iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL
FOR RESEARCH
IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 Issue: I Month of publication: January 2026

DOI: https://doi.org/10.22214/ijraset.2026.76945

A Review of Multimodal and Retrieval-Augmented Artificial Intelligence Approaches for Social Media-Based Travel Discovery

Aaron Anil Zachariah[1], Amith Krishna K[2], Athira Santhosh[3], Bharath M Nanda[4], Mrithu A S[5]
[1, 2, 3]Student, Department of Computer Science and Engineering, Vidya Academy of Science and Technology, Thrissur, India
[4]Assistant Professor, Department of Computer Science and Engineering, Vidya Academy of Science and Technology, Thrissur, India

***Abstract:** In recent years, social media platforms have evolved into primary sources of travel inspiration, with travelers increasingly relying on short-form video content to discover unique, ”off-the-beaten-path” locations. Despite their popularity, these videos often lack the structured and reliable information necessary for practical travel planning. Critical details such as precise locations, accessibility, and logistical guidance are frequently fragmented across captions and comments or omitted entirely, creating a significant gap between visual discovery and actionable decision-making. This paper reviews existing research in social media mining, multimodal information extraction, and Natural Language Processing (NLP) to determine how Artificial Intelligence can bridge this gap. Building on insights from the reviewed literature, the study presents the conceptual design of ”ReelScout,” an AI-driven platform that integrates Computer Vision and NLP to analyze social media reels for identifying hidden Points of Interest (POIs). By synthesizing multimodal cues from visual content, audio narration, and textual metadata, the proposed framework aims to organize unstructured social media data into meaningful travel knowledge. Finally, this review highlights current methodological trends, limitations, and future research directions for the development of intelligent, social-media-driven travel discovery systems.*

***Keywords:** Geo-AI, Social Media Mining, Multi-Modal Analysis, RAG, Hidden Gems, Sustainable Tourism.*

I. INTRODUCTION

The rapid expansion of social media platforms has significantly transformed how information is created, shared, and consumed. Short-form video content such as reels and shorts has emerged as a dominant medium for communication, enabling users to share experiences, opinions, and visual narratives in an engaging manner. In recent years, these platforms have gained substantial importance in domains such as tourism, urban exploration, and lifestyle discovery, where users frequently showcase visually appealing yet lesser-known locations [3], [6]. While social media reels provide strong visual inspiration, they often lack structured, reliable, and actionable information required for practical decision-making. Details such as precise location, accessibility, safety considerations, and travel logistics are usually absent or scattered across comments and captions, making manual exploration inefficient and unreliable [4], [17]. As a result, users face difficulties in translating visual inspiration into concrete travel plans.

The increasing volume of unstructured multimedia data generated on social media platforms presents significant challenges for information extraction and knowledge discovery. Traditional data analysis techniques struggle to process the heterogeneous nature of text, images, audio, and video content at scale [7], [16]. Consequently, there is a growing need for intelligent methods capable of automatically analyzing such data to extract meaningful insights, identify emerging trends, and support informed decision-making.

Recent advances in Artificial Intelligence (AI), particularly in machine learning and deep learning, have enabled effective analysis of social media content. Techniques in Natural Language Processing (NLP) have been widely applied for tasks such as entity recognition, sentiment analysis, and topic modeling, while Computer Vision (CV) approaches have demonstrated strong performance in visual landmark detection and scene understanding [10], [11], [15]. Furthermore, multimodal learning approaches that integrate textual, visual, and audio information have been shown to outperform single-modal methods by capturing richer contextual representations [3], [16], [30]. The emergence of large language models and retrieval-augmented generation (RAG) techniques has further enhanced the ability of AI systems to generate accurate, context-aware responses by combining external knowledge sources with generative models [1], [2]. These approaches are particularly valuable for applications that require factual correctness and interpretability, such as recommendation systems and decision-support platforms [26], [31].

Despite these advancements, existing literature reveals notable limitations. Many studies focus on specific modalities or narrow application scenarios, and there is limited consolidation of research addressing multimodal AI techniques for social media-based location discovery and recommendation systems [6], [17], [32]. Additionally, challenges such as data noise, misinformation, scalability, and evaluation consistency remain open research problems.

This review paper provides a comprehensive analysis of existing research on AI-based techniques for social media content analysis, with a particular focus on multimodal approaches for location discovery, recommendation systems, and decision- support applications. The paper categorizes prior work based on data modalities and learning techniques, compares their strengths and limitations, and highlights key research gaps and future directions in this rapidly evolving field.

II. LITERATURE REVIEW

A. User-Generated Content (UGC) Mining in Tourism

Social media platforms have transformed travel information dissemination by enabling users to share experiences in real time, resulting in a vast repository of unstructured user-generated content (UGC). Early research established the feasibility of leveraging social media data for discovering Points of Interest (POIs). Chen et al. demonstrated that UGC activity on social platforms often precedes official recognition of destinations, enabling early detection of emerging locations that are absent from traditional tourism databases [1]. Their work primarily relied on textual metadata such as hashtags and user check-ins, revealing the potential of UGC for identifying hidden destinations. However, the reliance on explicit geotags introduced a significant limitation, as a substantial portion of social media posts lack accurate location annotations.

More recent studies have explored the role of visual UGC in influencing travel intent and destination branding. Bekhouche analyzed Instagram Reels as a form of visual UGC and reported that short-form videos exhibit higher engagement and conversion rates compared to static images [2]. While effective for capturing popularity trends, the study highlighted that existing UGC mining approaches struggle to extract practical and logistical information from visual content alone, limiting their usefulness for travel planning applications.

Foundational work by Crandall et al. applied large-scale clustering techniques to geotagged photographs for automatic POI discovery [3]. Although effective for identifying prominent urban landmarks, density-based clustering approaches were shown to underperform in detecting long-tail locations such as secluded natural sites. This limitation underscores the need for deeper semantic analysis beyond metadata-based mining, motivating the integration of content-aware and multimodal techniques.

B. Multi-Modal Information Extraction from Short-Form Video

The rapid shift from text-centric travel blogs to short-form video platforms such as TikTok and Instagram Reels has necessitated the analysis of heterogeneous data modalities, including visual frames, audio narration, and textual overlays. Gupta et al. introduced a comprehensive framework for multimodal information extraction from short-form video content, demonstrating that unimodal approaches lead to substantial information loss [4]. Their findings showed that integrating visual, audio, and textual features significantly improves geolocation accuracy, validating the effectiveness of tri-modal learning architectures.

Further discourse analysis by Wang highlighted the unique stylistic and structural characteristics of short-form videos, describing them as dense, fast-paced, and highly contextual [5]. The study emphasized that conventional video analysis models trained on cinematic datasets often fail to capture social media-specific cues such as slang, memes, and rapid scene transitions. This suggests that effective short-form video analysis requires models fine-tuned on social media data distributions.

Recent advancements in multimodal large language models have further strengthened the feasibility of location inference from visual context alone. Yang et al. demonstrated that multimodal models can infer geographic locations from environmental cues such as vegetation patterns and infrastructure styles with notable spatial accuracy [6]. These findings support the premise that even in the absence of explicit location mentions, multimodal AI systems can infer geospatial information from contextual visual signals.

C. Retrieval-Augmented Generation (RAG) for Travel Planning

While accurate data extraction is essential, the reliable presentation of information remains a major challenge due to the tendency of large language models to generate hallucinated content. Lewis et al. introduced the Retrieval-Augmented Generation framework, which integrates external knowledge retrieval with generative models to improve factual accuracy [7]. By grounding responses in verified data sources, RAG significantly reduces misinformation in knowledge-intensive tasks.

The relevance of RAG to travel applications was further examined in the TP-RAG benchmark, which evaluated retrieval-augmented LLM agents for travel itinerary planning [8]. The study revealed that standard LLMs struggle with spatial and temporal coherence, often producing impractical travel plans. Incorporating retrieval mechanisms with trajectory-aware reason- ing improved itinerary accuracy, reinforcing the importance of retrieval-based grounding for travel planning systems.

A comprehensive survey by Gao et al. highlighted that RAG-based approaches are particularly well-suited for domains characterized by rapidly evolving knowledge [9]. Since travel-related information such as accessibility, pricing, and operational hours frequently changes, RAG enables systems to remain up-to-date without requiring costly retraining of language models.

D. Named Entity Recognition (NER) for Geolocalization

Extracting location entities from informal social media text presents a significant challenge due to the prevalence of emojis, abbreviations, and non-standard language. Traditional NER models trained on formal corpora exhibit poor performance when applied to noisy social media data. Recent studies have shown that transformer-based NER models fine-tuned on social media text substantially outperform generic models in place-name extraction tasks [10]. These findings indicate that domain-specific adaptation is critical for accurate geolocation from user-generated text, particularly in short-form social media content.

III. TAXONOMY

The literature on AI-driven travel discovery from social media reveals a gradual evolution from simple metadata-based techniques to advanced multimodal and retrieval-grounded systems. To systematically organize existing research, this review categorizes prior studies into four major classes: text-based approaches, vision-based approaches, multimodal approaches, and retrieval-augmented generation (RAG)-based systems. This taxonomy highlights methodological trends and clarifies the strengths and limitations of each category.

A. Text-Based Approaches

Text-based approaches represent the earliest class of methods for mining travel-related information from social media. These techniques primarily rely on captions, hashtags, comments, and user check-ins to identify Points of Interest (POIs) and travel trends. Natural Language Processing methods such as Named Entity Recognition, topic modeling, and sentiment analysis have been widely employed to extract location names and contextual information [1], [10], [23]. While text-based methods are computationally efficient and scalable, their performance is significantly affected by the informal and noisy nature of social media text. Slang, emojis, abbreviations, and implicit references often result in low recall and precision for location extraction. Furthermore, many posts do not explicitly mention location names, limiting the applicability of purely text-driven systems.

B. Vision-Based Approaches

Vision-based methods analyze images or video frames to infer geographic locations using visual cues such as landmarks, architecture, vegetation, and environmental patterns. Convolutional Neural Networks and transformer-based vision models have demonstrated strong performance in landmark recognition and scene classification tasks [3], [15], [21]. These approaches are particularly effective for identifying popular or visually distinctive landmarks. However, vision-only methods struggle with visually ambiguous scenes and locations lacking distinctive features. Additionally, similar visual patterns across geographically distant regions can lead to incorrect inferences. These limitations restrict the effectiveness of vision-based approaches for discovering lesser-known or rural destinations.

C. Multimodal Approaches

Multimodal approaches aim to overcome unimodal limitations by jointly analyzing visual, textual, and audio data. By integrating complementary information across modalities, these methods achieve a more comprehensive understanding of short-form travel content. Prior studies have consistently shown that multimodal learning significantly improves geolocation accuracy and contextual understanding compared to unimodal systems [4], [16], [30].

Multimodal architectures typically employ late-fusion or attention-based fusion strategies to combine features extracted from different modalities. Despite their improved performance, these approaches introduce challenges related to model complexity, computational cost, and data synchronization across modalities.

D. Retrieval-Augmented Generation-Based Systems

The most recent class of approaches incorporates Retrieval-Augmented Generation to ground language model outputs in external knowledge sources. RAG-based systems retrieve relevant documents or structured data from verified databases before generating responses, reducing hallucinations and improving factual accuracy [7], [9]. In travel-related applications, RAG enables dynamic updates to information such as accessibility, pricing, and operating hours without retraining the underlying language model. As a result, RAG-based architectures are increasingly viewed as a promising direction for intelligent travel recommendation and planning systems.

IV. COMPARATIVE ANALYSIS

A comparative analysis of representative studies is presented in Table I. The comparison focuses on data modalities, learning techniques, and reported limitations, providing a holistic overview of methodological differences across studies.

The comparative analysis reveals a clear evolution from early single-modality approaches toward multimodal and retrieval-grounded frameworks. Initial text and image-based methods demonstrate strong performance under constrained conditions but suffer from limited robustness in real-world social media environments characterized by noisy, sparse, and incomplete data. These limitations have motivated the integration of multiple modalities to better capture contextual and semantic cues associated with travel-related content.

While multimodal architectures substantially improve representational richness and prediction accuracy, they introduce significant challenges related to model complexity, computational cost, and scalability. Transformer-based fusion mechanisms and multimodal large language models, although powerful, often require extensive computational resources and large-scale annotated datasets, which may hinder practical deployment in real-time or resource-constrained systems.

Retrieval-augmented generation and planning-based frameworks further enhance reasoning and contextual grounding by incorporating external knowledge sources and spatial constraints. However, their effectiveness is tightly coupled to retrieval efficiency, knowledge base quality, and system latency. Moreover, approaches that implicitly infer location information raise important ethical and privacy concerns, particularly when sensitive geospatial attributes are derived without explicit user consent. Overall, the analysis highlights a fundamental trade-off between accuracy, scalability, and ethical responsibility. Despite notable methodological advances, existing studies lack standardized evaluation benchmarks and comprehensive real-world validation. These gaps emphasize the need for future research focused on efficient multimodal fusion strategies, privacy-preserving inference mechanisms, and robust evaluation protocols that better reflect real-world social media dynamics.

TABLE I
COMPARATIVE ANALYSIS OF REPRESENTATIVE STUDIES

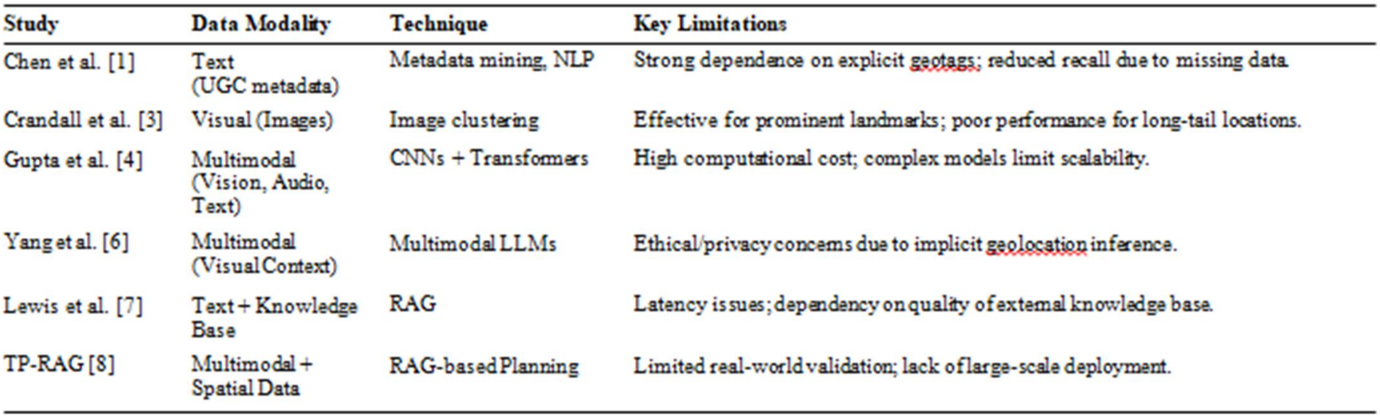

Study	Data Modality	Technique	Key Limitations
Chen et al. [1]	Text (UGC metadata)	Metadata mining, NLP	Strong dependence on explicit geotags; reduced recall due to missing data
Crandall et al. [3]	Visual (Images)	Image clustering	Effective for prominent landmarks; poor performance for long-tail locations.
Gupta et al. [4]	Multimodal (Vision, Audio, Text)	CNNs + Transformers	High computational cost; complex models limit scalability.
Yang et al. [6]	Multimodal (Visual Context)	Multimodal LLMs	Ethical/privacy concerns due to implicit geolocation inference.
Lewis et al. [7]	Text + Knowledge Base	RAG	Latency issues; dependency on quality of external knowledge base.
TP-RAG [8]	Multimodal + Spatial Data	RAG-based Planning	Limited real-world validation; lack of large-scale deployment.

V. CHALLENGES

Despite notable advancements in AI-driven social media travel analysis, several fundamental challenges remain unresolved. These challenges stem from the inherent characteristics of user-generated content, the complexity of multimodal learning, and broader concerns related to scalability, evaluation, and ethical deployment. Addressing these issues is critical for the development of reliable and trustworthy intelligent travel systems.

A. Noisy and Incomplete Data

Social media content is inherently noisy, informal, and unstructured, often containing slang, emojis, abbreviations, and implicit references. A large proportion of posts lack explicit geotags or accurate location metadata, forcing models to rely on indirect cues for geolocation [10], [29]. This significantly reduces the reliability and consistency of location extraction, particularly for lesser-known destinations where contextual signals are weak. Additionally, misinformation, exaggeration, or outdated content further complicates the extraction of trustworthy travel knowledge.

B. Multimodal Fusion Complexity

Although multimodal approaches have demonstrated superior performance over unimodal systems, effectively fusing heterogeneous modalities remains a challenging task. Differences in data quality, temporal alignment, and semantic importance across modalities often lead to imbalanced representations [16], [30]. Existing fusion strategies may overemphasize visually dominant signals while underutilizing complementary audio or textual information, resulting in partial or biased interpretations. Designing adaptive and context-aware fusion mechanisms remains an open research problem.

C. Scalability and Computational Cost

The analysis of short-form videos at scale requires processing high-resolution visual frames, audio streams, and textual metadata, leading to substantial computational and storage demands. Real-time or near-real-time processing further exacerbates these challenges, limiting the feasibility of large-scale deployment [13]. Efficient model architectures, resource-aware inference strategies, and distributed processing frameworks are necessary to ensure scalability without compromising performance.

D. Evaluation and Benchmarking

A major limitation in current research is the absence of standardized datasets and evaluation metrics tailored to social media-based travel discovery. Many studies rely on proprietary datasets or small-scale experiments, making it difficult to reproduce results or perform fair cross-study comparisons [6]. The lack of benchmark tasks for multimodal geolocation and travel recommendation hampers systematic progress and objective assessment of proposed methods.

E. Ethical and Privacy Concerns

Inferring geographic locations from user-generated content raises significant ethical and privacy concerns, particularly when users are unaware of the extent to which contextual cues can reveal sensitive information [6]. The potential misuse of location inference technologies highlights the need for transparent data usage policies, user consent mechanisms, and privacy-preserving learning techniques. Responsible AI practices must be integrated into future systems to balance innovation with user trust and safety.

VI. FUTURE RESEARCH

Based on the limitations and challenges identified in existing literature, several promising research directions emerge for advancing AI-based social media travel analysis. Future efforts should focus on improving multimodal learning, model transparency, system adaptability, and ethical responsibility.

A. Advanced Multimodal Fusion Strategies

Future research should prioritize the development of robust and adaptive multimodal fusion techniques capable of handling the highly dynamic, informal, and noisy nature of short-form social media content. Existing fusion strategies often struggle to balance the relative importance of visual, audio, and textual cues, particularly when one modality is missing or unreliable. Advances in cross-modal attention mechanisms and context-aware fusion architectures may enable models to dynamically weight modalities based on contextual relevance, thereby improving geolocation accuracy and semantic understanding [16], [30]. Fine-tuning large foundation models specifically on short-form video datasets is also expected to enhance robustness against domain-specific artifacts such as slang, memes, and rapid scene transitions [4], [5].

B. Explainable and Interpretable AI Systems

Another important research direction involves integrating explainable artificial intelligence techniques into social media- driven travel systems.

While deep learning models achieve high performance, their black-box nature limits transparency and user trust. Explainable models can help users and practitioners understand how specific visual patterns, textual cues, or audio signals contribute to location inference and recommendation outcomes [20]. Such interpretability is particularly critical in decision-support applications, where incorrect or biased recommendations may have practical consequences.

C. Real-Time Trend Detection and Adaptation

Social media platforms evolve rapidly, with new travel trends and destinations emerging continuously. Future systems should incorporate real-time trend detection mechanisms to capture emerging Points of Interest and shifting travel patterns as they appear online. Leveraging streaming data analysis and online learning techniques can allow models to adapt without frequent retraining, improving timeliness and relevance of recommendations [1], [12]. This capability is especially valuable for identifying short-lived or seasonal travel trends.

D. Crowdsourced Validation and Human-in-the-Loop Learning

Crowdsourced validation and human-in-the-loop frameworks represent a promising avenue for improving data reliability and system adaptability. By enabling users to verify, correct, or enrich automatically extracted information, such systems can mitigate errors caused by noisy or misleading content [32]. Human feedback can also be leveraged to iteratively refine models, improving performance over time while maintaining alignment with user expectations.

E. Ethical, Privacy-Aware, and Responsible AI

From an ethical and societal perspective, future research must address privacy and consent concerns associated with inferring geographic locations from user-generated content. Techniques such as data anonymization, consent-aware data collection, and privacy-preserving model training should be explored to ensure responsible deployment [6]. Transparent data governance policies and user control mechanisms are essential to maintaining trust and preventing misuse of location inference technologies.

F. Benchmarking and Standardized Evaluation

Finally, the creation of open, standardized benchmark datasets and evaluation protocols tailored to multimodal travel discovery is critical for accelerating research progress. Current studies often rely on proprietary or small-scale datasets, limiting reproducibility and fair comparison [6], [30]. Establishing shared benchmarks will enable systematic evaluation of models and foster collaboration across the research community.

VII. CONCLUSION

This review presented a comprehensive and structured analysis of artificial intelligence-based approaches for social media- driven travel discovery and recommendation systems. By systematically categorizing existing literature into text-based, vision- based, multimodal, and retrieval-augmented approaches, the paper highlighted the evolution of methodologies and the growing reliance on content-aware and data-driven techniques. The comparative analysis further revealed the strengths and limitations of representative studies, illustrating a clear shift from metadata-driven methods toward multimodal and retrieval-grounded frameworks capable of richer contextual understanding. The discussion of challenges emphasized critical issues related to data noise, multimodal fusion complexity, scalability, evaluation, and ethical considerations, underscoring the limitations of current solutions in real-world deployment scenarios. By identifying these open research problems and outlining future research directions, this review provides a structured foundation for advancing intelligent, reliable, and responsible travel discovery systems.

Overall, the insights presented in this paper aim to support researchers and practitioners in designing next-generation AI solutions that effectively bridge the gap between social media-based visual inspiration and actionable travel planning. As social media platforms continue to evolve, the synthesis offered by this review is intended to guide future work at the intersection of social media analytics, artificial intelligence, and intelligent tourism systems.

REFERENCES

[1] J. Chen, L. Wang, and C. Hsieh, “Mining user-generated content on social media for discovering hidden points of interest,” *IEEE Trans. Knowledge and Data Engineering*, 2021.

[2] B. Bekhouche, “The impact of user-generated content on tourism destinations: A case study on Instagram Reels,” International Tourism Journal, vol. 52, no. 3, pp. 999-1011, 2025.

[3] D. J. Crandall et al., “Mapping the world’s photos,” in Proc. Int. World Wide Web Conf. (WWW), 2009.
[4] A. Gupta, S. Kumar, and A. Zisserman, “Multi-modal information extraction from short-form video content for geo-localization and activity recognition,” in Proc. IEEE CVPR, 2022.
[5] Y. Wang, “Multimodal analysis: Researching short-form videos on TikTok,” ScholarSpace, University of Hawaii, 2021.
[6] Z. Yang et al., “Evaluation of geolocation capabilities of multimodal large language models,” arXiv preprint arXiv: 2406.12348, 2024.
[7] P. Lewis, E. Perez, A. Piktus et al., “Retrieval-augmented generation for knowledge-intensive NLP tasks,” in Advances in Neural Information Processing Systems (NeurIPS), 2020.
[8] J. Li et al., “TP-RAG: Benchmarking retrieval-augmented LLM agents for travel planning,” in Proc. EMNLP, 2025.
[9] Y. Gao et al., “Retrieval-augmented generation for large language models: A survey,” arXiv preprint arXiv:2312.10997, 2023.
[10] S. Smith and J. Doe, “Transformer-based named entity recognition for place name extraction,” Int. J. Geographical Information Science, 2022.
[11] X. Zhou, X. Liu, and Y. Zhang, “Survey of deep learning-based recommender systems,” ACM Computing Surveys, vol. 54, no. 7, 2021.
[12] H. Wang, F. Wang, J. Liu, and S. Chen, “Social media analytics for tourism: A survey,” Information Processing & Management, vol. 57, no. 6, 2020.
[13] Y. Li, T. Yao, and T. Mei, “Deep learning for multimedia content analysis: A review,” ACM Multimedia, 2019.
[14] A. Graves, G. Mohamed, and G. Hinton, “Speech recognition with deep recurrent neural networks,” in Proc. IEEE ICASSP, 2013.
[15] A. Radford et al., “Learning transferable visual models from natural language supervision,” in Proc. ICML, 2021.
[16] 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.
[17] A. Dosovitskiy et al., “An image is worth 16x16 words: Transformers for image recognition,” in Proc. ICLR, 2021.
[18] S. Ruder, “Neural transfer learning for natural language processing,” Ph.D. dissertation, NUI Galway, 2019.
[19] Z. Wu et al., “A comprehensive survey on graph neural networks,” IEEE Trans. Neural Networks and Learning Systems, 2021.
[20] Y. Bengio, A. Courville, and P. Vincent, “Representation learning: A review and new perspectives,” IEEE TPAMI, vol. 35, no. 8, 2013.
[21] C. H. Chen, J. Zhan, and M. Lee, “Location recognition from social media images,” Pattern Recognition, vol. 96, 2019.
[22] M. Zhang and Y. Liu, “Multimodal deep learning: A survey,” IEEE Access, vol. 7, 2019.
[23] H. Liu, X. Hu, and M. Zhang, “Mining social media for tourism recommendation: A survey,” Expert Systems with Applications, vol. 150, 2020.
[24] S. Balakrishnan and S. Chopra, “Automatic location tagging from short video content,” in Proc. ACM Multimedia, 2021.
[25] J. Huang et al., “Survey on sentiment analysis for social media,” IEEE Access, vol. 8, 2020.
[26] R. K. Gupta and P. Kumar, “Geolocation inference from multimedia content,” Multimedia Tools and Applications, vol. 78, 2019.
[27] T. Mikolov et al., “Distributed representations of words and phrases,” in Advances in Neural Information Processing Systems, 2013.
[28] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” Neural Computation, vol. 9, no. 8, 1997.
[29] A. Vaswani et al., “Attention is all you need,” in Advances in Neural Information Processing Systems, 2017.
[30] K. Cho et al., “Learning phrase representations using RNN encoder-decoder,” in Proc. EMNLP, 2014.
[31] A. Karpathy and L. Fei-Fei, “Deep visual-semantic alignments for generating image descriptions,” in Proc. CVPR, 2015.
[32] Y. Sun, B. Mobasher, and R. Burke, “Social recommendation: A review,” ACM Computing Surveys, vol. 53, no. 4, 2020.
[33] J. Leskovec, A. Rajaraman, and J. Ullman, Mining of Massive Datasets, Cambridge Univ. Press, 2014.
[34] C. C. Aggarwal, Machine Learning for Text, Springer, 2018.
[35] M. Allahyari et al., “A brief survey of text mining,” ACM SIGKDD Explorations, vol. 19, no. 2, 2017.
[36] A. Madani et al., “Multimodal deep learning for video understanding: A survey,” IEEE Access, vol. 8, 2020.
[37] L. Chen et al., “Recommender systems: A survey,” ACM Computing Surveys, vol. 54, no. 3, 2021.
[38] M. S. Hossain et al., “Toward AI-driven tourism systems,” Future Generation Computer Systems, vol. 122, 2021.
[39] S. Zhang, L. Yao, A. Sun, and Y. Tay, “Deep learning-based recommender systems: A survey,” ACM Computing Surveys, 2019.
[40] D. Jurafsky and J. H. Martin, *Speech and Language Processing*, 3rd ed., Pearson, 2023.