

Photo Sleuth: Combining Human Expertise and Face Recognition to Identify Historical Portraits

Vikram Mohanty

Department of Computer Science, Virginia Tech
Blacksburg, VA, USA
vikrammohanty@vt.edu

Sneha Mehta

Department of Computer Science, Virginia Tech
Arlington, VA, USA
sudo777@vt.edu

David Thames

Department of Computer Science, Virginia Tech
Blacksburg, VA, USA
davidct@vt.edu

Kurt Luther

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

ABSTRACT

Identifying people in historical photographs is important for preserving material culture, correcting the historical record, and creating economic value, but it is also a complex and challenging task. In this paper, we focus on identifying portraits of soldiers who participated in the American Civil War (1861-65), the first widely-photographed conflict. Many thousands of these portraits survive, but only 10–20% are identified. We created Photo Sleuth, a web-based platform that combines crowdsourced human expertise and automated face recognition to support Civil War portrait identification. Our mixed-methods evaluation of Photo Sleuth one month after its public launch showed that it helped users successfully identify unknown portraits and provided a sustainable model for volunteer contribution. We also discuss implications for crowd-AI interaction and person identification pipelines.

CCS CONCEPTS

• **Human-centered computing** → Collaborative and social computing systems and tools; • **Computing methodologies** → Computer vision tasks; • **Applied computing** → Arts and humanities;

KEYWORDS

Crowdsourcing, online communities, face recognition, person identification, crowd-AI interaction, history

ACM Reference Format:

Vikram Mohanty, David Thames, Sneha Mehta, and Kurt Luther. 2019. Photo Sleuth: Combining Human Expertise and Face Recognition to Identify Historical Portraits. In *24th International Conference on Intelligent User Interfaces (IUI '19)*, March 17–20, 2019, Marina del Rey, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3301275.3302301>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IUI '19, March 17–20, 2019, Marina del Rey, CA, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6272-6/19/03...\$15.00

<https://doi.org/10.1145/3301275.3302301>

1 INTRODUCTION

Identifying people in historical photographs provides significant cultural and economic value. From a cultural perspective, it can help recognize contributions of marginalized groups, as in the recent social media campaign to identify Sheila Minor Huff, the only female African American scientist visible in a group portrait of attendees at a 1971 biology conference [18]. Identification can also correct the historical record, as in the case of James Bradley, author of *Flags of Our Fathers*, who was convinced by visual evidence that his father was not pictured in the iconic photo of US Marines at Iwo Jima during World War II, as he once believed [46]. Identification can also create significant economic value, as when a photo purchased at flea market for \$10 was estimated to be worth millions of dollars following its identification as a circa-1875 portrait of American outlaw Billy the Kid [17].

Despite this cultural and economic value, identifying people in historical photos is complex and challenging, and researchers lack adequate technological support. The current research practices employed by historians, antiques dealers, and collectors for identifying portraits are largely manual and often time-consuming. These practices involve manually scanning through hundreds of low-quality photographs, military records, and reference books, which can often be tedious and frustrating, and lacks any guarantee of success. Automated face recognition algorithms can support this effort, but are not widely used by historical photo experts, and are often insufficient for solving the problem on their own. Many studies have compared face recognition algorithms to a human baseline, with mixed results [7, 9, 23, 60]. Further, historical photographs add unique challenges as they are often achromatic, low resolution, and faded or damaged, which might result in loss of useful information for identification.

In this paper, we present Photo Sleuth¹, a web-based platform that combines crowdsourced human expertise and automated face recognition to support historical portrait identification. We introduce a novel person identification pipeline in which users first identify and tag relevant visual clues in an unidentified portrait. The system then suggests filters based on these tags to narrow down search results of identified reference photos. Finally, the user can carefully inspect the narrowed search results, sorted using automatic face recognition, to make a potential identification. This

¹<http://www.civilwarphotosleuth.com>

pipeline also bootstraps crowdsourced user contributions to grow the site's database of reference images in a sustainable way, increasing the likelihood of a potential match in the future. Photo Sleuth initially focuses on identifying portraits from the American Civil War (1861–65), the first major conflict to be extensively documented through photographs. An estimated three million soldiers fought in the war and most of them had their photos taken at least once. After 150 years, many thousands of these portraits survive in museums, libraries, and individual collectors, but the identities of most have been lost.

We publicly launched Photo Sleuth in 2018 and conducted a mixed-methods evaluation of its first month of usage, including interviews with nine active users, content analysis of uploaded photos, and expert review of user identifications. We found that the system transformed users' research practice and helped them identify dozens of unknown portraits. Additionally, Photo Sleuth's pipeline encouraged users to voluntarily add hundreds of identified portraits to aid future research, suggesting a sustainable model for long-term participation. Our primary contributions are:

- a novel person identification pipeline combining crowdsourcing and face recognition
- a web-based tool and online community, Photo Sleuth, demonstrating this approach
- a mixed-methods evaluation of Portrait Sleuth after one month of deployment with real users

We also discuss implications for crowd-AI interaction and person identification pipelines.

2 RELATED WORK

2.1 Person Identification in Photographs

In recent years, commercial computer vision-based face recognition algorithms are finding use in many real-world applications, such as Uber using Microsoft's Cognitive Face API to verify their drivers [39] and C-Span using Amazon's Rekognition service to index their videos by who is speaking or who is on camera [2].

Kumar et al. [27] propose the use of generalizable visual attributes (i.e., labels to describe the appearance of an image) for the face, such as gender, age, jaw shape, and nose size, to search faces and verify whether two faces show the same person. Some deep learning approaches like DeepFace [52], DeepID2 [51], FaceNet [47] have shown near-perfect face verification accuracy on the Labeled Faces in the Wild (LFW) dataset. Schroff et al. [47] also propose a method to automatically cluster all faces of a particular person. Photo Sleuth, on the other hand, does not depend on a training set. Instead, it exploits the strengths of existing face recognition algorithms in a hybrid pipeline by integrating additional relevant information from visual clues in a photograph into the search process to enhance accuracy.

Multiple studies have compared face recognition algorithms to human baselines, and some show that human performance is superior [7, 9, 23, 60]. A recent study shows state-of-the-art face recognition algorithms performing in the range of professional face examiners and suggests optimal face recognition achieved by fusing human and machines [43]. However, these algorithms have also been seen failing to filter out false positives. Recently, Welsh police wrongly identified people as criminals 92% of the time at a

soccer game relying on face recognition technology [3]. Amazon's Rekognition wrongly identified 28 members of Congress as people charged with a crime [49]. The workflow of Photo Sleuth prevents face recognition *per se* from making the final decision, instead deferring to human judgment.

Crissaff et al. [14] propose an image manipulation system called ARIES for organizing digital artworks, allowing users to compare images in complex ways and use feature-matching to explore visual elements of interest. Bell & Ommer [5] use computer vision algorithms to retrieve similar images for a query search image of a historical painting. Srinivasan et al. [50] propose using automated face recognition techniques for addressing ambiguities in portrait subjects and understanding an artist's style. Google released an app [19] in which users could find their painting doppelgangers from museums worldwide. Inspired by these recent efforts, Photo Sleuth helps users retrieve the identities of unknown photos of soldiers from the Civil War era by building and searching a digital archive of historical photographs.

Civil War portrait identification has not yet been studied through an HCI or AI lens, but a survey of historical scholarship [13, 53, 55], practitioner articles [32–34], and media accounts [37, 45, 56] offers some insight into the key tasks and challenges. It is estimated that at least four million Civil War-era portraits survive today, of which 10–20% are already identified [1]. Civil War portrait identification or “photo sleuthing” typically requires extensive skill and domain expertise, from identifying obscure uniform insignia and weapons [37], to weighing probabilities [34], to consulting a wide range of reference works [33], to systematically reviewing thousands of potential matches [32]. Photo Sleuth attempts to ease the sleuthing process by bringing together a large repository of soldier portraits and military service records, and the visual clues one would typically use in this process, in a workflow designed for both novices and experts.

2.2 Crowdsourced History and Image Analysis

2.2.1 Crowdsourced History. Research on crowdsourcing systems with applications to historical research has largely been limited to transcription projects (e.g., [11, 20, 59]). While person identification is a more complex task than text transcription and requires more historical domain knowledge, we draw inspiration from the approaches these projects take to designing interfaces that help crowd workers visually inspect historical primary sources.

A smaller body of research considers how members of online communities can work together to synthesize complex historical information and even conduct original research. Rosenzweig [44] contrasted the solitary tradition of professional historical research and the collaborative nature of Wikipedia articles about history. Willever-Farr et al. [57] found that genealogists on Ancestry.com are more likely to engage in cooperative research (sharing data) and not collaborative instructions (sharing techniques). A follow-up study [58] of Ancestry.com and Find A Grave showed that contributors are conscious about information quality and inaccurate information, and show skepticism towards open editing practices. These studies drew our attention to the complexities of facilitating original historical research in a public online platform and guided

us to design a pipeline that foregrounded attribution and accountability to reward high-quality contributions and discourage the spread of misinformation.

2.2.2 Crowdsourced Image Analysis. The bulk of the projects involving crowdsourced image analysis often usually focuses on identifying everyday objects, transcribing text, or other tasks requiring only basic knowledge. Researchers have shown to yield impressive results by leveraging crowdsourced visual analysis in a well-defined layout where workers know what to look for e.g. identifying everyday objects [8, 10, 40], analyzing video data [28–30] or performing tasks at scale with speed [6, 26]. Investigating photographs, however, requires crowds to make sense of unfamiliar historical and cultural contexts without any prior idea about objects of interest in the photos, and thus such tasks warrant a different approach.

Different techniques are employed to use crowds for analyzing unfamiliar visual material in a systematic way such as crowds being combined with computer vision to annotate bus stops and sidewalk accessibility issues in Google Street View images [21, 22], tutorials being provided to non-expert volunteer crowds for analyzing scientific imagery in GalaxyZoo, a Zooniverse project [31] and volunteer crowds comparing photos of missing and found pets to reunite them with their owners after a disaster [4].

Platforms like Flock [12] and Tropel [42] use crowdsourcing to build hybrid crowd-machine learning classifiers. Due to scale and complexity issues, a person identification task cannot be seen as multi-label classification problem. Since these approaches required a user to define the prediction task and example labeled data, they cannot be directly applied to a person identification task.

3 SYSTEM DESCRIPTION

Photo Sleuth is an online platform we developed to identify Civil War-era portraits. The website allows users to upload photos, tag them with visual clues, and connect them to profiles of Civil War soldiers with detailed records of military service. This person identification problem can be seen as *finding a needle in the haystack*. Our novel pipeline (see Figure 1) has three components - *a) building the haystack, b) narrowing down the haystack and c) finding the needle in the haystack*.

3.1 Building the Haystack

3.1.1 System Database: Photo Sleuth's initial reference database contains over 15000 identified Civil War soldier portraits from public sources like the US Military History Institute [54], as well as other private sources. This is just a small proportion of the 4 million photos that might exist [1]. Therefore, a more comprehensive archive with more reference photos and identities would boost Photo Sleuth's goal of identifying a soldier, and therefore necessitates *building a haystack*.

3.1.2 Photo Upload and Primary Sources. A user begins the identification process by uploading a photograph with a mandatory front view and an optional back view. The user is also encouraged to provide the original source of the photo. We use Microsoft's Cognitive Services Face API [38] to detect a face in the photograph at the time of uploading. Photo Sleuth does not yet support photos with multiple faces.

3.1.3 Photo Metadata. Next, the user tags metadata related to the photograph, if available, such as the photo format, inscriptions on the front and back view of the photo, and the photographer's name and location. This metadata can offer insights into the subject's hometown, military unit, or name, both improving the search filters and providing useful source material for researchers.

3.1.4 Visual Tags. Our system then gathers information about visual evidence e.g., *Coat Color, Chevrons, Shoulder Straps, Collar Insignia, or Hat Insignia*. These visual tags are mapped on to the soldier's military service information, which qualifies as a useful search parameter. More tags imply looking at a more accurate candidate pool, and thus reduce the number of false positives.

3.1.5 Bootstrapping and Ownership. Photo Sleuth adds the photo along with this information into the reference database, irrespective of identity, while displaying authorship credentials to the user. These photos enrich the database for potentially identifying future uploads. Previous work suggests attribution is an important incentive for crowds conducting original research [35, 36]. By storing this information, a future feature of the platform would be to inform the users when their uploaded photos are identified by some other users.

3.2 Narrowing down the Haystack

3.2.1 Search Filters. A major challenge in person identification tasks is the size of the candidate pool. Larger pools mean greater possibilities for false positives. Photo Sleuth reduces the likelihood of wrong identifications by generating search filters based on the visual evidence tagged by the user. These search filters are based on military service details that would otherwise be unknown to a novice user and are built using domain expertise. The military records used by the filters come from a variety of public sources, including the US National Park Service Soldiers and Sailors Database [41]. We scraped the full military service record for every identified soldier portrait in our database, along with, in many cases, vital records and biographical details. This allows for users to filter by visual clues that would only be applicable for a snapshot of a soldier's career.

For example, if the user tagged *Hat Insignia* with a hunting horn, the system would recommend the "Infantry" branch filter, whereas *Shoulder Straps* with two stars would suggest the "Major General" rank filter. These filters narrow down the search pool to all soldiers who might ever have held these positions, including promotions, demotions, and transfers. Our system shows all search filters to the users, allowing expert users to make manual refinements. Photo Sleuth's interface also scaffolds domain knowledge to prevent users from applying search filters that might contradict each other.

3.2.2 Facial Similarity. Photo Sleuth augments the above search filters with facial similarity filtering via Microsoft's Cognitive Services Face API [38]. Our tests with gold standard Civil War photos have shown that this API yields near-perfect recall at a 0.50 similarity confidence threshold; i.e., retrieved search results at this level almost always include the correct results. However, its poor precision means many other similar-looking photos also show up in the search results.




Figure 1: System Workflow. (i) The user uploads a Civil War soldier portrait. All uploaded photos, identified or not, are added to the reference database for future searches. (ii) The system automatically detects the face in the uploaded photo. (iii) The user looks for visual clues in the photo (e.g., uniforms, insignia) and tags them. (iv) The user tags the photo for metadata, such as original source, photo format, and inscriptions. (v) Photo Sleuth converts the user-tagged visual clues into search filters for matching military service records and other biographical details. (vi) The system runs face recognition on the narrowed candidate pool from the previous step to find similar-looking soldiers with matching military records, sorting the results by facial similarity. (vii) The user can browse the search results and make a careful assessment, considering all relevant context, before deciding on a match.

The search filters create a reduced search space in which face recognition now looks for similar-looking photos of the query image. This complementary interaction between military records and facial similarity ensures that the most accurate information is retained in the search space.

3.3 Finding the Needle

3.3.1 Search Results. The search results page displays all the soldier portraits who satisfy the search filters and have a facial similarity score of 0.50 and above with the query photo, sorted by similarity. The user has the option to hide as-yet unidentified photos. The search results show military record highlights next to the names and photos. The user can then closely investigate the most promising search results before making the final decision of the soldier's identity. The user can also add new names and service records to the database if that soldier has not yet been added. In order to prevent misinformation being spread and promote cross-verification, all users are made to follow the entire workflow, even in the case of photos whose identities they believe they already

know. In such cases, the user is asked to provide the source of identification.

3.3.2 User Review. Users who find a potential match among the search results can closely inspect the two photos via a "Comparison" interface. The interface provides separate zoom/pan controls and also displays the service records of the reference photo to provide a broader context of who the soldier might be. Notably, the system hides the facial similarity confidence scores for verifying two faces to avoid biasing the user. If the user is confident about the photo being a match, they can click on an "Identify" button to link the query photo to the soldier's profile and receive "identifier" attribution. The user can also undo these identifications, if desired.

3.4 Implementation Details

Civil War Photo Sleuth is a web app built on the Python/Django framework with a PostgreSQL database for data storage and Amazon S3 for image storage. It is hosted on the Heroku cloud platform. The site also provides a public RESTful API to facilitate interchange with the digitized collections of libraries and museums.

4 EVALUATION

We released Photo Sleuth to the public on August 1, 2018. We recruited users via a launch event at the National Archives Building in Washington, DC and advertising in history-themed social media groups. Within one month of its launch, 612 users registered free accounts on the website. A majority (360) registered in the first three days of the launch, followed by a steady stream of 5–10 registrations per day. During that period, users uploaded 2012 photos, with 931 photos added in the first three days, followed by an average upload of 30 photos per day. As of January 2019, the site has over 4400 registered users and over 25000 photos, of which over 8000 have been added by users.

4.1 Log Analysis

We examined website logs for user-uploaded photos between August 1 and September 1, 2018. Users categorized uploads into front and back views. Uploads that did not have a face detected or with only a back view were excluded from our analysis. We then separated the remaining photos into *identified* and *unidentified* ones.

We further analyzed the logs to identify users who had uploaded or identified at least one photo. We also analyzed uploaded photos for which users had associated one or more visual tags, and identified the most commonly tagged categories for these photos. We give details of these log analyses below.

4.1.1 Categorizing Identified Photos. From the logs, we found that users performed 691 soldier identifications in the first month, and matched 850 photos to these identities. To clean the data, we first excluded accidental duplicate uploads. Next, we checked all photos for duplicate identities (i.e. different photos of the same soldier under the same name but saved as separate identities) and grouped them together as a single identity. Lastly, all the photos that did not have a full name but had some demographic or military information were separated out as *partial identities*. The final pool consisted of 648 photos (560 uploaded by users, 88 already in the system) sharing 479 soldier identities between them.

Our pipeline does not distinguish whether the identities of soldiers in photos are known prior to uploading or not. We therefore categorized these identified photos into two categories:

Pre-identified : Photos uploaded by users with their identities known prior to uploading

Post-identified : Photos matched by users to an existing identified photo in the database using Photo Sleuth's photo matching workflow

To determine *pre-identified photos*, we considered soldier identities with only one photo, since they had not been matched to any other photo in the database. We also grouped together all soldier identities matched with multiple photos in this category if none of the photos for an identity came from Photo Sleuth's reference archive. The remainder of the photos, i.e., soldier identities with multiple photos where at least one photo came from Photo Sleuth's archive of reference photos, were labeled *post-identified photos*.

4.1.2 Grouping Unidentified Photos. We filtered the unidentified photos (with faces detected) by removing 28 duplicate uploads. Photos with no names and no military information from the previous

filtering process were also added to the original set of unidentified photos.

4.2 Content Analysis

Based on the above-mentioned categories, we performed a more targeted, in-depth analysis of how users identified the photos using Photo Sleuth.

4.2.1 Sources of Identification. We first analyzed the sources of information users drew upon when adding identified photos. We analyzed both front and back views of all pre-identified photos for the presence of a Civil War-era inscription or autograph of the matched soldier's name. If no name inscription was present, we checked if the user had provided an alternative source of identification.

4.2.2 Supporting Face Recognition. We considered two factors to understand the extent to which face recognition supported a user's identification decision. One was the presence of prior name *inscriptions* in the front or back views of the photo (see Figure 3), as this would prompt an easy decision on the user's behalf to match the photo with a search result displaying the same name.

The second consideration was the possibility of an exact duplicate. One of the most popular photo formats during the 1860s was the *carte de visite*, where a subject would receive a dozen or more identical copies of their portrait on small paper cards they could collect in albums and exchange with friends and family. If multiple copies survive today, it is possible one of them is already identified, and a user could upload an unidentified version of the same photo that may differ only slightly due to cropping or age-related damage. We refer to such photos as *replicas* (see Figure 2). In such cases, we would expect face recognition to return search results featuring an identified reference copy of the photo with a high similarity score, making it a top result for the user to quickly recognize.

Considering these factors, we analyzed front and back views of all photos in the post-identified category for the presence of the soldier's name inscriptions, similar to our analysis of pre-identified photos. Then, we examined whether any of the user-uploaded photos was a *replica* of an identified reference photo of the matched soldier.

Based on our findings, we divided the soldier identities with post-identified photos into four sub-categories: *a) inscription and replica*, *b) inscription but no replica*, *c) replica but no inscription*, and *d) no replica and no inscription*. For example, if Capt. John Smith had five user-uploaded photos matched to his identified reference photo, and any one of the user-uploaded photos had a name inscription and none of them was a replica, we grouped Capt. John Smith in the *inscription but no replica* category. Similarly, if none of the photos was a replica and none of them had an inscription, we would place the identity in the *no replica and no inscription* category, and so on for the other categories.

4.2.3 Backtracing User Behavior. For a randomly chosen small sample in each of the above defined sub-categories, we backtraced (reconstructed) the identification workflow to re-match a post-identified photo. Backtracing helped us visualize the user's experience when posed with the search results under the original conditions.

4.3 User Interviews

We also conducted in-depth, semi-structured interviews [48] with nine Photo Sleuth users. These participants were active contributors to the site, each adding at least 10 photos to the site during the first month. They also had extensive prior experience identifying Civil War photos (mean=20 years, min=8, max=40), representing a mix of collectors, dealers, and historians. Eight participants were male (one female) and the average age was 54 (min=25, max=69). We anonymized participants with the identifiers P1–P9. All interviews were conducted over phone/video calls and were audio-recorded, fully transcribed, and analyzed with respect to the themes described in Section 5.

4.4 Expert Review

In order to assess the quality of user-generated identifications, an expert Civil War photo historian (and a co-author of this paper) reviewed all post-identified photos added by users and evaluated them whether they were correctly identified or not. We establish ground truth in terms of whether a Soldier X's photo was identified as Soldier X or some other Soldier Y. The expert used the same four sub-categories as defined above to provide a fine-grained assessment of users' identifications. We captured the expert's responses using a 4-point Likert scale (1 = definitely not, 2 = probably not, 3 = possibly yes, and 4 = definitely yes).

5 FINDINGS

Using the methods above, we evaluated Photo Sleuth along three themes: *adding photos*, *identifying photos*, and *tagging photos*.

5.1 Adding Photos

5.1.1 Users added photos with both front and back views.

From our logs analysis, we found 2012 photos uploaded in the first month, of which 1632 photos were front views and 380 photos were back views. Of the 612 users who had registered for the website in the first month, 182 users (excluding the authors) uploaded at least one photo to the system.

There were three power users who each uploaded more than 200 photos, while 11 users (excluding the authors) uploaded more than 30 photos each. On average, a Photo Sleuth user uploaded 13 photos (median = 3 photos) to the website.

5.1.2 Users added both identified photos and unidentified photos.

Our log analysis showed that the number of identified photos (560) is similar to unidentified ones (602). There were also 121 partially identified photos. If we consider only identified photos, 441 were pre-identified (i.e., their identities were already known by the uploader), whereas 119 photos were post-identified (i.e., their identities were discovered using Photo Sleuth's workflow). These post-identified photos were matched to 88 identities with a prior photo in the reference archive.

Additionally, 107 users added at least one *unidentified* photo, while 105 users had added at least one *identified* photo. Fifty-three users added both identified and unidentified photos.

Interviewees expressed a variety of motivations for adding pre-identified photos. Most commonly, participants mentioned trying to help other users identify their unknown photos, but they recognized

this generosity could also help themselves. P2 felt it was only fair to contribute, given the identifications he was able to make from others' contributions: "As a way of giving back, I think I'm obligated to now." For P6, the motivation was anticipated reciprocity: "I'm just trying to help other people out like I want me to be helped out." P8 was motivated by curiosity to learn more about his own images: "I just uploaded to see if maybe there's a collector out there that had the same image maybe of a different pose or a different backdrop, different uniform." Some participants made an intentional choice to add identified photos first, waiting to add their unidentified ones later. P4 explained: "As your database of identified people increases, then there's a greater [chance] ... later on when I [upload] an unidentified image then I'll get a hit, where if I do that today, my odds are much less."

Other interviewees explained why they did *not* add pre-identified photos. One concern was bootlegging, i.e., unscrupulous users printing out scans from Photo Sleuth and reselling them as originals. P8 said, "Look on eBay. Look at all the fakes ... Look at the Library of Congress. You can download a file format ... the TIF format where it's high resolution. Then if you have a good printer, you print it out and you can make fake easy as that." A second concern was reuse without attribution. P3 said, "I have a lot of identified images that probably would help other people identify some of their guys. But I'm worried about putting them on there, only because I don't want them using my stuff unless they get permission from me first."

5.1.3 Users provided attribution for most identified photos.

Users matched 441 pre-identified photos to 386 unique soldier identities. Based on our content analysis, we found that users, while adding pre-identified photos, generally referred either to the name inscription on the photos (173 cases) or to the original source of identity (177 cases), to support their identity claims about a soldier's photograph. Users did not attribute a source in only 36 cases.

5.2 Identifying Photos

5.2.1 Users identified unknown photos using the website's search workflow.

Based on our log and content analysis, we found that users successfully used the system's search workflow to identify unknown photos. In the first month, 119 user-uploaded post-identified photos were matched to 88 existing soldier identities with a prior photo in the database. In some cases, more than one photo was matched to an identity.

Participants who added unknown portraits to the site described their success rates in enthusiastic terms. P1 remembered, "I was a half dozen in, and all of a sudden I got a hit on one of them." P5 described his experience: "I started running that whole pile of images that I had trying to find IDs on 'em, and I wanna say I found maybe 10 to 15% hits on images that I had squirreled away, that [Photo Sleuth] were able to compare to and bring up either the exact same image or an alternative that was clearly the same person." P2 noted, "Out of those 30 or 40 or 50 that I posted on there, I've successfully identified I think at least three. That's a pretty good success rate considering there were hundreds [of] thousands of people fighting in the war."

Participants also favorably compared Photo Sleuth to traditional research methods. P5 lamented that US state archives often lacked

searchable databases or digitized imagery, and aside from Photo Sleuth, "there's really nothing else out there as far as trying to find identifications for unidentified images." P8 emphasized that Photo Sleuth "saves a ton of time because now I don't have to just go through every single picture that's available ... When I first get an image, that's usually what I do — books, go online, search different areas, old auction houses ... But I kind of don't have to do that anymore because Photo Sleuth helps a lot."

Participants also recognized how the public nature of the system would affect their future collecting positively and negatively. P5 used the metaphor of a double-edged sword: "If I can find a match, it's good for me, but then it also may give somebody else that match, and then it becomes a bidding war whether I'm gonna pay more for it on eBay than that person is."

5.2.2 Users decided on a match based on additional clues in the photo beyond face recognition.

Based on our content analysis, we found that the post-identified photos included additional information that supported identification beyond face recognition. Of the 88 soldier identities that users matched during the first month, a significant proportion had additional helpful clues, such as the presence of an inscription (see Table 1). Additionally, participants told us that they considered other contextual information besides facial similarity, such as military service records, when making an identification. In P9's words, "Without more information besides the face, I'm not gonna say it's 100%."

Table 1: Types of Post-Identified Photos

Categories with at least one photo (per identity) satisfying this condition	Soldier Identities
Inscription and Replica	17 (17 positive)
Inscription but No Replica	21 (20 positive)
Replica but No Inscription	13 (13 positive)
No Replica and No Inscription	37 (25 positive)

5.2.3 Users checked multiple search results carefully before confirming a match.

During the backtracing process for post-identified photos, we observed that the matched identity did not always appear as the top search result (see Figure 4). Out of 119 post-identified photos matched, 11 did not have their identities in the top 50 search results, while 19 had their identities in the top 50 but not the top search result. This suggests users confirmed a match only after carefully analyzing the search results beyond the top few ones.

Interviewees compared the automated face recognition to their own capabilities. P5 and P1 noted that, as human researchers, they were more likely to be distracted by similarities and differences in soldiers' facial hair, whereas the AI focused on features that remained constant across facial hairstyles. P1 also gave an example of how the AI challenged his assumptions by finding a matching soldier from a location he had not initially included: "I'm convinced I never would have figured that one out without the site."

Some participants mentioned drawbacks in the face recognition AI. P3 and P4 emphasized the differentiating value of ear shape, a




Figure 2: Search results for an identity with both an inscription and replica. The uploaded photo can be considered a replica of the reference archive version displayed as the top search result.




Figure 3: Search results for an identity with only an inscription and no replica. The inscription on the photo says "DR Roys", which might have prompted the user to match "David R Roys" from the search results.

feature the AI does not consider. P8 observed that the AI often failed to recognize faces in profile (side) views, whereas he had no trouble. P4 felt he could outperform the AI on individual comparisons, but fatigue limited the number of images he could consider: "I still think that my eye could make the match better, but you just lose energy about it."

Participants also expressed a desire to solicit a second opinion from the community on the possible matches. We saw many examples of users posting screenshots of potential matches on social media and requesting feedback from fellow history enthusiasts. One potential benefit, described by P2, was consensus: "If a person posts a photograph and it's supposedly identified, you'll sort of see Facebook's hive mind kind of spin into action and in the comments,

and if there's some dissent then I think there reasonably is doubt. But if everybody just says, 'Yes, duh,' that's the person." Another potential benefit, P8 said, was noticing details one might have missed: "It's always better to have a second opinion or a second pair of eyes to point out things that maybe you were focused on that you didn't really see."




Figure 4: Search results in which the top result does not show the matched identity. This photo was correctly matched to "Orlando W Dimick."

5.2.4 Users are generally good at identifying unknown photos.

The expert analyzed all 88 identities matched with the post-identified photos and provided responses assessing identifications done by the users using a 4-point Likert scale for all 117 photos matched. Based on the expert's response, we consider the matches to be either *positive matches* (Likert-scale ratings of 3–4) or *negative matches* (ratings of 1–2).

As shown in Table 1, for the first and third categories in post-identified photos, i.e., when at least one replica was present for an identity, the expert validated responses for all 30 identities to be positive matches. For the second category, in which there is an inscription but no replica, only one out of 21 identities was validated as a negative match. We considered the final category of identities that did not have any inscriptions nor replicas to be the most difficult one. Out of 37 identities in this category, the expert assigned 12 identities to be negative matches and 25 as positive matches. Thus, the expert considered the vast majority of the identifications done by users in all categories to be positive matches.

5.3 Tagging Photos

5.3.1 Users tagged both unidentified and identified photos.

Based on the logs, we found that users had provided one or more tags for at least 401 of the 602 unidentified photos they added to the

website. Out of the 560 identified photos (both pre-identified and post-identified) added by users, 445 photos had one or more tags associated with them. Further, 115 of the 182 user who uploaded photos also tagged a photo with at least one or more tags.

Because adding tags was optional, we asked participants why they did or did not provide tags. Some participants (P5, P3, P6) added tags because they believed the tags would help retrieve more relevant search results). For this reason, P8 skipped tags that were not linked to search filters: "If he's a straight-up civilian and there's nothing to go off of, I'll just bypass [tagging] and just hoping the face recognition brings something." In contrast, P2 thought the overhead was minimal: "It's probably just about as easy to put in the correct information as not." Other participants, like P4, added tags because they thought they would help future users, but not necessarily themselves.

5.3.2 Users added uniform tags more often than others.

From the logs, we observed that users on an average added 5 tags per (tagged) photo, which was also the median count. We found that users provide tags related to both the photo's metadata (*Photo Format*, *Photographer Location*, etc.) and the visual evidence in the photos like (*Coat Color*, *Shoulder Straps*, etc.). *Coat Color* and *Shoulder Straps* were the most commonly tagged visual evidence, which the system uses to reduce search results by filtering military records by army side and officer rank, respectively.

6 DISCUSSION

6.1 Fostering Original Research while Preventing Misinformation

Prior work pointed to problems with misinformation in online history communities [57, 58], a concern also voiced by our study participants. In Photo Sleuth, we made design decisions explicitly to support accuracy and limit the spread of misinformation. One such decision was to give users the option to provide the original sources of identification to add credibility to their identification claims. Although this feature was optional, users took advantage of it for all but 36 of 386 pre-identified photos.

A second design decision to promote accuracy was requiring all users to go through the entire pipeline, even if they believed they already knew the pictured soldier's name. A third, related design decision was asking users to separate the visual clues they could actually observe in the image (e.g., tagging visible rank insignia) from their interpretation of the clues (e.g., activating search filters for certain ranks). Both of these design decisions encouraged tagging of more objective visual evidence, with 401 of 602 unidentified photos and 445 of 560 identified photos receiving tags. These interfaces allowed for clearer delineations between fact and opinion, and left room for reasonable disagreement.

In the first month, users post-identified 75 unknown historical portraits, including 25 in the most difficult category (no inscription and no replica). This is promising evidence of the success of our approach — in P6's words, traditionally, "it's really rare that you can identify a non-identified image." However, 13 of the 88 post-identifications were judged by our expert as negative matches, indicating potential misinformation. In future work, we are exploring allowing users to express more nuanced confidence levels in

their identifications, based on the the expert's 4-point Likert scale, as well as capturing user disagreements.

6.2 Building a Sustainable Model for Volunteer Contributions

We observed substantial volunteer contributions to Photo Sleuth in its first month, even without typical incentive mechanisms like points and leaderboards. In interviews, participants described a variety of motivations for adding and tagging both unidentified and identified photos, ranging from making money to preserving history. Observing the usage numbers on our website, we are optimistic that we have built a sustainable model for volunteer contribution.

Our workflow leverages network effects so that the more people use it, the more beneficial it becomes to all. Users, when uploading and tagging known and unknown photographs, are enhancing the reference archive. These photos, along with their visual tags and metadata, are bootstrapped into the system for future searches and identifications, allowing the website to continuously grow. These users are also publicly credited for their contributions. Designing crowdsourcing workflows that align incentive mechanisms for enriching metadata and performing searches, as well as publicly recognizing contributions, can help build a sustainable participation model.

6.3 Combining the Strengths of Crowds and AI

We deliberately decided not to allow the Photo Sleuth system *per se* to directly identify any photos. Although this feature is one of our most persistent user requests, examples from popular media show the danger of a fully automated approach [3, 49]. Instead, the system suggests potential matches largely driven by objective user tagging, and hides quantitative confidence levels. The face recognition algorithm influences results in a more subtle way, by filtering out low-confidence matches and sorting the remainder. We believe this approach improves accuracy, but at the cost of increased requirements for human attention per image. Because Photo Sleuth helps users quickly identify a much more relevant set of candidates compared to traditional research methods, participants did not seem to view this attention requirement as a major drawback.

This human-led, AI-supported approach to person identification is further emphasized in our design decision to attribute individual users as responsible for particular identifications. This approach aims to promote accountability through social translucence [16], and to recognize the achievements of conducting original research, as recommended by prior work [35, 36]. It also aligns with traditions of expert authentication in the art and antiquarian communities [15].

Unexpectedly, we saw and heard about users posting screenshots of Photo Sleuth on social media to solicit second opinions from the community. This suggests a potential benefit of the wisdom of crowds not yet supported by our system, but also potential dangers of groupthink. In future work, we are exploring ways to capture discussions directly within Photo Sleuth's Comparison interface, drawing inspiration from social computing systems supporting reflection and deliberation around contentious topics [24, 25].

6.4 Enhancing the Accuracy of Person Identification

Prior work on person identification has mostly been limited to studies of face recognition algorithms. These studies often focus on face verification evaluations, i.e., comparing two photos and providing a confidence score about how similar or different they are. The algorithm gives the final verdict on a potential match based on a confidence threshold. Such approaches are usually evaluated on fixed datasets, and are therefore prone to false positives. Even though human-machine fusion scores are shown to outperform individual human or machine performances, none of these systems propose a hybrid pipeline where human judgment complements that of a machine or vice versa.

Photo Sleuth addresses accuracy issues in person identification by enhancing face recognition with different layers of contextual information, such as visual clues, biographical details, and photo metadata. Users provide visual clues along with the face, which help the system in generating search filters based on military service records. This ensures that the facial recognition runs on a plausible subset of soldiers satisfying the clues. We also show how users consider photo metadata like period inscriptions and historical primary sources to correctly match a person with an identity. Since the final decision of identification is reserved for the users, they can make an informed decision based on the contextual information along with facial similarity.

In future work, this pipeline could be adapted for other historical or modern person identification tasks by incorporating a domain-specific database and tagging features in a context-specific manner. For example, to identify criminal suspects in surveillance footage or locate missing persons from social media photos, an initial seed database of identified portraits with biographical data could be fed to the system. The user interface could be tuned, with the guidance of subject matter experts, to support tagging relevant photo metadata and visual clues like distinctive tattoos, clothing styles, and environmental features. These tags could similarly be linked to search filters to narrow down candidates after face recognition. Especially in high-stakes domains like these examples, where both false positives and false negatives can have life-altering impacts, it would be critical for experts in law enforcement or human rights investigation to oversee the person identification process.

7 CONCLUSION

Photo Sleuth attempts to address the challenge of identifying people in historical portraits. We present a novel person identification pipeline that combines crowdsourced human expertise and automated face recognition with contextual information to help users identify unknown Civil War soldier portraits. We demonstrate this approach by building a web platform, Photo Sleuth, on top of this pipeline. We show that Photo Sleuth's pipeline has enabled identification of dozens of unknown photos and encouraged a sustainable model for long-term volunteer contribution. Our work opens doors for exploring new ways for building person identification systems that look beyond face recognition and leverage the complementary strengths of human and artificial intelligence.

ACKNOWLEDGMENTS

We wish to thank Ron Coddington, Paul Quigley, Nam Nguyen, Abby Jetmundsen, Ryan Russell, and our study participants. This research was supported by NSF IIS-1651969 and IIS-1527453 and a Virginia Tech ICTAS Junior Faculty Award.

REFERENCES

- [1] 2013. <https://www.battlefields.org/learn/articles/military-images-magazine>
- [2] Amazon. 2018. Amazon Rekognition Customers - Amazon Web Services (AWS). <https://aws.amazon.com/rekognition/customers/>
- [3] Press Association. 2018. Welsh police wrongly identify thousands as potential criminals. *The Guardian* (May 2018). <https://www.theguardian.com/uk-news/2018/may/05/welsh-police-wrongly-identify-thousands-as-potential-criminals>
- [4] M. Barrenechea, K. M. Anderson, L. Palen, and J. White. 2015. Engineering Crowdsourcing for Disaster Events: The Human-Centered Development of a Lost-and-Found Tasking Environment. In *2015 48th Hawaii International Conference on System Sciences*. 182–191. <https://doi.org/10.1109/HICSS.2015.31>
- [5] P. Bell and Björn Ommer. 2016. *Digital Connoisseur? How Computer Vision Supports Art History*. Artemide, Rome.
- [6] Michael S Bernstein, Joel Brandt, Robert C Miller, and David R Karger. 2011. Crowds in two seconds: Enabling realtime crowd-powered interfaces. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*. ACM, 33–42.
- [7] L. Best-Rowden, S. Bisht, J. C. Klontz, and A. K. Jain. 2014. Unconstrained face recognition: Establishing baseline human performance via crowdsourcing. In *IEEE International Joint Conference on Biometrics*. 1A–58. <https://doi.org/10.1109/BTAS.2014.6996296>
- [8] Jeffrey P Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samuel White, et al. 2010. VizWiz: nearly real-time answers to visual questions. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology*. ACM, 333–342.
- [9] Austin Blanton, Kristen C Allen, Timothy Miller, Nathan D Kalka, and Anil K Jain. 2016. A comparison of human and automated face verification accuracy on unconstrained image sets. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 161–168.
- [10] Erin Brady, Meredith Ringel Morris, Yu Zhong, Samuel White, and Jeffrey P Bigham. 2013. Visual challenges in the everyday lives of blind people. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2117–2126.
- [11] Tim Causer and Melissa Terras. 2014. Crowdsourcing Bentham: Beyond the Traditional Boundaries of Academic History. *International Journal of Humanities and Arts Computing* 8, 1 (April 2014), 46–64. <https://doi.org/10.3366/ijhac.2014.0119>
- [12] Justin Cheng and Michael S Bernstein. 2015. Flock: Hybrid crowd-machine learning classifiers. In *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*. ACM, 600–611.
- [13] Ronald S. Coddington and Michael Fellman. 2004. *Faces of the Civil War: An Album of Union Soldiers and Their Stories* (1st edition edition ed.). Johns Hopkins University Press, Baltimore.
- [14] Lhaylla Crisaffi, Louisa Wood Ruby, Samantha Deutch, R Luke DuBois, Jean-Daniel Fekete, Juliana Freire, and Claudio Silva. 2018. ARIES: enabling visual exploration and organization of art image collections. *IEEE computer graphics and applications* 38, 1 (2018), 91–108.
- [15] David Cycleback. 2017. *Authenticating Art and Artifacts: An Introduction to Methods and Issues*. lulu.com, S.I.
- [16] Thomas Erickson and Wendy A. Kellogg. 2000. Social translucence: an approach to designing systems that support social processes. *ACM Trans. Comput.-Hum. Interact.* 7, 1 (March 2000), 59–83. <https://doi.org/10.1145/344949.345004>
- [17] Jacey Fortin. 2018. A Photo of Billy the Kid Bought for \$10 at a Flea Market May Be Worth Millions. *The New York Times* (Jan. 2018). <https://www.nytimes.com/2017/11/16/us/billy-the-kid-photo.html>
- [18] Jacey Fortin. 2018. She Was the Only Woman in a Photo of 38 Scientists, and Now She’s Been Identified. *The New York Times* (Mar 2018). <https://www.nytimes.com/2018/03/19/us/twitter-mystery-photo.html>
- [19] Google. 2018. Google App Goes Viral Making An Art Out Of Matching Faces To Paintings. <https://www.npr.org/sections/thetwo-way/2018/01/15/578151195/google-app-goes-viral-making-an-art-out-of-matching-faces-to-paintings>
- [20] Derek L. Hansen, Patrick J. Schone, Douglas Corey, Matthew Reid, and Jake Gehring. 2013. Quality control mechanisms for crowdsourcing: peer review, arbitration, & expertise at familysearch indexing. In *Proceedings of the 2013 conference on Computer supported cooperative work (CSCW '13)*. ACM, New York, NY, USA, 649–660. <https://doi.org/10.1145/2441776.2441848>
- [21] Kotaro Hara, Shiri Azenkot, Megan Campbell, Cynthia L Bennett, Vicki Le, Sean Pannella, Robert Moore, Kelly Minckler, Rochelle H Ng, and Jon E Froehlich. 2015. Improving public transit accessibility for blind riders by crowdsourcing bus stop landmark locations with google street view: An extended analysis. *ACM Transactions on Accessible Computing (TACCESS)* 6, 2 (2015), 5.
- [22] Kotaro Hara, Vicki Le, and Jon Froehlich. 2013. Combining crowdsourcing and google street view to identify street-level accessibility problems. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 631–640.
- [23] Ira Kemelmacher-Shlizerman, Steven M Seitz, Daniel Miller, and Evan Brossard. 2016. The megaface benchmark: 1 million faces for recognition at scale. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4873–4882.
- [24] Travis Kriplean, Caitlin Bonnar, Alan Borning, Bo Kinney, and Brian Gill. 2014. Integrating On-demand Fact-checking with Public Dialogue. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '14)*. ACM, New York, NY, USA, 1188–1199. <https://doi.org/10.1145/2531602.2531677>
- [25] Travis Kriplean, Michael Toomim, Jonathan Morgan, Alan Borning, and Andrew Ko. 2012. Is This What You Meant?: Promoting Listening on the Web with Reflect. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 1559–1568. <https://doi.org/10.1145/2207676.2208621>
- [26] Ranjay A. Krishna, Kenji Hata, Stephanie Chen, Joshua Kravitz, David A. Shamma, Li Fei-Fei, and Michael S. Bernstein. 2016. Embracing Error to Enable Rapid Crowdsourcing. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 3167–3179. <https://doi.org/10.1145/2858036.2858115>
- [27] Neeraj Kumar, Alexander Berg, Peter N Belhumeur, and Shree Nayar. 2011. Describable visual attributes for face verification and image search. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33, 10 (2011), 1962–1977.
- [28] Gierad Laput, Walter S Lasecki, Jason Wiese, Robert Xiao, Jeffrey P Bigham, and Chris Harrison. 2015. Sensors: Adaptive, rapidly deployable, human-intelligent sensor feeds. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 1935–1944.
- [29] Walter S Lasecki, Mitchell Gordon, Danaï Koutra, Malte F Jung, Steven P Dow, and Jeffrey P Bigham. 2014. Glance: Rapidly coding behavioral video with the crowd. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*. ACM, 551–562.
- [30] Walter S Lasecki, Mitchell Gordon, Winnie Leung, Ellen Lim, Jeffrey P Bigham, and Steven P Dow. 2015. Exploring privacy and accuracy trade-offs in crowd-sourced behavioral video coding. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 1945–1954.
- [31] Chris J Lintott, Kevin Schawinski, Anže Slosar, Kate Land, Steven Bamford, Daniel Thomas, M Jordan Raddick, Robert C Nichol, Alex Szalay, Dan Andreescu, et al. 2008. Galaxy Zoo: morphologies derived from visual inspection of galaxies from the Sloan Digital Sky Survey. *Monthly Notices of the Royal Astronomical Society* 389, 3 (2008), 1179–1189.
- [32] Kurt Luther. 2015. Blazing a Path From Confirmation Bias to Airtight Identification. *Military Images* 33, 2 (2015), 54–55. <http://www.jstor.org/stable/24864385>
- [33] Kurt Luther. 2015. The Photo Sleuth’s Digital Toolkit. *Military Images* 33, 3 (2015), 47–49. <http://www.jstor.org/stable/24864403>
- [34] Kurt Luther. 2018. What Are The Odds? Photo Sleuthing by the Numbers. *Military Images* 36, 1 (2018), 12–15. <http://www.jstor.org/stable/26240155>
- [35] Kurt Luther, Scott Counts, Kristin B. Stecher, Aaron Hoff, and Paul Johns. 2009. Pathfinder: An Online Collaboration Environment for Citizen Scientists. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 239–248. <https://doi.org/10.1145/1518701.1518741>
- [36] Kurt Luther, Nicholas Diakopoulos, and Amy Bruckman. 2010. Edits & credits: exploring integration and attribution in online creative collaboration. In *CHI '10 Extended Abstracts on Human Factors in Computing Systems*. ACM, 2823–2832.
- [37] Ramona Martinez. 2012. Unknown No More: Identifying A Civil War Soldier. <http://www.npr.org/2012/04/11/150288978/unknown-no-more-identifying-a-civil-war-soldier>
- [38] Microsoft. 2018. Face API - Facial Recognition Software | Microsoft Azure <https://azure.microsoft.com/en-us/services/cognitive-services/face/>
- [39] Microsoft. 2018. Uber boosts platform security with the Face API, part of Microsoft Cognitive Services. <https://customers.microsoft.com/en-us/story/uber>
- [40] Jon Noronha, Eric Hysen, Haoqi Zhang, and Krzysztof Z Gajos. 2011. Platamate: crowdsourcing nutritional analysis from food photographs. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*. ACM, 1–12.
- [41] NPS. 2018. Soldiers and Sailors Database - The Civil War (U.S. National Park Service) <https://www.nps.gov/subjects/civilwar/soldiers-and-sailors-database.htm>
- [42] Genevieve Patterson, Grant Van Horn, Serge Belongie, Pietro Perona, and James Hays. 2015. Tropel: Crowdsourcing Detectors with Minimal Training. In *Third AAAI Conference on Human Computation and Crowdsourcing*.

- [43] P Jonathon Phillips, Amy N Yates, Ying Hu, Carina A Hahn, Eilidh Noyes, Kelsey Jackson, Jacqueline G Cavazos, Géraldine Jeckeln, Rajeev Ranjan, Swami Sankaranarayanan, et al. 2018. Face recognition accuracy of forensic examiners, superrecognizers, and face recognition algorithms. *Proceedings of the National Academy of Sciences* (2018), 201721355.
- [44] Roy Rosenzweig. 2006. Can History Be Open Source? Wikipedia and the Future of the Past. *Journal of American History* 93, 1 (June 2006), 117–146.
- [45] Michael E. Ruane. 2014. Facebook helps identify soldiers in a forgotten Civil War portrait. *Washington Post* (March 2014). https://www.washingtonpost.com/local/facebook-helps-identify-soldiers-in-a-forgotten-civil-war-portrait/2014/03/07/a4754218-a47a-11e3-8466-d34c451760b9_story.html
- [46] Michael S. Schmidt. 2018. “Flags of Our Fathers” Author Now Doubts His Father Was in Iwo Jima Photo. *The New York Times* (Jan 2018). <https://www.nytimes.com/2018/05/04/us/iwo-jima-marines-bradley.html>
- [47] Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 815–823.
- [48] Irving Seidman. 2006. *Interviewing As Qualitative Research: A Guide for Researchers in Education And the Social Sciences* (3 ed.). Teachers College Press.
- [49] Natasha Singer. 2018. Amazon’s Facial Recognition Wrongly Identifies 28 Lawmakers, A.C.L.U. Says. *The New York Times* (Jul 2018). <https://www.nytimes.com/2018/07/26/technology/amazon-aclu-facial-recognition-congress.html>
- [50] Ramya Srinivasan, Conrad Rudolph, and Amit K Roy-Chowdhury. 2015. Computerized face recognition in renaissance portrait art: A quantitative measure for identifying uncertain subjects in ancient portraits. *IEEE Signal Processing Magazine* 32, 4 (2015), 85–94.
- [51] Yi Sun, Yuheng Chen, Xiaogang Wang, and Xiaoou Tang. 2014. Deep learning face representation by joint identification-verification. In *Advances in neural information processing systems*. 1988–1996.
- [52] Yaniv Taigman, Ming Yang, Marc’Aurelio Ranzato, and Lior Wolf. 2014. Deepface: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1701–1708.
- [53] Alan Trachtenberg. 1985. Albums of War: On Reading Civil War Photographs. *Representations* 9 (1985), 1–32. <https://doi.org/10.2307/3043765>
- [54] USAHEC. 2018. MOLLUS-MASS Civil War Photograph Collection <http://cdm16635.contentdm.oclc.org/cdm/landingpage/collection/p16635coll12>.
- [55] Sarah Jones Weicksel. 2014. To Look Like Men of War. *Clio* No 40, 2 (2014), 137–152. https://www.cairn-int.info/article-E_CLIO1_040_0137--to-look-like-men-of-war.htm
- [56] Charlie Wells. 2012. UNKNOWN SOLDIER in famed Civil War portrait identified. <http://www.nydailynews.com/news/national/unknown-soldier-famed-library-congress-civil-war-portrait-identified-article-1.1142297>
- [57] Heather Willever-Farr, Lisl Zach, and Andrea Forte. 2012. Tell me about my family: A study of cooperative research on Ancestry. com. In *Proceedings of the 2012 iConference*. ACM, 303–310.
- [58] Heather L Willever-Farr and Andrea Forte. 2014. Family matters: Control and conflict in online family history production. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. ACM, 475–486.
- [59] A. C. Williams, J. F. Wallin, H. Yu, M. Perale, H. D. Carroll, A. F. Lamblin, L. Fortson, D. Obbink, C. J. Lintott, and J. H. Brusuelas. 2014. A computational pipeline for crowdsourced transcriptions of Ancient Greek papyrus fragments. In *2014 IEEE International Conference on Big Data (Big Data)*. 100–105. <https://doi.org/10.1109/BigData.2014.7004460>
- [60] Wenyi Zhao, Rama Chellappa, P Jonathon Phillips, and Azriel Rosenfeld. 2003. Face recognition: A literature survey. *ACM computing surveys (CSUR)* 35, 4 (2003), 399–458.