When AI Designs Alike

: How the Midjourney Shock Standardized Trademark across Firms



Suhyoung An, Hyo Kang

Strategy and International Management
The Graduate School of Business
Seoul National University

Abstract

Suhyoung An
Strategy and International Management
The Graduate School of Business
Seoul National University

The study examines how the diffusion of generative artificial intelligence reshapes firms' visual identity formation and induces market-level convergence in trademark design. Theoretically, the paper argues that generative image models alter the structure of creative production by shifting design processes from manual construction to generation-selection workflows, while diffusion-based sampling biases systematically pull outputs toward high-probability aesthetic regions. These algorithmic forces interact with cognitive mechanisms—such as processing fluency and conformist selection—to create a self-reinforcing feedback loop that amplifies visual homogenization.

Empirically, the study uses the universe of U.S. trademark applications filed between 2019 and 2024 and measures visual similarity using CLIP-based image embeddings. Exploiting the public release of Midjourney in July 2022 as a plausibly exogenous technological shock, the analysis employs a repeated cross-sectional design to trace changes in the distribution of trademark designs. The

results show a significant increase in visual similarity among trademarks after the diffusion of generative AI. Moreover, the homogenization effect is more pronounced for intent-to-use filings and in highly regulated industries, consistent with mechanisms of algorithmic convergence and institutional constraint.

This research contributes to the strategy and innovation literature by providing large-scale evidence that generative AI reshapes creative output distributions at the market level rather than merely improving individual productivity. By integrating technological and institutional mechanisms, the paper highlights how AI adoption can unintentionally erode visual differentiation and alter the dynamics of symbolic competition among firms.

Keyword : Generative AI, Trademark Design, Visual Homogenization, Algorithmic Convergence, Institutional Legitimacy, Creative Diversity
Student Number : 2024-24572

Table of Contents

1. Introduction 1

2. Theoretical Background and Hypotheses 5

2.1 Visual Homogenization in Trademark Design 6

2.2 Technological Affordances 9

2.3 Algorithmic Convergence and Regression to the Mean..... 12

2.4 Illusion of Fluency and Algorithmic Conformity..... 14

2.5 Moderating Role of Intent-to-Use (ITU) Filings 17

2.6 Moderating Role of Trademark Regulation 19

3. Methodology 21

3.1 Data Sources and Samples..... 21

3.2 Variables 23

3.3 Emperical Strategy 29

4. Result.....30

4.1 Descriptive Statistics and Correlations 32

4.2 Baseline Effect of Generative AI 38

5. Discussion and Contribution 38

6. Limitations and Future Research 41

References.....45

국문 초록 47

1. Introduction

Generative artificial intelligence has rapidly transformed how visual artifacts are produced, diffused, and evaluated. Text-to-image systems such as Midjourney, Stable Diffusion, and DALL · E 2 allow users to generate high-quality images from natural language prompts at negligible marginal cost, fundamentally lowering the barriers to visual production. Since its public beta release in July 2022, Midjourney has diffused rapidly via Discord-based communities, with estimated annual image generation volumes exceeding 260 million by late 2023, reflecting an unprecedented scale of AI-assisted visual production (Nielsen, 2025). These tools are increasingly used in branding, marketing, and design-intensive industries, where firms rely on AI-generated images for early-stage ideation, visual identity development, and campaign assets (Ngoc Le, 2025). As a result, the visual landscape of markets is increasingly shaped by algorithmic generation rather than human-centered construction processes.

While prior research emphasizes the productivity and creativity benefits of generative AI at the individual level, its aggregate implications for market-level design diversity remain less understood. Existing studies show that generative AI can improve individual creative performance while simultaneously increasing

similarity across outputs at the collective level, suggesting a potential tension between individual creativity and systemic homogenization (Doshi & Hauser, 2024). Yet, empirical evidence on whether generative AI leads to convergence in real-world creative markets—where differentiation is strategically valuable—remains scarce.

This study investigates whether the diffusion of generative image AI increases visual homogenization in trademark design. Trademarks represent a critical empirical setting for studying creative convergence because they are designed to signal firm identity, differentiation, and brand positioning. Unlike transient marketing visuals, trademarks are legally registered and persist over time, providing a systematic record of firms' design choices. At the same time, trademarks are increasingly created with AI-assisted tools as firms seek to reduce design costs and accelerate branding processes (Ngoc Le, 2025). The public release of Midjourney in July 2022 constitutes a major technological shock to the cost structure and accessibility of visual design generation, enabling an unprecedented scale of AI-assisted trademark creation (Nielsen, 2025).

This study conceptualizes visual homogenization as a market-level distributional shift in the space of trademark images. Rather than focusing on direct copying, the analysis emphasizes structural

convergence, where many independent firms arrive at visually similar designs because their search processes are guided by shared algorithmic priors and reference pools. Drawing on diffusion-model theory, exploration-exploitation frameworks, and processing fluency theory, this study develops a multi-stage mechanism. First, generative AI reduces the gulf of execution and democratizes design production, expanding entry into the design space. Second, diffusion models disproportionately generate high-probability aesthetic patterns, which biases the creative search process toward exploitation and regression to the mean (Nielsen, 2025). Third, users are more likely to select fluent and polished AI outputs, leading to early termination of exploration and reinforcing convergence through feedback loops (Ngoc Le, 2025).

To test these arguments, this study constructs a comprehensive dataset of U.S. trademark applications filed between 2019 and 2024. Trademark images are embedded using CLIP-based computer vision representations, and visual homogenization is measured as similarity-to-centroid in the embedding space. The empirical design exploits the introduction of Midjourney in July 2022 as a plausibly exogenous shock to generative image production. Using repeated cross-sectional regressions with time controls, this study examines

whether trademark designs become more visually similar after the diffusion of generative AI.

Beyond the baseline effect, this study examines two theoretically motivated boundary conditions. First, the analysis investigates intent-to-use (ITU) filings, which capture early-stage branding contexts where design decisions are made under uncertainty and resource constraints. Second, the study examines highly regulated industries, where institutional pressures and legal constraints narrow the feasible design space. Building on institutional theory and exploitation-exploration frameworks, the study proposes a "double compression" mechanism in regulated industries, where trademark regulatory constraints narrow feasible designs and diffusion-model sampling biases narrow realized designs.

This study makes several contributions. First, it provides the first large-scale empirical evidence on how generative image AI affects market-level creative diversity in trademark design. Second, it integrates algorithmic mechanisms with organizational and cognitive theories to explain why generative AI can simultaneously democratize creativity and induce homogenization. Third, it contributes to strategy and innovation literatures by showing how shared production technologies shape the evolution of symbolic assets and brand identity. Finally, the findings have implications for

managers, policymakers, and intellectual property scholars concerned with the systemic consequences of AI-driven creative production.

2. Theoretical Background and Hypotheses

Before generative AI, high execution costs filtered design production toward experts and increased the value of deliberate experimentation. After generative AI, execution costs fall, entry expands, and production shifts from manual construction to generation followed by selection. Diffusion-model generators disproportionately sample high-probability regions of the design distribution, making typical styles easier to reach than tail styles (Ho et al., 2020; Shumailov et al., 2024). Selection further amplifies this tendency because users often satisfice on fluent, polished outputs and terminate search early (March, 1991; Dell' Acqua et al., 2023). At the aggregate level, such individual-level search and selection dynamics can generate systemic convergence in creative outputs (Doshi & Hauser, 2024). Together, these technological and behavioral forces predict a post-July 2022 increase in trademark visual similarity following the widespread accessibility of Midjourney.

2.1. Visual Homogenization in Trademark Design

Trademarks are designed to signal firm identity through distinctive visual forms, yet distinctiveness at the firm level does not guarantee diversity at the market level. Prior research shows that markets can converge toward a narrower set of recurring solutions when many independent actors rely on shared technologies, reference pools, and search heuristics, even in the absence of coordination (March, 1991; DiMaggio & Powell, 1983). In creative domains, common templates, conventions, and production tools can induce stylistic isomorphism, leading decentralized producers to gravitate toward similar visual conventions (Raffaelli et al., 2019). Algorithmic generation further amplifies this tendency because shared model priors and recursive feedback loops can compress the distribution of outputs and increase similarity across independently generated artifacts (Shumailov et al., 2024; Doshi & Hauser, 2024).

A distributional framing clarifies the mechanism. Trademark designs can be represented as points in a high-dimensional visual feature space, and a market becomes more homogenized when this cloud of points contracts such that extreme or unconventional designs become rarer while designs closer to the prevailing center become disproportionately represented. This intuition parallels formal results in generative modeling research on model collapse,

which show that when generative models are trained on recursively generated data, the tails of the original distribution disappear and probability mass concentrates around a narrower region (Shumailov et al., 2024). Although this evidence concerns training dynamics, the analogy is informative for creative markets because outputs become inputs through social reuse: AI-shaped images circulate as references, shape expectations, and influence what appears legitimate or aesthetically acceptable.

Consistent with this logic, empirical evidence suggests that generative AI can increase homogeneity in real-world outputs. In professional tasks, AI assistance improves quality while pushing outputs toward more standardized forms, particularly when users rely on default system suggestions (Dell' Acqua et al., 2023). In digital marketing, a natural experiment exploiting the Italy-wide ChatGPT ban shows that content similarity decreases when AI access is removed, implying that AI access increases homogenization under normal conditions (Liu et al., 2025). Although these studies do not focus on trademarks, they support the broader claim that generative AI can induce market-level convergence even when individual users pursue high-quality or creative outcomes rather than explicit imitation (Dell' Acqua et al., 2023; Liu et al., 2025).

To operationalize market-level homogenization, this study uses a similarity-to-centroid statistic that directly corresponds to the distributional definition of convergence. Each trademark image is first mapped into a continuous embedding vector that captures high-level visual semantics. The market's "center" is then defined as the centroid—the mean embedding across the relevant set of applications in a given period. Similarity-to-centroid is computed as the cosine similarity between each image's embedding and this centroid, so the measure increases precisely when designs accumulate around the market's central tendency and decreases when designs populate the tails.

This approach targets distributional compression rather than pairwise imitation. Homogenization in this study is not defined as copying between specific firms, but as a market-level reallocation of probability mass toward prototypical regions of the design space. A centroid-based metric is therefore appropriate because it detects whether the entire cloud of designs is moving inward, even when no two marks are near-duplicates. Moreover, the measure is naturally aligned with embedding-based operationalizations of collective convergence in recent creativity research, where similarity of a focal output to an average embedding is used to quantify whether outputs drift toward a shared "typical" form (Doshi & Hauser, 2024).

The centroid logic also matches the mechanism developed in this chapter. If diffusion-based generation and selection routines systematically privilege high-density patterns, then convergence should appear as a rising proximity to the market average, not merely as a rising frequency of direct lookalikes. In other words, the theoretical prediction is a tightening of the distribution around a common center. Similarity-to-centroid is designed to capture exactly that movement.

2.2 Technological Affordances

Generative image systems based on diffusion models fundamentally change how visual designs are produced by altering the structure of the creative process. Diffusion models create images through an iterative denoising process learned from large-scale training data, thereby generating samples that approximate the underlying data distribution (Ho et al., 2020). In practice, this shifts the workflow away from step-by-step construction in graphics software, toward a process in which users specify intent in natural language, receive candidate outputs, and iteratively refine prompts through selection and modification. As a result, design production moves from a construction-centric paradigm to a generation-selection paradigm; creative effort is increasingly allocated to choosing among machine-generated alternatives rather than

manually constructing artifacts.

Midjourney represents a salient inflection point in this transition since it became broadly accessible through a low-friction interface in mid-2022, when the system opened to the public in beta mode (PCWorld, 2022). This timing motivates the empirical use of July 2022 as a shock window for the widespread diffusion of text-to-image generation in consumer-facing design contexts.

In pre-AI design environments, translating conceptual intent into a finished visual artifact required substantial software expertise and iterative execution effort. Human-computer interaction research characterizes this friction as a *gulf of execution*, defined as the gap between what users intend to achieve and what the system allows them to do easily (Norman, 1988). Text-to-image systems compress this gulf by allowing users to express design intent through semantic instruction rather than procedural construction, thereby substituting linguistic description for technical manipulation.

Lowering the execution burden broadens participation in creative production by reducing the effective skill threshold required to generate high-quality outputs. Experimental evidence in writing tasks shows that access to generative AI increases productivity and output quality while disproportionately benefiting lower-skill participants, thereby compressing the performance distribution (Noy

& Zhang, 2023). Related evidence from customer-support contexts similarly shows that AI assistance yields larger productivity gains for novice or lower-skill workers (Brynjolfsson et al., 2023). Together, these findings support the argument that generative AI democratizes production by reducing the effective role of specialized expertise, which implies that a larger and more heterogeneous set of actors can produce visually competent designs.

A defining technological affordance of text-to-image systems is immediacy combined with stochastic variation. Users can generate multiple high-quality candidates within seconds, observe outputs, and iterate through prompt revision. Midjourney's interface reinforces this process by producing grids of variations for a single request, encouraging rapid comparison and selection (PCWorld, 2022). Since diffusion models incorporate probabilistic sampling during denoising, repeated generations from the same prompt yield multiple variants (Ho et al., 2020). While this immediacy and variation expand the apparent exploration space, they do not guarantee realized diversity, because the system can generate many outputs that remain concentrated around a narrow aesthetic region.

The generation-selection workflow also accelerates and broadens early-stage creative search. Rather than committing to a single concept and refining it through costly revisions, users can

evaluate multiple candidates and converge on an acceptable direction, which changes the economics of experimentation and convergence. Creativity experiments show that AI-generated suggestions can improve individual-level outcomes while simultaneously increasing similarity across outputs, because users anchor on the same classes of machine-generated ideas (Doshi & Hauser, 2024). Applied to trademark design, this implies that even when firms seek differentiation, reliance on shared generative tools can shape the pool of plausible starting points in similar ways across organizations.

2.3 Algorithmic Convergence and Regression to the Mean

Beyond altering production costs and workflows, diffusion-based generative systems shape the geometry of creative outputs through their algorithmic properties. Diffusion models are designed to generate samples that are probable under the learned data distribution, which creates a systematic pull toward high-density regions of the learned space unless users explicitly impose constraints that push the system toward atypical outputs (Ho et al., 2020).

This mechanism can be framed through the exploration-exploitation lens. March (1991) argues that learning involves a tradeoff between exploring novel possibilities and exploiting reliable

knowledge. In practice, diffusion-based generation pushes the creative process toward exploitation because default samples reflect dominant patterns that are easiest to learn and most likely under the training distribution. When users adopt early candidates that already appear polished, subsequent prompts and selections tend to reinforce similar aesthetic targets, thereby embedding exploitation dynamics into the generation-selection loop.

Repeated exploitation implies that outputs increasingly cluster around the center of the learned distribution. Model-collapse research formalizes this logic by showing that when synthetic outputs are recursively reused, the tails of the original distribution disappear and probability mass contracts toward the center (Shumailov et al., 2024). The implication is not merely that outputs appear similar, but that extreme variation becomes structurally rarer. In creative markets, a similar contraction can occur through social reuse, because AI-shaped aesthetics become widely visible, influence prompt conventions, and shape expectations of what constitutes a plausible or desirable design. This creates a market-level analogue of recursive sampling, in which generated outputs influence subsequent generation and selection processes.

Algorithmic convergence becomes a market phenomenon when many creators rely on a small set of dominant models and

interfaces. Even if each creator acts independently, shared algorithmic priors can induce collective convergence in outputs. Creativity research provides consistent evidence that AI assistance can raise individual-rated creativity while reducing collective diversity, because outputs become more similar across users (Doshi & Hauser, 2024). Applied to trademarks, this implies that the distribution of designs should shift toward a tighter center after the diffusion of a widely used system such as Midjourney.

H1. The introduction of Midjourney (July 2022) increases visual similarity among firms' trademarks.

2.4 Illusion of Fluency and Algorithmic Conformity

Algorithmic convergence explains why outputs drift toward the center of the distribution, but it does not explain why users repeatedly select such outputs even when differentiation is strategically valuable. The missing link lies in judgment and attention. When generative tools produce polished images quickly, selection decisions are often made under limited cognitive resources, which may bias users toward outputs that are easy to process.

Processing fluency theory posits that stimuli that are processed more fluently tend to be evaluated more positively, including in aesthetic judgments (Reber et al., 2004). Mean-regressed images often exhibit stable composition, familiar visual

grammar, and predictable structure, which increase processing fluency. Users may therefore misinterpret fluency as a signal of quality, even when fluency is not informative about distinctiveness or strategic fit. As a result, users may terminate search early because the first acceptable output appears sufficiently complete and professional.

When fluency becomes a quality cue, search behavior changes. Users terminate exploration earlier because the first acceptable output feels sufficient, which is consistent with evidence that AI-generated suggestions may anchor creators and increase similarity across outputs (Doshi & Hauser, 2024). In co-creative settings, the model's initial outputs define the perceived search space, which may anchor subsequent revisions to local adjustments rather than global shifts in aesthetic direction.

Fluency-based selection also creates a feedback loop. Users repeatedly select outputs that are easy to like and easy to justify, which increases the visibility of average-like designs as examples of "good" AI output. This visibility shapes prompt habits and reference pools, increasing the share of average-like designs circulating in the market and reinforcing algorithmic convergence. This feedback logic parallels the core intuition in model-collapse research: when outputs increasingly reflect outputs, diversity erodes

and the center dominates (Shumailov et al., 2024).

This feedback logic also resonates with recent work on self-consuming generative systems. Alemohammad et al. (2024) introduce the concept of model autophagy disorder (MAD), showing that when generative models increasingly train on or reference their own synthetic outputs, diversity collapses and the distribution contracts toward dominant patterns. In creative markets, algorithmically generated images may similarly become social reference points that shape human evaluation criteria and prompt conventions. This creates a socio-technical autophagic loop in which AI-shaped aesthetics recursively influence what humans select and what algorithms subsequently generate. Together with processing fluency biases, this recursive loop accelerates the contraction of the design distribution and amplifies market-level visual homogenization.

2.5 Moderating Role of Intent-to-Use (ITU) Filings

The baseline mechanism implies an average increase in visual similarity following the diffusion of Midjourney; however, the magnitude of this effect may vary systematically across organizational contexts that differ in their reliance on algorithmic defaults.

Under U.S. trademark practice, applicants may file on an intent-to-use basis when the mark has not yet been used in

commerce but there is a bona fide intention to use it in the future, as specified in Lanham Act Section 1(b) and USPTO guidance (USPTO, n.d.). ITU filings therefore correspond to an early-stage design context in which branding elements are provisional and subject to revision, and where organizations have not yet accumulated extensive market feedback or internal evaluation routines.

Organizational search theory suggests that actors facing high uncertainty and limited experience rely more heavily on heuristics and external knowledge sources rather than internalized routines (March, 1991; Levinthal & March, 1993). In such contexts, generative AI constitutes a salient external knowledge source that can substitute for specialized design capabilities. Consistent with this view, evidence from generative AI adoption shows that productivity gains are largest when skill or experience is lower, suggesting that organizations with fewer specialized resources rely more heavily on AI-generated defaults rather than expert-driven construction (Noy & Zhang, 2023; Brynjolfsson et al., 2023). Recent experimental evidence further reinforces this mechanism, showing that while generative AI lowers the barrier to entry, it can lead to "design fixation" among less experienced users who tend to anchor on the initial, often prototypical, suggestions provided by the model (Hou et al., 2025). Since diffusion-based defaults reflect high-density

regions of the learned distribution, heavier reliance on these defaults systematically increases convergence toward the distributional center.

Importantly, ITU contexts may also lack strong organizational routines for evaluating and rejecting algorithmic suggestions, which implies that algorithmic priors may exert a larger influence on realized design outcomes. In contrast, marks that are already in use typically undergo repeated validation through market feedback and internal governance processes, which may counteract convergence toward algorithmic averages. Taken together, these arguments imply that generative AI constitutes a homogenizing technological shock whose impact is amplified when organizations rely more heavily on algorithmic defaults and have weaker mechanisms for evaluating or correcting them. Thus, ITU contexts represent a setting where algorithmic convergence mechanisms are least likely to be counteracted by organizational routines and governance structures. Therefore, ITU filings should exhibit a larger post-Midjourney increase in similarity-to-centroid.

H2. The effect of Midjourney on trademark visual similarity is more pronounced in ITU applicants.

2.6 Moderating Role of Trademark Regulation

A second boundary condition concerns the institutional environment in which trademark design decisions are evaluated. Trademark design is not only a creative choice but also a legal signal subject to institutional scrutiny by trademark offices and potential opposition by incumbents. In highly regulated industries, firms face heightened legal and legitimacy pressures in trademark registration, which shape both the feasible and preferred set of design choices.

Institutional theory predicts homogeneity within organizational fields because organizations seek legitimacy by conforming to established norms and expectations (DiMaggio & Powell, 1983). In trademark contexts, legitimacy is assessed through formal examination criteria such as distinctiveness, likelihood of confusion, and compliance with industry-specific naming and labeling norms. Organizations facing intense scrutiny therefore have strong incentives to adopt conventional and defensible trademark forms rather than experimenting with atypical designs (Suchman, 1995). Regulation thus narrows the feasible trademark design space by making unconventional signals costly, risky, or legally contestable.

Generative AI introduces an additional technological source of convergence. Diffusion models tend to return high-probability mainstream patterns by default, which narrows the realized design

space toward statistically typical forms (Ho et al., 2020). When combined with trademark regulatory constraints, this creates a double compression in which both the legally feasible and algorithmically realized design spaces contract toward conventional templates.

Organizational research further suggests that oversight and standardization pressures bias organizations toward exploitation rather than exploration (Benner & Tushman, 2003). In trademark registration contexts, deviation from conventional forms may be interpreted as legal or legitimacy risk, strengthening incentives to rely on AI-generated templates that approximate statistically typical and legally defensible designs. In this sense, institutional pressures in trademark law and algorithmic priors operate as complementary forces that jointly reinforce convergence toward the center of the design distribution. Regulated industries therefore provide a setting in which both legal-institutional and algorithmic forces jointly bias organizations toward exploitation and conformity.

H3. The effect of Midjourney on trademark visual similarity is more pronounced in industries subject to stronger trademark institutional constraints.

3. Methodology

3.1 Data Sources and Sample

This study draws on the universe of U.S. trademark applications from the United States Patent and Trademark Office (USPTO) Trademark Case Files. The dataset covers all trademark applications filed between 2019 and 2024, allowing observation of the trademark design distribution before and after the widespread diffusion of generative image AI.

The empirical setting exploits the introduction of Midjourney as a plausibly exogenous technological shock to creative production. Midjourney opened public access in July 2022, substantially lowering the cost of generating high-quality visual designs (PCWorld, 2022). Following prior work on technology diffusion shocks, this study defines the post-AI period as 2022Q3 onward, which captures the initial large-scale adoption of text-to-image systems in consumer-facing creative workflows. Although diffusion-based image generation systems existed prior to Midjourney's release, Midjourney's public beta in July 2022 marked the first large-scale diffusion of high-quality text-to-image generation to non-expert users through a low-friction interface. Industry accounts document a rapid acceleration in user adoption and image generation volume

following this release, suggesting a discrete increase in exposure rather than a smooth trend. This timing therefore provides a plausibly exogenous shock to the cost and accessibility of visual design generation in branding contexts.

The unit of analysis is a trademark application-image pair. For each application, the associated trademark image is retrieved and processed using computer vision embeddings. Applications without image files, purely word marks, or missing filing dates are excluded. The resulting sample constitutes a repeated cross-section of trademark filings, which allows tracing changes in the market-level distribution of visual designs over time.

Although firms may file multiple trademarks over time, each application represents a distinct entry into the market-level design distribution. The empirical objective is to track changes in the distribution of designs rather than within-firm design evolution. A repeated cross-sectional design is therefore appropriate and avoids conditioning on post-treatment firm behavior.

3.2 Variables

- **Dependent Variable: Trademark Visual Homogenization**

  Measuring visual homogenization requires an embedding that captures semantic and stylistic regularities across heterogeneous images, including minimalist logos,

emblems, icons, and mixed graphical marks. This research uses CLIP-based image embeddings for two reasons. First, CLIP is explicitly designed to learn transferable visual representations through contrastive training on large-scale image-text pairs, producing an embedding space in which cosine similarity reflects high-level visual-semantic relatedness rather than low-level pixel overlap (Radford et al., 2021). This property matters for trademarks because convergence often takes the form of "similar in structure" rather than "identical in detail." A representation that is sensitive to compositional and semantic features—rather than exact pixel matching—fits the conceptual definition of market-level homogenization developed in Chapter 2.

Second, and more directly tied to logos/trademarks, recent work in computer vision demonstrates that image-text pretraining in the CLIP paradigm is effective for logo recognition, precisely because logos require representations that generalize across extreme variation in typography, abstraction, and context. For example, Hubenthal and Kumar (2023) develop an image-text pretraining approach for logo recognition and show that CLIP-style pretraining is well-suited to capturing logo-level identity features beyond pixel

similarity, supporting the use of CLIP-family embeddings when the empirical task is to compare logos in a shared latent space (Hubenthal & Kumar, 2023). Together, these studies provide concrete methodological grounding for treating CLIP embeddings as a credible basis for measuring visual similarity among trademark images in a way that aligns with the theoretical mechanism-distributional convergence toward prototypical aesthetic regions rather than direct duplication (Radford et al., 2021; Hubenthal & Kumar, 2023). Similarity-to-centroid captures distributional convergence rather than pairwise duplication. Alternative dispersion-based measures such as average pairwise similarity or embedding variance are conceptually related but less directly map to the theoretical mechanism of probability mass contraction toward a central tendency. The centroid-based metric therefore provides a parsimonious and theory-aligned measure of market-level homogenization.

- Independent Variable: Generative AI Adoption

The key independent variable is a Post-AI indicator that equals one for applications filed in 2022Q3 or later and zero otherwise. This variable captures the diffusion of

Midjourney and related diffusion-based text-to-image tools into mainstream creative practice.

The identifying assumption is that July 2022 represents a structural break in the cost and accessibility of high-quality visual generation, rather than a gradual trend. The identifying assumption is that July 2022 represents a structural break in the cost and accessibility of high-quality visual generation rather than a gradual trend. This assumption is consistent with contemporaneous industry accounts documenting Midjourney's rapid public diffusion (PCWorld, 2022), and therefore provides a plausibly exogenous shock to visual design production in branding contexts. Time fixed effects and flexible time trends are included to account for secular changes unrelated to generative AI. Therefore, it provides a plausibly exogenous shock to the cost and accessibility of visual design generation in branding contexts. Time fixed effects and flexible time trends are included to account for secular changes unrelated to generative AI.

- Moderators

  Intent-to-Use (ITU) Indicator

The ITU indicator equals one if a trademark application is filed on an intent-to-use basis and zero otherwise. Under the Lanham Act, ITU filings indicate that a mark has not yet been used in commerce but that the applicant intends to use it in the future. This status captures early-stage branding contexts in which design decisions are made under uncertainty and before market validation.

Theoretically, ITU contexts should exhibit greater reliance on generative AI defaults because applicants face tighter resource constraints and weaker internal design capabilities. As discussed in Section 2, heavier reliance on generation-selection workflows implies stronger convergence toward algorithmically typical styles. Therefore, ITU status is modelled as a moderator of the Post-AI effect.

Trademark Institutional Regulation Indicator

Trademark institutional regulation is operationalized using Nice classification codes associated with each trademark. Highly regulated industries are identified based on classes corresponding to pharmaceuticals, medical devices, financial services, alcoholic beverages, and other sectors subject to intensive regulatory oversight (e.g., Nice Classes 005, 010, 033, 036). These classes are coded as a binary

indicator for high regulation. Certain product and service categories are subject to heightened trademark examination and opposition risks due to public safety, consumer protection, and statutory restrictions on naming and labeling (e.g., pharmaceuticals, medical devices, financial services, and alcoholic beverages). These sectors face greater legal scrutiny from trademark offices and incumbent firms, which constrains the feasible space of trademark signals. Consistent with institutional theory, these classes are treated as settings with stronger trademark institutional constraints.

In Section 2, this mechanism is conceptualized as "double compression", where regulatory constraints narrow the feasible design space and diffusion-model sampling biases narrow the realized design space.

- Control Variables

Several control variables account for alternative explanations of visual similarity.

First, the number of Nice classes associated with a trademark is included to capture scope and complexity of the mark, which may correlate with design sophistication.

Second, attorney representation is measured using a Pro Se indicator. Pro Se filings may reflect lower

sophistication and could correlate with reliance on AI-generated designs.

Third, flexible time controls are included, including year and quarter fixed effects, to capture secular trends in visual design unrelated to generative AI.

3.3 Empirical Strategy

The empirical design uses a repeated cross-sectional framework to estimate how the distribution of trademark designs changes after the diffusion of generative AI. This approach is appropriate because trademark applications represent new entries into the design distribution each period rather than longitudinal observations of the same entities.

The baseline specification estimates the association between the Post-AI indicator and similarity-to-centroid, controlling for time trends and application characteristics. Interaction terms between Post-AI and the moderator variables test heterogeneous effects for ITU filings and regulated industries. Event-study specifications are estimated to examine dynamic treatment effects and assess the presence of differential pre-trends. The absence of significant pre-treatment differences would support the interpretation of the Post-AI coefficient as capturing a structural break associated with generative AI diffusion.

This design exploits the sharp timing of Midjourney's diffusion while controlling for secular trends in design practices, branding strategies, and legal environments. Since direct measures of firm-level AI usage are unavailable, the analysis exploits generative AI diffusion as a market-level technological shock. The estimates therefore capture the equilibrium effect of widespread exposure to generative tools rather than individual adoption decisions, reducing concerns about firm-level selection into AI usage.

4. Results

4.1 Descriptive Statistics and Correlations

Table 1 reports descriptive statistics and the correlation matrix for the main variables. The dependent variable, *Trademark Similarity*, has a mean of 0.754 (SD = 0.088), indicating meaningful dispersion in visual similarity across trademark filings. Approximately 41.2% of applications were filed in the post-AI period, and 43.5% were filed on an intent-to-use (ITU) basis. About 24.9% of applications were filed without attorney representation, and 19.6% involved multiple Nice classes.

Table 1. Descriptive Statistics and Correlation Matrix

	Variables	Mean	SD	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(1)	Trademark Similarity	0.7541	0.0882								
(2)	Post-AI	0.412	0.4922	**0.036*** **							
(3)	Intent-to-Use	0.4349	0.4957	**0.047*** **	0.023***						
(4)	Regulatory Constraint	0.1315	0.6036	**0.041*** **	0.007*	0.002					
(5)	Number of Classes	1.3953	1.6291	0.041***	0.001	0.067***	0.049***				
(6)	Pro Se Filing	0.2496	0.4328	-0.137***	-0.074***	-0.040***	-0.113***	-0.089***			
(7)	Experience	0.0595	0.2366	0.036***	0.008**	0.006	0.041***	0.090***	-0.124***		
(8)	Multi-class Filing	0.1957	0.3967	0.077***	0.024***	0.142***	-0.012***	0.492***	-0.152***	0.123***	

Pairwise correlations provide preliminary support for the hypotheses. Post-AI filings are positively correlated with trademark similarity ($\rho = 0.036$, $p < 0.01$), suggesting an increase in convergence after the diffusion of generative AI. ITU filings are also positively correlated with similarity ($\rho = 0.047$, $p < 0.01$), consistent with stronger convergence in early-stage contexts. Regulatory constraint is positively correlated with similarity ($\rho = 0.041$, $p < 0.01$), indicating greater homogeneity in regulated industries. In contrast, pro se filings are negatively correlated with similarity ($\rho = -0.137$, $p < 0.01$), suggesting lower standardization in visual branding among unrepresented applicants.

4.2 Baseline Effect of Generative AI

Table 2. Regression Analysis

	(1) Baseline Model	(2) ITU Moderation	(3) Trademark Regulation Moderation
Constant	0.7567***	0.7569***	0.7569***
	-0.0012	-0.0011	-0.0012
Post-AI (Dummy)	**0.0047*** **	-0.0028***	0.0051***
	-0.0008	-0.001	-0.0008
Post-AI × ITU		**0.0169*** **	
		-0.0015	
ITU Applicant		0.0001	
		-0.001	
Post-AI × Regulation			**0.0031** **
			-0.0013
Regulation			0.0023***
			-0.0008
Number of Classes	0.0016**	0.0014**	0.0015**
	-0.0007	-0.0007	-0.0007
Pro Se (No Attorney)	-0.0269***	-0.0270***	-0.0264***
	-0.0009	-0.0009	-0.0009
Observations	59717	59717	59717
R^2	0.0202	0.024	0.0209
Adjusted R^2	0.0202	0.0239	0.0208
F Statistic	410.9185***	293.5416***	255.1641***

Table 2 presents the regression results testing the baseline effect of generative AI diffusion and the proposed moderating mechanisms. Column (1) reports the baseline specification. The Post-AI indicator is positive and statistically significant ($\beta = 0.0047$, $p < 0.01$), indicating that trademark designs filed after the diffusion of Midjourney are more visually similar to the market centroid. This

finding supports H1 and is consistent with the prediction that diffusion-based generative models bias creative search toward high-probability regions of the design space, leading to regression toward prototypical visual forms.

Among control variables, the number of Nice classes is positively associated with similarity, suggesting that broader-scope trademarks tend to adopt more conventional and standardized visual forms, potentially reflecting greater branding complexity and legal scrutiny. Pro se filings are associated with significantly lower similarity, consistent with greater idiosyncrasy and lower standardization in branding practices among unrepresented applicants. The model explains a modest share of variance (R^2 = 0.020), which is expected given the high-dimensional and heterogeneous nature of creative visual design outcomes.

Figure 1. Trend in Trademark Visual Similarity, 2019-2024

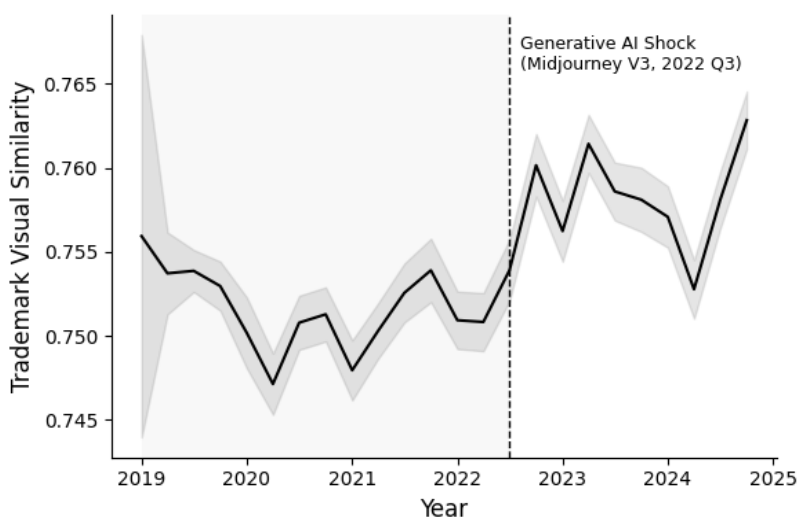


Figure 1 shows the mean trend in trademark visual similarity from 2019 to 2024, with the vertical dashed line marking the introduction of Midjourney in 2022Q3. The figure shows a relatively stable pre-AI trend followed by a discrete upward shift in similarity after the AI diffusion period, suggesting a structural break rather than a continuation of pre-existing trends. This visual pattern corroborates the identification assumption that generative AI diffusion represents a discrete technological shock to visual design production. Although the estimated coefficient is numerically small, it represents a meaningful shift in a high-dimensional embedding space. The baseline Post-AI effect corresponds to approximately 5.7% of the pre-AI standard deviation in visual similarity. Given the large volume and persistence of trademark filings, even modest shifts imply substantial contraction of the design distribution at the market level.

Column (2) introduces the interaction between Post-AI and Intent-to-Use (ITU) status. The interaction term is positive and highly significant ($\beta = 0.0169$, $p < 0.01$), indicating that the increase in visual similarity after generative AI diffusion is substantially stronger for ITU filings. This result supports H2. The main effect of ITU status is not significant once the interaction is included, suggesting that ITU filings are not inherently more homogeneous but

become more homogeneous specifically after the diffusion of generative AI. This pattern is consistent with organizational search theory, which predicts greater reliance on external heuristics under uncertainty and limited experience. In early-stage branding contexts, generative AI likely substitutes for specialized design capabilities, anchoring search around algorithmically typical outputs and amplifying convergence. The ITU interaction effect corresponds to approximately 20.6% of the pre-AI standard deviation, indicating a substantively large amplification of algorithmic convergence in early-stage branding contexts.

Figure 2. Trademark Visual Similarity for ITU vs Non-ITU

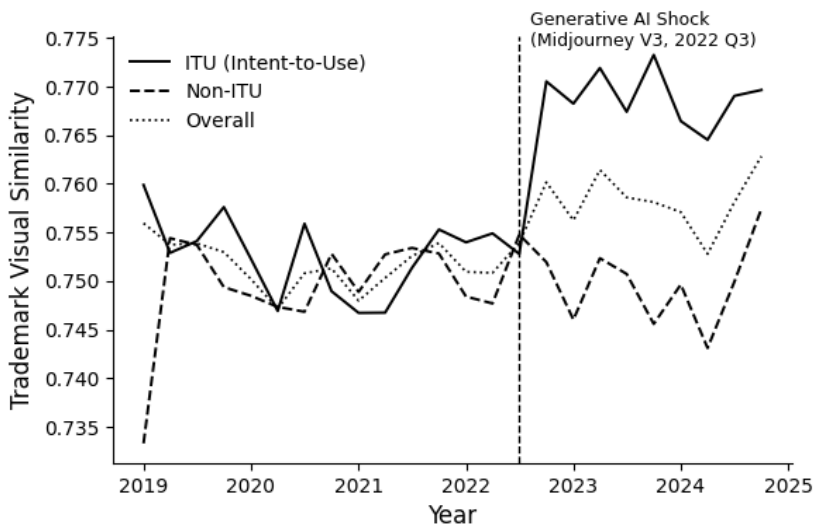


Figure 2 shows similarity trends separately for ITU and non-ITU filings. The divergence after 2022Q3 shows a steeper post-AI increase for ITU filings, visually confirming the interaction effect and

reinforcing the interpretation that generative AI disproportionately homogenizes early-stage trademark designs.

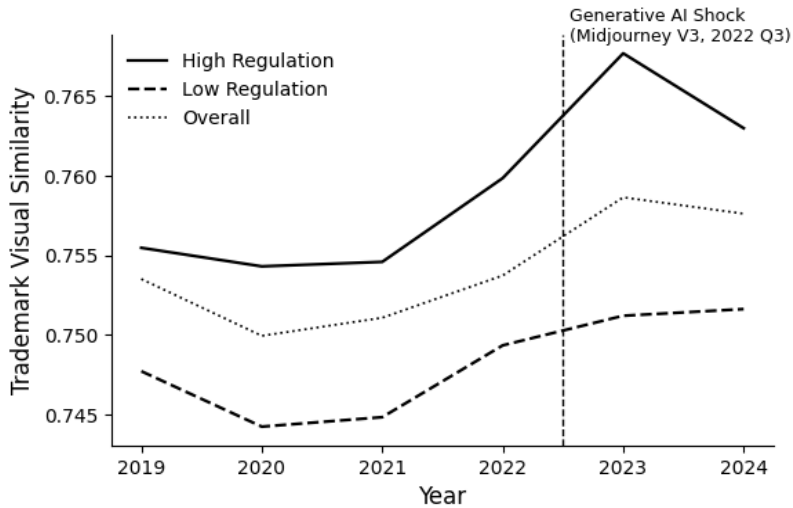
Column (3) presents the interaction between Post-AI and Regulatory Constraint. The interaction term is positive and statistically significant ($\beta = 0.0031$, $p < 0.05$), indicating that generative AI diffusion leads to greater visual convergence in industries subject to stronger trademark institutional constraints. This finding supports H3. The main effect of regulatory constraint is also positive and significant, suggesting that regulated industries exhibit higher baseline similarity even before AI diffusion.

Together, these findings are consistent with the "double compression" mechanism developed in Chapter 2. In highly regulated industries, trademark law and institutional scrutiny constrain the feasible set of signals, while diffusion-based generative models bias realized outputs toward statistically typical forms. The interaction suggests that institutional and algorithmic selection environments operate as complementary forces, jointly accelerating convergence toward conventional trademark designs.

Figure 3 shows trademark similarity trends by regulatory status. Regulated industries exhibit both higher baseline similarity and a steeper post-AI increase, consistent with institutional selection

pressures that narrow feasible design variation even before AI diffusion.

Figure 3. Trademark Visual Similarity by Trademark Regulation



Overall, the results indicate that generative AI diffusion affects not only the mean level of visual similarity but also the shape of the trademark design distribution. The upward shift in similarity-to-centroid implies a contraction of the design space, where extreme or unconventional visual forms become less prevalent and probability mass concentrates around prototypical aesthetic regions. This pattern is consistent with the socio-technical feedback mechanisms discussed in Chapter 2 and suggests that shared generative technologies systematically reshape the evolution of symbolic assets in markets.

5. Discussion and Contributions

This study investigates whether the diffusion of generative image AI leads to visual homogenization in trademark design and identifies organizational and institutional boundary conditions under which this effect is amplified. The results provide consistent evidence that the public diffusion of Midjourney is associated with increased convergence in trademark visual designs. Moreover, this convergence is significantly stronger for intent-to-use filings and in highly regulated trademark domains, supporting the multi-level theoretical framework developed in Chapter 2.

The research makes several contributions to the literatures on generative AI, creativity, and strategic management. First, it contributes to the emerging literature on AI and creativity by providing large-scale field evidence of a market-level homogenization effect. Prior experimental research shows that generative AI can enhance individual creativity while reducing collective diversity by anchoring creators on common machine-generated ideas (Doshi & Hauser, 2024; Dell' Acqua et al., 2023). This study extends these insights to a real-world institutional setting and demonstrates that AI diffusion systematically shifts the distribution of symbolic artifacts toward prototypical forms. By

documenting distributional convergence in a persistent and legally consequential creative domain, this study shows that generative AI can function as a structural force shaping the evolution of creative markets, rather than merely a productivity-enhancing tool at the individual level.

Second, this study contributes to the strategy literature on symbolic assets and identity by conceptualizing trademarks as points in a high-dimensional visual feature space and operationalizing market-level homogenization using similarity-to-centroid. This distributional approach moves beyond binary imitation and captures subtle shifts in the geometry of design spaces, offering a scalable framework for studying convergence in cultural and symbolic assets (Rindova et al., 2006; Gioia et al., 2013). The framework extends research on technological standardization and institutional isomorphism into the domain of visual identity and branding.

Third, this research integrates technological, cognitive, and institutional mechanisms into a unified theoretical framework. It shows how diffusion-model sampling biases generate algorithmic convergence, how processing fluency and anchoring biases shape human selection behavior, and how institutional pressures shape feasible design spaces. Together, these forces generate what this study terms a "double compression" effect, in which both the

feasible and realized design distributions contract toward dominant patterns. This integration bridges insights from machine learning, cognitive psychology, and institutional theory, offering a novel socio-technical perspective on how digital technologies reshape the evolution of symbolic assets (DiMaggio & Powell, 1983; Suchman, 1995).

Finally, the findings contribute to strategic management theories of exploration and exploitation by showing that generative AI systematically shifts creative search toward exploitation at the population level. While AI democratizes access to design production, it simultaneously biases the collective trajectory of creative search toward statistically typical solutions, reshaping the balance between exploration and exploitation in markets (March, 1991; Levinthal & March, 1993). This highlights a new tension in digital innovation: technologies that expand individual creativity may induce collective convergence, altering competitive differentiation dynamics.

The findings also have important managerial and policy implications. For firms, reliance on generative AI tools may inadvertently reduce distinctiveness in visual identity, potentially weakening brand differentiation and long-term competitive advantage. Managers should therefore implement governance mechanisms that counteract algorithmic convergence, such as prompt

diversification protocols, structured exploration phases, or human-in-the-loop review processes that explicitly reward atypical design directions. Firms may also need to develop organizational capabilities to audit AI-generated design pipelines to ensure strategic differentiation is preserved.

For intellectual property institutions and policymakers, the results indicate that generative AI tools can systematically reshape the visual landscape of trademarks. Increased baseline similarity across marks may complicate trademark examination and enforcement by raising the background level of visual convergence and potentially increasing disputes over likelihood of confusion. Policymakers may need to reconsider how distinctiveness and similarity standards are evaluated in an era where design production is increasingly mediated by shared generative algorithms.

6. Limitations and Future Research

Despite its contributions, this study has several limitations that point to important directions for future research. First, while this study exploits the public release of Midjourney as a plausibly exogenous technological shock, the empirical design relies on a repeated cross-sectional framework and cannot fully test parallel

trends in the absence of generative AI. Although event-study analyses show no sharp pre-trends, the counterfactual trajectory of trademark design convergence remains unobserved. Future research could leverage staggered platform adoption, firm-level AI integration data, or natural experiments exploiting differential access constraints to strengthen causal identification. Such designs would allow researchers to distinguish AI-induced convergence from secular trends in design aesthetics and institutional standardization, thereby sharpening causal claims about algorithmic homogenization.

Second, while this study documents convergence in trademark design outcomes, it does not examine downstream consequences for consumers, firms, or legal institutions. Trademarks are strategic signals, and increased visual similarity may affect consumer confusion, brand differentiation, and litigation risk. Future research could link AI-driven homogenization to consumer perception data, brand performance metrics, or trademark opposition and litigation records. Such work would illuminate whether algorithmic convergence merely alters aesthetic distributions or fundamentally reshapes competitive dynamics and intellectual property enforcement.

Third, this study conceptualizes homogenization as distributional convergence in visual embedding space, but it does not

observe how firms strategically respond to algorithmic convergence. For example, firms may deliberately counteract AI defaults by investing in human-driven differentiation, or conversely, may strategically exploit convergence to mimic dominant visual categories. Future research could examine organizational responses to AI-induced convergence, such as investments in design capabilities, governance structures, or deliberate aesthetic deviation strategies. This would extend theories of technological search by showing how organizations adapt to algorithmically shaped opportunity landscapes.

Taken together, this study suggests that generative AI is not merely a productivity-enhancing design tool but a technology that reshapes the structure of symbolic markets. By showing that shared generative infrastructures can systematically bias creative search toward prototypical forms, this research highlights how algorithmic systems can influence not only what organizations produce, but also what becomes typical, legitimate, and distinctive at the market level. More broadly, the findings point to a new role of digital technologies in organizational and institutional dynamics: as infrastructures that shape collective search trajectories and the evolution of symbolic assets. As generative AI continues to diffuse across creative and strategic domains, understanding how algorithmic systems interact

with organizational choices and institutional environments will be critical for theories of innovation, differentiation, and competitive advantage.

References

Alemohammad, S., Casco-Rodriguez, J., Luzi, L., Humayun, A. I., Babaei, H., Le Jeune, D., Siahkoohi, A., & Baraniuk, R. G. (2024). Self-consuming generative models go MAD. arXiv preprint arXiv:2402.07061.

Benner, M. J., & Tushman, M. L. (2003). Exploitation, exploration, and process management: The productivity dilemma revisited. Academy of Management Review, 28(2), 238–256.

Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). Generative AI at work. National Bureau of Economic Research Working Paper No. 31161.

Dell'Acqua, F., Eling, M., Fabius, V., von Krogh, G., Lakhani, K. R., & Levitt, R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. Harvard Business School Working Paper.

DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. American Sociological Review, 48(2), 147–160.

Doshi, A., & Hauser, O. (2024). Generative AI enhances individual creativity but reduces collective diversity. Science Advances, 10(6), eadn7637.

Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. Advances in Neural Information Processing Systems, 33, 6840–6851.

Hou, J., Wang, L., Wang, G., Wang, H. J., & Yang, S. (2025). The double-edged roles of generative AI in the creative process: Experiments on design work. *Information Systems Research*

Hubenthal, A., & Kumar, A. (2023). Logo recognition with multimodal contrastive pretraining. Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops.

Le, T. N. (2025). Artificial intelligence as a methodological innovation in the training of future designers: Midjourney tools. Information Technologies and Learning Tools, 97(5), 203–215.

Liu, Y., Wang, Z., & Yang, Y. (2025). The homogenization effect of generative AI on digital content: Evidence from the Italy ChatGPT ban. Working paper.

Levinthal, D. A., & March, J. G. (1993). The myopia of learning. *Strategic Management Journal, 14*(S2), 95–112.

March, J. G. (1991). Exploration and exploitation in organizational learning. Organization Science, 2(1), 71–87.

Nielsen, D. L. (2025). The protection of “style” under copyright and its application to generative artificial intelligence. Emory Law Journal, 74(4), 1009–1054.

Norman, D. A. (1988). The psychology of everyday things. Basic Books.

Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. Science, 381(6654), 187–192.

PCWorld. (2022, July). Midjourney opens beta access to everyone. PCWorld.

Raffaelli, R., Glynn, M. A., & Tushman, M. L. (2019). Frame flexibility: The role of frame ambiguity in the evolution of nascent categories. Strategic Management Journal, 40(3), 384–416. https://doi.org/10.1002/smj.2983

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., & Sutskever, I. (2021). Learning transferable visual models from natural language supervision. Proceedings of the 38th International Conference on Machine Learning (ICML).

Reber, R., Schwarz, N., & Winkielman, P. (2004). Processing fluency and aesthetic pleasure: Is beauty in the perceiver's processing experience? Personality and Social Psychology Review, 8(4), 364–382.

Shumailov, I., Shypula, A., Zhao, Y., Balestriero, R., Biroli, G., & Gal, Y. (2024). AI models collapse when trained on recursively generated data. Nature, 631, 755–759.

Suchman, M. C. (1995). Managing legitimacy: Strategic and institutional approaches. Academy of Management Review, 20(3), 571–610.

Tsidylo, I. M., & Sendra, C. E. (2023). Artificial intelligence as a methodological innovation in the training of future designers: Midjourney tools. Information Technologies and Learning Tools, 97(5), 203–215.

United States Patent and Trademark Office. (n.d.). Trademark Manual of Examining Procedure (TMEP).

국문 초록

생성형 AI 도입이 상표의 시각적 유사성에 미치는 영향: 알고리즘적 수렴과 제도적 압착의 역할

Suh Young An

경영학과 전략/국제경영전공

서울대학교 경영대학원

본 연구는 생성형 인공지능(Generative AI)의 도입이 기업의 시각적 정체성 형성 과정에 어떠한 구조적 변화를 초래하는지 분석한다. 특히 텍스트-이미지 생성 모델의 확산이 상표 디자인의 시각적 다양성을 축소하고 시장 수준의 시각적 동질화를 유발하는지를 실증적으로 검증한다. 이론적으로, 본 연구는 생성형 AI 의 기술적 어포던스가 디자인 생산 과정을 '구축(construction)' 중심에서 '생성-선택(generation-selection)' 중심으로 전환시키며, 확률적 샘플링 구조가 알고리즘적 수렴(algorithmic convergence)을 유발한다는 메커니즘을 제시한다. 또한 처리 유창성(processing fluency)에 기반한 인지적 편향과 순응적 선택 행동이 이러한 수렴 과정을 강화하는 자기증폭 피드백 루프를 형성함을 논의한다.

실증 분석을 위해 2019 년부터 2024 년까지 미국 특허청(USPTO)의 전수 상표 출원 데이터를 활용하고, CLIP 기반 이미지 임베딩을 사용하여 상표 이미지 간 시각적 유사성을 측정한다. Midjourney 의 2022 년 7 월 공개를 외생적 기술 충격으로 활용한 반복 횡단면 연구 설계를 통해, 생성형 AI 도입 이후 상표 이미지의 시각적 유사성이 유의미하게 증가함을 발견한다. 또한 Intent-to-Use 출원과

고규제 산업에서 이러한 동질화 효과가 더욱 강하게 나타나는 이질적 효과를 확인한다.

본 연구는 생성형 AI 가 창의적 다양성을 촉진할 수 있다는 기존 담론을 넘어, 디지털 창작 도구가 시장 수준의 디자인 분포를 어떻게 재편하는지를 실증적으로 규명한다. 나아가 알고리즘적 수렴과 제도적 압착이라는 이중 메커니즘을 제시함으로써, AI 도입이 조직의 전략적 차별화와 시각적 정체성에 미치는 장기적 함의를 확장한다.

주요어: 생성형 AI, 상표 디자인, 시각적 동질화, 알고리즘적 수렴, 제도적 정당성, 창의적 다양성
학번 : 2024-24572