A Multimodal Visualization System for Exploring Creative Fatigue in Digital Advertising

Satomi Hasegawa
Ochanomizu University
Tokyo, Japan
g2020545@is.ocha.ac.jp

Hayato Ohya
Septeni Japan, Inc.
Tokyo, Japan
hayato.oya@septeni.co.jp

Takayuki Itoh
Ochanomizu University
Tokyo, Japan
itot@is.ocha.ac.jp

***Abstract*—In digital advertising, repeated exposure to the same creative can lead to performance decline, a phenomenon known as creative fatigue. To support the analysis and mitigation of this issue, we developed an interactive visualization system that integrates both textual and visual features of advertisements. The system uses a pre-trained Japanese RoBERTa model to extract text embeddings and LPIPS to evaluate image similarity. Dimensionality reduction with t-SNE and clustering enables interpretable grouping of creatives. We define an effectiveness score based on ad frequency changes after introducing new creatives, and visualize fatigue trends using heatmaps. This system allows advertisers and researchers to intuitively explore how creative attributes relate to ad fatigue, supporting data-driven decisions for creative optimization.**

***Index Terms*—Creative fatigue, digital advertising, visualization system, t-SNE, LPIPS, heatmap, text clustering.**

I. Introduction

In recent years, internet advertising has experienced remarkable growth, both globally and in Japan. The total advertising expenditure in Japan reached ¥7.673 trillion, with internet advertising accounting for approximately ¥3.652 trillion, representing 47.6% of the total expenditure [1] in 2024. Among the various formats, video advertisements have shown the most rapid expansion, with a 23.0% year-on-year increase, largely driven by vertical video content on social media platforms.

As digital advertising continues to expand, advertisers face a growing challenge known as *ad fatigue* (i.e., *creative fatigue*). This phenomenon refers to the decline in advertising performance caused by promoters being repeatedly exposed to the same creative elements—such as identical images or videos—over time. More precisely, *creative fatigue* occurs when an end-user sees the same creative repeatedly, leading to decreased engagement, lower click-through and conversion rates, and rising costs per action [2]. It is distinct from broader audience saturation effects, as performance can often be restored simply by refreshing the creative asset.

Meta's internal analysis shows that repeated exposures to the same creative can reduce the likelihood of conversion by up to 45%, with creative fatigue being a statistically significant factor across various advertiser segments. This highlights the importance of monitoring creative wear-out and developing data-driven strategies for timely creative replacement.

To address this issue, we propose a web-based interactive visualization system that facilitates the analysis and interpretation of creative fatigue in digital ad campaigns.

Unlike prior systems that often focus on predictive modeling or single-modal analysis, our approach integrates multimodal feature analysis (text and image), temporal trend visualization, and interactive filtering within a unified dashboard.

In this study, we introduce several key technical components. We use t-distributed stochastic neighbor embedding (t-SNE) for dimensionality reduction, and the Learned Perceptual Image Patch Similarity (LPIPS) metric to quantify perceptual differences between images. To evaluate changes in ad delivery performance, we define a score (EFFECT), calculated based on changes in delivery frequency before and after the introduction of a new creative. Additionally, the advertising data includes structured identifiers such as AD SET ID (a campaign unit that groups multiple creatives) and AD ID (individual creative identifier).

Textual features are extracted from ad headlines using a pre-trained Japanese RoBERTa model. The system applies t-distributed stochastic neighbor embedding (t-SNE) [3] as a dimensionality reduction scheme for these high-dimensional embeddings to two dimensions, and then groups using K-means clustering. The system uses LPIPS (Learned Perceptual Image Patch Similarity) [3] for visual features to evaluate perceptual similarity between ad images.

The effectiveness of each creative is quantified using a metric called *EFFECT*, which captures changes in ad delivery frequency following the introduction of a new creative. Based on this, the system generates a series of heatmaps:

- A heatmap of average EFFECT values for each combination of text cluster and LPIPS similarity level
- A heatmap tracking changes in EFFECT over time for each text cluster
- A heatmap tracking changes in EFFECT over time for each LPIPS bin (groups divided by specific LPIPS value ranges)

The system allows users to isolate specific creative patterns and gain insight into how fatigue manifests in various segments by supporting filtering by demographic targets and content categories. This enables marketers and researchers to derive actionable strategies for creative refresh and campaign optimization.

DOI 10.1109/IV68685.2025.00029

The remainder of this paper is structured as follows: **Section 2** reviews related work, **Section 3** describes the *proposed visualization system*, **Section 4** presents the visualization results and interpretation, and **Section 5** presents the conclusion and future work.

II. Related Work

Creative fatigue, a phenomenon where repeated exposures to the same ad creative lead to declining user engagement, has attracted growing attention in the fields of online advertising and recommendation systems.

Silberstein et al. [4] proposed a novel approach to combating ad fatigue by incorporating *frequency-over-recent-intervals (FoRI)* features into CTR prediction models. Their method emphasizes the recency and frequency of user-ad interactions by allocating these interactions to unevenly distributed time intervals, placing higher weight on recent impressions. They achieved simultaneous improvements in user experience (e.g., a 15% decrease in repeated ads) and revenue metrics by integrating these features into a deep factorization machine model. Their study also introduced two novel evaluation metrics—Ad Novelty (AN) and Ad Repetitiveness (AR)—to quantify the extent of fatigue and diversity in ad exposure.

Guo et al. [5] extended the classical Nerlove-Arrow advertising response model to account for consumer ad fatigue by incorporating an attenuation factor that modulates advertising effectiveness. Their revised model leads to an inverted-U-shaped relationship between advertising spend and goodwill, reflecting diminishing and even negative returns under excessive exposure. Using optimal control theory, they showed that the optimal expenditure should decrease proportionally across all time periods when fatigue is present, although the direction of the expenditure trajectory remains unchanged.

Aharon et al. [6] proposed an in-model solution for addressing user fatigue in Yahoo Gemini's native advertising platform. Their method, Soft Frequency Capping (SFC), integrates user-ad frequency features directly into a collaborative filtering model (Offset), enabling frequency-sensitive click probability predictions. A/B testing showed a significant 7.3% lift in revenue compared to hard-capped systems, and further analysis revealed a stronger fatigue effect among older users.

While these studies offer valuable perspectives on the modeling and mitigation of creative fatigue, most of them focus on improving CTR prediction accuracy or optimizing budget allocation strategies. In contrast, our study emphasizes the exploratory visualization of creative fatigue using a web-based interactive system.

Our system leverages multimodal features—textual embeddings extracted via a Japanese RoBERTa model and visual similarity scores computed by LPIPS—to provide interpretable representations of ad creatives. By applying dimensionality reduction (t-SNE), K-means clustering, and visualizing performance trends (EFFECT) through heatmaps, our approach enables the discovery of fatigue-related patterns over time.

III. The Visualization System for Exploring Creative Fatigue

To explore and mitigate the impact of creative fatigue in digital advertising, we developed an interactive visualization system that enables analysis of ad content and performance trends. The system allows users to investigate which features in ad creatives are associated with performance degradation and which contribute to fatigue mitigation.

The system comprises two major components: (1) data preprocessing, including feature extraction, dimensionality reduction, and quantification of fatigue; and (2) visualization components that present the processed data in interactive visual formats.

A. System Purpose

This visualization system supports the analysis of creative fatigue by uncovering patterns in both textual and visual characteristics of ad creatives. It provides actionable insights for advertisers to design more engaging creatives and complements predictive modeling approaches by offering an intuitive, interpretable framework for data exploration.

B. Data Preprocessing

1) Dataset Overview: We used anonymized advertising log data provided by Septeni Japan, Inc. The original dataset includes 36,450,438 ad set records and 9,061,137 individual ad creatives. Each creative contains text information such as `HEADLINE` and `DESCRIPTION`, along with image data and performance metrics including frequency and other indicators.

An *ad set* refers to a unit within an advertising campaign that defines targeting conditions such as audience demographics, schedule, and budget. Within each ad set, multiple *ads* (creatives) can be deployed, each consisting of a specific combination of visuals and text. The dataset contains both AD SET IDs and AD IDs, enabling analysis at both group and individual creative levels.

2) Text Feature Extraction and Dimensionality Reduction: To analyze semantic similarity among ad creatives, we extracted textual features from the `HEADLINE` and `DESCRIPTION` fields, where `HEADLINE` represents the main title or slogan of the ad, and `DESCRIPTION` provides supplementary explanation or details about the ad content. We used a pre-trained Japanese language model (`nlp-waseda/roberta-base-japanese`[1]) to encode each field into a 768-dimensional vector. We then applied t-distributed stochastic neighbor embedding (t-SNE) as a dimensionality reduction scheme to the concatenated vectors to two dimensions, followed by K-means clustering to group semantically similar ads. In this study, the number of clusters was empirically set to 6 as a representative example.

3) Quantification of Creative Fatigue: To quantify creative fatigue, we define the effectiveness score EFFECT E as:

$$E = \frac{F_1 - F_2}{F_1} \tag{1}$$

[1] Available at https://huggingface.co/nlp-waseda/roberta-base-japanese

where F_1 is the frequency on the day a new creative was introduced, and F_2 is the average frequency over the following three days. Here, frequency refers to the number of times the same advertisement was shown to the same user.

The use of a three-day average for F_2 was based on a practical business consideration: in many advertising operations, performance evaluation and optimization are conducted in short cycles, typically within a few days. A three-day window strikes a balance between reflecting early performance trends and avoiding excessive delay in decision-making for creative refresh.

Positive E values indicate that the frequency decreased after the introduction of a new creative, which may suggest mitigation of creative fatigue. Conversely, negative E values imply that the frequency increased, indicating continued or worsened fatigue.

We excluded instances where multiple creatives were introduced on the same day from the analysis to isolate the effect of each creative. Additionally, we removed outliers in the E values using the Hampel identifier [7], resulting in a final dataset of 26,125 creatives used for visualization.

C. Visualization Techniques

Figure 1 provides an overview of the proposed visualization system. We designed the system to assist advertisers in identifying patterns of creative fatigue across different ad creatives by enabling interactive exploration of both textual and visual features. Based on the updated implementation, the user interface comprises five main components, described as follows:

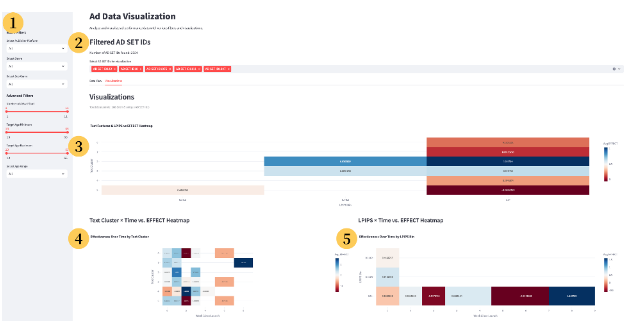

Fig. 1. Overview of the visualization system interface.

① **Filter Controls**: The interface provides filtering functionality based on publisher platform, genre, sub-genre, the number of ads already present in the ad set at the time a new creative is introduced, and the minimum, maximum, and range of target age groups.

② **AD SET ID Selection Panel**: After filters are applied, the system displays a list of corresponding AD SET IDs. Users can select up to five IDs to focus their analysis on a subset of creatives.

③ **Text Cluster × Visual Similarity Heatmap**: Using t-SNE and K-Means clustering, ad creatives are grouped by text similarity. This heatmap visualizes the average EFFECT scores across combinations of textual clusters and visual similarity bins (LPIPS), revealing interactions between copy and image similarity.

④ **Text Cluster × Time Heatmap**: This heatmap captures the average EFFECT score changes over time (in weeks) for each text cluster, helping to detect how fatigue trends vary depending on text similarity.

⑤ **Visual Similarity × Time Heatmap**: This heatmap tracks how average EFFECT scores evolve over time, grouped by LPIPS similarity bins, providing insight into visual fatigue over time.

These components are linked in an interactive dashboard, allowing users to apply filters and compare selected creatives across views. We implemented all components using Plotly and Altair, and computations such as t-SNE, clustering, and LPIPS binning are preprocessed for fast rendering.

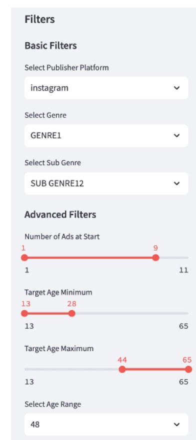

Fig. 2. Filtering Interface for Interactive Selection

1) Interaction and Filtering: Users can filter data by publisher platform, genre, sub-genre, ad count at the time of creative introduction, target age range, and age group (see Figure 2). They can select specific AD SET IDs to compare visualizations across creatives. All views are linked through these filters to support exploratory analysis.

2) Implementation Details: We implemented the system in Python (version 3.9.6) using the Streamlit framework(version 1.42.2). Dimensionality reduction using t-distributed stochastic neighbor embedding (t-SNE, from scikit-learn version 1.3.0), clustering with K-Means (from scikit-learn version 1.3.0), and visualization (Plotly version 5.19.0, Altair version 5.5.0) are performed using precomputed text embeddings and LPIPS scores. The interface supports interactive zooming, brushing, and linked filtering across views.

IV. Visualization Results and Interpretation

This section presents the visualization results using the system introduced in the previous section, and discusses the insights derived from the results. By analyzing patterns across textual and visual feature combinations, as well as differences across platforms and ad sets, we aim to reveal factors contributing to creative fatigue and performance variation in digital advertising campaigns.

A. Cluster-specific Effects in AD SET ID 89

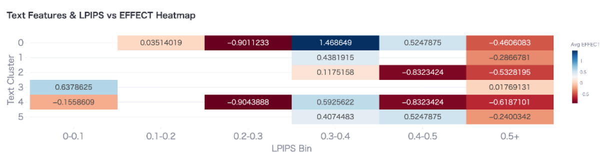

Fig. 3. Heatmap of Text Cluster × Visual Similarity for AD SET ID 89

Figure 3 shows the interaction between text clusters and visual similarity bins (LPIPS) in AD SET ID 89, with colors indicating the average effectiveness (EFFECT) for each combination. The LPIPS values range from 0.0 to 1.0, with lower values indicating higher visual similarity between images and higher values representing greater perceptual dissimilarity.

The following observations can be drawn from Figure 3:

Cluster 0: Highly sensitive to visual changes: Cluster 0 exhibited large variability in performance depending on LPIPS. A moderate change (LPIPS 0.3–0.4) led to a very high EFFECT of +1.46, while a slightly smaller change (0.2–0.3) resulted in a sharp drop to −0.90. Slight changes (0.1–0.2) yielded near-neutral results (+0.03), while larger changes (0.5+) reduced performance (−0.46). This cluster is highly sensitive to visual tuning, with a narrow optimal range that must be carefully targeted.

Cluster 2: Relatively stable but unresponsive: Cluster 2 maintained moderate EFFECT scores across all bins, ranging from −0.5 to +0.1. It exhibited no major spikes or dips, suggesting weak interaction between visual variation and performance.

Cluster 4: Consistently low performance with visual change: Cluster 4 showed consistently low EFFECT across all LPIPS bins. Specifically, values dropped to −0.90 (LPIPS 0.2–0.3), −0.83 (0.4–0.5), and −0.62 (0.5+), indicating that visual modification fails to improve effectiveness.

Overall observations on Visual Similarity: Moderate changes (LPIPS 0.3–0.4) were most effective in Cluster 0, while minimal changes (0.0–0.1) were optimal for Cluster 3. Intermediate changes (0.2–0.3, 0.4–0.5) often led to declines, and extreme visual variation (0.5+) was ineffective across clusters.

Strategic implications: These findings support differentiated creative strategies:

- **Cluster 0:** Focus on LPIPS 0.3–0.4, with precise visual tuning.
- **Cluster 2:** Consider reworking both text and design.
- **Cluster 4:** Improve textual content; design changes are unlikely to help.

B. Platform-based Comparison: AD SET ID 89 on Instagram vs. Facebook

To investigate how platform context affects the relationship between text clusters and visual similarity, we analyzed the same set of creatives from AD SET ID 89 by isolating those shown on Instagram and Facebook, respectively. Figures 4 and 5 show the heatmaps for each platform.

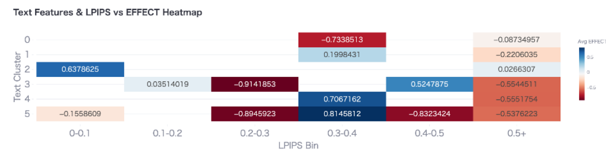

Fig. 4. Heatmap of Text Cluster × Visual Similarity for AD SET ID 89 on Instagram

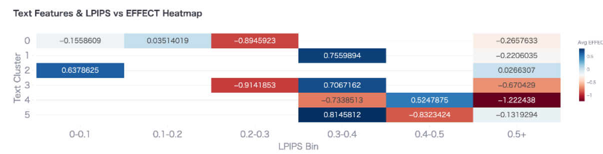

Fig. 5. Heatmap of Text Cluster × Visual Similarity for AD SET ID 89 on Facebook

Result of Instagram: Cluster 2 achieves the highest EFFECT (+0.71) at LPIPS 0.3–0.4, indicating strong synergy with moderate visual updates. Clusters 3 and 5 also perform well in this range (+0.52 to +0.81), but decline at higher LPIPS values. LPIPS 0.2–0.3 consistently yields negative scores, suggesting that subtle visual changes may confuse or fatigue users. Overall, Instagram favors moderate but noticeable visual variation.

Result of Facebook: Cluster 3 also peaks at LPIPS 0.3–0.4 (+0.71), similar to Instagram. However, overall scores are lower, especially at LPIPS 0.5+, where many clusters fall below −0.5. Cluster 2 remains stable across bins, performing best at 0.0–0.1 (+0.64). These results imply that on Facebook, strong text matters more than visual novelty, and users prefer minimal design changes.

A summary of these results is shown in Table I.

TABLE I
QUALITATIVE COMPARISON OF PLATFORM-SPECIFIC TENDENCIES

Aspect	Instagram	Facebook
High-performing LPIPS Bin	0.3–0.4	0.3–0.4, but less consistently
Poor-performing LPIPS Bin	0.2–0.3, 0.5+	0.5+ (most clusters)
Role of text	Synergistic with visuals	Text dominates performance
User reaction	Sensitive to subtle visual change	More stable, less reactive
Creative recommendation	Moderate novelty works best	Focus on copy effectiveness

Strategic Implications: On Instagram, creatives should be visually refreshed within a controlled LPIPS range (0.3–0.4), particularly for text clusters such as 2, 3, and 5. Excessive or insufficient visual change tends to decrease performance. On Facebook, restrained LPIPS values (0.0–0.3) are preferable, and textual strength—especially as found in Cluster 2—appears to be more important. These platform-specific behaviors indicate the need for tailored design and copy strategies depending on the delivery environment.

C. Temporal Effectiveness Trends in AD SET ID 89

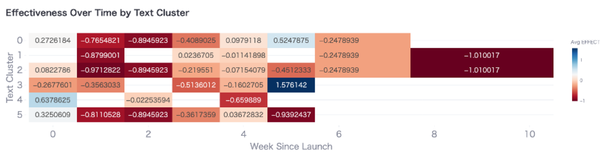

Fig. 6. Effectiveness Over Time by Text Cluster for AD SET ID 89

Figure 6 presents a heatmap showing the weekly changes in average effectiveness (EFFECT) by text cluster for AD SET ID 89. The vertical axis indicates text clusters, and the horizontal axis represents the number of weeks since ad launch. The color intensity corresponds to the average EFFECT score, with blue indicating high performance and red indicating fatigue.

Cluster 3: Strong late-stage performance: Cluster 3 shows a delayed but significant improvement in effectiveness, peaking at +1.57 in Week 6—the highest among all clusters. This suggests that the associated creatives are slow-burn in nature, performing better with prolonged exposure and suitable for brand awareness or long-term campaigns.

Clusters 0, 1, 2: Early effectiveness, rapid fatigue: Clusters 0 through 2 exhibit high or moderate performance in early weeks, but all converge to severe fatigue by Weeks 8 and 10, with EFFECT scores dropping to −1.00. Particularly, Cluster 1 shows a poor starting score (−0.87) from Week 0, indicating short-lived impact. These clusters are more suitable for short-term delivery cycles.

Cluster 5: Early promise, rapid decline: Cluster 5 begins with a moderately positive score (+0.32) in Week 0, but steadily declines to −0.99 by Week 6. This trend implies that while the initial response is positive, the content quickly wears out, and replacement or redesign should be considered within 1–2 weeks.

Cluster 4: Balanced and stable: Cluster 4 starts with a high score (+0.64), then gradually decreases to −0.66 by Week 4. The moderate drop-off suggests that this cluster is effective for medium-term campaigns—strong in the beginning and reasonably resilient thereafter.

Cross-week Insights: From a temporal perspective, Week 0–1 shows a wide range of performance across clusters. By Week 3–4, variation narrows. From Week 6 onward, polarization becomes evident: Cluster 3 remains highly effective, while others transition into fatigue.

Strategic Implications:

- **Cluster 3:** Long-term campaigns. Peaks late and resists fatigue.
- **Cluster 4:** Effective in early to mid-term.
- **Clusters 0–2:** Short-term use only. Rotate or replace after Week 3.
- **Cluster 5:** Moderate initial impact. Best used within the first two weeks.

This heatmap highlights when each cluster is likely to peak and when it begins to fatigue, providing a foundation for timing-optimized creative scheduling.

D. Effectiveness Over Time by Visual Similarity: AD SET ID 89

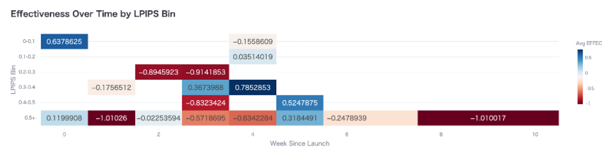

Fig. 7. Temporal EFFECT Trends by LPIPS Bin for AD SET ID 89

Figure 7 illustrates how the average effectiveness (EFFECT) changes over time depending on the visual similarity of ad creatives, as measured by LPIPS Bins. The horizontal axis shows the week since campaign launch, while the vertical axis shows the LPIPS bin representing the degree of visual difference. Color intensity reflects average effectiveness, with blue indicating high effectiveness and red indicating fatigue.

LPIPS 0.0–0.1: High initial performance, fast fatigue: Creatives with minimal visual change performed well initially (+0.64 at Week 0), but their effectiveness declined by Week 4 (−0.15), suggesting susceptibility to fatigue. This range is suitable for short-term reminders but not long-term deployment.

LPIPS 0.2–0.3: Consistently poor performance: This bin showed strong negative EFFECT values across time (e.g., −0.89 at Week 2), making it the least effective range. Slight visual variation may confuse users or signal repetitiveness, accelerating fatigue.

LPIPS 0.4–0.5: Inconsistent effects: This bin produced unstable results: severe underperformance at Week 2 (−0.83), followed by improvement at Week 4 (+0.52). Performance may depend on specific creative context or user expectations.

LPIPS 0.5+: High initial impact, sharp fatigue: Major visual changes yielded strong initial EFFECT (+1.19 at Week 0) but deteriorated rapidly, reaching −1.00 by Week 8. This bin is effective for burst campaigns but not for sustained use.

Strategic Implications:

- **0.0–0.1:** Use for short-term reminder creatives.
- **0.2–0.3:** Avoid. Associated with poor performance.
- **0.3–0.4:** Primary target range for visual design. Balanced and effective.
- **0.4–0.5:** Test cautiously. May require A/B testing.
- **0.5+:** Suitable for time-limited, high-impact campaigns.

This heatmap highlights how creative fatigue is influenced by the degree of visual novelty over time. Proper LPIPS tuning can improve both short-term impact and long-term performance management.

E. Temporal Effectiveness by Text Cluster: AD SET ID 182

Figure 8 visualizes how the effectiveness (EFFECT) of each text cluster evolves over time for AD SET ID 182. The vertical

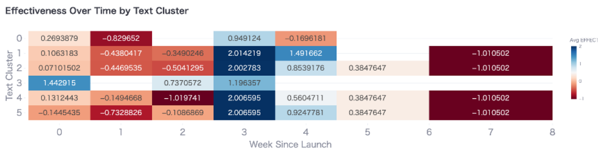

Fig. 8. Effectiveness Over Time by Text Cluster for AD SET ID 182

axis represents clusters of ad copy, while the horizontal axis shows the week since launch. Color intensity corresponds to the average EFFECT score.

Cluster 3: Long-lasting effectiveness: Cluster 3 demonstrates consistently high performance, maintaining EFFECT scores above +0.7 throughout the campaign. With +1.44 at Week 0 and +1.19 at Week 3, it is the most stable and fatigue-resistant cluster.

Clusters 2, 4, 5: Burst-type performance: Clusters 2, 4, and 5 peak sharply at Week 3 (EFFECT = +2.00) but quickly decline to -1.01 by Week 6 or 8. These clusters are highly effective in the short term but suffer from rapid fatigue, making them suitable for limited-time promotions or burst campaigns with early replacement strategies.

Clusters 0 and 1: Initially weak, temporarily recover: These clusters perform poorly in the first two weeks (EFFECT ≈ -0.8 to -0.3), but recover around Week 3 (+0.9 to +2.0), only to deteriorate again by Week 6. This suggests that with proper optimization—such as design or copy refinement—they can yield temporary improvements.

Strategic Implications:

- **Cluster 3:** Use in long-term, brand-focused campaigns.
- **Clusters 2, 4, 5:** Short-lived but powerful—ideal for burst promotions. Replace after Week 3.
- **Clusters 0, 1:** Candidates for copy or design testing. May perform after adjustment.

This visualization clarifies when and how fatigue manifests for different copy types, offering practical insights into optimal campaign timing and creative rotation strategies.

V. Conclusion and Future Work

A. Conclusions

This study proposed a multimodal visualization system to support the analysis of creative fatigue in digital advertising. The system provides an interactive environment for interpreting how creative characteristics affect audience response over time, by integrating textual embeddings, LPIPS-based image similarity, and time-series ad performance data.

Through a series of visualization components, such as text cluster, LPIPS heatmaps and temporal trend heatmaps, we demonstrated that different ad types exhibit distinct patterns of effectiveness and fatigue. Our findings highlight that **certain creative strategies, such as moderate visual variation (LPIPS 0.3–0.4) or stable messaging (Text Cluster 3), lead to more sustainable engagement**. Conversely, some combinations result in rapid fatigue or inconsistent performance, indicating the need for fine-grained creative tuning.

While our study did not identify universal patterns of creative fatigue across all digital advertising campaigns, we successfully extracted meaningful patterns within individual AD SETs.

B. Future Work

The visualization system proposed in this study offers a valuable approach for understanding the structural dynamics of creative fatigue. However, there remain several opportunities for future improvement and extension.

One important direction is the integration of visualization with predictive modeling techniques. By combining exploratory heatmaps with machine learning techniques, the system could support automatic recommendations for creative refresh timing and early detection of fatigue symptoms.

Furthermore, we plan to deploy the system in real advertising environments and conduct user studies with advertising professionals to evaluate its utility as a decision-support tool. Through practical use and user feedback, we aim to assess its impact on campaign performance and incorporate the findings to refine the interface, improve interpretability, and enhance overall usability.

These extensions will allow us to build a comprehensive, practical system that bridges the gap between algorithmic analysis and creative strategy, ultimately contributing to more effective, data-driven campaign planning.

References

[1] Dentsu, 'Advertising Expenditures in Japan in 2024: Internet Advertising Media Expenditures Detailed Analysis," Feb. 2025. [Online]. Available: https://www.dentsu.co.jp/news/release/2025/0227-010853.html [Accessed: Mar. 27, 2025]. (in Japanese)

[2] Analytics at Meta, 'Creative Fatigue: How Advertisers Can Improve Performance by Managing Repeated Exposures," *Medium*, May 2023. [Online]. Available: https://medium.com/@AnalyticsAtMeta/creative-fatigue-how-advertisers-can-improve-performance-by-managing-repeated-exposures-e76a0ea1084d [Accessed: Mar. 27, 2025].

[3] L. van der Maaten and G. Hinton, 'Visualizing Data Using t-SNE," *J. Mach. Learn. Res.*, vol. 9, pp. 2579–2605, 2008.

[4] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, 'The Unreasonable Effectiveness of Deep Features as a Perceptual Metric," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 586–595, 2018.

[5] N. Silberstein, O. Shoham, and A. Klein, 'Combating Ad Fatigue via Frequency-Recency Features in Online Advertising Systems," in *Proc. ACM Int. Conf. Inf. Knowl. Manag. (CIKM)*, pp. 4822–4828, 2023.

[6] R. Guo and Z. Jiang, 'Optimal Dynamic Advertising Policy Considering Consumer Ad Fatigue," *Decis. Support Syst.*, vol. 187, p. 114323, 2024.

[7] M. Aharon, Y. Kaplan, R. Levy, O. Somekh, A.Blanc, N. Eshel, A. Shahar, A. Singer, and A. Zlotnik, 'Soft Frequency Capping for Improved Ad Click Prediction in Yahoo Gemini Native," in *Proc. ACM Int. Conf. Inf. Knowl. Manag. (CIKM)*, pp. 2793–2801, 2019.

[8] F. R. Hampel, 'The Influence Curve and Its Role in Robust Estimation," *J. Amer. Statist. Assoc.*, vol. 69, pp. 383–393, 1974.