

FITMag: A Framework for Generating Fashion Journalism Using Multimodal LLMs, Social Media Influence, and Graph RAG

Jinda Han¹, Mengyao Guo², Shanghao Li³, Kailash Thiyagarajan⁴, Zhinan Cheng⁵, and Li Zhao⁶

¹University of Illinois at Urbana Champaign, Champaign, USA

²Shenzhen International School of Design, Harbin Institute of Technology, Shenzhen, China

³University of Illinois at Chicago, Chicago, USA

⁴Anna University Chennai, Chennai, India

⁵University of California San Diego, San Diego, USA

⁶University of Missouri Columbia, Columbia, USA

ABSTRACT

As generative artificial intelligence (AI) reshapes media and communication, its integration with social media creates new opportunities for human-centered content creation. We present FITMag, a framework that combines multimodal large language models (LLMs) with real-time social media influence and graph data to generate fashion articles approaching professional quality. FITMag ingests diverse multimodal inputs, including fashion influencer metadata, tweets, hashtags (e.g., #NYFW, #bohemian), images, and network graphs of follows, retweets, and mentions, to drive text and image generation using models such as ChatGPT, Claude, DeepSeek, and LLaMA, enhanced with graph databases for Retrieval-Augmented Generation (RAG). FITMag outputs three article formats: event based, niche community based, and trend based. In a qualitative user study with 15 fashion professionals, AI-generated text frequently matched human writing in authenticity, coherence, and style. However, image generation lagged in contextual accuracy, highlighting current limitations in multimodal modeling. These findings illustrate the potential of generative AI in fashion journalism and emphasize the need for stronger vision-language alignment in professional fashion media contexts.

Keywords: Fashion journalism, Social media influence, Generative AI, Large language models, Human-centered AI, Graph RAG, Multimodal LLMs

INTRODUCTION

Fashion journalism has long been shaped by elite publications and industry gatekeepers. However, with the proliferation of social media platforms such as X (formerly Twitter), Instagram, and TikTok, the dynamics of influence have shifted. Traditional fashion editors are now joined by independent influencers, bloggers, and social media personalities who drive discourse and

set trends in real time (Kim et al., 2021; Sharma et al., 2024; Vogue Business, 2024). This shift has led to an unprecedented democratization of fashion media, where individuals can shape public perception and consumer behavior through their online presence.

At the same time, advancements in AI, particularly generative models such as ChatGPT, Claude, LLaMA, and DeepSeek (Achiam et al., 2023; Anthropic, 2024; Guo et al., 2025; Touvron et al., 2023), have introduced new possibilities for content creation. These models can generate high-quality text, mimicking human writing styles and producing articles, summaries, and reports at scale. Within the fashion domain, generative AI vision models such as Stable Diffusion, DALL-E, and Midjourney (Podell, 2024; Nikolenko, 2023) have the potential to streamline the production of editorial content, provide design inspiration, and support journalists in crafting compelling narratives (Harreis et al., 2023; Iksula et al., 2023; Słomka et al., 2024).

Existing research has focused on analyzing social media influence within the fashion industry (Smith et al., 2024; Jones et al., 2023; Garcia et al., 2024; Whitaker et al., 2024; McCann et al., 2024). For instance, FITNet identified key influencers by constructing a fashion-specific subgraph of Twitter, leveraging mentions, retweets, and following relationships to map the most impactful voices in fashion (Han et al., 2021). FITViz further provided interactive visualization tools to explore these influencer networks, demonstrating how fashion influence spreads across social media (Han et al., 2025). However, these works have yet to explore how such influence can be harnessed to inform and generate high-quality fashion articles in the context of journalism.

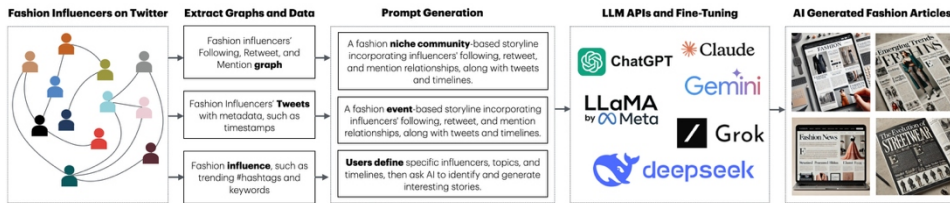


Figure 1: The FITMag workflow for generating fashion articles, from extracting influencer data and social graphs on Twitter to prompt construction, LLM processing, and AI-generated content.

In this paper, we introduce FITMag (see Figure 1), a novel framework that integrates influencer metadata, extracted social media trends, and fashion subgraphs with Graph RAG-enhanced multimodal LLMs to generate fashion articles approaching professional quality. FITMag builds on the foundational datasets of FITNet and FITViz, leveraging influencers' text and image content along with their interaction graphs, including follow, retweet, and mention relationships, as structured graph data. It also incorporates timeline-based trends and topical keywords and hashtags (e.g., circular fashion, #NYFW) as inputs to LLMs enhanced by Graph RAG (Peng, 2024). By combining these components, FITMag aims to close the gap between AI-generated content

and human-written journalism, ensuring that the resulting articles reflect the tone, style, and depth expected in the fashion industry.

To evaluate the effectiveness of FITMag, we compare AI-generated articles produced by models such as ChatGPT, Claude, LLaMA, and DeepSeek, along with images generated using Stable Diffusion, against articles published by leading fashion magazines and media outlets (e.g., Vogue, Elle, WWD, Fashionista, Harper's Bazaar). We then conduct a human-centered evaluation through interviews with 15 fashion professionals, assessing the authenticity, stylistic accuracy, and credibility of the AI-generated content. From this study, we draw preliminary conclusions about whether AI-driven fashion journalism can approximate human-level quality.

Our research contributes to the broader discourse on AI in fashion journalism, social media influence, and human-computer interaction (HCI) by exploring the viability of AI-generated writing in professional fashion media. We also examine whether human readers can reliably distinguish between AI- and human-authored content. By leveraging real-world social media data and large language models, FITMag seeks to advance the discussion around AI-assisted journalism and its implications for the future of fashion journalism.

RELATED WORKS

Research at the intersection of generative AI, social media, and fashion journalism spans multiple disciplines, including natural language processing, media studies, and human-computer interaction. As AI-generated content grows more sophisticated, recent studies have examined its editorial quality, credibility, and reader perception. In fashion, social media influencers play a central role in shaping trends, with works like FITNet and FITViz mapping digital influence. Human-centered evaluations have explored trust in AI-generated content, while multimodal AI research has expanded into fashion image generation, retrieval, and captioning. Building on this work, our study offers a unified framework that combines social media signals, multimodal generation, and expert evaluation to explore AI's role in fashion journalism.

AI in Fashion Journalism and Content Creation

The intersection of generative AI, social media influence, and fashion journalism has been explored across multiple research domains. The rise of AI-assisted journalism has driven advancements in natural language generation, with models such as GPT-4o, Claude, and DeepSeek R1 increasingly used to produce high-quality editorial content. Studies have examined the feasibility of AI-generated news articles, with findings suggesting that readers often struggle to distinguish between AI-written and human-authored content (Thasler et al., 2024; Simon, 2022; Krings, 2024). For example, a study published in *Computers in Human Behavior* found that readers rated AI-generated news articles as less informative, yet had difficulty accurately identifying their source (Sorab et al., 2023). Similarly, research from Pennsylvania State University highlights the growing challenge

of differentiating between AI- and human-generated content as AI technology continues to evolve (Fetzer, 2024).

However, fashion journalism presents unique challenges, as credibility depends not only on factual accuracy but also on stylistic consistency, creativity, and deep industry expertise (Glam Observer, 2024; Karen, 2023). According to the Fashion Retail Academy, fashion journalists must possess strong writing and communication skills to keep pace with the industry's fast-moving and ever-evolving nature (Fashion Retail Academy, 2024). Creativity is essential across all roles in the fashion sector, enabling professionals to innovate and contribute meaningfully to a dynamic media landscape (Glam Observer, 2024). In addition, a strong understanding of fashion history, current trends, and cultural relevance is crucial for producing high-quality, engaging content that resonates with readers (Karen, 2023).

Social Media Influence in Fashion

Research on social media influence in fashion has primarily focused on identifying and analyzing influencer networks (Rosário et al., 2025; Nissa et al., 2024; Heuritech, 2025). Studies have examined how fashion influencers impact consumer behavior and brand engagement, emphasizing the importance of authenticity and the challenges influencers face in consistently producing engaging content. For example, Rosário et al. (2025) highlight the need for influencers to craft narratives that resonate with their audiences. Nissa et al. (2024) explore how influencer marketing strategies shape consumer purchasing decisions, while Heuritech (2025) investigates the role of influencer marketing in driving fashion trends and consumer insights.

Works such as FITNet and FITViz have mapped social media interactions to identify key fashion influencers, demonstrating how trends propagate across digital platforms (Han et al., 2021; Han et al., 2025). These studies highlight the importance of influencer engagement, retweet patterns, and hashtag trends in shaping online fashion discourse. This foundation informs our approach, allowing us to leverage these dynamic influence patterns and graph structures to generate AI-powered fashion articles that reflect real-world trends and audience engagement.

Human-Centered Evaluation of AI-Generated Content

Human-centered evaluations of AI-generated content have emerged as a critical area of research. Prior work in HCI has explored how users perceive the credibility, trustworthiness, and engagement of AI-generated text across various domains (Huschens et al., 2023; Cohn et al., 2024; Sharma et al., 2024). For example, Huschens et al. (2023) found that participants attributed similar levels of credibility to both human- and AI-generated content, and even rated AI-generated text as clearer and more engaging. Cohn et al. (2024) demonstrated that anthropomorphic cues in AI interfaces, such as combining speech and text, can enhance users' trust and perceived accuracy. Sharma et al. (2024) showed that providing explanations alongside AI-generated responses significantly increases user trust, particularly when

users can compare different responses. Building on this work, our study conducts a blind evaluation with fashion professionals to assess the feasibility of AI-driven fashion journalism in comparison to human-written content.

Multimodal AI in Fashion

Finally, the integration of multimodal AI models, which combine text generation with vision capabilities, has gained increasing attention in fashion research (Han et al., 2022; Mirchandani et al., 2022; Baldrati et al., 2024). These models support a range of applications, including fashion image retrieval, captioning, and editing. For instance, Han et al. (2022) introduced FashionViL, a vision-and-language representation learning framework tailored for fashion tasks, demonstrating improved performance across several benchmarks. Similarly, Mirchandani et al. (2022) proposed FaD-VLP, a pretraining framework designed to unify retrieval and captioning tasks within the fashion domain. Baldrati et al. (2024) explored multimodal-conditioned latent diffusion models for fashion image editing, enabling human-centric image generation guided by diverse prompts.

While models like Fashion-GPT focus primarily on vision-based tasks such as automated design recommendations, there remains a significant gap in ensuring text-image coherence for fashion journalism (Chen et al., 2023; Alikhani et al., 2023). Fashion-GPT integrates ChatGPT with fashion-specific AI modules to enhance multimodal search and retrieval (Chen et al., 2023), but it does not fully address the alignment between textual narratives and corresponding visual content. Research on image-text coherence emphasizes the importance of modeling the relationship between images and language to improve the performance and interpretability of multimodal systems (Alikhani et al., 2023). This highlights the need for further development of models capable of integrating coherent visual-textual narratives, especially within editorial contexts.

Our work addresses this gap by incorporating vision models alongside large language models (LLMs), enhanced with additional graph interaction data. This approach ensures that AI-generated fashion articles not only meet high textual standards but also align with appropriate and contextually relevant fashion imagery. By synthesizing insights from social media influence, generative AI, and human-centered evaluation, our study aims to advance the discourse on AI in fashion journalism through a comprehensive multimodal framework, reinforced by timely influence signals and Graph RAG.

METHODOLOGY

The FITNet dataset is designed to identify and analyze fashion influencers on Twitter by mapping social media interactions to understand how fashion influence spreads. We use this dataset not only to extract textual and visual content but also to identify fashion trends, niche communities, and topics. The dataset includes graph data capturing influencer relationships through follows, mentions, and retweets to analyze influence dynamics.

Dataset

FITNet was constructed using a fashion classifier trained on 11.5k manually labelled accounts, achieving 92% accuracy. A re-weighted random walk sampling method was used to estimate a fashion subgraph of 260k accounts. The full FITNet-300k dataset includes 300k fashion-related accounts, with the top 10k influencers selected using PageRank and categorized into influencers (42%), brands (20%), retailers (14%), and media (31%). Tweets, retweets, and mentions were collected from January 1, 2018, to February 1, 2019, revealing patterns such as homophily in influencer interactions, niche community formation (e.g., #plussize, #sustainability), and geographic clustering around fashion hubs.

To track engagement, we constructed three types of interaction graphs—mentions, retweets, and follows—focused on both event-based trends (e.g., #NYFW, #PFW, #StreetStyle, #OOTD), niche communities (e.g., #plussizefashion, #sustainability), and emergent fashion trends (e.g., Boho Chic in Suede). We also monitored key fashion influencers such as Kylie Jenner, Wisdom Kaye, Paige Lorenze, Chiara Ferragni, and Gabbriette, who play significant roles in shaping digital fashion discourse. These structured datasets were then used to generate three categories of fashion articles: event-based, niche community-based, and trend-based.

Model Selection and Implementation

To support fashion article generation, we implemented FITMag using a combination of LLMs and image generation models informed by graph, text, image, and extracted event, niche community, and trend data. For text generation, we selected ChatGPT (GPT-4o), Claude 3.5, DeepSeek R1, and LLaMA 3.3 due to their strong natural language generation performance and API or open-source availability. These models were integrated into a Graph RAG pipeline that utilized fashion influencer subgraphs derived from real-time Twitter data. The graph structure encoded follow, retweet, and mention relationships and was used to retrieve relevant prompts and context. For multimodal generation, we paired the LLMs with Stable Diffusion to produce corresponding visual content. Hashtags, trending topics, and influencer metadata were transformed into structured graph queries that informed both text and image generation. The system was modular, with separate components for LLM invocation, graph retrieval, prompt formatting, and image rendering, allowing for comparative evaluation across models.

EXPERIMENT DESIGN

We conducted an interview-based evaluation to assess the quality of FITMag-generated content in comparison to existing fashion journalism. The study involved 15 fashion professionals, including models, stylists, content creators, bloggers, public relations specialists, and fashion researchers. Participants were selected to reflect a diverse cross-section of professional roles within the fashion media industry. Each participant was compensated \$60 and spent approximately 1.5 to 2 hours completing the session.

Each participant reviewed three randomized set of 14 fashion articles in each category: 10 were generated by FITMag using different LLM and graph input configurations (GPT-4o API, GPT-4o API + Stable Diffusion, Claude 3.5 API + Stable Diffusion, DeepSeek R1 (70B) + Stable Diffusion, and LLaMA 3.3 (70B instruct-q3_K_M) + Stable Diffusion), two were curated from top-tier fashion publications (e.g., Vogue, Elle, Harper's Bazaar) matching the timeframe of the FITNet dataset, and two additional articles were generated from the most recent tweets (GPT-4o API + Stable Diffusion) to mitigate bias from LLMs trained on past data. Across the three categories of event-based, niche community-based, and trend-based, participants evaluated a total of 52 short articles (estimated reading time of 2–3 minutes per article).

Task Design

Participants were asked to complete three primary tasks for each article using a 5-point Likert scale style format (Joshi, 2015; Herbold, 2024). The evaluation followed a 3×3 task design applied to seven different model scenarios. Each of the three rating dimensions-authenticity, coherence, and style-was assessed using two articles from each article category: event-based, niche community-based, and trend-based. Participants assessed every article across these three key dimensions.

1. **Authenticity Rating:** Evaluate how believable and credible the article reads, on a 5-point Likert scale style ranging from “Does not trust at all” to “Very trustworthy.”
2. **Coherence Rating:** Assess the logical flow and structure of the article, on a 5-point Likert scale from “Totally disconnected” to “Extremely well connected.”
3. **Style Rating:** Judge how well the writing style aligns with professional fashion journalism, on a 5-point Likert scale from “Unprofessional” to “Highly professional.”

Participants were unaware of the origin (AI or human) of each article to ensure a blind review process. After completing all ratings, they were asked to guess whether each article was AI- or human-authored, followed by open-ended feedback.

Evaluation Metrics

The evaluation focused on three core metrics: authenticity, coherence, and style. Authenticity reflects the believability and journalistic credibility of each article. Coherence captures the logical structure and clarity of the narrative. Style evaluates how well the tone and phrasing align with professional fashion writing. These dimensions were used both for participant scoring and for assessing perceived credibility. In addition to these three metrics, we collected identification accuracy, which measures participants' ability to correctly classify content as human- or AI-generated, and qualitative feedback offering insights into participant perceptions, model strengths and weaknesses, and

stylistic suggestions. Together, these metrics provide a holistic view of how AI-generated content is perceived in professional fashion contexts.

Table 1: Average ratings of fashion event-based articles by model and dimension (authenticity, coherence, style).

Models Scenarios	Authenticity	Coherence	Style
Two human-authored fashion articles from the FITNet period	4.73	4.77	4.77
GPT-4o API + Stable Diffusion (FITNet data + two articles)	4.17	4.23	4.77
Claude 3.5 API + Stable Diffusion (FITNet data + two articles)	4.13	4.13	4.27
GPT-4o API + Stable Diffusion (New data + two articles)	4.10	3.90	4.40
DeepSeek R1 (70B) + Stable Diffusion (FITNet data + two articles)	4.03	4.13	4.07
LLaMA 3.3 (70B) + Stable Diffusion (FITNet data + two articles)	3.83	3.93	3.83
GPT-4o API (FITNet data + two articles)	3.50	3.93	4.07

RESULTS AND DISCUSSION

Demographically, 15 participants had an average of 6.5 years of experience in fashion, with roles ranging from PR and modeling to content strategy. Most reported using Instagram, TikTok, and WGSN for staying updated, indicating that AI tools must align with current social media practices and industry-standard resources.

Preliminary results from event-based articles indicated (see Table 1) that models such as GPT-4o, Claude 3.5, DeepSeek R1, and LLaMA 3.3 produced promising content. Based on evaluations from 15 fashion professionals, human-written articles unsurprisingly outperformed all AI-generated content across dimensions, particularly in authenticity (4.73) and coherence (4.77).

Among the AI models, GPT-4o with Stable Diffusion using FITNet data achieved the highest average across authenticity (4.17), coherence (4.23), and style (4.77)—equivalent to human-written articles in stylistic quality. This was closely followed by Claude 3.5 with FITNet data and GPT-4o using newly scraped data.

Participants frequently misidentified AI-generated articles as human-written, especially those by GPT-4o and Claude, highlighting their strong textual realism. However, participants also noted that image-text coherence remained a limitation. While the articles themselves were clear and well-structured, accompanying images occasionally failed to match the article's context or tone. Notably, when using GPT-4o for image generation, the system declined to depict real influencers due to copyright restrictions, which affected perceived authenticity.

DeepSeek R1 and LLaMA 3.3 received slightly lower scores in stylistic alignment, suggesting further refinement is needed to meet fashion journalism's tone and branding conventions. Despite this, coherence scores

were relatively consistent across AI models, indicating that structural flow was generally robust.

We then investigated niche community-based articles. Besides the human-authored content, Claude 3.5 + Stable Diffusion performed best among the evaluated models. The results indicate coherent narratives and an appropriate tone for sensitive topics (e.g., #plussize). However, it was slightly weaker in stylistic polish compared to GPT-4o.

In the trend-based category, the top performer was GPT-4o with new data. The model benefited from better contextual freshness derived from recent tweets, which helped enhance style and realism. However, coherence slightly decreased due to the emergent and less structured nature of prompt data in this category.

These findings validate the textual capabilities of FITMag's generative pipeline while highlighting opportunities to enhance visual supplements and fine-tune stylistic expression. With improvements to multimodal alignment, AI-generated fashion journalism has the potential to meet professional editorial standards.

CONCLUSION AND FUTURE WORK

FITMag demonstrates the potential of combining social media influence, graph-based interaction data, and generative AI to produce fashion journalism that approaches professional standards. By integrating fashion influencer subgraphs, Twitter interactions, and trending topics with multimodal LLMs, our framework generates articles that closely mirror human writing in style, authenticity, and topical relevance. Preliminary feedback from fashion professionals suggests that AI-generated text is often indistinguishable from human-authored content. However, the inclusion of visual elements, such as images, reveals ongoing challenges in vision-language alignment and contextual accuracy.

Looking ahead, we plan to incorporate more advanced vision models to improve text-image coherence, expand our datasets to include larger volumes of text, images, and graph data, and scale human-centered evaluations through a broader, quantitative study involving a wider pool of fashion experts. We also aim to further examine the ethical, editorial, and industry-specific implications of AI-driven journalism. As generative AI continues to evolve, understanding its role in content creation, influencer engagement, and media credibility will be essential to shaping the future of fashion journalism.

REFERENCES

- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altschmidt, J., Altman, S., Anadkat, S. and Avila, R. (2023) GPT-4 technical report. arXiv preprint arXiv:2303.08774.
- Alikhani, M., Khalid, B. and Stone, M. (2023) Image-text coherence and its implications for multimodal AI. *Frontiers in Artificial Intelligence*, 6, p. 1048874.
- Anthropic (2024) The Claude 3 Model Family: Opus, Sonnet, Haiku. Anthropic.

- Baldrati, A., Morelli, D., Cornia, M., Bertini, M. and Cucchiara, R. (2024) Multimodal-conditioned latent diffusion models for fashion image editing. arXiv preprint arXiv:2403.14828.
- Bhonsle, P., Soni, N. and Singh, C. M. (2024) Fashion in the Digital Age: Social Media Marketing's Influence on the Apparel Market. *ShodhKosh: Journal of Visual and Performing Arts*, 5(6).
- Chen, Q., Zhang, T., Nie, M., Wang, Z., Xu, S., Shi, W. and Cao, Z. (2023) November. Fashion-GPT: integrating LLMs with fashion retrieval system. In *Proceedings of the 1st Workshop on Large Generative Models Meet Multimodal Applications*.
- Cohn, M., Pushkarna, M., Olanubi, G. O., Moran, J. M., Padgett, D., Mengesha, Z. and Heldreth, C. (2024) Believing anthropomorphism: Examining the role of anthropomorphic cues on trust in large language models. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (pp. 1–15).
- Fashion Retail Academy (2024) *Guide to Fashion Journalism: Skills, Trends, and Industry Insights*. Fashion Retail Academy.
- Fetzer, M (2024) *Increasing Difficulty Detecting AI Versus Human Writing*. Information Sciences and Technology. The Pennsylvania State University.
- Garcia, L. and Patel, A. (2024) The Impact of Social Media on Communication and Popularity in the Fashion Industry. *Fashion Communication Review*.
- Glam Observer (2024) *How Creativity Powers Every Career in Fashion*. Glam Observer.
- Guo, D., Yang, D., Zhang, H., Song, J., Zhang, R., Xu, R., Zhu, Q., Ma, S., Wang, P., Bi, X. and Zhang, X. (2025) Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. arXiv preprint arXiv:2501.12948.
- Han, J., Chen, Q., Jin, X., Xu, W., Yang, W., Kumar, S., Zhao, L., Sundaram, H. and Kumar, R. (2021) FITNet: Identifying fashion influencers on twitter. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1), pp. 1–20.
- Han, J., Guo, M., Ding, W., Thiyagarajan, K. and Zhao, L. (2025) FITViz: An Interactive Visualization Framework for Demonstrating Fashion Influence through Social Media Influencer Interactions. In *Companion Proceedings of the 2025 ACM International Conference on Supporting Group Work* (pp. 68–75).
- Han, X., Yu, L., Zhu, X., Zhang, L., Song, Y. Z. and Xiang, T. (2022) FashionVil: Fashion-focused vision-and-language representation learning. In *European conference on computer vision* (pp. 634–651). Cham: Springer Nature Switzerland.
- Harreis, H., Koullias, T., Roberts, R. and Te, K. (2023) *Generative AI: Unlocking the Future of Fashion*. McKinsey & Company.
- Herbold, S., Trautsch, A., Kikteva, Z. and Hautli-Janisz, A. (2024) Large Language Models can impersonate politicians and other public figures. arXiv preprint arXiv:2407.12855.
- Heuritech (2025) *Zoom on Social Media: Influencer Marketing & Fashion Trends*. Heuritech x Hivency.
- Huschens, M., Briesch, M., Sobania, D. and Rothlauf, F. (2023) Do You Trust ChatGPT?—Perceived Credibility of Human and AI-Generated Content. arXiv preprint arXiv:2309.02524.
- Iksula (2023) *Harnessing Generative AI: The Next Evolution in Fashion*. Iksula Newsletter on LinkedIn.
- Joshi, A., Kale, S., Chandel, S. and Pal, D. K. (2015) Likert scale: Explored and explained. *British journal of applied science & technology*, 7(4), p. 396.

- Jones, M. and Taylor, S. (2023) Exploring the Influence of Social Media on Sustainable Fashion Consumption. *Journal of Fashion and Sustainability*, 12(3).
- Karen, G. R. (2023) A Guide to Writing for Fashion Magazines. Fashiony City. Guy Vernes.
- Kim, M. and Lee, J. (2024) The Influence of Fashion Influencers on Consumer Behavior and Trend Dissemination. *Fashion and Textiles*.
- Krings, M (2024) Study finds readers trust news less when AI is involved, even when they don't understand to what extent. KU News Service.
- McCann, A. M. M. (2024) *Social Media: One of Fast Fashion's Biggest Influencers*. *Maine Law Journal of Intellectual Property and Law*, 15(2).
- Mirchandani, S., Yu, L., Wang, M., Sinha, A., Jiang, W., Xiang, T. and Zhang, N. (2022) Fad-vlp: Fashion vision-and-language pre-training towards unified retrieval and captioning. arXiv preprint arXiv:2210.15028.
- Nissa, K., Chauhan, P., Asim, M. and Chaudhary, P. (2024) The Effect of Influencer Marketing on Consumer Behavior of Fashion Style in the Age of Social Media. *International Journal of Tourism and Hospitality in Asia Pacific*, 7(3), pp. 360–378.
- Nikolenko, S. (2023) Generative AI & VI: Stable Diffusion, DALL·E 2, and Midjourney. Synthesis AI.
- Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., M'uller, J., Penna, J. and Rombach, R. (2024) Sdxl: Improving latent diffusion models for high-resolution image synthesis. ICLR 2024 spotlight.
- Peng, B., Zhu, Y., Liu, Y., Bo, X., Shi, H., Hong, C., Zhang, Y. and Tang, S. (2024) Graph retrieval-augmented generation: A survey. 2024 Conference on AI, Science, Engineering, and Technology (AIXSET).
- Rosário, A. M. and da Silva, J. M. (2025) Social Media Influencers in Fashion: Challenges and Opportunities. In: *Social Media and Online Consumer Decision Making in the Fashion Industry*, pp. 243–266.
- Sharma, Priya and Gupta, Rahul (2024) The Role of Social Media Influencers in Fashion Trends and Consumer Engagement. *International Journal of Novel Research and Development*, 9(2), pp. 15–27.
- Simon, F. M. (2022) 'Uneasy Bedfellows: AI in the News, Platform Companies and the Issue of Journalistic Autonomy', *Digital Journalism*.
- Słomka, D. (2024) Generative AI in Fashion Industry: How It Transforms the Market. Miquido.
- Sharma, M., Siu, H. C., Paleja, R. and Peña, J. D. (2024) Why Would You Suggest That? Human Trust in Language Model Responses. arXiv preprint arXiv:2406.02018.
- Sorab, G. (2023) Can Humans Distinguish Between Machine-Created and Human-Written Content? Medium.
- Th'asler-Kordonouri, S., Thurman, N., Schwertberger, U. and Stalpf, F. (2024) Too many numbers and worse word choice: Why readers find data-driven news articles produced with automation harder to understand. *Journalism: Theory, Practice & Criticism*.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M. A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F. and Rodriguez, A. (2023) Llama: Open and efficient foundation language models. arXiv preprint arXiv:2302.13971.
- Vogue Business Editorial Team (2024) Want to Be More Sustainable? Develop a Better Sense of Style. Vogue Business.
- Whitaker, K. L. (2024) The Effects of Social Media in the Fashion Industry. Lindenwood University.