

Find Your Twin in History: Exploring Ethical Design Challenges in Facial Recognition




Figure 1: Results page for Civil War Twin. Showing the top 4 twins for an uploaded selfie in a reference database of Civil War-era soldiers and civilians.

**Vikram Mohanty, Marx Wang,
Manisha Kusuma, Kurt Luther**
Virginia Tech, USA
{vikrammohanty,boyuan,manishak,
kluther}@vt.edu

David Thames
Google, USA
davidc.thames@gmail.com

ABSTRACT

Facial recognition systems pose numerous ethical challenges, yet little guidance is available for designers. We explore these challenges in a three-step design process to create Civil War Twin, an educational web application where users can discover their lookalikes from the American Civil War era while learning more about both history and facial recognition.

KEYWORDS

Human-AI interaction, facial recognition, digital history, ethical design

INTRODUCTION

Recent advancements in AI and machine learning have opened up exciting possibilities for enabling novel, beneficial forms of human-AI interaction. However, AI's complexity, unpredictability, and over-reliance on data poses numerous challenges for designing ethical and effective AI-infused applications [9]. To address these challenges, HCI researchers have proposed multiple sets of guidelines for designing human-AI interaction [1, 15], AI fairness checklists [6, 12] and toolkits [2]. While these resources are valuable for addressing ethical issues in designing human-AI systems generally, relatively little guidance is available for designing facial recognition systems in particular.

We recently wrestled with these challenges while creating an educational web application in collaboration with the American Battlefield Trust (ABT), a Washington, DC-based non-profit organization focused on preserving battlefields in the United States. The application, Civil War Twin (CWT), would

Find Your Twin in History

Civil War Twin

Behind The Scenes

Now, we will search for your twins in a database provided by Civil War Photo Sleuth, a free, public website with over 30,000 Civil War-era portraits. Most of these photos come from public collections like the Library of Congress and the National Archives, but thousands have been contributed by individual users with interests in history and photography — people like you. Because there are so many potential matches, you can customize the search results that you want to see.

We use primary sources (such as military and medical records) and secondary sources (such as books published by historians) to determine whether the historical persons in our database served in the military or were civilians.

WE PROTECT YOUR PRIVACY AND DO NOT SAVE YOUR PHOTO OR FACEPRINT AFTER YOU LEAVE THE WEBSITE. [READ OUR PRIVACY POLICY HERE.](#)

Step 2 of 3

APPLY SEARCH PREFERENCE

Military

Gender

Race

91.1%

Union Civilian 143M

Confederate 12M

Civilian 10M

What results would you like to see?

You can decide whether you want to see matches from soldiers, civilians, or both.

 ☒ Union

☒ We have 143M Union soldiers in our database

There were approximately 2.1 million enlistments in the Union Army.

Read more: [ABT: Who fought?](#)

 ☐ Confederate

☐ We have 12M Confederate soldiers in our database

Due to destroyed and missing records, estimates vary, but approximately 750,000 to 1 million soldiers fought for the Confederacy.

Read more: [ABT: Who fought?](#)

 ☐ Civilian

☐ We have 10M civilians in our database

About 13.7 million men and most of the 10.3 women in the US remained civilians during the war. Men did not serve in Civil War armies for a variety of reasons, including being too young or old, disabilities, objections to the war, and even being substitutes to serve in their place. Some men were drafted into the Union or Confederate armies against their wishes.

Back

Next

Figure 2: Military search preferences page.

Civil War Twin

Behind The Scenes

We use primary sources (such as military and medical records) and secondary sources (such as books published by historians) to identify the gender of historical persons in our photo database.

Research shows that face recognition algorithms are less accurate for female faces compared to male faces, due to biases in their training datasets and the algorithms themselves.

Some face recognition algorithms try to guess the gender of a face using visual clues. Since that technology has problems, we don't use it here.

In the American Civil War era, most people identified as men or women, so our photo database is limited to these two genders.

WE PROTECT YOUR PRIVACY AND DO NOT SAVE YOUR PHOTO OR FACEPRINT AFTER YOU LEAVE THE WEBSITE. [READ OUR PRIVACY POLICY HERE.](#)

Step 2 of 3

APPLY SEARCH PREFERENCE

Military

Gender

Race

91.1%

Union Civilian 143M

Female 1M

Male 12M

What results would you like to see?

You can decide whether you want to see matches from male, female, or both.

 ☐ Male

☒ We have 12M men in our database

About 14 million males lived in the US during the Civil War, including about 2 million enslaved African American men. Of the 5.6 million white men of military age, about 3 million served in the military as volunteers or draftees.

 ☒ Female

☐ We have 1M women in our database

Women experienced the Civil War in numerous roles, as nurses, spies, slaves, abolitionists, civil rights advocates and promoters of women's suffrage.

Over 400 women dressed as men and served in the army. Many of them, like Sarah Rosetta Wakeman who served in the 133rd New York Infantry as Lyris Wakeman were not revealed to be women until after the war was over.

African American women like [Harriet Tubman](#) often took on especially dangerous roles, operating behind Confederate lines as Union scouts.

Learn more: [ABT: Female soldiers in the Civil War](#), [Women in the Civil War](#)

Back

Next

Figure 3: Gender search preferences page.

allow users to upload a selfie and, using facial recognition, discover their “Civil War twins,” i.e., photographs of people from the American Civil War era (1840s–1870s) who look like them. This “hook” was used to engage members of the public, who would then be presented with educational material about the historical figures, as well as information about the ABT’s mission. CWT builds on Civil War Photo Sleuth (CWPS), a free website we previously launched that combines crowdsourcing and facial recognition to identify unknown Civil War photos, which has attracted over 10,000 registered users who have identified hundreds of previously unknown photos [14].

Even though the workflow of CWT may appear simple, we quickly encountered a number of underlying ethical issues that required careful attention, including data privacy and transparency, gender and racial bias, and limitations of the historical archive. For example, prior work has shown facial recognition technologies have substantial accuracy and bias problems, especially on faces of women and dark-skinned people [3, 16], and also have been used as a surveillance tool that poses risks to individual privacy [8]. Despite these findings, majorities of Americans believe that “facial recognition can effectively identify individual people, as well as classify them by gender and race” [18]. This mismatch suggested to us a second educational opportunity: using CWT to educate the public not only about history, but also about facial recognition’s strengths and limitations. To accomplish these goals, we employed a three-step ethical design process: i) addressing known issues with facial recognition; ii) consulting experts in race, gender, and history; and iii) iterating on feedback from representative user groups. Below, we briefly describe the current system prototype, followed by a discussion of our design process and some key ethical challenges we have encountered and addressed.

SYSTEM DESCRIPTION: CIVIL WAR TWIN

Civil War Twin is a web-hosted Python/Django application built on the CWPS API. To begin, a user begins the matching process by uploading a selfie to Civil War Twin. We use Microsoft Azure’s Face API [13] to detect whether a face is present or not. If a face is detected, the user then provides their preferences for the type of search results they would like to see. There are three categories, one on each page: *Military* (Union, Confederate, and/or Civilian, see Figure 2), *Gender* (Man or Woman, see Figure 3), and *Race* (White, African American, Native American, Hispanic, and Asian, see Figure 4). Users can select multiple choices for each of these categories. On a sidebar, a “Behind the Scenes” section explains in layperson’s terms how the technology works, and a “What Could Go Wrong?” section describes potential shortcomings of the technology and/or historical records. Above the selection, a real-time interactive visualization shows how each user preference affects the search pool. Finally, facial recognition searches for the top four similar-looking candidates (“twins”) in this filtered pool based on similarity confidence scores, and displays the selected twin as a “baseball card” artwork, which the user can printed or share on social media (see Figure 1). The user can also learn more about their twins by reading their biographical details and/or clicking on their CWPS profile links.

Find Your Twin in History

Civil War Twin

Behind The Scenes

We use primary sources (such as military and medical records) and secondary sources (such as books published by historians) to identify the gender of historical persons in our photo database.

What Could Go Wrong?

Research shows that face recognition algorithms are more accurate for white faces compared to non-white faces, due to biases in their training datasets and the algorithms themselves.

Some face recognition algorithms try to guess the race of a face using visual clues. Since that technology has problems, we don't use it here.

Race in the American Civil War era was categorized by a complex mix of social and legal forces. For example, free mixed-race people were considered legally "white" if they had less than 1/16 African ancestry, depending on the state. Therefore, historical records may not reflect the racial identity that some historical persons preferred to use.

WE PROTECT YOUR PRIVACY AND DO NOT SAVE YOUR PHOTO OR RECORD ANY INFORMATION YOU LEAVE THE WEBSITE. READ OUR [PRIVACY POLICY](#) HERE.

Step 2 of 3

APPLY SEARCH PREFERENCE

Military

Gender

Race

91.1%

0.01%

0%

Union Soldier: 143%

Female: 14

White + African American: 0

What results would you like to see?

You can click the checkbox next to each race to see results from White, African American, Native American, Asian or Hispanic.

☒ White

We have 134 white people in our database

About 27 million white people lived in the US in 1860. About 13% of white Americans were immigrants, primarily from European countries like Germany and Ireland. One in four Union soldiers was an immigrant, compared to relatively few Confederate immigrants.

Learn more: [TIME: Immigrants in the Civil War](#)

☒ African American

We have 143 African Americans in our database

About 4.4 million African Americans lived in the US in 1860, of which about 90% were enslaved.

African American men were authorized by the Emancipation Proclamation to fight for the Union beginning in 1863. Over 180,000 African American soldiers and 20,000 sailors enlisted. Enslaved African Americans worked for the Confederacy as servants, laborers, cooks, and in other support roles, but not as soldiers.

Learn more: [ABT: Black soldiers in the Civil War](#) Kevin Levin, Searching for Black Confederates

☐ Native American

We have 1 Native American in our database

More than 20,000 Native Americans fought in the Civil War on both sides of the conflict.

☐ Asian

We have 1 Asian in our database

Several hundred men from Asia and the Pacific Islands fought in the Civil War as Union or Confederate soldiers.

Learn more: [NPS: Asians and Pacific Islanders in the Civil War](#)

☐ Hispanic

We have 1 Hispanic in our database

More than 20,000 Hispanics fought on both sides in the Civil War, primarily in the Southwest. Some had been Americans for generations, while others were immigrants from Spain, Cuba, Mexico, and other Latin American countries.

Learn more: [NPS: Hispanics and the Civil War](#)

Back

Next

Figure 4: Race search preferences page.

DESIGN PROCESS

1. *Addressing Known Issues.* We began our design process with a literature of academic and popular accounts of ethical concerns with facial recognition, and designed the CWT user experience to ameliorate them as much as possible.

Data Privacy and Transparency: In recent years, facial recognition has drawn criticism as a surveillance tool for law enforcement agencies [8]. A particular concern for individual privacy is that the training database for such tools often includes face images collected from CCTV camera footage, social media profiles and posts, and driver licence databases, often without the subjects' knowledge or consent [10].

To support privacy and user control, we designed our facial recognition system so that it does not permanently store the user's selfie photo on any web server (ours or Microsoft's). Their selfie is not used for training the facial recognition model, either. We also explicitly inform the user that the photo will be deleted as soon as the session is over. This design required a trade-off between user privacy and usability, as we could not create persistent URLs for individual users to find or share their twin search results. Instead, we encourage users to download the "baseball card" graphic for saving locally, printing, or sharing; we also email the user a message with links to their twins.

To increase transparency, we added the aforementioned "Behind the Scenes" sidebar panels to describe how the facial recognition algorithm uses the uploaded selfie to detect the facial landmarks, generate a template for the face, and search the reference database for similar-looking candidates. We further emphasize the conditions that may affect the face detection process.

Determining Search Results: Another key challenge was providing appropriate search results for each user. CWT targets a demographically diverse user base in the general public. One approach might be to pick twins for a user by matching automatically inferred demographic characteristics. For example, if facial recognition determined the user's selfie showed an African American male, then CWT could show only twins who were also African American males. However, we found this approach problematic for several reasons. First, automated techniques might perform badly when trying to guess the race or gender of the selfie. Multiple studies have shown facial recognition has accuracy and bias problems, such as for dark-skinned or transgender faces [3, 16, 17]. The chances of misgendering the user or inferring the wrong race are unacceptably high. Second, even if the user's demographics could be automatically inferred with high accuracy, it is hard to generalize what types of twins users will want to see. For example, some users might prefer to see twins of a different gender, or Confederate twins, while other users might find these matches offensive.

To address these issues, CWT gives users complete control of what type of search results (twins) they want to see by specifying search preferences, instead of automatically detecting or even requesting their demographic information. Users have the option of filtering the search results to particular

1. How do you feel about the overall premise, motivation and goals of the project? What value (if any) do you see in this project?
2. How can this project address rising ethical concerns related to facial recognition?
3. How might African American users respond to the way this project addresses the topic of Civil War history?
4. How might the focus on facial recognition software be perceived by people of color?
5. What are some issues regarding race, in the context of the Civil War era, this project may have overlooked?
6. What are some ethical considerations when asking users for their search preferences related to race?
7. What are some concerns regarding gender identities and gender roles this project may have overlooked?
8. How can we use this project to encourage contribution of historical photos to the CWPS database?
9. How can this website potentially be misused? Should aspect(s) of the platform be modified to prevent this?
10. How should we address the lack of photos for certain genders and ethnic groups in the CWT database?

Table 1: Sample questions for expert feedback.

genders, races/ethnicities, and military categories. These categories reflect the historical dataset and thus, the practices of the era. For example, soldiers' races are inferred based on their segregated military units, and gender is limited to the male–female binary. Instead of trying to hide these categories by mapping them onto modern-day labels, we employ seamful design [4] by presenting the historical context behind these categories and inviting the user to consider their own relationship to them. We also selected the Microsoft Azure Face API because among commercial facial recognition services, it has shown some of the most substantial gains in reducing race and gender bias.

Limits of the Archive: Due to historical circumstances, ranging from a Union naval blockade to discrimination, some groups, especially Confederate soldiers, women, and people of color, have fewer surviving photos [5]. These historical biases were echoed in the composition of the original CWPS database, which was seeded by public collections like the US Army's MOLLUS-Mass collection, which primarily contains portraits of white Union officers. One negative consequence is that users searching for twins in the other categories run the risk of seeing the same set of twins due to fewer photos in the reference database. Another negative consequence is that extant photos in these categories lack diversity. For example, in 2017, Google Arts & Culture released an app [7] which allowed users to upload a selfie and, via facial recognition, discover matching faces in digitized artworks from museums around the world. Although the app went viral, it drew criticism from the Asian community as the search results contained a disproportionately high number of Japanese geishas [11].

We attempt to mitigate these problems in several ways. First, we completed two targeted database enrichment projects to increase the number of photos of two underrepresented categories: African Americans and women. Consequently, our collection of African American Union soldier portraits, although small, is believed to be the largest digital collection in existence. Second, we took this opportunity to educate users about how the composition of reference databases affects the performance of facial recognition algorithms in CWT's sidebar paragraphs. Third, on CWT's twin results page, we invite users to contribute additional photos (e.g., from their personal collections or new sources) to further enrich underrepresented categories in the CWPS database.

2. Consulting Experts. After the initial design, we plan on consulting experts in gender studies, race studies, and Civil War history to critique our design decisions, specifically on how our design engages with some of the sensitive issues discussed above, and to have a better understanding of the societal implications of our proposed tool (see Table 1). Based on their feedback, we plan to iterate on our design and develop a high-fidelity prototype.

3. Collecting Prospective User Feedback. We also plan on getting feedback on the prototype through focus groups conducted with participants from a diverse range of demographics who are interested in history and may have little or no knowledge about facial recognition. After incorporating their feedback, we will plan for a public launch of Civil War Twin.

ACKNOWLEDGEMENTS

This research is supported by NSF IIS-1651969 and a Virginia Tech ICTAS Junior Faculty Award.

REFERENCES

- [1] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, et al. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–13.
- [2] R. K. E. Bellamy, K. Dey, M. Hind, S. C. Hoffman, S. Houde, K. Kannan, P. Lohia, J. Martino, S. Mehta, A. Mojsilović, S. Nagar, K. N. Ramamurthy, J. Richards, D. Saha, P. Sattigeri, M. Singh, K. R. Varshney, and Y. Zhang. 2019. AI Fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias. *IBM Journal of Research and Development* 63, 4/5, 4:1–4:15.
- [3] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. 77–91.
- [4] Matthew Chalmers, Ian MacColl, and Marek Bell. 2003. Seamful design: Showing the seams in wearable computing. (2003).
- [5] Ronald S. Coddington. 2008. *Faces of the Confederacy an album of Southern soldiers and their stories*. Johns Hopkins University Press.
- [6] Henriette Cramer, Jean Garcia-Gathright, Sravana Reddy, Aaron Springer, and Romain Takeo Bouyer. 2019. Translation, tracks & data: an algorithmic bias effort in practice. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–8.
- [7] Google Arts Culture. 2017. Art Selfie <https://artsandculture.google.com/camera/selfie>.
- [8] Clare Garvie, Alvaro Bedoya, and Jonathan Frankle. 2016. The Perpetual Line-Up | Unregulated Police Face Recognition in America <https://www.perpetuallineup.org/>.
- [9] Kelly A. Gates. 2011. *Our Biometric Future: Facial Recognition Technology and the Culture of Surveillance*. NYU Press, New York.
- [10] Kashmir Hill. 2020. The Secretive Company That Might End Privacy as We Know It <https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html>. *The New York Times* (2020).
- [11] Rachael Krishna. 2018. Asian People Are Not Impressed With Their Matches On Google’s Museum Selfie Feature <https://www.buzzfeednews.com/article/krishrach/asian-people-are-not-impressed-with-their-matches-googles>. *BuzzFeed News* (2018).
- [12] Michael A Madaio, Luke Stark, Jennifer Wortman Vaughan, and Hanna Wallach. 2020. Co-Designing Checklists to Understand Organizational Challenges and Opportunities around Fairness in AI. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [13] Microsoft. 2018. Face API - Facial Recognition Software | Microsoft Azure <https://azure.microsoft.com/en-us/services/cognitive-services/face/>.
- [14] Vikram Mohanty, David Thames, Sneha Mehta, and Kurt Luther. 2019. Photo sleuth: Combining human expertise and face recognition to identify historical portraits. In *Proceedings of the 24th International Conference on Intelligent User Interfaces*. 547–557.
- [15] Google PAIR. 2019. People + AI Guidebook <https://pair.withgoogle.com/guidebook/>.
- [16] Inioluwa Deborah Raji and Joy Buolamwini. 2019. Actionable auditing: Investigating the impact of publicly naming biased performance results of commercial ai products. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 429–435.

Find Your Twin in History

- [17] Morgan Klaus Scheuerman, Jacob M Paul, and Jed R Brubaker. 2019. How computers see gender: An evaluation of gender classification in commercial facial analysis services. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–33.
- [18] Aaron Smith. 2019. More Than Half of U.S. Adults Trust Law Enforcement to Use Facial Recognition Responsibly <https://www.pewresearch.org/internet/2019/09/05/more-than-half-of-u-s-adults-trust-law-enforcement-to-use-facial-recognition-responsibly/>. *Pew Research Center Internet Technology* (2019).