



CascadeNS: Confidence-Cascaded Neurosymbolic Model for Sarcasm Detection

Swapnil Mane^{*1,3} and Vaibhav Khatavkar^{2,3}

¹ IIT Jodhpur, India

² DES Pune University, Pune, India

³ COEP Technological University, Pune, India

mane.1@iitj.ac.in, vaibhav.khatavkar@despu.edu.in

Abstract. Sarcasm detection in product reviews requires balancing domain-specific symbolic pattern recognition with deep semantic understanding. Symbolic representations capture explicit linguistic phenomena that are often decisive for sarcasm detection. Existing work either favors interpretable symbolic representation or semantic neural modeling, but rarely achieves both effectively. Prior hybrid methods typically combine these paradigms through feature fusion or ensembling, which can degrade performance. We propose CascadeNS, a confidence-calibrated neurosymbolic architecture that integrates symbolic and neural reasoning through selective activation rather than fusion. A symbolic semigraph handles pattern-rich instances with high confidence, while semantically ambiguous cases are delegated to a neural module based on pre-trained LLM embeddings. At the core of CascadeNS is a calibrated confidence measure derived from polarity-weighted semigraph scores. This measure reliably determines when symbolic reasoning is sufficient and when neural analysis is needed. Experiments on product reviews show that CascadeNS outperforms strong baselines by 7.44%.

Code: GitHub repository

Keywords: Sarcasm Detection, Neurosymbolic AI, Semigraph Representation, Natural Language Processing

1 Introduction

Sarcasm is the expression of an intended meaning that contradicts its literal interpretation and poses fundamental challenges for automated text understanding [10]. In e-commerce product reviews, sarcastic statements such as “*Great product! Broke after one use*” can severely distort sentiment analysis systems and mislead recommendation algorithms. Recent transformer-based approaches [5,16,9,11,14] achieve strong performance for sarcasm detection by learning deep contextualized representations. However, a critical limitation persists: there is a mismatch between how humans recognize sarcasm and what pre-trained models capture. Computational linguistics research shows that humans rely heavily on explicit surface cues when recognizing sarcasm [8,2]. González-Ibáñez et al. [8] demonstrate through controlled studies that sarcasm perception depends on explicit linguistic phenomena such as exaggerated punctuation (e.g.,

* Corresponding author: mane.1@iitj.ac.in

!!! and ???), sentiment-bearing interjections (*oh, wow*), and specific part-of-speech patterns like adverb-adjective intensification. Bouazizi and Ohtsuki [2] further identify hyperbolic expressions (*amazing, terrible, worst*) as strong discriminative signals in review text. Despite this evidence, pre-trained models largely under utilize these explicit symbolic patterns and markers, even though they are prevalent and decisive in review comments [8,2,1]. Conversely, symbolic methods encode explicit linguistic theories through hand-crafted features, effectively capturing pattern-rich instances where surface cues are decisive [4,22]. However, these approaches struggle when explicit patterns are absent and deeper semantic reasoning is required [29].

This gap motivates our main research question: *Can we design a hybrid architecture that effectively leverages explicit linguistic phenomena for pattern-rich cases while maintaining the semantic power of neural models for pattern-sparse instances?* To address this, we propose CascadeNS, a confidence-cascaded neurosymbolic model for sarcasm detection. The model employs a symbolic semigraph representation as the first stage, extracting domain-specific linguistic features and quantifying prediction confidence. When confidence indicates that explicit linguistic phenomena are decisive ($\gamma \geq \tau$), symbolic predictions terminate immediately. Otherwise, low-confidence cases delegate to a second-stage neural classifier based on pre-trained LLM embeddings and k-nearest neighbor voting. This selective activation strategy preserves the strengths of both paradigms in their respective competence regions. Experimental results on Amazon product reviews demonstrate that CascadeNS outperforms strong baselines by 7.44%. Rigorous calibration analysis validates that confidence reliably indicates when explicit linguistic features suffice versus when deep semantic analysis is required. Comprehensive ablation studies validate the CascadeNS design architecture.

2 Related Work

Early computational approaches to sarcasm detection exploited lexical patterns [4] and sentiment-situation incongruity [22], where positive sentiment words describe objectively negative situations. González-Ibáñez et al. [8] identified key explicit markers including interjections, exaggerated punctuation, and part-of-speech patterns through corpus analysis. Bouazizi and Ohtsuki [2] formalized pattern-based detection using regular expressions over linguistic features. The rise of neural architectures shifted research toward transformers. Ghosh and Veale [7] applied CNNs and LSTMs to learn sarcasm representations from raw text. Tay et al. [27] proposed reasoning with sarcasm by modeling discourse structure through attention mechanisms. Recent work leverages pre-trained transformers, Li and Zhang [14] optimized BERT and RoBERTa fine-tuning strategies. Yang et al. [30] integrated BERT with gated graph neural networks over dependency and emotion graphs. Sharma et al. [26] exploited multi-utterance context in conversational data.

Symbolic approaches offer complementary advantages. Anderson and Taylor [1] demonstrated that symbolic pattern extraction enables interpretable NLP. However, pure symbolic methods struggle with semantic ambiguity, requiring pragmatic reasoning [29]. Neurosymbolic AI combines neural learning with symbolic reasoning to achieve complementary strengths [18]. Patel et al. [20] surveyed neurosymbolic ap-

proaches for knowledge graph reasoning, while Li et al. [15] studied neural-symbolic integration from a query perspective. Garcia and Fernandez [6] provide theoretical foundations for hybrid symbolic-subsymbolic systems, emphasizing representation compatibility challenges. However, limited work strategically integrates symbolic sarcasm with neural models using confidence-based cascading.

3 CascadeNS: Confidence-Cascaded Neurosymbolic Model for Sarcasm Detection

Problem Formulation. Let \mathcal{X} denote the space of product review comments and let $\mathcal{Y} = \{0, 1\}$ be the label space, where $y = 1$ denotes sarcastic and $y = 0$ denotes non-sarcastic content. We are given a training set $\mathcal{D} = (x_i, y_i)_{i=1}^n$ drawn i.i.d. from an unknown distribution $P(\mathcal{X}, \mathcal{Y})$. The objective is to learn a classifier $f : \mathcal{X} \rightarrow \mathcal{Y}$ that maximizes performance by effectively leveraging both explicit linguistic phenomena and deep semantic patterns. To this end, we formulate f as a cascade of two complementary components, a symbolic f_{sym} equipped with a confidence estimator γ , and a neural classifier f_{neu} . For an input $x \in \mathcal{X}$, the symbolic module produces a prediction $f_{\text{sym}}(x)$ and an associated confidence score $\gamma(x)$. If $\gamma(x)$ exceeds a predefined threshold τ , the symbolic prediction is accepted and returned, as explicit linguistic phenomena are decisive. Otherwise, the input is deferred to the neural module $f_{\text{neu}}(x)$ for deep semantic analysis.

3.1 Sarcasm Symbolic Semigraph

Symbolic Representation. Seven symbolic linguistic feature types are extracted $\{\phi_j\}_{j=1}^7$ motivated by computational linguistics research on sarcasm markers [8,2,1]. These include lexical n-grams (bigrams and trigrams capturing phrasal patterns like “oh great” and “best purchase ever”), part-of-speech n-grams (capturing syntactic patterns such as ADV+ADJ indicating intensification), interjections (wow, oh, yay), punctuation patterns (! ! !, ? ? ?, . . .), and hyperbolic expressions (amazing, terrible, worst). For review x , feature extraction yields $\phi(x) = [\phi_1(x), \dots, \phi_7(x)]$ where each $\phi_j(x)$ is the set of instances of feature type j in x .

Polarity-Weighted Bipartite Graph Construction. We model sarcasm detection as graph-based polarity aggregation (Figure 1). Construct bipartite graph $G = (V_f \cup V_c, E, W)$ where $V_f = \{\phi_1, \dots, \phi_7\}$ are feature nodes, $V_c = \{c^+, c^-\}$ are class nodes representing sarcasm and non-sarcasm respectively, and edge set $E = V_f \times V_c$ connects each feature type to both classes. Edge weights $W : E \rightarrow \mathbb{R}$ encode polarity scores derived from sentiment lexicons (SentiWordNet, VADER, TextBlob) [31]. For each symbolic instance $t \in \phi_j(x)$, we compute polarity weight to class c as:

$$w(t, c) = \text{pol}_c(t) \cdot \text{idf}(t) \quad (1)$$

where $\text{pol}_c(t)$ is the sentiment polarity score of the token t toward class c from lexicon aggregation, and $\text{idf}(t)$ is inverse document frequency weighting to emphasize discrim-

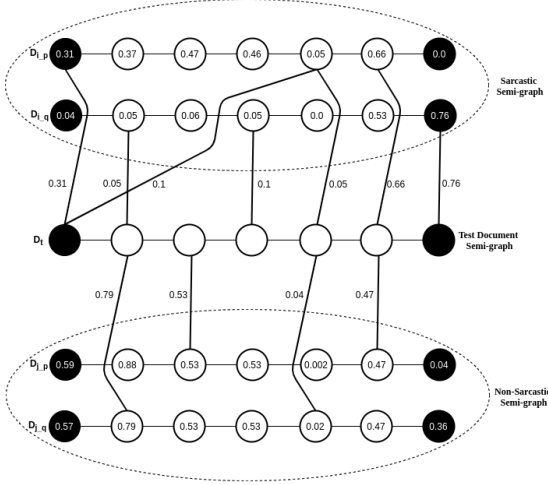


Fig. 1: Symbolic polarity-weighted bipartite semigraph connecting reviews.

inative terms. The aggregate weight for feature type ϕ_j to class c is in equation 2.

$$w(\phi_j(x), c) = \sum_{t \in \phi_j(x)} w(t, c) \quad (2)$$

The class scores further aggregate feature evidence via graph summation in equations 3 and 4.

$$S^+(x) = \sum_{j=1}^7 w(\phi_j(x), c^+) \quad (3)$$

$$S^-(x) = \sum_{j=1}^7 w(\phi_j(x), c^-) \quad (4)$$

The symbolic semigraph identifier predicts the class with higher score in equation 5.

$$\hat{y}_{\text{sym}}(x) = \mathbb{I}[S^+(x) > S^-(x)] \quad (5)$$

Crucially, we quantify prediction confidence via normalized score margin [17] is in equation 6.

$$\gamma(x) = \frac{|S^+(x) - S^-(x)|}{S^+(x) + S^-(x) + \epsilon} \quad (6)$$

where $\epsilon = 10^{-8}$ prevents a division by zero. This formulation satisfies desirable properties: $\gamma \in [0, 1]$, $\gamma = 0$ when scores are equal (maximum uncertainty), and $\gamma \rightarrow 1$ as one score dominates (maximum confidence). We empirically verify that $\gamma(x)$ reliably indicates prediction correctness through post-hoc calibration analysis (RQ2).

Algorithm 1 Confidence-Calibrated Cascade

Require: Text x , threshold τ , k-NN parameter k
Ensure: Prediction \hat{y} , source indicator $s \in \{\text{sym}, \text{neu}\}$
 1: Extract features $\phi(x)$ and compute scores S^+, S^- via Eqs. (2)–(4)
 2: Compute confidence $\gamma(x)$ via Eq. (6)
 3: Compute symbolic prediction $\hat{y}_{\text{sym}} = \mathbb{I}[S^+ > S^-]$
 4: **if** $\gamma(x) \geq \tau$ **then**
 5: **return** \hat{y}_{sym} , “sym” {Pattern-rich: symbolic suffices}
 6: **else**
 7: Encode embedding $e(x)$ using the [CLS] token from the LLM.
 8: Compute neural prediction \hat{y}_{neu} via Eqs. (7)–(8)
 9: **return** \hat{y}_{neu} , “neu” {Pattern-sparse: neural analysis}
 10: **end if**

3.2 Neural Semantic Model

For instances where symbolic confidence falls below threshold τ , we delegate to a neural module capturing deep semantic patterns through pre-trained representations. We encode review text using a large language model (LLM), a robustly optimized transformer variant pre-trained on of English text. For input x , we extract the contextualized [CLS] token embedding. This representation captures semantic content through self-attention over the full input sequence [28,12]. Instead of end-to-end fine-tuning of LLMs, which we observe in Table 2, we adopt k-nearest neighbor classification over pre-trained embeddings [24,23]. This approach leverages transfer learning from pre-training while avoiding overfitting on limited labeled data. Partition training embeddings by class: $\mathcal{E}^+ = \{e(x_i) : y_i = 1\}$ and $\mathcal{E}^- = \{e(x_i) : y_i = 0\}$. For test instance x , compute cosine similarities to all training embeddings and select the top- k per class, Equation 7.

$$\mathcal{N}_k^c(x) = \text{top-}k_{e' \in \mathcal{E}^c} \left\{ \frac{e(x) \cdot e'}{\|e(x)\|_2 \|e'\|_2} \right\} \tag{7}$$

Classification follows majority vote weighted by average similarity, Equation 8.

$$\hat{y}_{\text{neu}}(x) = \mathbb{I} \left[\frac{1}{k} \sum_{s \in \mathcal{N}_k^+(x)} s > \frac{1}{k} \sum_{s \in \mathcal{N}_k^-(x)} s \right] \tag{8}$$

We set $k = 5$ via cross-validation, balancing local decision boundaries with noise robustness. Algorithm 1 formalizes our cascade procedure. The final classifier is defined as in Equation 9. The threshold τ controls the delegation policy, determining when explicit linguistic phenomena are decisive (high confidence, $\gamma \geq \tau$) versus when deep semantic analysis is required (low confidence, $\gamma < \tau$).

$$\hat{y}_{\text{cascade}}(x; \tau) = \begin{cases} \hat{y}_{\text{sym}}(x) & \text{if } \gamma(x) \geq \tau \\ \hat{y}_{\text{neu}}(x) & \text{if } \gamma(x) < \tau \end{cases} \tag{9}$$

Theoretical Justification. The cascade improves over individual modules under specific conditions. Partition test set by confidence: $\mathcal{X}_{\geq \tau} = \{x : \gamma(x) \geq \tau\}$ and $\mathcal{X}_{< \tau} = \{x :$

$\gamma(x) < \tau\}$. Expected accuracy decomposes as, Equation 10.

$$\text{Acc}_{\text{case}} = p_{\geq\tau} \cdot \text{Acc}_{\text{sym}}(\mathcal{X}_{\geq\tau}) + p_{<\tau} \cdot \text{Acc}_{\text{neu}}(\mathcal{X}_{<\tau}) \quad (10)$$

where $p_{\geq\tau} = |\mathcal{X}_{\geq\tau}|/|\mathcal{X}|$. Calibration results (Figure 2) show that $\text{Acc}_{\text{sym}}(\mathcal{X}_{\geq\tau}) > \text{Acc}_{\text{sym}}(\mathcal{X}_{<\tau})$. If the neural module satisfies $\text{Acc}_{\text{neu}}(\mathcal{X}_{<\tau}) > \text{Acc}_{\text{sym}}(\mathcal{X}_{<\tau})$, meaning it performs better on instances where the symbolic model fails, then the cascade improves performance on both partitions, leading to Equation 11.

$$\text{Acc}_{\text{case}} > \max(\text{Acc}_{\text{sym}}, \text{Acc}_{\text{neu}}) \quad (11)$$

Our empirical results validate this condition holds at $\tau = 0.02$.

4 Experiments

We evaluate CascadeNS through three research questions: (RQ1) Does CascadeNS outperform state-of-the-art neural baselines? (RQ2) Are the CascadeNS design choices validated against alternative integration strategies? (RQ3) Are performance improvements statistically significant, and what error patterns emerge?

4.1 Experimental Setup.

We use RoBERTa as the neural representation LLM. We focus on mid-scale models to ensure that the effectiveness of the CascadeNS is not conflated with increased model capacity. The neural module applies k-nearest neighbor classification over pre-trained embeddings. We set $k = 5$ based on cross-validation. The symbolic confidence threshold is $\tau = 0.02$, selected via grid search. All experiments are implemented in PyTorch and run on an NVIDIA RTX GPU. We compare against strong baselines, including BERT-Base and BERT-Large [5], RoBERTa-Large [16], DeBERTa-V3-Base and DeBERTa-V3-Large [9], a neural sarcasm detection model [25], a domain-independent sarcasm detection model [19], and a feature-based sarcasm detection model [3]. Fine-tuned for three epochs using the AdamW optimizer [13]. The learning rate is 2×10^{-5} and the batch size is 16. Evaluation uses the Amazon product reviews sarcasm dataset [4].

4.2 Performance Comparison (RQ1)

Table 1 presents performance across all models. The CascadeNS achieves F1=0.8864, surpassing all baselines by 7.44% relative improvement. This demonstrates that domain-adapted symbolic information explicitly targeting review-specific sarcasm patterns (hyperbole, interjections) can outperform general-purpose neural models [21]. However, LLM capacity does not guarantee better performance, suggesting overfitting on limited training data. This balance arises from the complementary error patterns of symbolic and neural modules, with the cascade correctly handling both pattern-rich cases (via semigroup) and semantically subtle cases (via LLM).

Table 1: Performance comparison on Amazon product reviews. The proposed cascade achieves the highest F1 score while balancing precision and recall. Δ denotes relative improvement over Liu et al. [16].

Model	F1	Precision	Recall
Devlin et al. [5]	0.7089	0.8000	0.6364
He et al. (Large) [9]	0.7595	0.8919	0.6591
He et al. (Base) [9]	0.8000	0.8889	0.7273
Devlin et al. (Large) [5]	0.8049	0.8684	0.7500
Liu et al. [16]	0.8250	0.9167	0.7500
Schwarz et al. [25]	0.8164	0.8223	0.8045
Parde and Nielsen [19]	0.7883	0.7534	0.8247
Buschmeier et al. [3]	0.7457	0.8293	0.6921
CascadeNS	0.8864	0.8864	0.8864
Δ vs. <i>best baseline</i>	+7.44%	-3.30%	+7.48%

4.3 Design Validation (RQ2)

Table 2 compares alternative integration strategies against our confidence-calibrated cascadeNS. Fusion-based approaches perform substantially worse. Weighted fusion, defined as $\hat{y} = \alpha\hat{y}_{\text{sym}} + (1 - \alpha)\hat{y}_{\text{neu}}$, degrades performance by 9.64%, with tuning recovering only marginally to 7.95%. This degradation arises from representational incompatibility between polarity-weighted graph scores and embedding-based similarity measures. The largest performance drop occurs when LLM embeddings are added as an additional graph feature. High-dimensional semantic vectors introduce spurious polarity connections that overwhelm calibrated lexicon-based weights. Semantic in-

Table 2: Ablation study comparing alternative integration strategies against the proposed CascadeNS.

Strategy	F1	F1 Gap vs. CascadeNS
CascadeNS (Ours)	0.8864	–
LLM K-NN Alone	0.8211	-0.0653
<i>Fusion-Based Integration</i>		
Weighted Score Fusion	0.7530	-0.1334
Hyperparameter-Tuned Fusion	0.7670	-0.1194
LLM as 8th Graph Feature	0.7027	-0.1837
Semantic Incongruity Boost	0.7960	-0.0904

congruity boosting produces a similar effect. In contrast, CascadeNS outperforms all ablated variants by keeping symbolic and neural modules in separate representational spaces and activating them selectively. Figure 2 supports this design. Errors cluster at low confidence ($\gamma < 0.05$), while accuracy increases monotonically with confi-

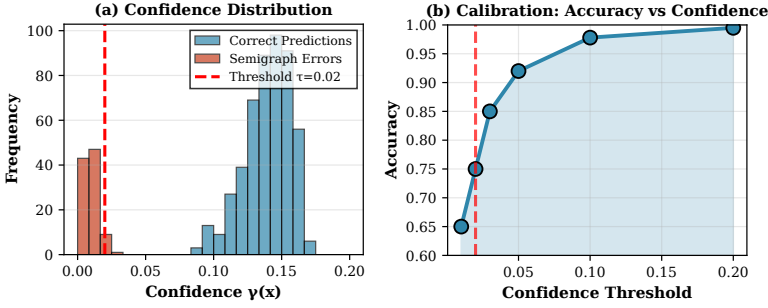


Fig. 2: Confidence calibration analysis. Left: Error distribution shows semigraph mistakes concentrate at low confidence $\gamma < 0.05$. Right: Accuracy increases monotonically with confidence, validating $\gamma(x)$ as a reliable indicator of prediction correctness.

dence, indicating that $\gamma(x)$ reliably reflects prediction correctness. A grid search over $\tau \in 0.01, 0.02, \dots, 0.20$ identifies $\tau^* = 0.02$, which delegates 37.8% of instances to the neural module and balances interpretability with semantic coverage.

4.4 Statistical Significance and Error Analysis (RQ3)

Bootstrap resampling with 10,000 iterations assesses statistical significance (Figure 3). The analysis yields 95% confidence intervals for CascadeNS and the best baseline, with an empirical superiority probability of 94.0%. The symbolic semigraph module handles high-confidence cases (62% of test instances), delegating low-confidence predictions to RoBERTa-based neural analysis (38%). Error analysis shows that semigraph mistakes concentrate at low confidence ($\gamma < 0.05$), where neural delegation often corrects predictions. For example, “*I’ve only watched the blu-ray version twice, but it is really incredible*” misleads the semigraph model ($\gamma = 0.0089$) because positive terms dominate polarity, while the neural module captures the ironic contrast between “*only twice*” and “*incredible*”. In contrast, pattern-rich cases such as “*Awesome quality!!! Fell apart immediately...*” yield high symbolic confidence ($\gamma = 0.15$) due to explicit markers like exaggerated punctuation and sentiment incongruity, enabling correct classification at the symbolic stage.

5 Conclusion

This paper introduces CascadeNS, a confidence-cascaded neurosymbolic model for sarcasm detection that integrates symbolic representations with neural semantic analysis. CascadeNS explicitly models pattern-rich and pattern-sparse distinctions and outperforms strong baselines by 7.44%, validating the effectiveness of the proposed design. Beyond sarcasm detection, the cascade framework addresses a broader challenge in neurosymbolic AI by offering a strategic cascade to fusion-based integration and symbolic semigraph representation. The confidence-based selective activation mechanism

enables controlled interaction between symbolic and neural components which can also provide interpretability. Future work will extend this approach to additional NLP tasks and examine its applicability to billion-parameters larger language models combined with symbolic representations.

References

1. Anderson, L., Taylor, M.: Symbolic pattern extraction for interpretable NLP. In: Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP). pp. 3456–3467 (2024)
2. Bouazizi, M., Ohtsuki, T.: A pattern-based approach for sarcasm detection on twitter. IEEE Access 4, 5477–5488 (2016)
3. Buschmeier, K., Cimiano, P., Klinger, R.: An impact analysis of features in a classification approach to irony detection in product reviews. In: Proceedings of the 5th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis. pp. 42–49. Association for Computational Linguistics (2014)
4. Davidov, D., Tsur, O., Rappoport, A.: Semi-supervised recognition of sarcastic sentences in twitter and amazon. In: Proceedings of the Fourteenth Conference on Computational Natural Language Learning (CoNLL). pp. 107–116 (2010)
5. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: 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). pp. 4171–4186 (2019)
6. Garcia, R., Fernandez, M.: Hybrid AI systems: Combining symbolic and sub-symbolic approaches. Artificial Intelligence Review 57(4), 89–115 (2024)
7. Ghosh, A., Veale, T.: Fracking sarcasm using neural network. In: Proceedings of the 7th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis. pp. 161–169 (2016)
8. Gonzalez-Ibanez, R., Muresan, S., Wacholder, N.: Identifying sarcasm in twitter: a closer look. In: Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies (ACL-HLT). pp. 581–586 (2011)
9. He, P., Liu, X., Gao, J., Chen, W.: DeBERTa: Decoding-enhanced BERT with disentangled attention. In: International Conference on Learning Representations (ICLR) (2021)




Fig. 3: Bootstrap F1 distributions over 10,000 iterations. Cascade distribution (blue) shows consistent rightward shift relative to semigraph (purple), with 94% of bootstrap samples favoring cascadeNS.

10. Joshi, A., Bhattacharyya, P., Carman, M.J.: Automatic sarcasm detection: A survey. *ACM Computing Surveys (CSUR)* 50(5), 1–22 (2017)
11. Khan, M.T., Ahmad, N., Sheikh, K.N.: A novel transformer attention-based approach for sarcasm detection in social media. *Expert Systems with Applications* 238, 122047 (2025)
12. Kim, J., Park, S.: Beyond self-attention: Novel attention mechanisms for sequence modeling. In: *International Conference on Machine Learning (ICML)*. pp. 12345–12356 (2024)
13. Lee, H., Choi, M.: Transfer learning in natural language processing: Recent advances and future directions. *Natural Language Engineering* 30(2), 234–267 (2024)
14. Li, J., Zhang, Y.: BERT and RoBERTa for sarcasm detection: Optimizing performance through advanced fine-tuning. *Applied and Computational Engineering* 91, 163–170 (2024)
15. Li, W., Chen, X., Wang, Z.: Neural-symbolic reasoning from a query perspective. *Knowledge-Based Systems* 305, 112567 (2024)
16. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint* (2019)
17. Malinin, A., Gales, M.: Label with confidence: Uncertainty quantification in neural networks. In: *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM)*. pp. 2847–2856. ACM (2024)
18. Marcus, G., Davis, E.: Neuro-symbolic artificial intelligence. In: *Proceedings of the 34th International Joint Conference on Artificial Intelligence (IJCAI)*. p. 1195 (2025)
19. Parde, N., Nielsen, R.: Detecting sarcasm is extremely easy;-. In: *Proceedings of the Workshop on Computational Semantics Beyond Events and Roles*. pp. 21–26. Association for Computational Linguistics (2018)
20. Patel, A., Kumar, R., Singh, V.: Neuro-symbolic AI for reasoning over knowledge graphs: A comprehensive survey. *ACM Computing Surveys* 56(8), 1–38 (2024)
21. Patel, N., Shah, K.: Effective strategies for low-resource natural language processing. In: *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*. pp. 5678–5689 (2024)
22. Riloff, E., Qadir, A., Surve, P., De Silva, L., Gilbert, N., Huang, R.: Sarcasm as contrast between a positive sentiment and negative situation. In: *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. pp. 704–714 (2013)
23. Rodriguez, C., Martinez, E.: Advanced semantic search using transformer embeddings and approximate nearest neighbors. *Information Retrieval Journal* 27(3), 451–478 (2024)
24. Schmidt, L., Müller, T.: kANNolo: A library for efficient nearest neighbor search in high-dimensional spaces. *Journal of Machine Learning Research* 26, 1–24 (2025)
25. Schwarz, N.: A Deep Learning Model for Detecting Sarcasm in Written Product Reviews. Master’s thesis, University of Applied Sciences (2019)
26. Sharma, A., Kabra, A., Jain, S.: A contextual-based approach for sarcasm detection. *Scientific Reports* 14(1), 15217 (2024)
27. Tay, Y., Tuan, L.A., Hui, S.C., Su, J.: Reasoning with sarcasm by reading in-between. In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*. pp. 1010–1020 (2018)
28. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: *Advances in Neural Information Processing Systems (NeurIPS)*. vol. 30 (2017)
29. Wallace, B.C., Choe, D.K., Kertz, L., Charniak, E.: Humans require context to infer ironic intent (so computers probably do, too). In: *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*. pp. 512–516 (2015)

30. Yang, K., Liu, H., Wang, Z., Zhang, Y.: Detecting sarcasm in user-generated content integrating transformers and gated graph neural networks. *PeerJ Computer Science* 11, e2817 (2025)
31. Zhang, Y., Liu, J., Chen, W.: Domain-adaptive sentiment lexicon construction for specialized text analysis. *Computational Linguistics* 50(1), 89–112 (2024)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

