

Trace-Aware Workflows for Co-Creating Branded Content with Generative AI

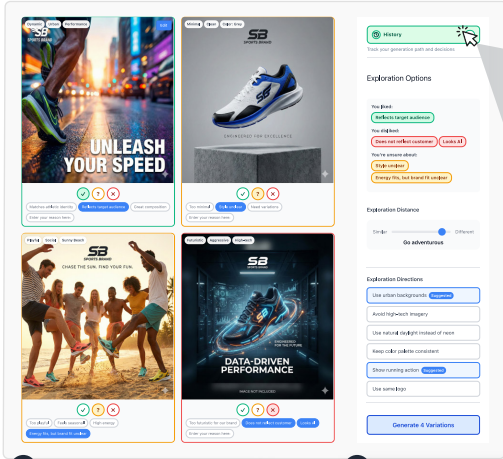
TAEHYUN YANG, University of Maryland, USA

EUNHYE KIM, KAIST, Republic of Korea

ZHONGZHENG XU, University of Maryland, USA

FUMENG YANG, University of Maryland, USA

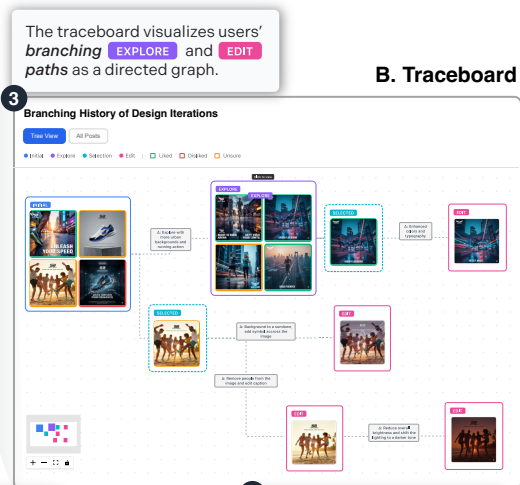
A. Exploration Workspace



1 Four candidate images: users provide **feedback** (like/dislike/unsure) and **rationale** to express their preferences.

2 Exploration control: users adjust a **novelty slider** and select **exploration directions** to steer the next set of image generations.

B. Traceboard



3 The traceboard visualizes users' **branching EXPLORE** and **EDIT** paths as a directed graph.

4 We also provide Δ summaries to help users make sense of the changes between image iterations.

Fig. 1. Overview of our prototype for branded content creation: (A) a creative exploration workspace, where users generate, evaluate, and refine branded images, and (B) a traceboard, which traces and visualizes the history of image generations and edits. In the exploration workspace, 1 users provide lightweight feedback (like/dislike/unsure along with short rationales) to express preferences towards individual images, and 2 use AI-guided exploration controls (suggested directions and a novelty slider) to steer the next set of image generations. The traceboard 3 visualizes users' branching exploration and edit paths as a directed graph while 4 change summaries (Δ) form the edges of the graph describing what was modified between iterations, helping users make sense of the changes between iterations.

Generative AI tools have lowered barriers to producing branded social media images and captions, yet small-business owners (SBOs) still struggle to create on-brand posts without access to professional designers or marketing consultants. Although these tools enable fast image generation from text prompts, aligning outputs with a brand's intended look and feel remains a demanding, iterative creative task. In this position paper, we explore how SBOs navigate iterative content creation and how AI-assisted systems can support SBOs' content creation workflow. We conducted a preliminary study with 12 SBOs who independently manage their businesses and social media presence, using a questionnaire to collect their branding practices, content workflows, and use of generative AI alongside conventional design tools. We identified three recurring challenges: (1) translating brand "feel" into effective prompts, (2) difficulty revisiting and comparing prior image generations, and (3) difficulty making sense of changes between iterations to steer refinement. Based on these findings, we present a prototype that scaffolds brand articulation, supports feedback-informed exploration,

Authors' Contact Information: Taehyun Yang, University of Maryland, College Park, USA, taeyang@umd.edu; Eunhye Kim, KAIST, Daejeon, Republic of Korea, gracekim027@kaist.ac.kr; Zhongzheng Xu, University of Maryland, College Park, USA, zxu169@umd.edu; Fumeng Yang, University of Maryland, College Park, USA, fy@umd.edu.

and maintains a traceboard of branching image iterations. Our work illustrates how traces of the iterative process can serve as workflow support that helps SBOs keep track of explorations, make sense of changes, and refine content.

CCS Concepts: • **Human-centered computing** → **Human computer interaction (HCI)**.

Additional Key Words and Phrases: Creativity support tools, Content creation, Exploration traces, Generative AI

1 Introduction

Small businesses comprise roughly 90% of businesses worldwide [10, 13]. For a local bakery owner, platforms like Instagram, Facebook, and LinkedIn have become the first place a potential customer encounters their brand, forms an impression, and decides whether to engage. The content shared on these social media platforms, such as promotional images, branded graphics, and visual assets, can directly influence brand reputation and influence purchasing [4, 9]. Therefore, a consistent, professional presentation is a competitive necessity for these social media contents.

Yet most small business owners (SBOs) struggle to meet this expectation. Unlike large enterprises with dedicated marketing teams, SBOs typically operate without the time, budget, or design expertise needed to produce professional visual content [1]. Tasks that experienced designers and brand strategists handle systematically fall entirely on the business owner themselves, including defining a coherent visual direction, maintaining aesthetic consistency across posts, and ensuring every asset feels “on brand.”

Recent generative AI tools such as Midjourney [7] and DALL-E [8], present both opportunities and challenges for SBOs. They have enabled non-experts to produce visual content through natural language prompts [5, 12, 14], showing promising results in generating complex, high-quality visuals, and many SBOs have adopted them hoping to reduce production costs and time [3, 6, 11]. However, creating effective visual content is rarely a one-shot process. Due to the inherent imprecision of generative AI outputs and the difficulty of articulating intents, SBOs typically iterate extensively. However, current generative AI tools offer little support for this iterative process, providing no structure for organizing explorations, externalizing brand preferences, or converging on a coherent visual identity over time. This very gap presents an opportunity for specialized interfaces that scaffold the iterative creation process and help SBOs articulate and refine their brand identity.

In this paper, we propose a system that preserves exploration history and makes iterations easy to compare. By making prior attempts and their outcomes visible, the design helps SBOs avoid losing promising directions and assess whether each step moves closer to their intended brand direction. We first conducted a questionnaire with 12 SBOs to understand the challenges they face during their process of creating social media content. Guided by the insights discovered, we introduce a prototype system that scaffolds brand articulation, preserves branching exploration histories, and makes generative changes interpretable through structured creativity traces. Through this work, we contribute a trace-aware approach to support users working with generative systems and highlight how interaction traces can be served as creative scaffolds for understanding iterative design practices.

2 Preliminary Study

We first conducted a formative study to understand the specific challenges SBOs face when creating social media content. We recruited 12 SBOs who independently operated their businesses and administered a questionnaire on their content workflows, definition of brand identity, and tool use with generative AI systems alongside conventional design tools (e.g., Illustrator, Figma, Canva, etc.). This study was approved by our IRB with participant demographics available

in Table 1 (Appendix A). After thematic analysis, we identify three challenges, presented as follows.

Participants struggled to translate implicit brand knowledge into effective prompts (C1). Business owners often had a well-formed sense of how their brand should feel but lacked the vocabulary or design expertise to express it in actionable terms (e.g., visuals, tone, composition). This knowledge was experience-based, and articulating it required additional effort within already time-constrained workflows. P12 shared: *“I didn’t know how to put that warmth into words, so I ended up not posting it.”* Because this translation gap prevents SBOs from effectively prompting AI tools, systems should provide scaffolds and examples that help convert experience-based intuitions into prompts.

Participants handled end-to-end content production as an iterative but poorly structured process (C2). SBOs repeatedly cycled through generating, revising, and evaluating content. Very often, their ideas evolved across tools but intermediate artifacts were easily lost, making prior directions difficult to revisit or compare. P3 explained: *“I move between three apps to see which version is better. It’s a hassle to try different captions and images across tools, and it’s hard to keep track.”* Because this lack of structure makes it difficult to evaluate progress or return to other design alternatives, systems should maintain persistent records of exploration.

Finally, **SBOs struggled to judge whether iterations were moving closer to their intended style and tone (C3).** Because they lacked formal design training and worked without experts, participants found it difficult to evaluate whether each new iteration moved closer to their intended style and tone. Small prompt edits could produce large changes in style or tone without clear explanations, making progress difficult to evaluate or steer. P4 noted: *“...the AI takes a giant leap when I only asked for a small step. I’ll change one word and the whole vibe flips. I can’t tell if I’m getting closer.”* As such, systems should make changes between iterations more explicit to support the content exploration and selection.

3 Prototype Features

Based on the findings, we discuss four features to better support SBOs in creating social media content with generative AI tools. We implemented these features in a prototype system (see Figure 1) using Gemini 3 Pro Image Preview (“Nano Banana”) for image generation, GPT-4o-mini for prompt auto-suggestions and edit recommendations, and GPT-4o for generating natural-language summaries of changes between iterations.

Feature 1. Scaffolding Brand Knowledge Through Structured Prompting. Users begin by entering core brand information, including the brand name, category, and a textual brand description. To help SBOs express their brand intuitions (C1) in prompts, the system provides AI-assisted autocomplete that helps SBOs externalize knowledge without requiring professional marketing knowledge (Figure B.1). Autocomplete operates in two modes. First, at sentence onset, the system detects missing brand dimensions using a schema inspired by Kapferer’s Brand Identity Prism [2] and offers scaffolding topics that frame brand knowledge in accessible terms (e.g., when information about uniqueness is absent, the system suggests “Something that makes [brand name] different from competitors is...”). Second, during mid-sentence pauses, the system offers context-sensitive phrase completions that help SBOs articulate implicit feelings (e.g., after “When customers interact with my brand, I want them to feel...,” the system proposes emotional descriptors based on prior context). A hoverable indicator visualizes which portions of the description correspond to specific identity facets, helping users identify gaps in articulation and surface overlooked aspects of their brand understanding.

Feature 2. Translating Brand Judgments into Visual Directions. Articulating brand knowledge, however, does not guarantee visually aligned outputs. To help SBOs translate brand descriptions into visuals that match their intended feel (C1), the system supports iterative evaluation and refinement of AI-generated images. After the brand description is completed, the system generates four candidate images conditioned on the provided information. Users evaluate each image through their existing brand knowledge, providing lightweight feedback (like, dislike, unsure;

see Figure 1A) accompanied by automatically inferred rationales from GPT-4o (e.g., “Too futuristic for our brand”). Users may revise these rationales using their own language. Selecting *Generate 4 Variations* generates new images that incorporate this feedback, with a novelty control allowing users to balance between diverse exploration and incremental refinement (Figure 1A). When users identify a promising direction, they can further refine images through prompt-based regeneration with suggested edit instructions or region-specific modifications.

Feature 3. Preserving Process Visibility Through Creativity Traces. To address the lack of structure in SBOs’ iterative workflows (C2), the system provides a centralized traceboard (Figure 1B) that maintains a record of the entire content development process. The traceboard records all generation and refinement actions (e.g., *Explore*, *Edit*, *Selection*) as creativity traces, visualized as a directed graph in which nodes represent artifacts (images) and edges represent transformations (prompt changes, feedback updates, edits). This structure preserves branching exploration trajectories from the initial set, allowing users to see how their content evolved over time rather than losing track across iterations. Users can select any node to revisit prior states, compare alternatives side-by-side, or resume exploration from earlier points without losing progress. When users identify a promising direction but later want to explore a different approach, they may revise feedback at any earlier node to regenerate variants from that branch while maintaining the full context of what they’ve already tried.

Feature 4. Supporting Solo Creative Exploration Through AI Guidance. To help SBOs working independently understand AI-generated changes and steer refinement (C3), the system offers creative suggestions and makes generative changes visible. First, the system proactively suggests exploratory directions in the *Edit* function, recommending targeted edit prompts (e.g., “try a warmer color palette”) that guide refinement without requiring design vocabulary (Figure B.2). These suggestions help users quickly specify refinement directions without needing design expertise or extensive prompt iteration. Second, the system makes each iteration understandable by revealing what changed between generations. For each generated image, the system extracts descriptive keywords that summarize visual attributes in natural language (Figure 1A). When users provide feedback, the system guides subsequent *Explore* generations toward preferred characteristics while avoiding undesired ones. The system then summarizes differences between successive outputs and attaches these natural language change descriptions to the corresponding traceboard nodes (Figure 1B), helping users evaluate brand fit.

4 Workshop Alignment and Future Work

Our work aligns with the workshop’s goal of advancing trace analysis by proposing how interaction traces for AI-assisted content creation can be captured beyond simple prompt histories. Tracing interactions that were previously unrecorded and transient, we represent SBOs’ iterative processes as structured traces—recording feedback (and rationales), edits, branching alternatives, and natural language summaries of change—to reflect how SBOs refine visual direction and assess brand fit over time. Our traceboard organizes these traces as a directed graph of generated artifacts and transformations, enabling inspection of exploration paths and iteration-to-iteration changes.

In doing so, we hope to provide a concrete trace representation that can serve as analyzable data for studying creative activity traces. The traceboard supports analysis at multiple levels, from local edits between single images to higher-level patterns across multiple image sets. In follow-up work, we plan to trace and analyze additional signals within these workflows (e.g., feedback rationales, edit types, and branching strategies) and examine how trace patterns (e.g., branching depth, revision frequency, and feedback granularity) relate to downstream outcomes such as perceived content quality and audience response to the resulting branded images.

References

- [1] Philip Alford and Stephen John Page. 2018. Marketing technology for adoption by small business. In *Social Media and Interactive Communications*. Routledge, 99–113. doi:10.1080/02642069.2015.1062884
- [2] Jean-Noël Kapferer. 2008. *The New Strategic Brand Management: Creating and Sustaining Brand Equity Long Term*. Kogan Page, London, UK.
- [3] Nimisha Karnatak, Adrien Baranes, Rob Marchant, Triona Butler, and Kristen Olson. 2025. ACAI for SBOs: AI Co-creation for Advertising and Inspiration for Small Business Owners. arXiv. arXiv:2503.06729 doi:10.48550/arXiv.2503.06729
- [4] Sascha Kraus, Johanna Gast, Moritz Schleich, Paul Jones, and Michael Ritter. 2019. Content is king: How SMEs create content for social media marketing under limited resources. *Journal of Macromarketing* 39, 4 (2019), 415–430. doi:10.1177/0276146719882746
- [5] Jie Li, Hancheng Cao, Laura Lin, Youyang Hou, Ruihao Zhu, and Abdallah El Ali. 2024. User experience design professionals’ perceptions of generative artificial intelligence. In *Proceedings of the 2024 CHI conference on human factors in computing systems*. 1–18. doi:10.1145/3613904.3642114
- [6] Xuye Liu, Annie Sun, Pengcheng An, Tengfei Ma, and Jian Zhao. 2025. Influencer: Empowering Everyday Users in Creating Promotional Posts via AI-infused Exploration and Customization. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–19. doi:10.1145/3706598.3713309
- [7] Midjourney. [n. d.]. Midjourney. <https://www.midjourney.com/>.
- [8] OpenAI. [n. d.]. DALL-E. <https://openai.com/dall-e>.
- [9] Heini Maarit Taiminen and Heikki Karjaluo. 2015. The usage of digital marketing channels in SMEs. *Journal of Small Business and Enterprise Development* 22, 4 (2015), 633–651. doi:10.1108/JSBED-05-2013-0073
- [10] United Nations. 2025. Micro-, Small and Medium-sized Enterprises Day. <https://www.un.org/en/observances/micro-small-medium-businesses-day> Accessed: 2026-02-07.
- [11] Verizon Business. 2025. 2025 State of Small Business Survey: Surge in AI, Cybersecurity and Social Media Demand. <https://www.verizon.com/about/news/2025-state-of-small-business-survey>. Accessed: 2026-02-06.
- [12] Justin D Weisz, Jessica He, Michael Muller, Gabriela Hoefler, Rachel Miles, and Werner Geyer. 2024. Design principles for generative AI applications. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–22. doi:10.1145/3613904.3642466
- [13] World Bank. 2025. Small and Medium Enterprises (SMEs) Finance. <https://www.worldbank.org/en/topic/sme/finance> Accessed: 2026-02-07.
- [14] J Diego Zamfirescu-Pereira, Richmond Y Wong, Bjoern Hartmann, and Qian Yang. 2023. Why Johnny can’t prompt: how non-AI experts try (and fail) to design LLM prompts. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–21. doi:10.1145/3544548.3581388

Appendices

A Participant Demographics

Table A.1. Participant demographics in our preliminary study. Business age indicates how long participants have operated their brand. Abbreviations: IG=Instagram, YT=YouTube, TT=TikTok, LI=LinkedIn.

ID	Role	Business age	Industry/Product	Platforms
P1	Aspiring entrepreneur	6 mo	Fashion	IG, YT
P2	Small-business owner	2 yr	Food Business	IG, Blog
P3	Side-business owner	3 yr	Publishing / branding	IG, TT
P4	Side-business owner	4+ yr	Design consulting	IG, LI
P5	Small-business owner	3 yr	Food Business	IG, Kakao
P6	Side-business owner	2 yr	Fashion (sportswear)	IG, YT
P7	Small-business owner	3 yr	Education	IG, Blog
P8	Small-business owner	1 yr	Fashion	IG, Blog, Kakao
P9	Side-business owner	1 yr	Beauty	IG, Kakao, Band
P10	Small-business owner	2 yr	Cafe	IG, Blog
P11	Freelancer	1 yr	Music production	IG
P12	Small-business owner	6 mo	Cafe	IG, TT

B Additional System Details

B.1 Prompt Autocomplete

The Prompt Autocomplete system scaffolds brand articulation during content setup. It suggests structured prompts based on missing brand dimensions and provides context-aware phrase completions to help users express brand attributes in actionable terms.

Tell us about your brand
Start with the basics, then describe it in your own words

Brand Name * Industry *

Brand Description * (Use AI autocomplete to help craft your story)

What to include:

- What your brand does and who it's for
- The problem you solve or need you address
- What makes your brand unique and special
- The feeling or values you want to convey

Sports Brand is a sports and wellness brand that creates functional, high-quality apparel designed to support active lifestyles in everyday settings.

→ Our audience includes athletes. → We solve issues with comfort.
→ Our mission is holistic health. → We empower everyone to move.

Great description! Ready to continue 150 characters

All requirements met [Next](#)

Fig. B.1. Prompt autocomplete system. The interface suggests structured brand dimensions and context-aware phrase completions to help users articulate brand intent.

B.2 Edit Interface

The Edit interface supports targeted refinement of generated images after an initial exploration round. Users can apply suggested edit prompts (e.g., color adjustments, stylistic shifts), modify text instructions directly, or regenerate specific regions of an image.

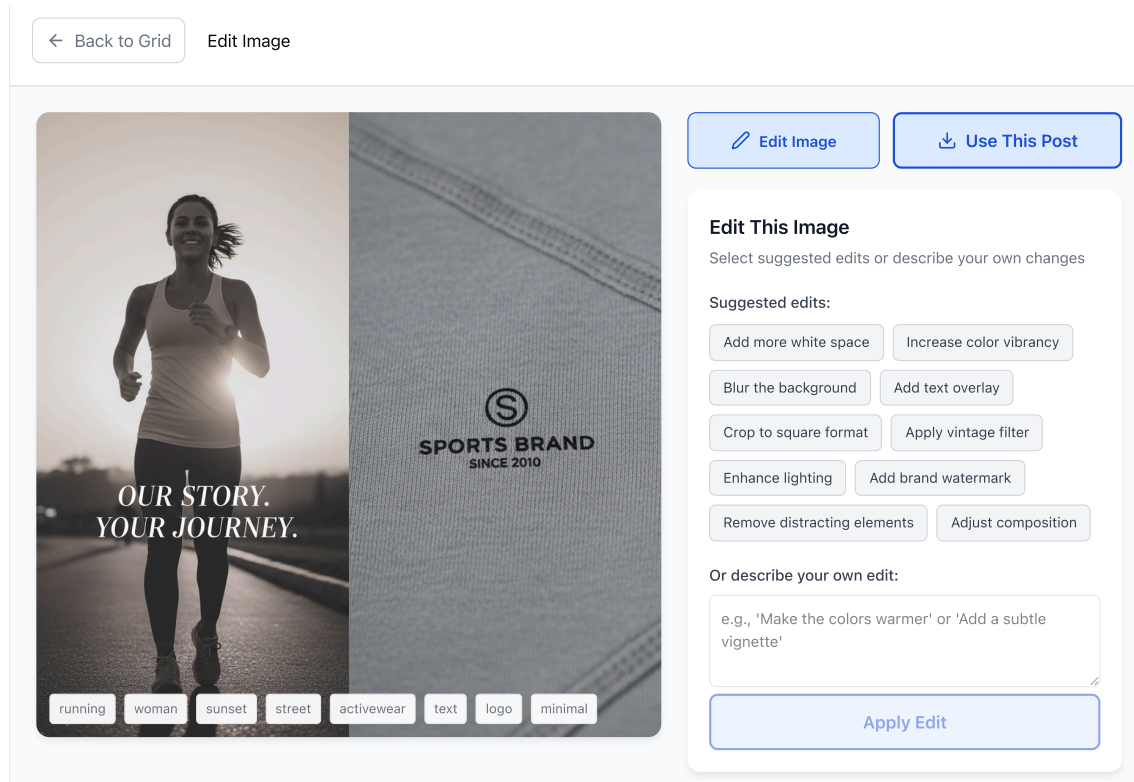


Fig. B.2. Edit interface for targeted refinement. Users can apply AI-suggested edit directions or modify prompts directly; each edit creates a new branch in the traceboard.