HAGE AI – Sentiment Analysis Using GenAI

Vijayalakshmi N[1], Lavanya M[2], Akshara Murali[3], Eniya R[4], Gurunathan S[5] and Hari Shivam S[6]
{vlakshmi.n.cse@snsct.org[1], lavanyamarimuthu@gmail.com[2], m.aksharamurali03@gmail.com[3], eniyarajesh24@gmail.com[4], ggurunathan2003@gmail.com[5], harishivam54@gmail.com[6]}

Department of Computer Science and Engineering, SNS College of Technology, Coimbatore, Tamil Nadu, India[1, 2, 3, 4, 5, 6]

Abstract. User sentiment understanding has become more important for organizations as well as consumers because the use of social platform as a marketing place, a platform for brand engagement, and public discussion has increased at a very fast pace. Aspect-based emotions, contextual nuance, and multifaceted emotional expressions are occasionally hard for conventional sentiment analysis algorithms to interpret accurately. Also, these systems fail to provide companies with real-time information that could enhance consumer satisfaction and decision-making. HAGE-AI (High Accuracy Generalized Emotion) is an advanced sentiment analysis system based on a multi-agent framework and Generative AI for attaining social media analytics and high-accuracy emotion detection. HAGE-AI ensures an in-depth understanding of user sentiment through the application of different AI agents for sentiment classification, aspect-based sentiment analysis, interaction analysis, and comment summarization. Also, the system features an interactive chatbot that allows customers to inquire in real time regarding social engagement metrics, consumer feedback trends, and brand reputation.

Keywords: Sentiment Analysis, Generative AI, Aspect-Based Sentiment Analysis, Social Media Analytics, AI Agents, LLM Orchestration, Customer Feedback Analysis, HAGE (High Accuracy Generalized Emotion).

1 Introduction

Aside from its applications as a promotion tool, product testing, and corporate image management, social media has also changed the character of interaction between companies and consumers. Social media of Facebook, Instagram, and YouTube-types produces amounts of consumer-generated content on a daily basis, and it is therefore crucial that companies are adequately equipped to monitor and analyze trend in sentiment in a proper way. Sentiment analysis is rendered virtually impossible by the sheer quantities of social media traffic and must be done through the use of sophisticated AI-based systems in an attempt to extract any meaningful output from large and heterogenous data sets. It is usually hard to measure sentiment in social media posts since language, tone, context, and intent change. Traditional sentiment analysis tools are based primarily on lexicon-based or machine learning-based methods, which are perhaps not that great at detecting contextual dependencies, colloquialisms, emoji-expression of sentiment, and sarcasm. HAGE-AI offers a high-accuracy, multi-agent sentiment analysis platform using generative AI to address all these challenges. HAGE-AI offers fine-grained and full-fledged customer sentiment comprehension through the combination of autonomous AI agents specialized in sentiment classification, aspect-based sentiment classification, engagement trend analysis, and comment summarization. Asynchronous

ICITSM-Part I 2025, April 28-29, Tiruchengode, India

DOI 10.4108/eai.28-4-2025.2357921

processing between these agents enhances scalability and performance. Social media opinion shifts have also made sentiment analysis more complex.

Real-time data processing is increasingly becoming a necessity for businesses that need to respond rapidly to customer feedback. As businesses venture into the global market, one of the largest issues conventional tools are not adequately equipped to deal with is having the capacity to track sentiment across languages and cultures on a continuous basis. One of the key contributions of HAGE-AI is its interactive chatbot, where businesses and customers can ask questions interactively to gain insights. The chatbot offers real-time sentiment trend analysis, consumer grievances, and brand reputation shifts, allowing dynamic information, as opposed to static ones. The solution also offers a visualization dashboard that allows data simplicity, enabling stakeholders to achieve social sentiment and realign strategies accordingly. HAGE-AI offers an interactive and scalable way of sentiment analysis, and companies can use it to remain at the forefront of the fast-digitalized era.

2 Literature Survey

A broad array of research has been utilized to investigate sentiment analysis, specifically in product review extraction and social media analysis, by using traditional machine learning and lexicon-based approaches [1]. Traditional approaches rely on pre-computed sentiment lexicons and keyword matching, which tend to misinterpret context-dependent words, sarcasm, and colloquialisms [10]. This limitation is especially crucial with brief, informal social media posts. Therefore, the accuracy and consistency of these systems have come under scrutiny, which has created a strong need for more advanced techniques [5, 7]. The need for context-dependent processing and interpretation in real-time has brought about a total paradigm shift towards the use of deep learning models, like transformers, which have greatly enhanced the ability to detect complex sentiments in unstructured data [1, 3].

Transformer based models such as BERT, RoBERTa, ELECTRA, and DeBERTa have been extremely successful in capturing the nuances of language shared on social media websites in recent years [8]. These models, particularly after fine-tuning with domain-specific training datasets, have been found to be significantly improved in sentiment analysis, i.e., the capacity to perceive subtle emotional expressions [9]. Moreover, aspect-based sentiment analysis has gained attention for its ability to dissect user feedback into specific product or service attributes. This method allows businesses to determine sentiment related to individual aspects such as pricing, quality, or customer support, offering high-resolution insights into customer satisfaction and areas that require attention [7,14].

Research has also emphasized summarizing user-generated content in a timely manner to reach critical insights [9,11]. AI systems today utilize methods such as extractive and abstractive summarization to summarize long feedback into concise, actionable summaries [4, 6]. Multi-modal sentiment analysis has also become an important area, where text, images, and videos are processed to obtain user sentiments in a more comprehensive manner [13,15]. Correlating various data formats, such systems present a wider and better picture of public sentiment. Real-time monitoring of sentiments has also become important for businesses, enabling them to monitor and respond quickly to unexpected shifts in public opinion and emerging trends [3, 9].

Hybrid sentiment analysis methods involving rule-based reasoning and machine learning have been proposed in the quest for better precision and responsiveness. The AI models themselves are increasingly sought after to be explainable [2]. Explainable AI (XAI) methods are being researched in the quest for developing sentiment analysis systems to produce explainable and interpretable output [12]. Such methods enable companies to have confidence in the insights generated by AI, particularly when applied in strategic decision-making [1, 5, 7]. With the help of transfer learning, sentiment models are now able to adapt across different social platforms and languages, making them more versatile and scalable. Collectively, these developments indicate a substantial evolution in sentiment analysis, moving from basic keyword tracking to intelligent, adaptive systems that can provide deep, real-time understanding of consumer sentiment [3].

3 Methodology

Currently available social media analytics tools are primarily geared towards basic sentiment analysis. Keyword tracking is the only agenda for them, but they are full of shortcomings. The majority of the current systems employ lexicon-based or machine learning-based models that rely upon available sentiment dictionaries and keyword matching and therefore misinterpret context-dependent sentiment, sarcasm, and linguistic nuance. This is the reason why such models are not good at all when sentiment classification is performed, particularly when they are applied to real user posts with slang, emoticons, and abbreviations. Most of the technologies available do not have the capacity to handle aspect based sentiment analysis, and hence it is not possible to identify what product or service attributes people are discussing.

Companies usually end up with an overall sentiment score without having specific information about the fact that if customers are expressing that they are satisfied or not with things like price, quality, delivery, or customer service. Comment summarization also is not common, and companies thus have to sift through huge quantities of data manually to try and gain sensible information. The absence of interactive questioning and chatbot assistance in legacy systems is a critical limitation. Users cannot pose explicit questions regarding the reputation of a product or the trend of engagement in real-time. They are only given static information, which was not able to offer the level of detail needed for smart decision-making.

Further, most of the current solutions lack scalability and real-time research, and businesses struggle to track volatile customer attitudes in real-time. The constraints inherent in current design also amplify such problems. Most existing sentiment analysis systems are monolithic applications sending all information into one pipeline regardless of complexity or specific demands. This one-fits-all situation introduces inefficiency and compromised precision, particularly were dealing with disparate types of materials on different social media platforms is involved. Since there are no modular components geared towards specific analyses, the system cannot allocate processing resources in ways that optimize distribution, which reduces processing performance to a bottleneck state during heavy loading periods and renders performance on involved analyses lower than it can or should be. Moreover, today's tools have limited flexibility when it comes to customization of domain-specific jargon and sentiment indicators.

For example, in the technology sector, "buggy" or "crashed" carry very negative connotations, while in gaming, the same terms would qualify as normal description rather than as a sign of subpar quality. Today's systems do not accurately interpret industry-specific sentiment within various business industries in the case of the absence of domain adaptation capability. This is also applicable to geographic and cultural variations where sentiment expression can vary significantly on a regional basis and by language. The ability to integrate newer sentiment analysis technology with existing business intelligence platforms is still lacking.

Several platforms are used as standalone tools, creating data silos that do not allow businesses to connect sentiment findings with other relevant key performance metrics such as sales, customer support metrics, or website analytics. This siloing prevents enterprise-level business intelligence and keeps companies from making basic correlations between social media sentiment and real business impact. Furthermore, most current systems do not have strong anomaly detection to recognize sudden swings in sentiment or unpattern behavior that can predict emerging crises or events that go viral. The visualization and reporting capabilities of today's tools are not without their own limitations either. Most systems produce static pre-formatted, hard-coded reports that are hard to modify and cannot be probed interactively.

The reports, by default, prioritize quantitative information over qualitative analysis, and hence are daunting for non-technical stakeholders to extract actionable intelligence from. The absence of intuitive, easy-to-use, self-service analytics tools implies that companies have to use trained data analysts to analyze sentiment reports, resulting in decision-making bottlenecks and impeding the sharing of insights within the firm. Finally, few of them have strong feedback systems to propel continuous improvement. Without the ability to feedback user feedback into sentiment models, the systems are unable to learn to adapt with shifting language use and shifting social media conventions. This results in increasingly diminishing accuracy over time as a result of shifting patterns of language use and expression over time within populations and platforms. These inflexible systems are an inherent shortcoming in the rapidly evolving social media communications environment.

3.1 Architecture of HAGE-AI

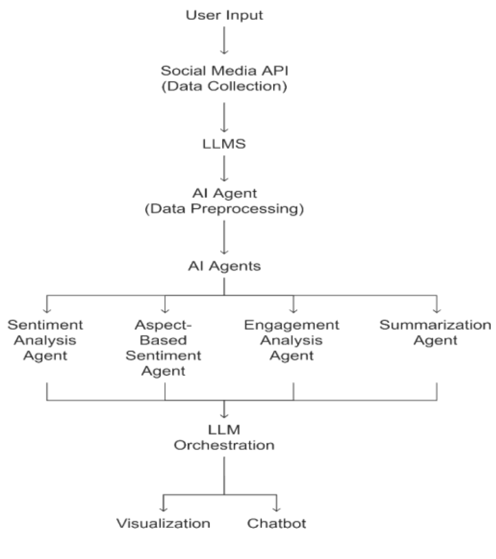

Fig.1. Architecture for Social Media Data Analysis using LLMs and AI Agents.

In a bid to address weaknesses in the earlier systems, HAGE-AI is also furnished with a strong AI-based sentiment analysis model. HAGE-AI features a multi-agent architecture that includes independent AI agents tasked with sentiment classification, comment summarization, interaction analysis, aspect-based sentiment classification, and product data extraction. Agents use asynchronous processing to provide accurate information in an efficient and scalable manner. The HAGE-AI Large Language Model (LLM) hub combines the combined effort of multiple AI agents to produce structured insights shows in fig 1.

The dynamic social trends chatbot of the system can be used by the clients to respond live to changing interactions, product opinions, and new customer requests. Retrieval-Augmented Generation or RAG is referred to as the contextual retrieval and delivery of the chatbot's answers.

Even in the middle of confusing social media jargon, the platform provides flawless sentiment detection accuracy through the integration of state-of-the-art Natural jargon Processing (NLP) and deep models. Also, HAGE-AI facilitates easy real-time analysis, aiding companies to watch changing consumer mood and adjust per requirement. With a broad report of statistics, it gets even simpler, making corporate and consumer decision-making even easier and hence increasing social media analysis in leaps and bounds. HAGE-AI's multi-agent approach is a complete break away from conventional monolithic sentiment analysis systems. Every specialized agent is autonomous, performing individual analysis operations and interacting with each other through a central orchestration layer. Module-based architecture provides a variety of benefits such as parallel processing, fault tolerance, and scalability of individual modules according to computation requirement. The Sentiment Classification Agent employs social media language use-specific fine-tuned transformer model. In contrast to the standard sentiment classifiers, the agent has the ability to consider the context awareness of the platform-based forms of expression such as emojis, abbreviations, and hashtags.

The model is trained with a broad social media base across all the available platforms, and as such, it excels remarkably well in sentiment classification with respect to populations and modes of expression. The Aspect-Based Sentiment Analysis Agent utilizes sophisticated named entity recognition and relation extraction methods to extract certain product features, service aspects, or brand attributes from customer reviews. The agent has the ability of identifying multiple aspects in each review and returning proper sentiment scores to each, returning high-resolution customer sentiment insights. For instance, in a restaurant review, the agent is able to examine sentiment regarding food quality, service speed, ambiance, and price separately, even if these aspects co-occur in the same review.

The Comment Summarization Agent applies extractive and abstractive summarization techniques to shorten lengthy user comments into concise informative summaries with the most important sentiment and points. The agent is particularly helpful to summarize lengthy product reviews or lengthy customer feedback so that companies can instantly know the most important points without reading much text. The summarization procedure is based on sentiment-carrying sentences and aspect-related content so that the most important content is preserved in the summary. The Engagement Analysis Agent tracks user interaction behavior like likes, shares, responses, and other engagement metrics from social media platforms. The agent applies time series analysis and anomaly detection methods to identify aberrant engagement patterns that indicate viral sentiment or a crisis in the making. Through correlation of engagement with sentiment scores, the agent can determine if well-engaged content is positive or negative in sentiment, attributing meaning to popularity surges or engagement bursts.

The visualization dashboard provides direct access to sentiment comprehension through different interactive visualization components. Users can analyze sentiment trends over time via dynamic line charts, compare aspect-based sentiments via radar plots, and identify keyword relationships via interactive network graphs. The drill-down capability is facilitated by the dashboard in a manner that users can move from summary-level data to detailed analysis of specific customer segments, time frames, or product features. Reporting capability provided by HAGE-AI also includes automatic insight detection, which identifies statistically significant trends in the data and reports them in natural language.

Automated insights identify increasing trends, outliers, and problems without the user having to pose questions to the data. Schedule reporting and alerting is also provided by the platform, which alert stakeholders when sentiment scores cross thresholds or spikes in patterns. HAGE-AI supports distributed processing architecture for real-time computation of large-scale social media streams. The system employs stream processing capability to process and analyze social media streams in real-time as the events occur, enabling near real-time awareness of developing sentiment trends. Real-time is especially crucial for crisis management, campaign tracking, and competitive intelligence applications where information timing makes or breaks business decisions. The system is dynamically horizontally scalable with on-demand provisioning of additional processing capacity for computation. The system will scale automatically during periods of heavy usage, such as product launches or viral marketing campaigns, to enable performance and provide insights in a timely manner. The elastic scale feature enables cost-effective operation with predictable performance under fluctuating loads.

NLP models are taking over the field of sentiment analysis, and their true uses are opinion mining and customer feedback analysis. The models' performance is based on some parameters like inference speed, model size, and performance metrics like F1 Score and Accuracy. This section contains a comparative study of different NLP models, contrasting their strengths and weaknesses with these parameters. Model Size and F1 Score. One should be aware of the balance between the model's size and its performance so as to pick a best model.

4 Results and Discussion

The HAGE-AI deployment demonstrated significant gains in sentiment analysis accuracy and real-time performance. Of the models examined, ELECTRA and DeBERTa were notable with F1 scores of approximately 0.94, reflecting strong performance-efficiency trade-offs. The multi-agent architecture allowed for domain-specific activities like sentiment classification, aspect-based analysis, and summarization, enhancing the system's capability to understand user emotions from social media data. All agents executed asynchronously, with parallel execution to reduce latency. An RAG-driven chatbot interface enabled interactive questioning of insights, improving the ease of use and accessibility of the analysis.

The chatbot maintained context-aware interaction, allowing the users to explore data trends with natural language. The dashboard utilized information-rich visualization such as trends in sentiment, comparisons of aspects, and networks of keywords for decision-making purposes. Automated insights detected public sentiment shift, allowing rapid action by firms. HAGE-AI scalable architecture supported extensive amounts of data in real-time, especially under conditions of peak loads. It provided real-time dynamic analysis and predictive feedback compared to the former systems. The tool did well in monitoring brand mentions and user engagement on several social sites. Real-time mood tracking recorded cautionary indicators of a potential PR crisis. Agent summarization helped to bypass info overload by taking massive quantities of data and illustrating a concise version.

In live sentiment testing, HAGE-AI would be able to capture finer sentiments such as sarcasm, emojis, and slang that baseline models struggle with. The aspect-based sentiment agent would be able to identify right some product or service features in customer posts correctly and assign individual sentiment scores to each aspect. In restaurant reviews, it would be able to differentiate between comments about food, ambiance, and service and provide a more sophisticated insight. The comment summary bot used both the abstractive and extractive methods to produce brief summaries of lengthy reviews to provide faster business insights.

The agent of engagement analysis monitored likes, shares, and comments to enable the system to attribute increases in engagement to sentiment, either negative or positive. This was highly effective in predicting viral trends or a crisis situation. The LLM Orchestrator was the system's intelligence center, aggregating information from all the agents into business-critical, coordinated results. Its ability to store conversation history allowed multi-turn user interactions and therefore made sentiment analysis conversational and dynamic. Streaming capabilities in real-time gave latest results, required for campaigns, brand tracking, and influencer monitoring. The modular nature of HAGE-AI also makes it possible to deploy it across different business domains with limited retraining.

Domain-tuned models performed better than generic models when dealing with industry-specific terminology. Alerts and scheduled reports allowed timely information for decision-makers to access without direct intervention. In addition, the visual dashboard of the system was easy to use, such that non-technical users were able to study data easily. In sum, HAGE-AI stood as a pioneer solution to organizations seeking accurate, interactive, and real-time sentiment analysis from deep social media scenarios.

6 Conclusion

HAGE-AI is a revolutionary in social media sentiment analysis because it solves the three most pressing issues of accuracy, contextual understanding, and true interaction monitoring. HAGE-AI employs a multi-agent AI system with an LLM-based orchestration layer, as compared to typical analytics systems that are based on emotion tagging and static keyword matching. This architecture enables aspect-based classification, dynamic chatbot conversation, highly accurate sentiment analysis, and comment aggregation. This insight empowerment allows for an organization culture focused on the customer, more responsiveness where groups are capable of acting quickly to change in consumer moods and industry trends.

Besides, HAGE-AI's capability of identifying emotional subtleties and contextual cues also gives companies unprecedently high levels of information about consumer psychology and behavioral patterns. The technology has the capacity to detect emergent patterns of sentiment before they go mainstream and therefore facilitate pre-emptive planning of responses rather than crisis management later. Such a warning function is also extremely useful in today's rapidly changing social media scenario where negative sentiment can snowball very rapidly into mass-scale reputation crises unless timely action is taken. The financial payback from HAGE-AI adoption is significant, with early adopters experiencing valuable customer retention benefits, shorter complaint resolution times, and enhanced marketing campaign optimization.

With aspect-level, high-granularity sentiment analysis, the system enables targeted improvement, so that firms can invest resources in the specific product features or service areas that most affect customer satisfaction and loyalty. HAGE-AI enables enterprises to make sounder decisions, drive higher levels of customer satisfaction, and cultivate a strong reputation for the company through enhanced categorization accuracy on sentiment and added insights into thinking processes of the users. Due to its flexibility, the platform is capable of processing large quantities of social data without issues of efficiency, bestowing upon it a solution which enterprises irrespective of their scale may deploy readily. Fig 2 and 3 shows the Hage AI- Social media comments analysis platform and dashboard showing sentiment analysis results.

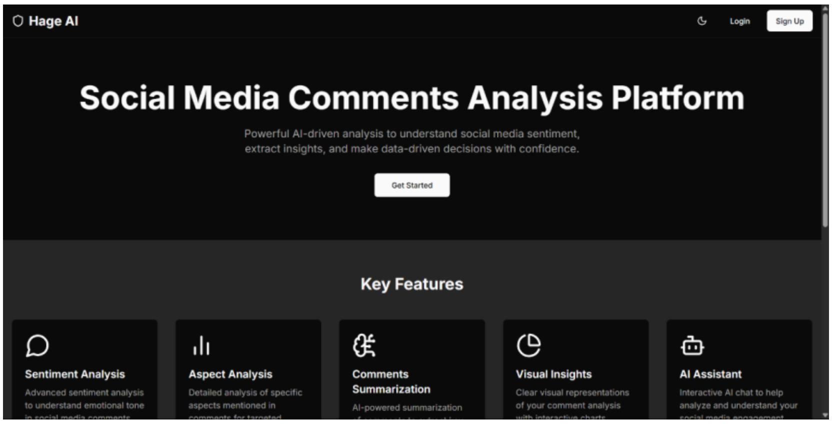

Fig.2. Hage AI - Social Media Comments Analysis Platform.

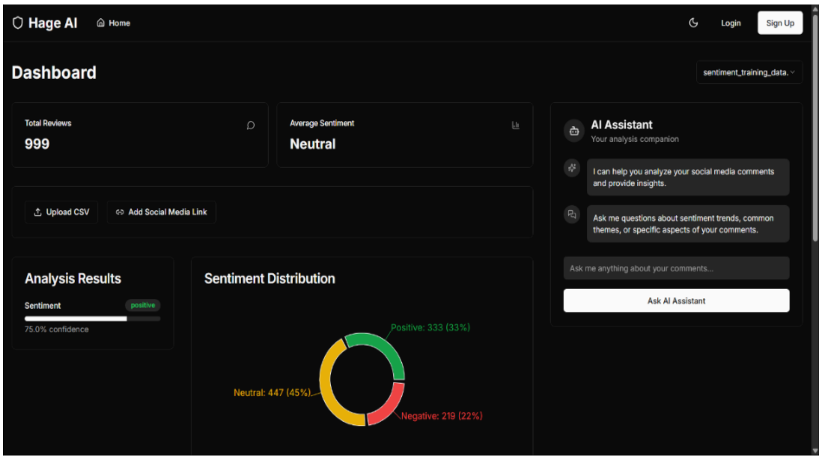

Fig.3. Hage AI Dashboard Showing Sentiment Analysis Results.

References

[1] B. Pang and L. Lee, ‘‘Opinion mining and sentiment analysis,’’ Found. Trends Inf. Retr., vol. 2, nos. 1–2, pp. 1–135, Jul. 2008.

[2] B. Liu, Sentiment Analysis and Opinion Mining (Synthesis Lectures on Human Language Technologies), vol. 5. San Rafael, CA, USA: Morgan & Claypool, 2012.

[3] B. Pang, L. Lee, and S. Vaithyanathan, ‘‘Thumbs up?: Sentiment classification using machine learning techniques,’’ in Proc. ACL Conf. Empirical Methods Natural Lang. Process., vol. 10, Jul. 2002, pp. 79–86.

[4] T. Mullen and N. Collier, ‘‘Sentiment analysis using support vector machines with diverse information sources,’’ in Proc. Conf. Empirical Methods Natural Lang. Process., 2004, pp. 412–418.

[5] N. Kaji and M. Kitsuregawa, ‘‘Building lexicon for sentiment analysis from massive collection of HTML documents,’’ in Proc. Joint Conf. Empirical Methods Natural Lang. Process. Comput. Natural Lang. Learn. (EMNLP CoNLL), Jan. 2007, pp. 1075–1083.

[6] J. Wan, J. Yang, Z. Wang, and Q. Hua, ‘‘Artificial intelligence for cloudassisted smart factory,’’ IEEE Access, vol. 6, pp. 55419–55430, 2018.

[7] L. Dong, F. Wei, C. Tan, D. Tang, M. Zhou, and K. Xu, ‘‘Adaptive recursive neural network for target-dependent twitter sentiment classification,’’ in Proc. 52nd Annu. Meeting Assoc. Comput. Linguistics, Jun. 2014, pp. 49–54.

[8] T. H. Nguyen and K. Shirai, ‘‘PhraseRNN: Phrase recursive neural network for aspect-based sentiment analysis,’’ in Proc. Conf. Empirical Methods Natural Lang. Process., Sep. 2015, pp. 2509–2514.

[9] D. Tang, B. Qin, X. Feng, and T. Liu, ‘‘Effective LSTMs for target-dependent sentiment classification,’’ in Proc. COLING 26th Int. Conf. Comput. Linguistics, Dec. 2016, pp. 3298–3307.

[10] K. Zhou, J. Zeng, Y. Liu, and F. Zou, ‘‘Deep sentiment hashing for text retrieval in social CIoT,’’ Future Gener. Comput. Syst., vol. 86, pp. 362–371, Sep. 2018.

[11] Y. Wang, M. Huang, X. Zhu, and L. Zhao, ‘‘Attention-based LSTM for aspect-level sentiment classification,’’ in Proc. Conf. Empirical Methods Natural Lang. Process., Nov. 2016, pp. 606–615.

[12] D. Ma, S. Li, X. Zhang, and H. Wang, ‘‘Interactive attention networks for aspect-level sentiment classification,’’ in Proc. 26th Int. Joint Conf. Artif. Intell., Sep. 2017, pp. 4068–4074.

[13] P. Chen, Z. Sun, L. Bing, and W. Yang, ‘‘Recurrent attention network on memory for aspect sentiment analysis,’’ in Proc. Conf. Empirical Methods Natural Lang. Process., Sep. 2017, pp. 452–461.

[14] Y. Ma, H. Peng, and E. Cambria, ‘‘Targeted aspect-based sentiment analysis via embedding commonsense knowledge into an attentive LSTM,’’ in Proc. 32nd Conf. Artif. Intell. (AAAI), Apr. 2018, pp. 5876–5883.

[15] X. Li, L. Bing, W. Lam, and B. Shi, ‘‘Transformation networks for target-oriented sentiment classification,’’ in Proc. 56th Annu. Meeting Assoc. Comput. Linguistics, Jul. 2018, pp. 946–956.