Behind the Canvas: A Human-AI Workflow for Tracing 19th-Century Photographers and Studio Backdrops

Vikram Mohanty, Jude Lim, Terryl Dodson, Kurt Luther

Department of Computer Science, Virginia Tech
Arlington, VA, USA
vikrammohanty@vt.edu, lkyoong428@vt.edu, tdodson3@vt.edu, kluther@vt.edu

Abstract

Painted backdrops in 19th-century photography can provide critical provenance clues that are of interest to historians, museum professionals, and collectors. Yet, the absence of specialized tools leaves researchers to rely on labor-intensive manual methods. To address this challenge, we introduce *BackTrace*, a human-AI collaborative system that helps users discover and organize historical photographs sharing similar painted backdrops. Our system allows users to iteratively improve search results by providing feedback to the system, and subsequently create public "collections" that can aid in future research.

Introduction

In the second half of the 19th century, photographers used painted backdrops (see Figure 1) to enhance aesthetics, convey status, and offer diverse settings within studio confines, reflecting both practical needs and the era's cultural zeitgeist (Brown 2018). Many of these backdrops are sufficiently unique that they can be attributed to specific photographers, providing important provenance clues for historical photo research (Fleischer and Kelbaugh 2020; Fleischer 2022; Zeitlyn 2010).

Historians, museum professionals, and collectors with interests in American Civil War photography, for instance, often utilize these painted backdrops to enrich metadata related to photographers and locations (Canberg 2012; Fleischer and Kelbaugh 2020; Fleischer 2022). Yet, the predominant methods for researching and cataloging these backdrops remain manual and labor-intensive (Keller 2021; Fulmer 1990; Fleischer 2022). This problem limits the scalability of individual efforts and poses challenges to consolidate and access the collective knowledge on the subject.

Computer vision (CV) techniques can help researchers to sift through vast image pools, by presenting a curated selection of relevant images for more detailed visual examination (Cai et al. 2019; Mohanty et al. 2019). Notably, Zeitlyn, Coto, and Zisserman (2021) analyzed the work of Jacques Touselle, a Cameroonian photographer, using facial recognition and pattern matching. This study highlighted the capabilities of computer vision for backdrop-focused investigations but lacked backdrop-specific features and categorization based on unique backdrops.

Figure 1: Examples of American Civil War painted backdrops, showing different painted features such as pillars, tents, foliage, and balustrades. (Bbreneman 2020; Morin 2018; Pomerantz 2021)

Despite the potential of such Content-based Image Retrieval (CBIR) approaches, there are no publicly available systems that specifically cater to painted backdrop research. Popular reverse image search engines like Google and Bing Images often fail to capture the nuances specific to this research task. Furthermore, fine-tuning an existing CV model for the task would require an extensive amount of manual annotation due to the absence of a tagged database of painted backdrops.

To address this research gap, we present BackTrace, a CBIR system designed to discover and organize painted backdrops in historical photos. In this work, we focus on the American Civil War (1861–1865), the first major conflict to be extensively documented via photography, underscoring the intertwined history of the Civil War and painted backdrops (mil 1996; Keller 2021). Given its vast repository of Civil War photos, we chose Civil War Photo Sleuth (CWPS),[1] an online community dedicated to identifying unknown Civil War portraits (Mohanty et al. 2019), as the base platform for developing BackTrace. Here, we specifically address the following research question: *How can we design a system that supports users in discovering and organizing historical photos with matching painted backdrops?*

Presented at the Works-in-Progress and Demonstrations track, AAAI HCOMP 2023.

[1]http://www.civilwarphotosleuth.com

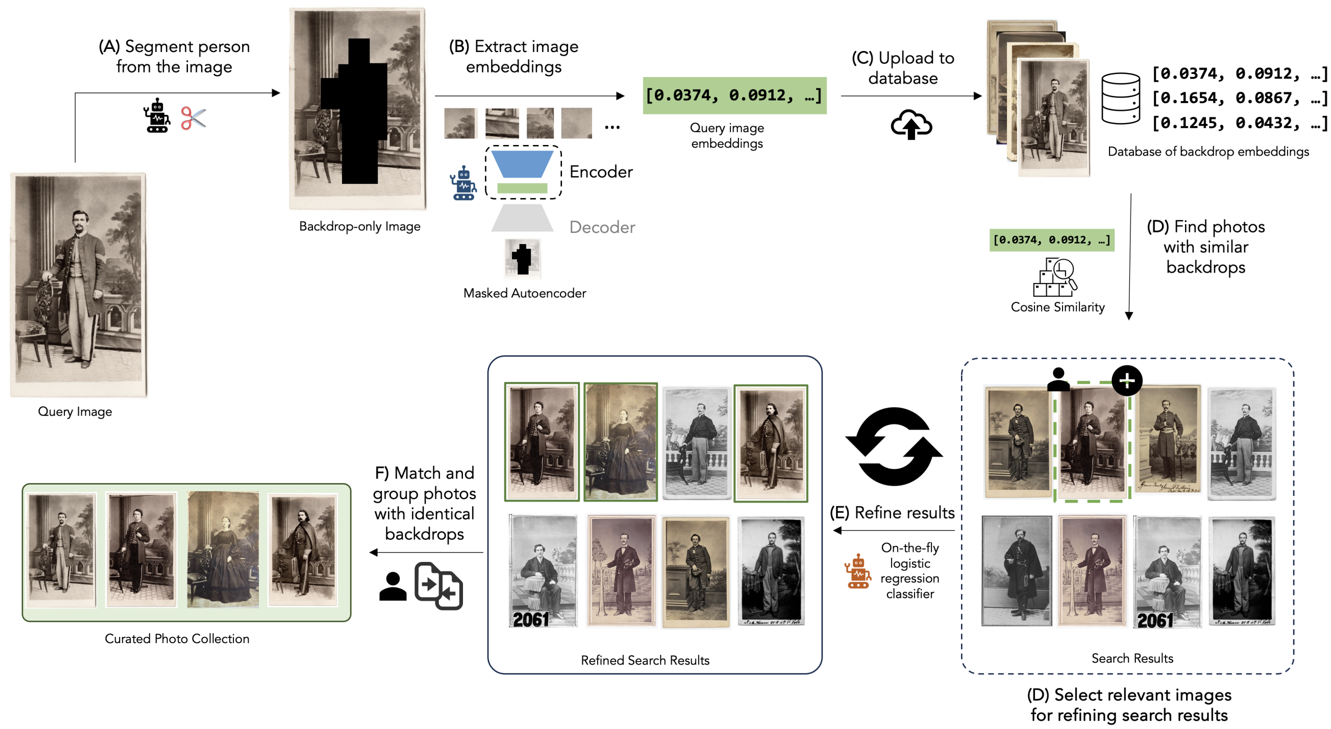

Figure 2: Overview of BackTrace's workflow for discovering and organizing photos with matching painted backdrops.

System Overview

Imagine Jill, an avid collector of historical photos, who acquires a 19th-century photograph featuring a distinctively painted studio backdrop. However, it lacks any imprint identifying the photographer's name or studio location. To learn more about the photographer, Jill searches for other photos featuring the same unique backdrop.

To assist Jill in finding photos with matching studio backdrops, we introduce BackTrace, an exploratory CBIR tool. BackTrace combines computer vision with human feedback to discover similar backdrops and supports users in curating collections of photos with matching backdrops (see Figure 2). These collections subsequently enable downstream investigative efforts, such as identifying the photographer or studio information.

Jill starts her investigation by uploading the query photo to BackTrace. She tags the photo's backdrop with descriptive terms, serving as search keywords for subsequent stages. Using PixelLib (Olafenwa 2021), BackTrace then segments the image to focus solely on the backdrop, removing any individual subjects present (see Figure 2A).

After isolating the backdrop, BackTrace uses MAE (Masked Autoencoder) (He et al. 2021), a transformer model based on Vision Transformers (ViT) (Dosovitskiy et al. 2021), to extract embeddings (see Figure 2B). Given MAE's capability to reconstruct missing sections of an image using sparse patches, the encoder's embeddings suggest a potential to recreate the full backdrop even without any individual present. Typically, MAE generates N $\times$ 1024-dimensional vectors, where N represents the varying number of patches based on the image. We adapted the MAE to our requirements, averaging patches to produce a uniform 1 $\times$ 1024-dimensional vector. This method was applied to all images in the CWPS database (the base platform), forming our search pool (see Figure 2C).

BackTrace then uses a K-Nearest Neighbors (KNN) algorithm to identify images with backdrops resembling Jill's query, ranked by cosine similarity (see Figure 2D). Recognizing that humans might emphasize specific features like flags or tents, BackTrace incorporates relevance feedback (Fogarty et al. 2008; Cai et al. 2019; Rui et al. 1998; Lee and Weld 2020). Jill analyzes the results and selects the images she deems relevant, prompting an on-the-fly training of a logistic regression classifier. This classifier, weighing positive examples (selected images and the query image) tenfold compared to negative ones (remaining KNN results), computes probability scores for the top-1000 KNN matches (i.e., likelihood of belonging in the same class as the positive examples). BackTrace then refreshes the search results based on these scores (see Figure 2E). Jill also has the option to input backdrop descriptions for keyword-based filtering, allowing her to iteratively refine results.

Jill then uses BackTrace's comparison interface to closely inspect search results she believes share the same backdrop as her query image. This tool offers a side-by-side comparison, and when she confirms a match, those photos are added to a collection. Each collection gets a dedicated page, showcasing member photos and related metadata. Here, Jill discovers another photo bearing a photographer's backmark, helping her identify the query photo's photographer.

Preliminary Results and Future Work

We conducted a mixed-methods evaluation involving 9 participants with varying expertise levels, encompassing observations, log analysis, and interviews. Our findings showed that BackTrace effectively assists users in matching photos with similar backdrops. Participants expressed satisfaction with BackTrace's AI performance, highlighting its speed and efficiency compared to their current methods. Further, participants successfully curated collections of photos with matching backdrops, while recognizing the potential of these collections as valuable repositories for future cross-referencing.

As part of future work, we aim to refine BackTrace based on user feedback, followed by a public release and a subsequent large-scale study. We also plan on conducting an in-person crowdsourcing event to create a repository of Civil War backdrops from the CWPS database. Longer term, we hope to extend BackTrace's capabilities to other historical and even modern photos, starting with testing our system on non-Civil War painted backdrops from the 19th century.

Acknowledgements

This research was supported by NSF IIS-1651969 and the GEM Fellowship Program.

References

1996. The Grand Illusion: Painted Backdrops in Civil War Photography. *Military Images*, 17(5): 20–29.

Bbreneman. 2020. Photo 39386.

Brown, T. 2018. Painted Backgrounds for Turn of the Century Photographers.

Cai, C. J.; Reif, E.; Hegde, N.; Hipp, J.; Kim, B.; Smilkov, D.; Wattenberg, M.; Viegas, F.; Corrado, G. S.; Stumpe, M. C.; et al. 2019. Human-centered tools for coping with imperfect algorithms during medical decision-making. In *Proceedings of the 2019 chi conference on human factors in computing systems*, 1–14.

Canberg, K. 2012. ENOCH LONG: Benton Barracks Photographer. *Military Images*, 31(6): 3–5.

Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. arXiv:2010.11929.

Fleischer, A. O. 2022. Origins, artistry and photographers: Reflections on Two Years. *Military Images*, 40(2 (220)): 75–77.

Fleischer, A. O.; and Kelbaugh, R. J. 2020. Origins, artistry and photographers: A Daguerreian Pioneer at the Rendezvous of Distribution. *Military Images*, 38(2 (212)): 72–75.

Fogarty, J.; Tan, D.; Kapoor, A.; and Winder, S. 2008. CueFlik: Interactive Concept Learning in Image Search. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, 29–38. New York, NY, USA: Association for Computing Machinery. ISBN 9781605580111.

Fulmer, R. 1990. BACKDROPS. *Military Images*, 11(5): 6–10.

He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2021. Masked Autoencoders Are Scalable Vision Learners. arXiv:2111.06377.

Keller, K. 2021. A Great Variety of New and Fine Designs: Advertisements for Painted Backgrounds, 1856–1903.

Lee, B. C.; and Weld, D. S. 2020. Newspaper Navigator: Open Faceted Search for 1.5 Million Images. In *Adjunct Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, UIST '20 Adjunct, 120–122. New York, NY, USA: Association for Computing Machinery. ISBN 9781450375153.

Mohanty, V.; Thames, D.; Mehta, S.; and Luther, K. 2019. Photo Sleuth: Combining Human Expertise and Face Recognition to Identify Historical Portraits. In *Proceedings of the 24th International Conference on Intelligent User Interfaces*, IUI '19, 547–557. New York, NY, USA: Association for Computing Machinery. ISBN 9781450362726.

Morin, D. 2018. Photo 24490.

Olafenwa, A. 2021. Simplifying object segmentation with pixellib library. *Online.(2021). https://vixra.org/abs/2101.0122*.

Pomerantz, H. 2021. Photo 44739.

Rui, Y.; Huang, T.; Ortega, M.; and Mehrotra, S. 1998. Relevance feedback: a power tool for interactive content-based image retrieval. *IEEE Transactions on Circuits and Systems for Video Technology*, 8(5): 644–655.

Zeitlyn, D. 2010. Photographic Props / The Photographer as Prop: The Many Faces of Jacques Tousselle. *History and Anthropology*, 21(4): 453–477.

Zeitlyn, D.; Coto, E.; and Zisserman, A. 2021. Bounding an archiving: assessing the relative completeness of the Jacques Toussele archive using pattern-matching and face-recognition. *Visual Studies*, 0(0): 1–25.