BackTrace: A Human-AI Collaborative Approach to Discovering Studio Backdrops in Historical Photographs

Jude Lim, Vikram Mohanty, Terryl Dodson, Kurt Luther

Center for Human-Computer Interaction and Department of Computer Science, Virginia Tech, USA
{lkyoong428, vikrammohanty, tdodson3, kluther}@vt.edu

Abstract

In historical photo research, the presence of painted backdrops have the potential to help identify subjects, photographers, locations, and events surrounding certain photographs. However, there are few dedicated tools or resources available to aid researchers in this largely manual task. In this paper, we propose BackTrace, a human-AI collaboration system that employs a three-step workflow to retrieve and organize historical photos with similar backdrops. BackTrace is a content-based image retrieval (CBIR) system powered by deep learning that allows for the iterative refinement of search results via user feedback. We evaluated BackTrace with mixed-methods evaluation and found that it successfully aided users in finding photos with similar backdrops and grouping them into collections. Finally, we discuss how our findings can be applied to other domains, as well as implications of deploying BackTrace as a crowdsourcing system.

Introduction

The issue of missing metadata is frequently cited as a key problem in the curation of rapidly growing cultural heritage image collections (Cordell 2020), with many computer vision (CV) techniques applied towards solving it (Mohanty et al. 2019; Chumachenko et al. 2020; Zeitlyn, Coto, and Zisserman 2021; Caraffa et al. 2020; Resig 2014). One aspect of this problem is inadequate artist or photographer information in visual collections. This information is frequently used as a means of cataloging these large collections, and serves as a step towards other forms of identification, such as dating or attribution to certain individuals.

While the organization and identification of these images are generally informed by the contents of the entire image, a focus on background information can yield other interesting findings. A prime example of this information can be found in the painted backdrops of the 19th century (Figure 1), which were widely used in portrait photography of the period (Keller 2021; Haidt 2011). Many of these backdrops are sufficiently unique that they can be attributed to specific photographers (Zeitlyn 2010; Fleischer 2022), contributing directly towards enriching photographer metadata in collections. However, current practice around researching and organizing these backdrops are largely manual and time-consuming (Keller 2021; Fulmer 1990; Fleischer 2022).

Figure 1: Examples of the “Benton Barracks” painted backdrop, showing common backdrop features such as a cannon, foliage, and a tent. These examples also showcase two photo formats of the era: the Tintype and the Carte-de-Visite (CDV) (Montgomery1861 2018; Karnes 2020; Kurtz 2021).

CV has great potential to aid in populating metadata by speeding up the analysis of visual data, while ensuring that inconsistencies in metadata do not obscure objects of interest. Prior works have successfully used pattern matching and iterative refinement to group similar photos (Caraffa et al. 2020; Crissaff et al. 2017; Lee and Weld 2020) and even photos with similar backgrounds (Zeitlyn, Coto, and Zisserman 2021). However, none of these approaches computationally isolates and analyzes backdrops as a distinct feature.

To address this research gap, we present BackTrace, a Content-Based Image Retrieval (CBIR) system that facilitates the discovery and organization of painted backdrops in historical photos. For this project, we focused on the American Civil War (1861–1865), the first major conflict to be extensively documented via photography and whose history is closely connected to that of painted backdrops (Military Images 1996; Keller 2021). 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 of our system due to its vast repository of Civil War photos.

We conducted a mixed-methods evaluation, including user studies, log analysis, and interviews. We found that BackTrace successfully aided users in matching photos with identical backdrops. We also found that users were success-

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

ful in creating and browsing backdrop collections. Our primary contributions are:

- proposing a workflow for historical photo backdrop identification,
- implementing a human-AI CBIR system, BackTrace, developed using this workflow,
- evaluating BackTrace with participants engaged in Civil War photo research, and
- documenting Civil War backdrop research current practices and challenges.

We also discuss how our findings and workflow can be applied to other forms of photography and imagery, as well as the workflow's potential as a crowdsourcing system.

Related Work

Historical Backdrops

Historical photo research typically aims to explore the connection between backdrops, the photographers who use them, and the subjects whose photographs are enriched by them. Most of this research focus on analyzing images as a composition of posing, clothing, and accessories, along with background objects. Painted backdrops frequently feature as a part of these research efforts, and play a significant role in conveying meaning and social norms (Burke 2001; Zeitlyn 2010; Haidt 2011).

Although many backdrops can share similarities, variations in execution and detail allow each backdrop to be reliably associated with a photographer, studio, or location (Fleischer and Kelbaugh 2020; Fleischer 2022; Zeitlyn 2010). This connection helps enrich photo metadata by allowing information such as the photographer or location to be inferred between photos sharing the same backdrop. A popular example is the "Benton Barracks" backdrop used by Enoch Long (Figure 1), which helps researchers localize portraits of soldiers from various states who passed through St. Louis, Missouri during the war (Ertzgaard 1989; Canberg 2012).

However, some complexities must be considered. Photographers frequently altered the lighting, position of backdrop, and other elements in the photo that could obscure the backdrop. Photographers can also own multiple backdrops or even modify them over time (Keller 2021). The process for photo investigations using a painted backdrop involves first searching for photos with similar backdrops from sources such as books, digital collections, and physical collections. If matches are found, the researcher investigates other visual clues in the photo, such as the subject's uniform, any notes or inscriptions on the photo itself, military service records of other soldiers pictured with the same backdrop, and biographies of photographers who might have taken the photo (Lindberg 2021). This process of seeking and creating connections requires specialized domain knowledge. Even if matching backdrops are found, successful identification of the photographer or subject is not guaranteed, paralleling observations from other types of Civil War photo investigations (Mohanty et al. 2019).

A central reference point has the potential to speed up the process of such research. However, there is no authoritative database or reference for painted backdrops. Furthermore, current processes are not suitable for organizing large scale repositories of backdrops due to their manual, time-consuming nature. Utilizing a large collection of photos as a base, we design a workflow to facilitate backdrop research and subsequently allow users to save their research efforts in a central database.

Human-AI Collaborative Approaches for Visual Investigation Tasks

The combination of computer vision (CV) and human computation has revolutionized investigative tasks by using algorithms to quickly extract meaningful information from visual data. This approach allows investigations to be performed at scale and with impressive accuracy.

Visual recognition techniques have been applied to various domains such as citizen science, where they have been combined with human intelligence to identify plant and animal species (Nugent 2018). CV's ability to recognize scenes have also been used to combat human trafficking (Stylianou et al. 2017, 2019) as well as assist visually impaired persons explore indoor environments (Afif et al. 2020).In the context of historical photo identification, facial recognition combined with crowdsourcing has successfully assisted in identifying soldiers in Civil War photos (Mohanty et al. 2019). We take inspiration from these application for BackTrace, utilizing CV to assist and empower users in painted backdrop research.

Another method to enable these investigations combines content based image retrieval (CBIR) (Datta et al. 2008) and relevance feedback (RF) (Rui et al. 1998). Also known as reverse image search, CBIR allows for searches to be based on visual features instead of potentially scarse metadata. RF complements CBIR by providing a framework for users to iteratively improve the accuracy of search results by providing feedback to the system.

While publicly available CBIR systems exist (Google 2023; Microsoft 2023), these are general-purpose solutions and do not capture the nuances specific to this research task. Specialized solutions are often required to achieve optimal task performance. While fine-tuning an existing CV model may be an option for our context, it requires an extensive amount of manual annotation due to the absence of a tagged database of painted backdrops. Often, a combination of human and artificial intelligence is utilized to tackle more complex search tasks such as ours. For example, SMILY focuses on hispathology images and allows users to "refine" their searches using three different types of feedback (Cai et al. 2019). This inspired BackTrace's visual refinement feature, which allows users to rerank search results by selecting images with visually similar backdrops.

Exploring and Curating Visual Cultural Heritage Datasets

With the rapid digitization of galleries, libraries, archives, and museums (GLAMs), there is a growing need for strategies to explore and organize these vast collections. While

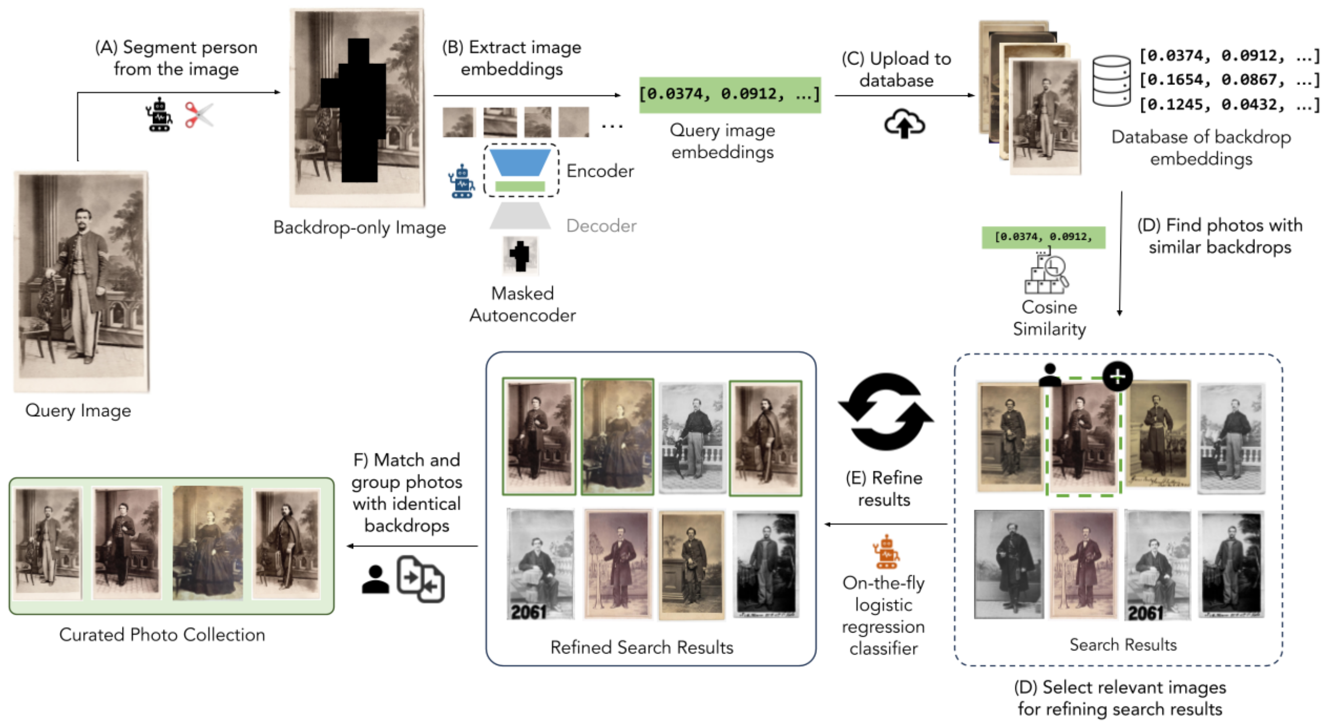

Figure 2: BackTrace's workflow. (A) Subjects are masked out of the photo. (B) Image embeddings are extracted from the masked photo. (C) Embeddings are stored in a database for future use. (D) After selecting a query photo, users are shown a list of photos with similar backdrops. Users can select photos from these results to use for refinement. (E) Search results are reordered to match photos selected for refinement. (F) Users can match and group photos with identical backdrops into collections.

design strategies that maximize exploration and accessibility (Whitelaw et al. 2015) have been proposed, the application of advanced computational methods remains key to achieving this goal (Cordell 2020; van Strien et al. 2022).

Distant viewing acts as a visual alternative to the more popular distant reading (Moretti 2005). While distant reading computationally extracts high-level concepts from text, distant viewing does so on visual content. These concepts are used to transform a collection into a form more suitable for closer inspection (Arnold and Tilton 2019). For example, Arnold and Tilton (2019) ran images through a compression algorithm to acquire a high-level representation of images before clustering similar representations together. This methodology heavily inspires our approach to pre-processing and presenting visual data in BackTrace.

Indeed, there has been a significant increase in attempts at applying computational analysis to these collections. Some use visual features to apply labels to historical photographs (Chumachenko et al. 2020), while others choose to minimize or even ignore textual metadata altogether (Resig 2014; Caraffa et al. 2020).

Following their previous research (Zeitlyn 2010), Zeitlyn, Coto, and Zisserman (2021) used a combination of pattern-matching and face-recognition to computationally analyze the work of Jacques Touselle. They successfully assessed the completeness of their own archives, using backdrops extensively as a unit of analysis. Rather than using pattern matching, which may include background imagery as well as other patterns such as clothing, our work seeks to computationally isolate and analyze backdrops as distinct features.

Conceptually, BackTrace is most similar to Newspaper Navigator, which analyzed a vast collection of digitized newspaper pages (Lee et al. 2020). The search interface launched alongside the database (Lee and Weld 2020) included a feedback system inspired by Cueflik (Fogarty et al. 2008) and shared similarities with SMILY (Cai et al. 2019). Users were able to train an underlying AI to refine their search results by selecting images of their preference.

We expand on these works by emphasizing the backdrop as a unit of analysis and focusing on the exploratory nature of backdrop research. By relying primarily on visual features, the dataset is made explorable even when metadata is scarce. We also provide tools for users to retain and make public the results of their research, laying the foundation for a central reference of backdrops.

System Description

We developed BackTrace (Figure 2), an exploratory CBIR tool designed to assist users in discovering photos that share the same painted studio backdrops using computer vision

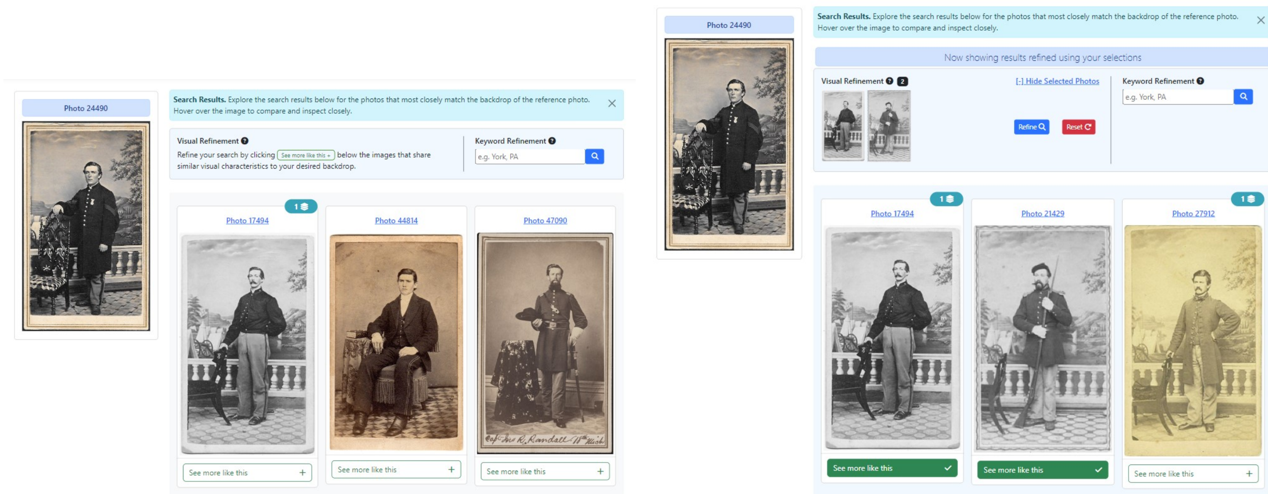

(a) Before refinement

(b) After refinement

Figure 3: Search result page before and after refinement.

and human feedback. Further, it also supports the organization of these photos into meaningful backdrop collections for aiding future photo investigations.

Consider a scenario where a user has a 19th-century photograph with a distinctive painted studio backdrop but no backmark bearing information about the photographer. In order to learn more about the photographer, the user wants to find other photos that share the same backdrop.

On BackTrace, a user goes through an *iterative discovery* process that comprises of three stages: a) *Isolating Backdrops*, b) *Exploring Matches*, and c) *Curating Collections*.

Isolating Backdrops

We initialize BackTrace's iterative discovery process by focusing attention on the backdrop region within the image.

Describing Backdrops: The user begins the process by first assigning descriptive tags for the photo's backdrop, serving as search keywords (e.g., tents, flags, cannons, etc.) in subsequent stages. Given the absence of a well-established taxonomy for backdrops, the user is provided with a free-form text input field, enabling them to input diverse and potentially unique descriptions.

Masking and Extracting Features: After the user tags the photo, BackTrace uses PixelLib (Olafenwa 2021), an image segmentation library, to segment the individuals present in the image. The portion of the image that remains, encompassing the backdrop and any remaining foreground elements (i.e., floor, props, etc.), is then used for generating image embeddings (Figure 2B).

BackTrace uses MAE (Masked Autoencoder) (He et al. 2021), a transformer model based on Vision Transformers (ViT) (Dosovitskiy et al. 2021), for extracting the embeddings. MAE's design is particularly suitable for our purpose as it allows for the reconstruction of missing pixels using only a quarter of the original image in the form of random patches. The assumption is that the encoder-derived embeddings carry enough significant information to ideally allow for a complete reconstruction of the backdrop, even when no individuals are present in the image.

We modified the MAE architecture to accept our segmented backdrop mask in place of the default random mask it usually employs. In its original form, MAE constructs an N (fixed number of patches from the unmasked region) $\times$ 1024-dimensional vector. Given that we are segmenting the person out of the image, the number of unmasked patches (N) can vary from image to image. To maintain the consistency of each vector, we compute the mean of all unmasked patches and generate a 1×1024-dimensional vector to represent the backdrop of a given image (Figure 2B), which will be used for comparing against the backdrops of other images in the search pool (Figure 2C).

Preparing the Search Pool: In order to ensure that BackTrace retrieves relevant images that contain painted backdrops, we trained a Support Vector Machine (SVM) classifier in an offline pre-processing step for filtering out images in the search pool which did not have any backdrops, such as images with portrait busts or plain white backgrounds (details about the training process in Appendix). We used the database provided by CWPS to extract the image embeddings (using the process described above) and set up the search pool for BackTrace.

Exploring Matches

Once the image embeddings and the search pool have been initialized, BackTrace retrieves similar-looking backdrops for the user to discover potential matches by iteratively refining the results.

Finding Similar Backdrops: BackTrace employs a K-Nearest Neighbors (KNN) algorithm to retrieve images from the search pool that resemble the backdrop of the user's query photo. The KNN algorithm compares the cosine similarity between the embeddings of the query image and those in the search pool to find the top-1000 closest matches (Figure 2D). These results, sorted according to the cosine similarity scores, are then presented to the user in a grid format on an ”overview” page, displaying each image in its entirety, allowing the user to make quick comparisons with the query image (Figure 3a).

Refining the Search Results: While the KNN algorithm is effective at retrieving similar matches based on image embeddings, the results may not always align with human perceptions of similarity. This is due to the fact that the algorithm takes into account the overall structure and features of the image, whereas a human observer might focus on specific key features such as flags, tents, or particular patterns. To bridge this gap, BackTrace draws from prior work on relevance feedback (Fogarty et al. 2008; Cai et al. 2019; Rui et al. 1998; Lee and Weld 2020) and allows users to iteratively refine the search results according to their needs.

For example, a user might find that the initial search results do not encompass certain key features present in the query photo. This could occur if the backdrop in the query photo contains a distinctive element, like a flag, absent in some of the retrieved images. Similarly, the KNN algorithm might rank images with identical backdrops differently due to variations in factors such as viewpoint, lighting, or other image characteristics. This could result in some images, which are relevant to the user's query, appearing lower in the search results, despite having the same backdrop. In these situations, BackTrace provides the user with the option to refine the search results such that the most relevant images appear higher.

To support this, BackTrace provides users with a “see more like this” button below each search result. The user can select images they deem relevant, prompting the system to refine the search results. This initiates an on-the-fly training of a logistic regression classifier, which uses the embeddings of the selected images and the query image as positive examples, while treating the embeddings of the remaining search results as negative examples. To emphasize the importance of the positive examples as well as address the imbalanced nature of the training data, the classifier assigns a weight that is ten times greater for the positive examples than the negative ones. This number was decided via trial and error.

The classifier calculates predicted probability scores, which estimates the likelihood that each image in the KNN's top-1000 potential matches belongs to the same class as the user's selected relevant images (i.e., the positive training examples) (Figure 2E). BackTrace then presents the user with a refreshed list of search results, re-ranked based on these probability scores (Figure 3b). The user also has the flexibility to update their selection of relevant images iteratively, allowing for continued refinement of the search results by re-training the classifier with each pass.

In addition to visual refinement, BackTrace offers the user an additional tool for refining results — a text box for keyword-based filtering. Here, the user can enter backdrop descriptions. The system then adjusts the results to prioritize images that match these keyword filters, pulling them to the top of the results list.

Curating Collections

Once the user is satisfied with the search results, they can start to closely examine the images for similar backdrops and organize them into a collection. Collections on BackTrace are groups of photographs that users curate based on their perception of backdrop similarity (Figure 2F).

Comparing Backdrops and Collections: BackTrace's comparison interface allows users to compare a target search result and the query image side-by-side for detailed inspection. This interface supports a two-step decision-making process. First, the user determines whether the two photos show the same backdrop by selecting one of three options (*Same backdrop*, *Not sure*, or *Different backdrop*) (Figure 4). Second, if the target backdrop is part of an existing collection, the user rates their confidence (on a five-point Likert scale, ranging from *Definitely No* to *Definitely Yes*) regarding the membership of the query photo in that collection. If the user expresses positive confidence, the query photo gets added to the target collection. The second step will be repeated if the target photo is part of multiple collections.

Building Collections: In case of comparisons where there is no target collection, BackTrace allows users to create a new collection with two or more photos that share the same backdrop. If the user votes on a target photo as having the same backdrop as the query photo, BackTrace automatically creates a “Work-in-Progress” collection, and queues up all target photos deemed by the user to have the same backdrop as the query photo.

This new “Work-in-Progress” collection, along with the other target collections that were updated with the query photo, are displayed on a sidebar. The user can then finalize the new collection by clicking the “Create Collection” button. BackTrace will then ask the user to provide some metadata for the collection, which includes a nickname and description (pre-populated from the original backdrop description tags). Once the user creates the collection, a dedicated page is created for the collection that shows all of its member photos and backdrop-related metadata. A separate gallery page lists all the user-curated collections.

Linking Collections: BackTrace restricts users from directly adding new images to existing collections, even if the user is confident about their membership or has discovered the images from the same grouping elsewhere. This is to ensure that every collection is created and updated through this iterative exploratory workflow using its two-step comparison interface, thus minimizing the potential for confirmation bias. All potential additions to a collection must go through the process of comparison and confidence rating to maintain the integrity and consistency of the collections.

This policy, however, can result in the creation of multiple collections associated with the same backdrop, stemming from different query images. As a result, there may be overlapping collections where some are supersets or subsets of others, or are interconnected through one or more

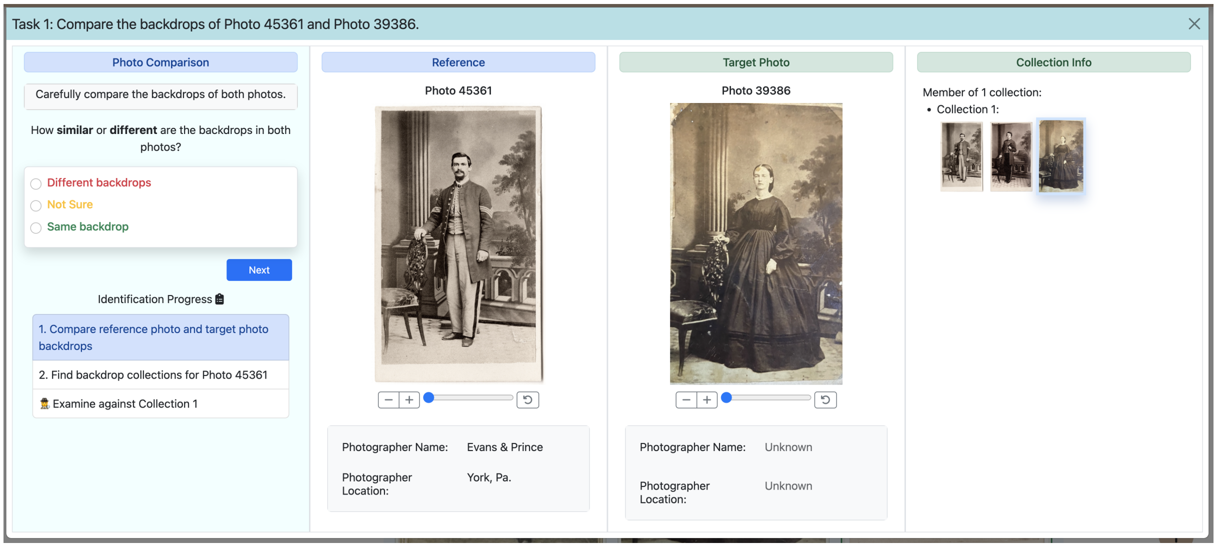

Figure 4: BackTrace's comparison interface showing two photos being compared.

shared images. To address the issue of fragmented collections, BackTrace, on the "Collection" page, also displays photos from other collections that have a shared membership under a "Related Photos" section. These links helps provide a comprehensive picture of all the photos associated with the same backdrop.

Evaluation

We conducted an exploratory, mixed-methods evaluation study of BackTrace to understand how well our system helps users 1) discover historical photos with matching painted backdrops and 2) organize them into meaningful collections. We also sought to understand how these outcomes can support the downstream task of historical photo identification.

Dataset

We manually curated a dataset of 153 photos that spanned 43 unique backdrops from multiple sources including scholarly articles, Pinterest, and inductive exploration of photos from CWPS using BackTrace. For our lab study, we randomly selected six photos from the dataset and three collections with unique backdrops. The number of matches for each photo ranged between one and seven; this variation is due to the random selection of test cases.

Based on the first match's rank in the search results (without any refinement), we divided these six photos into two broad categories: a) *Instant Matches* and b) *Challenging Matches*. Instant Matches, requiring minimal to no scrolling on the search results page to find the first match, had first match rankings of 1, 1, and 11. Conversely, Challenging Matches, requiring extensive scrolling, had first match rankings of 147, 224, and 293.

The collections selected each had a minimum of four photos sharing the same backdrop. We seeded each collection in BackTrace with just two photos, with the goal that participants would find and add the remaining two or more photos themselves.

Participants

We recruited nine participants from different Civil War and 19th-century photography Facebook groups and direct engagement with known Civil War photography experts. These participants had a diverse range of exposure to Civil War photography research (mean=20.9 years, min=4 years, max=40 years) and conducted backdrop research with a wide range of frequencies (mode=Daily, max=Daily, min=Never). Participants consisted of five men and four women, with an average age range of 44 to 53 years (min=18 to 30 years, max=61 to 70 years).

Procedure

The study was conducted virtually, with each participant joining a recorded Zoom session that lasted between 45 minutes and an hour, and was attended by at least one researcher. Participants started by filling out a consent form and a pre-survey detailing their demographics and experience with Civil War photography. The study was approved by the university's IRB protocol. Participants were compensated $20 for completing the study.

After watching an instructional video of how BackTrace works, each participant was asked to complete the following three tasks, using a think-aloud protocol:

Single Photo Matching Participants were asked to search matches for a given query photo without any particular objective.

Collection Enhancement Participants were asked to start a search from a given "Collection" page. They could select any photo, or photos, present in the collection and find other photos that are likely to belong (or be linked) to the collection.

	Single Match	Photo Difficulty	Timed Match	Photo Difficulty
P1	3 (7)	Instant	0 (6)	Challenging
P2	1 (6)	Challenging	3 (7)	Instant
P3	2 (7)	Instant	0 (6)	Challenging
P4	0 (1)	Challenging	6 (6)	Instant
P5	5 (6)	Instant	0 (1)	Challenging
P6	0 (1)	Challenging	4 (6)	Instant
P7	3 (4)	Instant	1 (3)	Challenging
P8	1 (3)	Challenging	4 (4)	Instant
P9	3 (4)	Instant	0 (3)	Challenging

Table 1: Number of identical backdrop matches identified by each participant for the single match and timed search tasks. The numbers in parentheses represent how many matching photos were present in the search results. Instant difficulty photos have their first match appear prior to 100 images while challenging difficulty photos have their first match appear after 100 images.

Timed Photo Matching Similar to single photo matching, participants were asked to find as many potential matches as possible for a specific query photo, adding these to a "Work-in-Progress" collection. However, this task was time-limited to three minutes. The three-minute time limit was included to create urgency and test the speed at which users can find matches in the system. Once the time limit was reached, they were asked to finalize their "Work-in-Progress" into a completed new collection.

For single and timed photo matching, participants were randomly assigned different photos from the pool of six photos in the test dataset. Similarly, for collection enhancement, they were randomly assigned to one of the three collections present in the dataset. We also ensured that each photo and collection was tested by at least three different participants. These tasks were carried out in the same sequence for all participants.

Upon completing these tasks, participants filled out a standard usability survey and answered a series of semi-structured questions about their experience using BackTrace.

Data Analysis

We carried out inductive thematic analysis on the qualitative data (interview transcripts, open-ended survey responses, observation notes) to uncover any common themes or behaviors between sessions with respect to the guiding themes in our research questions. In addition, we also analyzed the Likert-scale survey responses and system logs related to feature engagement, scroll amount in the search page, comparison votes and user confidences, interaction with the refinement feature, and collection creation.

Findings

Discovering Photos with Similar Backdrops

Participants successfully discovered new photos with similar backdrops using BackTrace's search workflow. A majority of participants (14 out of 19 search sessions) successfully identified at least one photo with a matching backdrop across both discovery tasks: *single photo match* and *timed photo match* (see Table 1). However, in the scenarios we categorized as "challenging" — where the first (ground truth) match was present beyond the initial 100 search results, and thus, required extensive scrolling — some participants faced difficulties. Excluding three specific cases, participants typically found multiple matching photos within the search results.

P3 highlighted the system's ability to discern not just direct matches but also contextual similarities, stating that "*. . . seeing not only all those obvious hits, but seeing similarities between the images, too, like that suggested that maybe they're all taken around the same time period, or were members from the same unit.*" This sentiment was echoed by P7, who found the experience enlightening: "*It was actually very freeing to realize that there are other examples out there with the same backdrop, similar backdrops, and that, you know, it can actually lead you to an answer.*" P9 underscored the tool's capability to enhance their search, particularly in situations where manual identification would be challenging: "*There's going to be times where the photograph just has a tiny fragment of the backdrop visible and your software is going to be able to perfectly match it to an example that you have, and it's gonna be one of those where I mean, if I were to look at, I would have never drawn that conclusion by myself.*"

Participants iteratively refined the search results for surfacing backdrop matches. Across all the search sessions, we recorded a total of 49 refinement instances (40 visual and nine text-based), which helped surface 39 matches for the participants. Twenty-one of these matches were a result of visual refinements (for nine distinct query photos), while the remaining were found with text refinement.

Participants found visual refinement feature helpful (mean = 4.11/5), with P8 stating, *"I think the the visual refining feature was probably the most memorable [feature], because [. . .] the ones [backdrops] that were similar to the Van Stavorens weren't even within the first 20 photos. And I finally found one down at the bottom. But when I clicked on the only other photo I saw with a balustrade, it brought me up photos with balustrades, and then within those there I found another Van Stavorens."* P6 echoed this sentiment, expressing a clear preference for visual refinement: *"I was seeing the matches visually, and it seemed better to match the visual images and search on those visual images as opposed to refining with the keyword, because I had them in front of me."* Of the 62 photos added for refinement, only 13 had different backdrops from that of the query photo. P8 stated that he only used non-identical backdrops for refinement because he *"didn't have an alternative."* Yet, this strategy proved effective, as two out of three successes in the "challenging" scenarios (P2, P8) resulted from refining with non-identical backdrops.

Conversely, text refinements had a narrower scope of success. All matches from text refinement were associated with one query photo and text query: the aforementioned "Benton

Barracks" backdrop well known in the Civil War photography community. Nonetheless, participants still found some value in this feature, giving it an average rating of 3.65 out of 5. P5 shed light on their typical research process and the advantages of combining both visual and text refinements: *"I usually look for photographer, and then I look for flag or backdrop, or whatever. And this [BackTrace], you know, just with the visual search, is so much faster."*

Participants valued the role of AI despite having mixed feelings about its performance. Participants appreciated the AI's capability to identify potential matches, with some favorably comparing it to existing image search engines (P3) and archival databases (P6), and others noting that it met or exceeded their expectations (P9). P4 highlighted its effectiveness, stating, *"The AI that's comparing the backgrounds... I think it's effective, because you're able to find some examples."* P3 further commended the system's attention to detail: *"What's cool is that it does search the reverse images. So, like the CDV is printed, showing you what it looks like straight on, but an ambrotype or a tintype would be reversed."*

However, while participants acknowledged the AI's strengths, they also encountered moments of confusion with some search results. This mixed experience was evident in the survey response to the statement, "I was confused by the search results" (mean = 2.5/5). P6's comment encapsulated this confusion: *"Interesting that these fellows here don't really have any of the similarities other than their standing. So I'm wondering why they've made the cut."*

Despite these challenges, participants emphasized the collaborative nature of the tool, suggesting that the AI's role is to assist, not replace, human judgment. P3 eloquently summarized this perspective, stating, *"If you get like three good hits and like 30 that don't match, I wouldn't say it's a three out of 30 accuracy, because you can visually still look through it. I'd rather have more misses than potentially miss something where the backdrop is cut off."*

Participants focused on a wide-range of visual cues beyond the overall backdrop while analyzing search results. Participants typically focus only on matching individual features in the backdrop instead of comparing them as a whole. For instance, P1 stated she was *"trying to decipher which trees have the right foliage in the right place to be part of this background"*. P6 even suggested looking for signs of damage on the backdrop as a means of identification.

This focus is not just limited to backdrop features however, as seven of our nine participants focused on features outside the backdrop in their search. Common examples include props (P1, P2, P5, P6, P7), furniture (P4, P5, P6), and flooring (P1, P2, P3, P4, P5, P6). P5 even considered some of these features more important than the backdrop, saying, *"It's usually the chair that I find, and not the backdrop that I find."*. P7 made a "same backdrop" comparison on two photos even though one did not show a backdrop at all, stating, *"I'm just gonna put same, because it's the same exact prop just doesn't have the painted backdrop on the back."*

Participants preferred BackTrace over traditional backdrop research methods. Participants overwhelmingly expressed optimism about BackTrace's potential to streamline their research process. They emphasized significant time and effort savings compared to traditional methods. P3 highlighted the system's efficiency in reducing the monotony of manual image comparison, remarking, *"It takes a lot of mind-numbing, flipping through images, staring at things until you get a headache. It takes a lot of that out of it."* P8 echoed this sentiment, predicting, *"Once it's developed out, I think it's gonna be an amazing time-saver for me."*

Survey results further underscored this enthusiasm. Participants not only enjoyed using BackTrace (mean = 4.89/5) but also rated it favorably over their existing backdrop research methods (mean = 4.45/5). P2 emphasized the uniqueness of the tool: *"I really don't know of any other way to have access to other photos that can do that type of thing."* P7 lauded its search capabilities: *"It was totally amazing. That's something you cannot Google. You cannot Google 'backdrop with column and wrought-iron railing'; that wouldn't even come up with anything."*

Organizing Photos with Matching Backdrops

Participants curated different photo collections with shared backdrops, foreseeing a range of downstream benefits. All participant successfully completed the "collection enhancement" task by linking a photo with a matching backdrop to their assigned collection, finding on average 75% of known matches by using existing collection members as query photos. Across all tasks, participants created a total of 19 collections, each averaging four images in size. Participants also added 15 photos to 10 existing collections. Participants appreciated the system's capability to store and display backdrop collections (mean = 4.22/5). P1 envisioned using the system as a repository to cross-reference new photos in the future. P3 highlighted the potential collaborative benefits of these collections: *"And so maybe with my collection, I found that reverse one. It might help somebody else. They might have a what they think is a positive image, and it's the one that was reversed and colorized right? or Tinted. Maybe that will spur something in their mind."* P2 liked that collections in tandem with related photos were able to give him a quick overview of a backdrop and emphasized the value of collections in drawing inferences: *"I like seeing the building of the collection, and then you can just see them all splayed out. Observing them collectively gives a sense of confidence, allowing one to infer, for instance, that these photos likely originated from the same studio."*

Participants selectively used the comparison interface for analyzing potential matches, while achieving high accuracy in their comparisons. Out of the 66 photo-to-photo comparisons made using the comparison interface, only five were inaccurate (Precision = 0.98, Recall = 0.92). In addition, 15 photo-to-collection comparisons were made with no mistakes. Participants found the comparison interface to be highly useful (mean = 4.77/5), with P9 commenting, *"If I only have a tiny portion of one backdrop visible on an image, and I'm comparing it to one where it's just not* *immediate clear that they match, I would have to compare it, and blow them up and sort of look at it".*

Almost no negative comparisons were made during the evaluation sessions. Of the 66 votes cast for photo-to-photo comparisons, only 5 were negative (Different Backdrop), while 4 votes were unsure. No negative comparisons were made for photo-to-collection comparisons. This suggests that participants saw limited value in making negative comparisons. As P5 stated, *"Unless I was told that voting no was helpful in the training of it, I don't think I would seek out different images to vote no on. I mean, if they were close, but no dice, I think I'd be more likely to vote no, but if it's a totally different backdrop, then I just don't wanna waste my time on it."* Additionally, participants only reserved the comparison interface for more complex comparisons, and were comfortable making most of their comparisons in the overview. On average, participants only compared 10 of the 186 photos that they viewed over each session. In the words of P9, *"I don't need to compare it because I am certain it's the same"*.

Participants employed diverse descriptions for collections, showing no unified taxonomy or nomenclature. Participants mainly used two different methods in naming their collections. Some participants assigned location-based nicknames such as "York PA" (P8), while others used object-based nicknames like "The Cliff and Tree Backdrop" (P9) for the same backdrop. Of the 19 created collections, nine adopted the latter method of naming based on backdrop contents, while four were named after the location where the photo was taken. Six collections were unnamed. This naming approach mirrored their choices in text-based refinement keywords. As P6 observed, *"The keywords performed better than anticipated. I was hoping to use terms like 'tree', 'lake', or 'flagpole' to refine the search more effectively."*

Broadly, participants provided descriptions that consistently describe the contents of the backdrop, with 6 (of 19) also containing references to props. However, the variation in terminology and taxonomy used by participants in descriptions was quite high, with several terms frequently being used to describe the same object. "Balustrade" was described using other words like "spindle" (P2) and "bannister" (P9). Some participants added more specific descriptors to the base word, such as "Sibley tent" (P3) instead of "tent" or "map of Missouri" (P9) instead of "map".

Participants expressed a need for enhanced flexibility and features in the collections feature. While participants were successful in creating different collections, we observed some users struggling with the collections feature. P1 expected a bulk-add feature and was slightly confused by repeating the sequence of first making comparisons followed by adding photos to collections. This sentiment was echoed by P8, who highlighted the potential need to incorporate 10+ images into a collection simultaneously. Another point of confusion were the restrictions placed on adding new images to existing collections. Additionally, several participants expressed disappointment at the lack of the ability to merge existing collections (P2, P3, P7).

Participants also showed confusion when asked to create duplicate or fragmented collections (P2, P3, P5, P9), preferring instead to add directly to a shared, singular collection. Knowledge that a photo is already in a collection also seemed to dampen motivation to compare a photo, with P4 stating, *"Let's see... he's already in that collection. So I feel like I'm not adding anything."*

Discussion

Enabling Backdrop Research at Scale

Participants valued BackTrace's potential to speed up the research process significantly. Much of this value is attributed to the layout of the search results, which allowed participants to choose a single salient background feature such as a tent, and use that as a point of comparison to rapidly compare the results while retaining the option of closer inspection where necessary, a kind of overview plus detail interface (Shneiderman 1996). However, this method of comparison could cause some matches to be overlooked, due to the variety of compositions and techniques employed by photographers in the period (Keller 2021).

At times, the chosen background feature may not be directly related to the backdrop. Certain props and studio features, such as distinctive floor tiles, are considered equally valuable as a means of photographer identification.

Our current organizational structure of one collection per search was found to be confusing to users, who generally prefer to have unified photo collections that contain all photos sharing the same backdrop. However, a new concern emerged: some users are hesitant to share the results of their research or photos for fear of having their research or photos appropriated without permission (Luther 2021). Even though Civil War-era portraits fall under public domain, community norms typically defer to the photo's current owner regarding permission to reproduce the image. Therefore, steps should be taken to ensure users are comfortable with uploading images to the platform. Technical restrictions or social guidelines around downloading content could be enacted to prevent violating community norms (Fiesler and Bruckman 2019).

We also note a varied taxonomy in participants' textual metadata, such as backdrop names and descriptions. These metadata adequately described the contents of the backdrops despite their unstructured nature (Kittur et al. 2014). Moving forward, these descriptions can be processed and converted into more uniform annotations to improve text search functionalities. While the vast majority of participants did provide photo descriptions during their sessions, our choice of making photo descriptions optional could reduce the amount of metadata generated in the wild, limiting their use for searching and informing future users.

Crowd-AI Collaboration

Textual search has long been considered the default when querying large datasets. However, in cases where text annotations are scarce, or if the user lacks domain knowledge required to form effective search terms, textual search becomes less effective. While approaches such as text-based generous interfaces (Whitelaw et al. 2015) and social autocomplete (Kittur et al. 2014) can help address this issue, they require either extensive manual annotation or an active user base, diminishing their utility in a cold start scenario.

To address the cold start problem, we use automatically generated image embeddings as the base of our system. This approach also allowed for extensive pre-processing of the data (Arnold and Tilton 2019), significantly reducing the users' workload of sorting through images without backdrops. However, this does not guarantee a match when a user initiates a search, or that matches will appear on top of the search results. As such, it is important to afford users control in these situations.

Our solution is to introduce relevance feedback (Rui et al. 1998), inspired by prior works that allow for iterative improvement to AI performance via the selection of examples matching a user's query (Cai et al. 2019; Fogarty et al. 2008; Lee and Weld 2020). Our combination of visual search and relevance feedback in the form of the visual refinement system proved effective in assisting backdrop research and was well received by users. Given the visual nature of the task, this comes as no surprise, as users are able to directly apply their learnings from visual inspection to improve search results, bypassing the step of describing a photo for textual search and removing the dependence on complete tagging or a unified taxonomy.

Despite experiencing inconsistent AI performance, participants remained satisfied with the system as a whole. This observation is consistent with Heer (2019), who posits that users who are allowed to maintain autonomy are more receptive to AI outputs, even if said outputs are unrelated to the users' input. However, improvements to AI performance improvements should still be made, as a faulty AI can mislead users, even experts, into making erroneous decisions (Tschandl et al. 2020).

Generalizing beyond Backdrops

The clearest extension to this work can seen in other cultural heritage collections. Due to the widespread use of painted backdrops in the 19th century and beyond (Keller 2021; Zeitlyn 2010), BackTrace can be adapted to other collections of portrait photography with minimal modifications required. Such an example can be seen in Zeitlyn, Coto, and Zisserman (2021)'s work, where backdrops were used extensively to catalog a collection of photos.

We propose that the BackTrace workflow can be expanded to other investigative tasks. For example, many citizen science tasks parallel that of painted backdrop identification, where users compare (Gould, Clulow, and Clulow 2021) and label (Saha et al. 2019; Wieck 2018; Simpson, Page, and De Roure 2014; Nugent 2018) images before sharing their findings with other users (Simpson, Page, and De Roure 2014; Nugent 2018). A workflow that helps users focus on only the most relevant images in a database, while providing advanced search tools has the potential to improve the outcomes and efficiency of these investigations.

Additionally, the BackTrace workflow could be adapted to modern photographs to assist journalists, law enforcement officials, and human rights investigators. Stylianou et al. (2017)'s work combating human trafficking specifically calls for the manual removal of the subject from the query photo when initiating a search. Another potential adaptation of our workflow is for geolocating modern-day photos and videos, where the subject of a photo is commonly considered secondary to the background information present (Higgins 2014; Venkatagiri et al. 2019).

Limitations and Future Work

While diverse in terms of domain experience, our pool of participants only numbered nine due to difficulties in the recruitment process. As such, quantitative data from this evaluation is limited, and should be considered supplementary findings to our qualitative analyses.

Furthermore, the test cases, while randomly selected, only represent a fraction of the overall photo database. Although we achieved theoretical saturation (Guest, Bunce, and Johnson 2006) over our nine participants, some potential user behaviors and findings might have been overlooked due to our constrained photo and participant pool.

Despite reasonable performance, we did not fine-tune our embedding extraction model for the purpose of this application. Additional work could also be done to fine-tune the refinement and backdrop classifiers. For example, negative user feedback can also be leveraged to either train the refinement classifier or exclude images from search results.

Conclusion

The study of 19th-century painted backdrops can yield useful clues in supporting the identification of persons, events, and places in the American Civil War. However, technological support for this task is lacking, with researchers generally having to rely on manual methods for their research. We present a three-step workflow that combines novel computer vision techniques and relevance feedback (RF) into a Content-Based Image Retrieval (CBIR) system for identifying and clustering historical backdrops. We developed a web-based software tool called BackTrace and showed that this workflow effectively supported users in discovering and organizing backdrops in a large, sparsely tagged photo dataset. Our work can guide the future design of human-AI collaboration systems for exploring large photo archives, providing interaction requirements and empirically validated workflows to enable effective discovery and organization of photos across the ages.

Acknowledgements

We wish to thank Benjamin C.G. Lee, Vishnu Dutt Sharma, and Adam Ochs Fleischer for early discussions. This research was supported by NSF IIS-1651969 and the GEM Fellowship Program.

References

Afif, M.; Ayachi, R.; Said, Y.; and Atri, M. 2020. Deep learning based application for indoor scene recognition. *Neural Processing Letters*, 51: 2827–2837.

Arnold, T.; and Tilton, L. 2019. Distant viewing: analyzing large visual corpora. *Digital Scholarship in the Humanities*, 34(Supplement_1): i3–i16.

Burke, P. 2001. *Eyewitnessing: The uses of images as historical evidence*. Cornell University Press.

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.

Caraffa, C.; Pugh, E.; Stuber, T.; and Ruby, L. W. 2020. PHAROS: A digital research space for photo archives. *Art Libraries Journal*, 45(1): 2–11.

Chumachenko, K.; Männistö, A.; Iosifidis, A.; and Raitoharju, J. 2020. Machine learning based analysis of finnish world war ii photographers. *IEEE Access*, 8: 144184–144196.

Cordell, R. 2020. Machine learning and libraries: a report on the state of the field. Library of Congress.

Crissaff, L.; Ruby, L. W.; Deutch, S.; DuBois, R. L.; Fekete, J.-D.; Freire, J.; and Silva, C. 2017. ARIES: enabling visual exploration and organization of art image collections. *IEEE computer graphics and applications*, 38(1): 91–108.

Datta, R.; Joshi, D.; Li, J.; and Wang, J. Z. 2008. Image retrieval: Ideas, influences, and trends of the new age. *ACM Computing Surveys (Csur)*, 40(2): 1–60.

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.

Ertzgaard, J. 1989. The Elusive Photographer of BENTON BARRACKS. *Military Images*, 10(6): 24–25.

Fiesler, C.; and Bruckman, A. S. 2019. Creativity, Copyright, and Close-Knit Communities: A Case Study of Social Norm Formation and Enforcement. *Proc. ACM Hum.-Comput. Interact.*, 3(GROUP).

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.

Google. 2023. Google Images. https://images.google.com/. Online; accessed 27 August 2023.

Gould, J.; Clulow, J.; and Clulow, S. 2021. Using citizen science in the photo-identification of adult individuals of an amphibian based on two facial skin features. *PeerJ*, 9: e11190.

Guest, G.; Bunce, A.; and Johnson, L. 2006. How many interviews are enough? An experiment with data saturation and variability. *Field methods*, 18(1): 59–82.

Haidt, R. 2011. Commodifying place and time: photography, memory, and media cultures around 1850. https://hdl.handle.net/11299/202471. Online; accessed 27 August 2023.

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

Heer, J. 2019. Agency plus automation: Designing artificial intelligence into interactive systems. *Proceedings of the National Academy of Sciences*, 116(6): 1844–1850.

Higgins, E. 2014. A beginner's guide to geolocating videos. https://www.bellingcat.com/resources/how-tos/2014/07/09/a-beginners-guide-to-geolocation/. Online; accessed 27 August 2023.

Karnes, S. 2020. Photo 34619. https://www.civilwarphotosleuth.com/photos/view/34619/. Online; accessed 27 August 2023.

Keller, K. 2021. A Great Variety of New and Fine Designs: Advertisements for Painted Backgrounds, 1856–1903. https://rshare.library.torontomu.ca/articles/thesis/A_Great_Variety_of_New_and_Fine_Designs_Advertisements_for_Painted_Backgrounds_1856--1903/14662737. Online; accessed 27 August 2023.

Kittur, A.; Peters, A. M.; Diriye, A.; and Bove, M. 2014. Standing on the schemas of giants: socially augmented information foraging. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*, 999–1010.

Kurtz, W. 2021. Photo 40060. https://www.civilwarphotosleuth.com/photos/view/40060/. Online; accessed 27 August 2023.

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.

Lee, B. C. G.; Mears, J.; Jakeway, E.; Ferriter, M.; Adams, C.; Yarasavage, N.; Thomas, D.; Zwaard, K.; and Weld, D. S. 2020. The Newspaper Navigator Dataset: Extracting And Analyzing Visual Content from 16 Million Historic Newspaper Pages in Chronicling America. arXiv:2005.01583.

Lindberg, M. 2021. Finding clues in Civil War Photographs: Picture this. https://blogs.loc.gov/picturethis/2021/08/finding-clues-in-civil-war-photographs/. Online; accessed 27 August 2023.

Luther, K. 2021. Real-life accounts on the research trail. *Military Images*, 39(2 (216): 8–11.

Microsoft. 2023. Bing Images. https://www.bing.com/images/feed. Online; accessed 27 August 2023.

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

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.

Montgomery1861. 2018. Photo 27371. https://www.civilwarphotosleuth.com/photos/view/27371/. Online; accessed 27 August 2023.

Moretti, F. 2005. *Graphs, maps, trees: abstract models for a literary history*. Verso.

Nugent, J. 2018. inaturalist. *Science Scope*, 41(7): 12–13.

Olafenwa, A. 2021. Simplifying object segmentation with pixellib library. https://vixra.org/abs/2101.0122. Online; accessed 27 August 2023.

Resig, J. 2014. Using computer vision to increase the research potential of Photo Archives. https://journalofdigitalhumanities.org/wp-content/uploads/2014/07/Using-Computer-Vision-to-Increase-the-Rese-John-Resig.pdf. Online; accessed 27 August 2023.

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.

Saha, M.; Saugstad, M.; Maddali, H. T.; Zeng, A.; Holland, R.; Bower, S.; Dash, A.; Chen, S.; Li, A.; Hara, K.; and Froehlich, J. 2019. Project Sidewalk: A Web-Based Crowdsourcing Tool for Collecting Sidewalk Accessibility Data At Scale. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, 1–14. New York, NY, USA: Association for Computing Machinery. ISBN 9781450359702.

Shneiderman, B. 1996. The eyes have it: a task by data type taxonomy for information visualizations. In *Proceedings 1996 IEEE Symposium on Visual Languages*, 336–343.

Simpson, R.; Page, K. R.; and De Roure, D. 2014. Zooniverse: observing the world's largest citizen science platform. In *Proceedings of the 23rd international conference on world wide web*, 1049–1054.

Stylianou, A.; Schreier, J.; Souvenir, R.; and Pless, R. 2017. TraffickCam: Crowdsourced and Computer Vision Based Approaches to Fighting Sex Trafficking. In *2017 IEEE Applied Imagery Pattern Recognition Workshop (AIPR)*, 1–8.

Stylianou, A.; Xuan, H.; Shende, M.; Brandt, J.; Souvenir, R.; and Pless, R. 2019. Hotels-50k: A global hotel recognition dataset. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 726–733.

Tschandl, P.; Rinner, C.; Apalla, Z.; Argenziano, G.; Codella, N.; Halpern, A.; Janda, M.; Lallas, A.; Longo, C.; Malvehy, J.; et al. 2020. Human–computer collaboration for skin cancer recognition. *Nature Medicine*, 26(8): 1229–1234.

van Strien, D.; Bell, M.; McGregor, N. R.; and Trizna, M. 2022. An Introduction to AI for GLAM. In Kinnaird, K. M.; Steinbach, P.; and Guhr, O., eds., *Proceedings of the Second Teaching Machine Learning and Artificial Intelligence Workshop*, volume 170 of *Proceedings of Machine Learning Research*, 20–24. PMLR.

Venkatagiri, S.; Thebault-Spieker, J.; Kohler, R.; Purviance, J.; Mansur, R. S.; and Luther, K. 2019. GroundTruth: Augmenting expert image geolocation with crowdsourcing and shared representations. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW): 1–30.

Whitelaw, M.; et al. 2015. Generous interfaces for digital cultural collections. Online; accessed 27 August 2023.

Wieck, L. P. 2018. Building Inspector. *Journal of American History*, 105(2): 477–478.

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.