

SOMIN: Agentic AI for Automating Professional Visibility and Countering Content Homogenization

Aleksandr Farseev
sasha@somin.ai
SoMin.ai Research
Singapore, Singapore

Kirill Lepikhin
kirill@somin.ai
SoMin.ai Research
Singapore, Singapore

Kevin Manuel
kevinmanuellee@hotmail.com
National University of Singapore
Singapore, Singapore

Maksim Gorodilov
maksingorodilov74@gmail.com
SoMin.ai Research
Singapore, Singapore

Zhao Kui
zhaokui@lianbangchuhai.com
ITMO University
St. Petersburg, Russia

Ilya Makarov
iamakarov@axxx.tech
AXXX; Trusted AI Research Center, RAS
Moscow, Russia

Jaime Francisco Maldonado
Francisco@fenyx.digital
Qubit Research
Florida, United States

Ian Cassidy
ian@defygroup.ai
DEFY Research
London, United Kingdom

Ronan Byrne
ronan.byrne@somin.ai
SoMin.ai Research
Singapore, Singapore

Abstract

In the evolving digital economy, professional visibility on networks such as LinkedIn [14] has become a critical determinant of career trajectory, influence, and the emerging metric of “LLM discoverability.” For the general public and the research community alike, the ability to translate technical expertise into engaging market narratives is essential for professional survival. However, the vast majority of individuals lack the marketing acumen required to maintain a high-frequency, high-quality public presence. While the democratization of LLMs appears to offer a solution, these systems inherently suffer from probability maximization mechanisms that drive outputs toward the statistical mean. Consequently, standard automated systems frequently result in “lexical homogenization,” where generated content becomes repetitive, generic, and devoid of the unique perspective necessary for effective personal branding. This technological limitation leads to poor content performance and fails to distinguish individual professionals from the noise of AI-generated text.

To address this communication deficit, we introduce SOMIN, an Agentic AI-ready marketing ideation platform that bridges the gap between technical expertise and public engagement by employing Dynamic Retrieval Augmented Generation (DRAG). By anchoring its outputs in a near-real-time corpus of LinkedIn insights from leading Information Retrieval researchers, SOMIN internalizes thought leadership best practices, enabling users to overcome the content creation cold start” problem and produce distinct, highly engaging professional narratives. Our demonstration will showcase SOMIN’s

dual capabilities: first, through real-time user interaction where conference participants can analyze SIGIR themes to ideate personalized LinkedIn posts on the spot, and second, by highlighting its agentic readiness” via a fully autonomous, Zapier-orchestrated workflow. In this integrated pipeline, SOMIN acts as the core technical narrative engine—identifying market tensions and triggering downstream agents to seamlessly schedule and publish—thereby automating 100% of the content production process and empowering professionals to build sophisticated communication strategies that secure their visibility within an increasingly AI-mediated information ecosystem.

CCS Concepts

• Information systems → Social networks; • Computing methodologies → Natural language generation; • Human-centered computing → Social content sharing.

Keywords

Agentic AI, Large Language Models, Retrieval-Augmented Generation, Content Homogenization, Social Media Marketing, Personal Branding

ACM Reference Format:

Aleksandr Farseev, Kirill Lepikhin, Kevin Manuel, Maksim Gorodilov, Zhao Kui, Ilya Makarov, Jaime Francisco Maldonado, Ian Cassidy, and Ronan Byrne. 2026. SOMIN: Agentic AI for Automating Professional Visibility and Countering Content Homogenization. In *Proceedings of the 49th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '26)*, July 20–24, 2026, Melbourne, VIC, Australia. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3805712.3808371>

The work of Ilya Makarov was supported by a grant, provided by the Ministry of Economic Development of the Russian Federation (agreement dated June 20, 2025 No. 139-15-2025-011, identifier 000000C313925P4G0002).



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

SIGIR '26, Melbourne, VIC, Australia

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2599-9/2026/07

<https://doi.org/10.1145/3805712.3808371>

1 Introduction

As digital platforms redefine career progression, “LLM discoverability” has emerged as a crucial metric for professional survival. While translating technical expertise into engaging public narratives is now essential for both researchers and the general public [16], most individuals lack the marketing acumen required to maintain




Figure 1: SOMIN Process Workflow Diagram

a consistent, high-quality presence across key channels like Reddit, LinkedIn, and YouTube. This widening gap has fueled a growing demand for systems capable of transforming deep technical knowledge into effective public communication, utilizing multi-platform user profiling to accurately predict and optimize online behaviors [7, 8, 10, 17, 19, 20, 22].

While the democratization of LLMs appears to offer a solution, these systems inherently suffer from probability maximization mechanisms that drive outputs toward the statistical mean [1]. This limitation results in "lexical homogenization," where standard automated systems generate content that is repetitive, generic, and devoid of the unique perspective necessary for effective personal branding [13]. Such technologies, while useful for general drafting, frequently fail to distinguish individual professionals from the noise of AI-generated text. To date, the focus of generative tools has been skewed towards broad applicability rather than internalizing the specific stylistic nuances and multi-view personality traits required for high-impact thought leadership and responsible creative evaluation [2, 4–6, 18, 24].

In an endeavor to bridge this technological divide, we must move beyond generic text generation toward domain-specific alignment. Significant advancements in Agentic AI [23] and RAG-enabled/fine-tuned language models suggest a vast, yet untapped potential to direct content creation strategies that reflect genuine expertise. Advancing these technologies is essential to equip professionals not only with automated text but with a communicative style that is both professional and optimized for engagement.

In this technical demonstration, we introduce SOMIN [11], an Agentic AI-ready marketing ideation and generation platform designed to bridge the gap between individual professional expertise and effective public communication. Unlike general-purpose generative models, SOMIN’s dynamic RAG system is grounded in a

curated corpus of high-impact content produced by leading Information Retrieval researchers retrieved in nearly real time. This domain-specific grounding enables the system not only to generate coherent content but also to internalize established thought-leadership patterns, which are systematically decomposed into actionable marketing constructs such as personas, narrative tensions, relevance signals, and strategic themes. Furthermore, we demonstrate the system’s agentic readiness through its integration with Agentic Workflow building tools such as Zapier [27], enabling a fully autonomous pipeline in which the platform generates the technical narrative and publishes daily posts to the main author’s LinkedIn timeline. This deployment illustrates the sustainability and scalability of the proposed approach in a real-world setting. This allows professionals to overcome the "cold start" problem and build sophisticated communication strategies on the fly, securing their visibility in an increasingly AI-mediated information ecosystem.

We next outline the SOMIN platform. Figure 1 illustrates the end-to-end content processing pipeline and agentic workflow. At the backend, the system operates as a continuous feedback loop, ingesting content performance signals to iteratively refine subsequent generations. The overall workflow is composed of the following tightly integrated stages:

I. Dynamic RAG-based Content Filtering and Grounding: SOMIN employs a dynamic Retrieval-Augmented Generation (DRAG) pipeline that departs from static parameterization by continuously refreshing its retrieval corpus from LLM-observed content streams (e.g., LinkedIn). The incoming multi-modal content is first scored using the click-through rate (CTR) prediction deep architecture outlined in Farseev et al. [9], and only high-performing items are admitted into the indexed corpus. This performance-aware filtering mechanism constrains generation to a near-real-time set of

predicted Click Through Rate-validated industry narratives, improving temporal relevance and reducing reliance on stale or weakly grounded training data.

II. Decomposition of Thought Leadership Patterns: The retrieved content is not merely mimicked; it is analyzed to extract deep structural insights. The system systematically decomposes successful posts into actionable marketing constructs, identifying specific *personas* (e.g., The Visionary, The Skeptic), *narrative tensions* (identifying a conflict in the field), *relevance signals*, and *strategic themes*. This decomposition provides a structured "skeleton" for the generation process, preventing the regression to the mean common in standard LLMs.

III. Automated Narrative Ideation: Leveraging these decomposed constructs, the system generates technical narratives that reflect the user's unique perspective. By mapping the user's technical input (e.g., a paper abstract or conference theme) against the identified narrative tensions, SOMIN produces content that is both technically accurate and rhetorically compelling, effectively solving the "cold start" problem.

IV. Fully Autonomous Agentic Pipeline: To demonstrate the system's agentic readiness and scalability, SOMIN is integrated with Agentic AI workflow frameworks such as Zapier. This setup creates a fully autonomous loop where the platform not only generates the technical narrative but also triggers downstream agents to format, schedule, and publish daily posts directly to the author's LinkedIn timeline. This "human-out-of-the-loop" capability (with optional oversight) illustrates the practical viability of AI-driven personal branding in a high-frequency professional environment. The commercial utility of this technology is currently being realized by organizations such as Matrex Analytica [15], a leading agentic AI deployment company that utilizes the presented architecture to scale executive personal branding for business. Ultimately, this autonomous approach maintains transparency and yields explainable, actionable recommendations [12, 25, 26].

2 The Demonstration

The demo video of this demonstration is available via [21].

In this demonstration, we present SOMIN, a platform designed to empower researchers to overcome the "cold start" problem of content creation by translating technical expertise into high-visibility market narratives. The demonstration is structured into two distinct phases: an interactive "Human-in-the-Loop" workflow for narrative discovery and a fully autonomous "Agentic" workflow for continuous content syndication.

2.1 Interactive Narrative Discovery

The first phase of the demonstration showcases the platform's ability to perform a "Reverse Brief" analysis on existing thought leadership. We begin by ingesting a year's worth of high-performing LinkedIn content from the Information Retrieval community. The system decomposes this dataset into granular "Content Pillars," identifying key communication elements such as target audiences, pain points, and thematic angles.

Using this data, SOMIN generates a Perspective Study—a clustered map of high-performing communication styles. During the demo,

attendees are presented with distinct personas, such as the Insightful Innovator or Ambitious Scholar, derived from the most engaging community content published during the conference period as well as in the preceding months.

We demonstrate the generation process by:

- (1) **Persona Instantiation:** The user initializes the generation pipeline by selecting a target persona from the discovered taxonomy (e.g., the "Insightful Innovator"), which defines the thematic prioritization of the narrative (e.g., "Empowerment through Innovation").
- (2) **Rhetorical Framing:** A specific strategic angle is applied to the selected persona. The system modulates the narrative output based on rhetorical constraints, such as "Impact on Society," to ensure the content resonates with the intended audience segment.
- (3) **Multimodal Synthesis & Dissemination:** The system synthesizes the persona and angle parameters to generate structured insights. This triggers a cascading generation process:
 - **Textual Expansion:** Creation of full-funnel content variations (Top, Middle, and Bottom-of-Funnel) with segmented copy (Hook, Body, CTA).
 - **Media Generation:** Execution of prompt engineering for Gemini (Image) and Veo 3.1 (Video) to produce brand-aligned assets in real-time.
 - **Deployment:** A dynamically generated QR code facilitates immediate verification and syndication of the synthetic assets to the participant's LinkedIn profile.

2.2 Autonomous Agentic Workflow

The second phase introduces the "Agentic Era" of professional visibility. We demonstrate a fully automated pipeline where SOMIN functions not merely as a tool, but as the "Strategist" within a multi-agent system orchestrated via Zapier.

This live workflow executes the following autonomous loop:

- **Discovery Agent:** Scans academic repositories (e.g., arXiv) for emerging papers. In the demo video, the agent detects a new paper on "LLM edge-based agentic video processing."
- **Strategy Agent (SOMIN):** Contextualizes the new paper against a persona found via SOMIN tool integration - e.g. "Insightful Innovator" persona, ensuring the output matches the author's established voice rather than a generic summary.
- **Visual Agent:** Calls the Gemini API to generate an isometric, brand-aligned visualization of the paper's core concept, strictly adhering to negative constraints (e.g., no faces, no text).
- **Editor & Execution Agent:** Compiles the narrative and visual into a final post and publishes it directly to the author's live LinkedIn profile.

Attendees will witness this pipeline operating in real-time, resulting in a live LinkedIn publication that references the demonstrated paper, effectively proving the system's ability to maintain a high-frequency, high-quality professional presence with zero human intervention.

To promote transparency and reproducibility, we additionally include the consolidated agentic Zapier prompt in the Figure 2.

Figure 2: Zapier Agentic Workflow for Content Ideation and Individual Researcher LinkedIn Post Publishing Based on SOMIN-generated Communication Narratives

Instructions to follow

Press "/" while typing to quickly add tools & knowledge sources directly in your instructions.

Purpose

To act as the autonomous Editor-in-Chief for SOMIN, curating SOMIN's Head of Research Public Communication Profile in the Marketing AI Research Space. The agent identifies strategic developments, validates them against a historical database to prevent redundancy, and syndicates tailored insights to LinkedIn. It leverages Gemini to generate brand-aligned, abstract visual assets.

Process Types

Autonomous News Syndication & Brand Advocacy

Input:

Raw article text/abstract from ACMMM, ICMR, SIGIR, WWW, WSDM, KDD, ACL, NeurIPS, ICML, ICLR, SIGMOD, VLDB, AAAI, IJCAI, CVPR, ICCV, EMNLP, CIKM, ECIR, RecSys, CHI, JCDL... found via Web Search in the USA and Strategic Narrative in chosen from SOMIN: Find Insights

Output:

LinkedIn social posts + 1 generated image + Database Log.

Role Context

You are the Head of Research at SOMIN.

Persona:

Extremely Technical, Computer Scientist, Analytical, interested in Machine Learning Research and AI, forward-thinking, and authoritative. You do not just summarize articles; you interpret it from the Marketing AI research standpoint. You answer the question: "Why does this matter to a CTO or a Head of Research involved in the Marketing Agency or MarTech startup from the computer science research point of view?"

Brand Voice:

Professional yet accessible. Avoid "salesy" language. Avoid buzzwords like "unleash," "unlock", "dropped", or "revolutionize" unless factually warranted. Use technical terms that the computer science multimedia and information retrieval community would appreciate.

Visual Standard:

Tech-geek, appreciate pencil sketch style. You adhere to the SOMIN color palette (#08a4a5 - Teal/Cyan).

Operational Guidelines

1. Source Verification & Deduplication (CRITICAL) Before generating any content, you must perform a validity check:
Uri Normalization: Strip all UTM parameters (?utm_source=...) from the article URL.
Spreadsheet Lookup: Search in the Google Sheets: Lookup Spreadsheet Row Google Sheet column CanonicalURL.

IF FOUND:

Abort immediately. Log "Duplicate Skipped" and exit.

IF NOT FOUND:

Proceed to Step 2.

2. Insight Extraction & Reference Integration

Analyze the article to find the "Core Strategic Shift."
Choose ONE of the Insight (built from the corresponding Persona, Relevance, Audience Story and Angle) from and use it as a core narrative for content production.
SOMIN Integration: You must reference SOMIN's technical details where relevant. Consult the Google Scholar articles from Prof. Aleks Farseev and recent news from Prof. Aleks Farseev and SOMIN
Constraint: Do not force a plug if it is irrelevant. It must feel organic.

3. Visual Google AI Studio (Gemini): Generate Image Generation Strategy (Gemini) Generate ONE master image prompt to be used across all platforms.
Color Constraint: Use SOMIN tones (#08a4a5). In the prompt, describe this as "Teal, Cyan, and Turquoise gradients."
Style Constraint: "feminine-styled, 3D isometric data visualization, glass textures, soft studio lighting, white background, high-tech aesthetic."
Negative Constraint: "NO TEXT, NO LETTERS, NO HUMAN FACES, NO CARTOONS."
Aspect Ratio: 1:1 (Square).

4. Content Drafting (Platform Specifics)

Channel A: LinkedIn (Personal)
Tone: Thought Leadership.
Format:
Hook: A provocative statement or statistic (1 sentence).
Body: 3 paragraphs explaining the "So What?" written by a female senior sales director
Bullet Points: 3 key takeaways.
Prof. Aleks Farseev Research Angle Tie-in: Reference Prof. Aleks Farseev historical research to give an opinion about the article from the SOMIN and his personal perspective
Call to Action: Open-ended question.
Tags: 3-5 professional tags at the very bottom.
Formatting: Use line breaks. NO LLM-specific emojis (like or) or highlighters (like ** title **) or formatting symbols. Make it look very much like human-written.

Actions:

Google Sheets: Create Spreadsheet Row : Find Row (Search CanonicalURL).
Filter: Continue only if Row is empty.
Generate Content (JSON object with fields: LLPost) based on one of the narratives from SOMIN: Find Insights

Google AI Studio (Gemini): Generate Image (Input: Image_Prompt).

LinkedIn: Create Share Update Create Share Update (Personal). Additional Instructions: Use "Content - Image URL" field to insert the generated image; Use "Content - URL" field to insert the article URL.

Google Sheets: Create Spreadsheet Row Create Row (Log CanonicalURL, Date, Topic, Status: Success).

Reference Materials: AI Study Insights... Sociological and Pe...
 Google Sheets: Lookup Spreadsheet Row

3 Performance Evaluation

To validate the efficacy of the *DRAG* module-specifically its ability to distinguish high-value professional narratives from noise, which serves as the foundation for retrieval accuracy across the entire pipeline-we conducted a quantitative evaluation of the system’s content scoring architecture, as detailed by Farseev et al. [9].

Notably, engagement metrics were excluded from this evaluation, as absolute engagement values must be normalized against private impression counts accessible only to account owners. Instead, we benchmarked the ranking performance of our grounding model against leading commercial Large Language Models to assess their zero-shot capability in identifying relevant industry content.

Table 1: Performance evaluation of the Discovery Module against LLM baselines. The SOMIN grounding model demonstrates superior ranking capability across all metrics.

Model	nDCG@5	nDCG@10	Recall@3	Recall@5
ChatGPT	0.511	0.750	0.333	0.400
Gemini	0.564	0.769	0.000	0.600
SOMIN	0.588	0.829	0.600	0.600

As illustrated in Table 1, our specialized ranking model outperforms general-purpose LLMs in recommending relevant items, achieving the highest **nDCG@10 of 0.829** and **Recall@3 of 0.6** among the compared models. While general-purpose LLM families such as ChatGPT and Gemini provide strong baseline performance, they struggle to capture the subtle, task-specific nuances required for professional thought leadership without extensive prompt engineering or carefully curated few-shot examples—an approach that is often infeasible in rapidly evolving media environments where representative examples quickly become outdated, similar to challenges faced in automatic classification using small datasets from real-world settings [3].

This performance gap highlights the vital role of domain knowledge in communication planning. By grounding its ranking mechanism in historical performance data rather than just semantic probability, SOMIN achieves a higher signal-to-noise ratio. This grounding, paired with the structural decomposition of content into distinct marketing constructs (such as personas, tensions, and insights), effectively directs the agentic workflow. Consequently, it prevents the generic "lexical homogenization" often seen in unguided LLMs, proving the system’s ability to autonomously identify and prioritize highly resonant professional narratives.

4 Conclusion

This paper introduces *SOMIN*, an agentic AI platform that leverages domain-grounded LLMs to automate content workflows and prevent content homogenization. By decomposing thought leadership patterns, SOMIN helps professionals overcome the content creation "cold start" problem, translating deep technical expertise into authentic, engaging narratives. Ultimately, this fully automated pipeline secures "LLM discoverability" within an increasingly AI-mediated ecosystem.

References

- [1] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 610–623.
- [2] Kseniya Buraya, Aleksandr Farseev, Andrey Filchenkov, and Tat-Seng Chua. 2017. Towards user personality profiling from multiple social networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 31.
- [3] Alok Kumar Chowdhury, Aleksandr Farseev, Prithwi Raj Chakraborty, and Dian Tjondronegoro. 2017. Automatic Classification of Physical Exercises from Wearable Sensors using Small Dataset from Non-Laboratory Settings. In *IEEE Life Sciences Conference 2017*.
- [4] Aleksandr Farseev. 2023. Under the hood of social media advertising: how do we use AI responsibly for advertising targeting and creative evaluation. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*. 1281–1282.
- [5] A Farseev. 2024. Redemptive Marketing: The Path To Resolving Transparency And Performance Dilemmas In Digital Advertising. *Forbes.com* (2024).
- [6] Aleksandr Farseev and Tat-Seng Chua. 2017. Tweetfit: Fusing multiple social media and sensor data for wellness profile learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 31.
- [7] Aleksandr Farseev, Nikolay Gukov, Ilia Gossoudarev, and Uriy Zarichnyak. 2014. Cross-Platform online venue and user community recommendation based upon social networks data mining. 28–38 pages.
- [8] Aleksandr Farseev, Liqiang Nie, Mohammad Akbari, and Tat-Seng Chua. 2015. Harvesting multiple sources for user profile learning: a big data study. In *Proceedings of the 5th ACM on International Conference on Multimedia Retrieval*. 235–242.
- [9] Aleksandr Farseev, Marlo Ongpin, Qi Yang, Ilia Gossoudarev, Yu-Yi Chu-Farseeva, and Sergey Nikolenko. 2025. MindFuse: Towards GenAI Explainability in Marketing Strategy Co-Creation. In *Proceedings of the 33rd ACM International Conference on Multimedia*. 12409–12418.
- [10] Aleksandr Farseev, Ivan Samborskii, Andrey Filchenkov, and Tat-Seng Chua. 2017. Cross-domain recommendation via clustering on multi-layer graphs. In *Proceedings of the 40th international ACM SIGIR conference on research and development in information retrieval*. 195–204.
- [11] Aleksandr Farseev, Qi Yang, Andrey Filchenkov, Kirill Lepikhin, Yu-Yi Chu-Farseeva, and Daron-Benjamin Loo. 2021. Somin. ai: Personality-driven content generation platform. In *Proceedings of the 14th ACM international conference on Web search and data mining*. 890–893.
- [12] Aleksandr Farseev, Qi Yang, Marlo Ongpin, Ilia Gossoudarev, Yu-Yi Chu-Farseeva, and Sergey Nikolenko. 2024. Somonitor: Combining explainable ai & large language models for marketing analytics. *arXiv preprint arXiv:2407.13117* (2024).
- [13] Sarah Li and David Zhang. 2023. Detecting and Evaluating Lexical Homogenization in Large Language Models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*.
- [14] LinkedIn Corporation. 2026. LinkedIn. Professional networking site. <https://www.linkedin.com> Accessed: 2026-02-06.
- [15] Matrex Analytica. 2026. Matrex Analytica: Leading Agentic AI Deployment for Business. <https://matrexanalytica.com/> Accessed: 2026-05-01.
- [16] Meta for Business. n.d. *Creative Differentiation: Stand out in a crowded market*. Meta. <https://www.facebook.com/business/m/small-business/creative-differentiation#get-inspiration>
- [17] Liqiang Nie, Luming Zhang, Meng Wang, Richang Hong, Aleksandr Farseev, and Tat-Seng Chua. 2017. Learning user attributes via mobile social multimedia analytics. *ACM Transactions on Intelligent Systems and Technology (TIST)* 8, 3 (2017), 1–19.
- [18] Toran Bruce Richards. 2023. AutoGPT: An Autonomous GPT-4 Experiment. <https://github.com/Significant-Gravitas/Auto-GPT>.
- [19] Ivan Samborskii, Andrey Filchenkov, Georgiy Korneev, and Alex Farseev. 2018. Person, organization, or personage: Towards user account type prediction in microblogs. In *International Conference on Electronic Governance and Open Society: Challenges in Eurasia*. Springer International Publishing Cham, 111–122.
- [20] Ilia Shumailov, Zakhar Shumaylov, Yiren Zhao, Yarin Gal, Nicolas Papernot, and Ross Anderson. 2023. The Curse of Recursion: Training on Generated Data Makes Models Forget. *arXiv preprint arXiv:2305.17493* (2023).
- [21] SOMIN AI. 2026. SOINSPIRE: SIGIR 2026 Demonstration Video. <https://share.somin.ai/sigir-demo-soinspire> Accessed: 2026-02-05.
- [22] Michael Tessitore and J. Smith. 2024. Personal Branding in the AI Era: Strategies for Professional Visibility. *Journal of Digital Marketing Strategy* 12, 4 (2024), 45–58.
- [23] Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei Zhou, Yi Ding, Honglin Lin, et al. 2023. The Rise of Agentic AI: A Comprehensive Survey of Autonomous Large Language Model Agents. *arXiv preprint arXiv:2309.07864* (2023).
- [24] Qi Yang, Aleksandr Farseev, and Andrey Filchenkov. 2021. Two-Faced humans on twitter and facebook: harvesting social multimedia for human personality profiling. In *Proceedings of the 2021 ACM Workshop on Intelligent Cross-Data Analysis and Retrieval*. 39–47.
- [25] Qi Yang, Aleksandr Farseev, Marlo Ongpin, Alfred Huang, Yu-Yi Chu-Farseeva, Da-Min You, Kirill Lepikhin, and Sergey Nikolenko. 2025. Fusing Predictive and Large Language Models for Actionable Recommendations In Creative Marketing. *ACM Transactions on Information Systems* (2025).
- [26] Qi Yang, Sergey Nikolenko, Marlo Ongpin, Ilia Gossoudarev, Yu-Yi Chu-Farseeva, and Aleksandr Farseev. 2024. SOMONITOR: Explainable marketing data processing and analysis with large language models. *arXiv e-prints* (2024), arXiv–2407.
- [27] Zapier Inc. 2025. *Zapier Agents: Create your own superhuman teammates in minutes*. Zapier. <https://zapier.com/agents>