local n-gram features relevant to sentiment. The introduction of the transformer architecture ^[6] and subsequently BERT constituted inflection points, enabling context-sensitive representations that capture long-range dependencies and polysemous semantics with unprecedented fidelity.

2.3 Influencer Marketing and Consumer Credibility

Influencer marketing has become an indispensable component of luxury brand digital strategy, yet its effects on consumer sentiment are highly variable. Djafarova and Rushworth ^[7] investigated the credibility of online celebrities' Instagram profiles in influencing the purchase decisions of young female consumers, finding that micro-level influencers—those with smaller but highly engaged follower bases—generated greater perceived authenticity and purchasing intent than mega-celebrities with mass followings. This parasocial intimacy effect is particularly salient in the luxury domain, where brand association is intertwined with aspirational identity construction.

For YSL and comparable luxury houses, the strategic deployment of influencers must balance reach with authenticity. The brand's positioning as an emblem of provocative French elegance demands influencer partnerships that can credibly embody the brand archetype without diluting the exclusive aura. Computational sentiment analysis provides a means to evaluate the aggregate sentiment impact of influencer activations in near real-time, enabling more agile and evidence-based partnership decisions.

3. Methodology

3.1 Data Collection and Preprocessing

The dataset for this study was assembled from three primary digital channels: Sina Weibo (China's dominant microblogging platform), Instagram, and YSL's official brand community portal. Data collection was conducted using platform-authorized APIs combined with structured web scraping protocols compliant with each platform's terms of service. The collection window spanned January 2020 to December 2024, capturing a five-year longitudinal arc that encompasses the COVID-19 pandemic disruption, the resurgence of offline retail, and the maturation of live-stream commerce in China.

Keyword-based queries included YSL's Chinese romanized brand names ('YSL', 'Saint Laurent', '圣罗兰'), flagship product lines (Black Opium, Libre, Loulou, Niki), and branded hashtags. After deduplication, bot-generated content filtering (using a rule-based classifier validated against manual inspection), and removal of posts with fewer than five characters, the final corpus comprised 127,840 items: 68,420 Weibo posts, 41,350 Instagram posts and comments, and 18,070 brand portal reviews. Language distribution was approximately 61% Simplified Chinese, 34% English, and 5% other languages (primarily French and Korean).

Preprocessing pipelines were language-specific. For Chinese text, the Jieba tokenizer was employed for word segmentation, supplemented by a custom luxury fashion lexicon of 2,847 domain-specific terms. For English text, SpaCy's `en_core_web_lg` model was used for tokenization, lemmatization, and named entity recognition. Across both language streams, URL removal, emoji standardization (emojis were retained as informative sentiment signals and mapped to textual descriptors), and HTML entity decoding were applied. A 10% stratified subsample of 12,784 items was reserved as the final held-out test set; the remaining 90% constituted the training and validation pool.

3.2 Annotation Schema and Labeling

The annotation schema was designed to support both document-level sentiment classification and aspect-level sentiment analysis (ABSA). At the document level, texts were assigned to one of five sentiment categories: Strongly Positive, Moderately Positive, Neutral, Moderately Negative, and Strongly Negative. For ABSA, six luxury-relevant aspects were defined a priori through thematic content analysis of a 2,000-post pilot sample: (1) Product Quality, (2) Brand Heritage and Identity, (3) Price and Value Perception, (4) Customer Service Experience, (5) Influencer and Ambassador Credibility, and (6) Packaging and Unboxing Experience.

A team of six trained annotators—three native Mandarin speakers and three native English speakers with backgrounds in luxury marketing—performed manual labeling on a 15,000-item annotation sample. Inter-annotator agreement was assessed

using Cohen’s kappa, yielding a mean kappa of 0.78 for document-level classification and 0.71 for ABSA, both exceeding the conventionally accepted threshold of 0.70 for substantial agreement^[8]. Disagreements were resolved through adjudication by a senior annotator. The 15,000 human-annotated items served as the gold-standard training data, augmented by a semi-supervised self-training procedure that extended labels to the full corpus using iterative model-assisted annotation with confidence thresholding.

3.3 Model Architectures

Three principal deep learning architectures were implemented and evaluated, alongside two conventional machine learning baselines for benchmarking purposes.

Baseline Models: A TF-IDF weighted Naive Bayes classifier and a linear Support Vector Machine (SVM) with RBF kernel were trained on unigram and bigram feature vectors as conventional benchmarks, representing the performance ceiling achievable without deep representation learning.

BiLSTM with Attention: A bidirectional LSTM with 256 hidden units per direction, coupled with a self-attention layer to weight the contribution of each hidden state to the classification decision. Pre-trained word embeddings (GloVe 300-d for English; Tencent AI Lab 200-d for Chinese) were used for initialization. Dropout rate of 0.3 was applied to recurrent connections to mitigate overfitting.

TextCNN: A multi-filter-size CNN (filter widths of 3, 4, and 5; 128 filters per size) over pre-trained word embeddings, with max-over-time pooling and a fully connected classification head. This architecture excels at extracting localized sentiment-bearing phrases, making it particularly effective for the short, emotionally direct posts typical of Weibo and Instagram.

Fine-Tuned BERT (BERT-FT): The BERT-base-multilingual-cased checkpoint was fine-tuned on the luxury UGC training set for document-level classification. The transformer architecture underlying BERT enables full bidirectional context modeling through self-attention, capturing semantic relationships that sequential models cannot. Fine-tuning was performed for 5 epochs with a learning rate of $2e-5$, batch size of 32, and AdamW optimizer with linear warmup and weight decay of 0.01.

4. Results and Analysis

4.1 Document-Level Sentiment Classification Performance

Table 1 presents the comparative performance of all models on the held-out test set across standard evaluation metrics: accuracy, macro-averaged precision, recall, and F1-score.

Table 1. Model Performance Comparison on Document-Level Sentiment Classification

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Naive Bayes (TF-IDF)	72.4	70.1	68.9	69.5
SVM (RBF)	76.6	75.3	73.8	74.5
TextCNN	83.7	82.9	81.4	82.1
BiLSTM + Attention	87.2	86.5	85.1	85.8
BERT-FT (Proposed)	91.3	90.8	90.1	90.4

The BERT-FT model achieves the highest performance across all metrics, recording an accuracy of 91.3% and a macro-F1 of 90.4%—a statistically significant improvement over the strongest baseline SVM ($p < 0.001$, McNemar’s test). The performance gap between BERT-FT and the BiLSTM+Attention model (4.1 percentage points in accuracy) underscores the superior contextual encoding capability of transformer self-attention relative to sequential recurrence. The self-attention mechanism in BERT allows every token to directly attend to every other token in the sequence, enabling the model to resolve long-range sentiment dependencies that sequential models process only indirectly through hidden state propagation.

Error analysis reveals that BERT-FT’s residual misclassifications are concentrated at the boundary between Moderately Positive and Neutral, and between Moderately Negative and Neutral—findings consistent with the inherent ambiguity of hedged luxury consumption language. Phrases such as ‘surprisingly reasonable for the quality’ or ‘could have been more

exclusive’ resist unambiguous polarity assignment even by human raters, as reflected in the kappa values reported in Section 3.2.

4.2 Aspect-Level Sentiment Analysis

The ABSA results reveal a differentiated sentiment landscape across YSL’s six brand aspects. Product Quality and Brand Heritage are the highest-frequency discussion topics (32.1% and 24.7% of aspect-level mentions, respectively) and attract the most positive mean sentiment polarity scores (0.68 and 0.71 on a -1 to +1 scale). These findings affirm that YSL’s core brand equity dimensions—craftsmanship and aesthetic heritage—continue to resonate strongly with digitally engaged consumers, serving as the bedrock of brand loyalty.

Price and Value Perception exhibits the most negative mean sentiment (-0.31), reflecting widespread consumer discourse around price increases—YSL has implemented multiple price adjustments globally since 2021—and perceptions of diminishing product differentiation relative to competing luxury houses at equivalent price points. Customer Service Experience (mean sentiment: 0.12) and Packaging Experience (0.58) occupy intermediate positions. Influencer and Ambassador Credibility presents the highest sentiment volatility (standard deviation: 0.49), reflecting the episodic nature of influencer controversies and the polarized reception of celebrity endorsements—a pattern aligned with the credibility research of Djafarova and Rushworth, who found that different types of celebrities elicit highly divergent trust and identification responses.

4.3 Temporal Dynamics and Campaign Analysis

Longitudinal analysis of monthly aggregated sentiment scores reveals several noteworthy temporal patterns. Overall positive sentiment exhibited a modest but consistent upward trend from 2020 to 2022 (from 58.3% to 64.7% of all classified posts), followed by a slight decline in 2023 coinciding with global inflationary pressures and social media controversies. A pronounced sentiment spike—the highest observed in the dataset—occurred in October 2021, corresponding to YSL’s 60th anniversary campaign, which generated over 2.3 million Weibo interactions in a two-week period and a sentiment positivity rate of 78.4%.

Conversely, a significant negative sentiment cluster emerged in March 2023 following an advertising campaign perceived by a segment of Chinese consumers as culturally insensitive. Within 72 hours of the campaign’s publication, negative sentiment posts constituted 61.2% of daily YSL-related discourse on Weibo—a reversal unprecedented in the five-year observation window. The brand’s subsequent retraction and public apology were followed by a 45-day sentiment recovery period before positivity rates returned to baseline. This episode vividly illustrates the reputational vulnerability of luxury brands in the hyper-connected digital environment and the critical importance of real-time sentiment monitoring.

5. Discussion

5.1 Theoretical Contributions

This study makes three principal theoretical contributions to the intersection of computational linguistics and luxury marketing scholarship. First, it demonstrates that fine-tuned multilingual BERT models, when trained on domain-specific luxury fashion corpora, achieve sentiment classification accuracy comfortably exceeding existing luxury NLP benchmarks, establishing a new methodological standard for this research domain. The superior performance derives from BERT’s capacity to model bidirectional contextual dependencies—a property enabled by the transformer’s self-attention mechanism—which is critical for interpreting the ironic understatement and aspirational hedging characteristic of luxury consumer discourse.

Second, the ABSA framework applied in this study enriches extant luxury brand equity theory by providing a computational operationalization of multi-dimensional brand perception. The finding that Brand Heritage sentiment significantly predicts overall brand positivity (Pearson’s $r = 0.74$, $p < 0.001$) corroborates Kapferer and Bastien’s theoretical proposition that heritage authenticity constitutes the primordial source of luxury value, while simultaneously demonstrating the measurability of this construct through UGC text analysis. This represents a methodological bridge between interpretive luxury marketing theory and quantitative computational social science.

Third, the longitudinal sentiment analysis contributes to the emerging literature on digital brand crisis management,

demonstrating the quantifiable rapidity of sentiment deterioration and recovery following reputational events. The 45-day recovery arc observed following the March 2023 controversy aligns with established communication theory's predicted timelines for credibility restoration through genuine remediation, validating computational sentiment dynamics as a real-time proxy for brand trust trajectories.

5.2 Managerial Implications and Strategic Framework

The findings carry concrete strategic implications for luxury brand managers navigating the digital marketing landscape. First, the robustness of Product Quality and Brand Heritage as sentiment drivers argues for content marketing strategies that foreground craft narratives, archival storytelling, and making-of documentation. YSL's most positively received digital content during the observation period (mean positivity: 76.3%) featured behind-the-scenes ateliers, artisan interviews, and heritage archival imagery. Godey et al. similarly found that content perceived as authentic and culturally rich generated stronger brand equity effects in luxury social media contexts, suggesting a convergence between large-scale computational findings and survey-based brand equity research.

Second, the pronounced negative sentiment associated with Price and Value Perception necessitates proactive communication strategies that justify premium pricing through transparent value articulation. This may include limited edition releases with artisanal framing, exclusive customer experience programs, and sustainability certifications that reframe price premiums within an ethical consumption narrative—Kapferer and Bastien's luxury strategy framework would characterize such moves as consistent with the anti-law principle that luxury must never justify its price, but can instead make price irrelevant by elevating the experiential and symbolic stakes.

Third, the high volatility of Influencer and Ambassador Credibility sentiment underscores the importance of rigorous KOL vetting. The data suggest that micro-influencers (10,000–100,000 followers) generate more consistently positive and authentic sentiment signals than mega-celebrities. Djafarova and Rushworth found that online celebrities perceived as more accessible and relatable generated higher credibility scores and stronger purchase intentions, particularly among younger female consumers—precisely the demographic that constitutes YSL's primary growth market.

Building upon these findings, this study proposes a Data-Driven Luxury Digital Marketing (DLDM) framework comprising three operational layers: (1) the Real-Time Sentiment Intelligence Layer, employing continuously updated deep learning models to monitor brand sentiment across all digital touchpoints and trigger alerts when sentiment deviates beyond defined thresholds; (2) the Aspect-Adaptive Campaign Layer, translating ABSA outputs into content strategy adjustments on a bi-weekly basis; and (3) the Predictive Crisis Management Layer, utilizing historical sentiment patterns and anomaly detection algorithms to identify emerging reputational risks before they escalate. Implementation of the DLDM framework requires investment in computational infrastructure, data science talent, and cross-functional integration between marketing, IT, and brand strategy functions.

5.3 Cross-Cultural Dimensions and Platform Specificity

A typical way in this study to address cross-cultural and cross-platform differences in sentiment expression is presented below. Chinese-language posts on Weibo show a systematic higher intensity of sentiment than English-language Instagram content for the same opinion valence, which can be attributed to differences in politeness norms, platform affordances, and the culturally specific emotional expressiveness of Chinese social media culture. BERT has introduced language-specific subword vocabularies to address the cross-lingual gap and improved the performance of multilingual models; however, its F1 score for Chinese ABSA (87.6%) is slightly lower than that for English ABSA (91.2%), and this is due to the greater morphological complexity of Chinese aspect boundary identification.

Based on the above analysis, Instagram comments are generally shorter, more positive, and contain relatively more emojis; Weibo posts are comparatively longer and more detailed, and they cover all kinds of emotions. Reviews in the brand portal account for a small proportion of the total, but they are rich in attribute-level information and show a strong correlation between expressed sentiment and subsequent purchase behavior; thus, they are highly beneficial sources of input for the ABSA layer of the proposed DLDM framework.

6. Limitations and Future Research Directions

Several limitations qualify the findings reported in this study. First, the dataset is confined to textual modalities; luxury brand discourse on platforms such as Instagram and Xiaohongshu is increasingly multimodal, with images, short-form videos, and live streams constituting substantial proportions of consumer communication. The BERT model analyzed here operates solely on text; future research should incorporate multimodal deep learning architectures that jointly encode visual and textual content to capture the full semantic richness of luxury digital discourse.

Second, the study focuses on a single luxury brand (YSL), limiting the generalizability of findings across the luxury sector's heterogeneous brand landscape. Comparative analyses spanning multiple luxury houses—for instance, contrasting YSL with Gucci, Prada, and Dior—would enable identification of brand-specific versus sector-wide sentiment dynamics and facilitate the development of more universally applicable analytical frameworks. The luxury social media marketing literature has demonstrated that different luxury brands elicit systematically different sentiment architectures even when operating in the same product categories and geographies.

Third, while this study establishes correlational linkages between sentiment dynamics and marketing events, causal inference from observational social media data remains methodologically challenging. Future studies should explore quasi-experimental designs—such as difference-in-differences analyses of controlled campaign interventions—to strengthen causal claims. Additionally, the ethical dimensions of large-scale social media data collection and automated sentiment profiling merit sustained scholarly attention, particularly regarding consumer consent and the potential for algorithmic targeting to exacerbate exclusionary consumption dynamics.

7. Conclusion

This paper presents all-round research on deep learning-based consumer sentiment analysis as the foundation for digital marketing of luxury brands and takes Yves Saint Laurent as a case study. The fine-tuned multilingual BERT model has reached a new high of 91.3% accuracy for sentiment classification of luxury fashion UGC and has greatly outperformed traditional machine learning methods. At the level of attributes, YSL's brand equity in the digital space is mainly supported by perceptions of "Product Quality" and "Brand Heritage", while "Price and Value" discourse accounts for the largest portion of negative sentiment and thus needs to be addressed through strategy.

According to the results of the longitudinal study, most people's attitudes towards luxury brands in the digital era have been stable for a long time and are well-supported by a robust system of brand equity and authentic heritage; however, they are also very sensitive to changes in specific marketing activities, price adjustments and cultural controversies at the micro level. The 2023 reputational crisis case study in this paper shows quantitatively that the speed at which sentiment declined and how long it took to restore credibility are relatively slow; therefore, proactive real-time monitoring systems are needed for luxury brand managers.

The new Data-driven Luxury Digital Marketing (DLDM) framework will organize the research results into an all-encompassing management plan for real-time sentiment analysis, aspect-adaptive campaign planning, and forward-looking crisis warning. With the continuous development of artificial intelligence and the expansion of applications for large language models and multimodal understanding, deep learning sentiment analysis will be introduced to help address problems related to change and diversity in global digital marketing for luxury brands.

The new scale at which we can listen is large, and by organizing the vast amount of consumer voices from various digital platforms into meaningful brand intelligence, a new form of knowledge for luxury marketing research is about to emerge. This paper offers a strict research method and a practical path for the above-mentioned transformation, adding to the body of scholarly work at the intersection of computational social science and luxury brand management.

Funding

No

Conflict of Interests

The authors declare that there is no conflict of interest regarding the publication of this paper.

Reference

- [1] Kapferer, J.-N., & Bastien, V. (2012). *The luxury strategy: Break the rules of marketing to build luxury brands* (2nd ed.). Kogan Page.
- [2] Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT 2019)* (pp. 4171–4186). Association for Computational Linguistics.
- [3] Godey, B., Manthiou, A., Pederzoli, D., Rokka, J., Aiello, G., Donvito, R., & Singh, R. (2016). Social media marketing efforts of luxury brands: Influence on brand equity and consumer behavior. *Journal of Business Research*, 69(12), 5833–5841.
- [4] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
- [5] Kim, Y. (2014). Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 1746–1751). Association for Computational Linguistics.
- [6] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems 30 (NeurIPS 2017)* (pp. 5998–6008). Curran Associates.
- [7] Djafarova, E., & Rushworth, C. (2017). Exploring the credibility of online celebrities' Instagram profiles in influencing the purchase decisions of young female users. *Computers in Human Behavior*, 68, 1–7.
- [8] Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33(1), 159–174.