

## Ghosts in the Machine

NINA DEWI TOFT DJANEGARA

Ph.D. Candidate, Department of Anthropology, Stanford University

---

### ABSTRACT

The National Institute of Standards and Technology's (NIST's) Facial Recognition Vendor Test is considered the gold standard for assessing the performance of facial recognition models. Since the test series began in 1994, it has been regarded as an objective benchmark for corporations and academic research groups to compare their models against one another. To evaluate the accuracy of facial recognition models, NIST draws upon a database of millions of photos of US visa applicants and border crossers. Without their knowledge or consent, immigrants' photos have become raw material for the refinement of facial recognition software. This essay discusses the afterlife of these images, photos of immigrants and non-citizens that have been repurposed for a function that goes far beyond their original intent.

---

## PART ONE: THE FACIAL RECOGNITION VENDOR TEST

### Introduction

It is difficult to overstate the importance of the US National Institute of Standards and Technology (NIST) to the world of facial recognition. This modest US government laboratory in Gaithersburg, Maryland administers the Facial Recognition Vendor Test (FRVT), a competition that evaluates and ranks facial recognition algorithms submitted by developers from all around the world. Since the test series began in 1994, it has been regarded as an objective benchmark for corporations and academic research groups to compare their models against one another. If you have ever come into contact with facial recognition, chances are your face was scanned by an algorithm that took part in the FRVT.

The FRVT is like the Olympics for facial recognition algorithms. It sets a standard goal— like the 100-meter sprint— and competitors are ranked by their ability to achieve that target. Because the goal remains the same year after year, analyzing the results of the FRVT allows us to trace how much computers have improved in their ability to recognize faces.

FRVT rankings and performance metrics are featured in the promotional materials of many major tech companies. Take, for instance, NEC, a Japanese corporation that supplies facial recognition services to companies and government agencies in seventy countries. In 2019, NEC issued a press release declaring “NEC Face Recognition Technology Ranks First in NIST Accuracy Testing.”<sup>1</sup> The press release went on to boast that “NEC’s technology ranked No. 1 in NIST testing for the fifth time, following its top placement in the face recognition testing for video in 2017. The high

performance of NEC’s technology is reflected in the test results, which placed the company significantly ahead of the runner-up.”<sup>2</sup> NEC’s thousands of clients include major entities like Delta Airlines and the London Metropolitan Police. As of 2019, NEC’s products were used by over one-third of all state police departments and law enforcement agencies in the United States, according to the company’s own marketing materials.<sup>3</sup> Notably, NEC performs facial recognition for the U.S. Customs and Border Protection, scanning the faces of over sixty million travelers as they cross the border into the United States.<sup>4</sup>

At the time of this writing, the most recent champion of the FRVT was the French company IDEMIA, which beat out 281 other algorithms to earn the top spot in 2021.<sup>5</sup> One of NEC’s primary competitors, IDEMIA is a facial recognition provider with customers around the world, from Singapore’s Changi Airport to the Massachusetts Department of Motor Vehicles. Like NEC, IDEMIA also proudly touts its FRVT results on its marketing materials. At a 2021 trade show for the travel industry, a banner on IDEMIA’s booth proclaimed “#1 Face Recognition Vendor as Tested by NIST.”



***Figure 1. IDEMIA’s booth at the “Future Travel Experience Global Conference and Exhibition,” a travel industry event held in Las Vegas, Nevada in December 2021. IDEMIA signage advertises two major achievements: its top performance on the FRVT, as well as its longstanding contract with the US Transportation Security Administration (TSA). Photo by the author.***

Even those who don't achieve first place choose to feature their FRVT results in their marketing campaigns. This will often require some inventive reframing, including certain caveats to boost their perceived ranking in the FRVT tables. For instance, CyberLink says their company is "ranked 18th globally and ranked 8th if China & Russia vendors are excluded."<sup>6</sup> Similarly, Paravision frames itself as the "#1 ranked vendor from the US, UK, and EU and the #4 ranked vendor globally in both "Identification" and "Investigation" modes with mugshot and webcam images."<sup>7</sup> Meanwhile, Innovatrics bills itself as a "top performer" on the FRVT. Only if you scroll further down on their webpage, will you find that they ranked: "13th from 145 vendors," which they clarify is "better than all [Automated Biometric Identification System] competitors who submitted to NIST FRVT 1:1 benchmark."<sup>8</sup> The effort put into these creative reformulations demonstrates the importance of a prominent FRVT ranking within the facial recognition community.

What is clear is that technology companies believe that their performance on the Facial Recognition Vendor Test is a major selling point for potential buyers. But how did we get here? What does this test do and why has it become so influential within the facial recognition industry? Exactly whose faces are being identified? To understand the stakes of the FRVT, one must trace the complex history behind the development of facial recognition technology, the arrangements of power that funded early research in this field, and the incentives that have shaped what is now a multi-billion dollar industry.

## **The Beginning: Standardizing Success**

In 1993, the US Department of Defense (DoD) was trying to find a way to assess the accuracy of facial recognition algorithms.<sup>9</sup> Under the umbrella of its "Counter Drug Technology Program," the DoD launched the Face Recognition Technology (FERET) program, with the goal of providing "an independent method of evaluating algorithms and assessing the state of the art in automatic face recognition."<sup>10</sup> To illustrate the importance of the FERET program, computer scientists Inioluwa Deborah Raji and Genevieve Fried mark it as a major turning point in the evolution of facial recognition.<sup>11</sup>

One of the first tasks of the program was developing a benchmark dataset so that all algorithms could be tested against the same collection of images. The conceit of the FERET dataset is simple. Prior to its development, there was no way to know whether any given facial recognition algorithm was better than another. This is because the accuracy of facial recognition algorithms will differ depending upon the images used to evaluate their performance. At this point in time, research groups self-reported their own accuracy measures in their journal publications. However, it was impossible "to accurately evaluate or compare face-recognition algorithms published in the literature" because "each researcher collected their own database under conditions relative to the problem they were investigating."<sup>12</sup> In other words, one algorithm may have appeared to be better than another based on the accuracy rates reported in scientific publications, but it was impossible to tell whether the higher accuracy was because the model had been tested on an "easier" set of images.


Imagine a sprinter running 100 meters on a standardized running track. Now imagine another sprinter running 100 meters on a pebbly beach. Afterwards, one wouldn't be able to say who was the better athlete just by looking at their race times, since they ran under significantly different conditions. The FERET dataset was built with the intention of creating something like an Olympic running track: standardized, predictable, and comparable.

Professor Harry Wechsler of George Mason University was contracted by the DoD to assemble the images for the dataset. Over the course of 15 sessions between August 1993 and July 1996, 1,199 people volunteered to have their photo taken, eventually resulting in a dataset of 14,126 facial images.<sup>13</sup> Each volunteer’s face appears at least eight times in the database, with each image exhibiting minute differences in expression and pose.

To build a facial recognition model, developers supply the model with millions of images so that it can learn what a human face looks like. If an image is poorly lit, if the person’s face is cast in shadow or turned at a strange angle, it can be difficult for the model to decipher the face portrayed in the image. As the NIST evaluators state:

The characteristics and quality of the images are major factors in determining the difficulty of the problem being evaluated. For example, the location of the face in an image can affect problem difficulty. The problem is much easier if a face must be in the center of image compared to the case where a face can be located anywhere within the image. (P. Jonathon Phillips et al., *Face Recognition Vendor Test 2002 Evaluation Report* [Arlington, VA: NIST, 2003], 10)<sup>14</sup>

Visual inspection of the FERET database reveals how its creators attempted to overcome these difficulties by providing high-quality images with uniform backgrounds, closely cropped around the subject’s face.



***Figure 2. A selection of facial photographs included in the FERET dataset. Images in the public domain.***

In addition to compiling the dataset of faces, the DoD administered the FERET test, thereby offering a third-party comparison of facial recognition algorithms conducted by “an organization that will not see any benefit should one algorithm outperform the others.”<sup>15</sup> Eventually, this test

evolved to become the Facial Recognition Vendor Test (FRVT), which today is considered the gold standard for assessing the performance of facial recognition models. The test works as follows: developers submit their algorithms to a trusted third party that will then measure how all the algorithms perform against the same pool of data. The FRVT is a closed evaluation, meaning that participating researchers do not have direct access to the images that comprise the test set. This limits their ability to game the system by designing an algorithm that performs perfectly on that one specific dataset. Because the test is administered by an independent arbiter, developers can submit their algorithms without fearing that proprietary knowledge will be compromised or that test results will be manipulated. This evaluation model has been widely accepted by the facial recognition industry, as evinced by the large number of developers who have elected to participate in the FRVT. While only five algorithms participated in the first FERET evaluation in 1994, the 2022 iteration of the FRVT assessed 742 algorithms, which were submitted by 271 unique developers.<sup>16</sup>

In 2002, the National Institute of Standards and Technology (NIST) took over responsibility for running the FRVT from the Department of Defense. Apart from the FRVT, NIST is charged with standardizing weights and measures, maintaining the atomic clock, and establishing standard references for a range of commercial products, from alloyed steel to peanut butter. One might find it somewhat curious that the Facial Recognition Vendor Test was transferred from the DoD to a bureaucratic organization housed within the US Department of Commerce,<sup>17</sup> but as a federal agency responsible for standardization, NIST was in fact well positioned to maintain a benchmark test that was created to harmonize facial recognition technology. The transfer of responsibilities from the Department of Defense to NIST, however, also suggests that by 2002 the federal government was starting to regard facial recognition as a commercial product, rather than a defense technology. In this sense, facial recognition followed a path paved by other technologies, like the internet or GPS, which were initially developed by the military before being introduced into a civilian life.

Under NIST's oversight, the FERET database was replaced by the High Computational Intensity Test (HCInt) dataset, which contained "121,589 operational images of 37,437 subjects," which were "provided from the U.S. Department of State's Mexican non-immigrant visa archive."<sup>18</sup> This substitution of FERET for a larger dataset of visa images was motivated in part by the desire to understand how database size affects recognition performance. Because the FERET dataset contained fewer than 2,000 subjects, the introduction of the HCInt dataset represented a major increase in the number of people that the algorithms were tasked with identifying.

Beyond the size of the dataset, the switch to the HCInt dataset represents a major conceptual and political shift. While the FERET dataset contained images of people who consented to have their photo taken for scientific experimentation, the HCInt database was made up of photos of unwitting visa applicants.<sup>19</sup> Their photos were taken from their original purpose as a form of bureaucratic documentation and converted into raw materials for the development of facial recognition. This was one of the first instances of a logic that continues to underpin the facial recognition industry: the idea that it is easier and inconsequential to appropriate existing images than to construct unique datasets.<sup>20</sup>

Creating the FERET dataset required locating volunteers and transporting them to the George Mason University campus or one of the Army Research Laboratories in Maryland or Virginia. Efforts were required to standardize the images: to maintain a degree of consistency throughout the database, the same physical setup is used in each photography session.<sup>21</sup> Even so, the researchers

noted that some variations were introduced due to the necessity of disassembling and reassembling the studio equipment between photo sessions.

Visa photos offered an attractive alternative to the costly and painstaking work of assembling the FERET dataset. They could be acquired in large numbers, since they were already in possession of the US government. Furthermore, their illumination, facial expression, and pose are already standardized, thereby meeting the optimal conditions for facial recognition developers. In its 2002 Evaluation Report, the NIST team described the value of the compiled dataset of visa photos:

The result is a set of well-posed (i.e., frontal to within 10 degrees) images of cooperative subjects usually with a mouth-closed neutral facial expression. The subject usually occupies a quarter of the image area. The top of the subject's shoulders is almost always visible. The background is universally uniform, with a white background (in most cases). (P. Jonathon Phillips et al., 2003, 15)


The organizers of the FRVT regarded the State Department's collection of visa photos as raw data, ripe for extraction.<sup>22</sup> In doing so, they naturalized the work that goes into making visa photos, which are a very particular sort of image.

### **Another Beginning: The Identity Photograph**

Visa photos have a distinct visual aesthetic. With a clear, white background, neutral facial expression, and full-frontal angle, a visa photograph abstracts the image of a person's face from underlying context or nuance. The simplicity of the image belies the labor that goes into creating it.

Visa and passport photographs must meet certain requirements to ensure that identifiable facial features are visible and that the photo is at an appropriate resolution to be digitized and algorithmically processed.<sup>23</sup> There are rules about proper lighting and pose, an effort to tame the untidy reality of human expression. During fieldwork at one of the many passport photo studios adjacent to the US embassy in Jakarta, I observed how photographers and photo subjects are made responsible for adhering to these standards. On one afternoon in August 2019, Halim,<sup>24</sup> a middle-aged photographer with carefully groomed facial hair, arranged me in the proper pose for a passport photo, adjusting the slant of my shoulders and instructing me to tuck my hair behind my ears. Later, I watched as he used photo editing software to remove any shadows from the background, transforming it into an empty white plane. On the wall was a wrinkled poster detailing the international standards for visa and passport photos; a red X appeared in the lower right-hand corner of certain photos that had been deemed inadequate—rejected because the subject was wearing glasses, smirking, showing too much of their hair, or none at all.

Paying attention to photographic practices in studios like Halim's illuminates how the work of making facial recognition feasible is distributed to actors beyond the government or the technology industry. Human labor goes into creating photographs that meet standards for lighting, resolution, and pose, images that are later used as benchmarks for testing facial recognition model performance. This work is often overlooked because of its remoteness (both geographic and conceptual) from the site of algorithmic calibration. Nevertheless, passport photographers have made an important contribution to the development of facial recognition technology by producing data that are legible to machines.



**Figure 3.** Guidelines for passport photos as specified by the UN International Civil Aviation Organization. Image in the public domain.

To my chagrin, when editing my passport photo, Halim took it upon himself to erase a constellation of pimples from my left cheek. The result is a portrait that manages to be both aesthetically flattering and somewhat insulting. This kind of airbrushing is relatively commonplace, though not

officially sanctioned. In fact, the U.S. State Department stipulates that “digital manipulation and/or retouching of the facial image on photographs is not acceptable. The use of beauty filters or other photo filter tools ... are also not acceptable.”<sup>25</sup> Halim’s efforts to beautify my photo are a reminder that photographers do not always adhere to the official guidelines for passport photos. There is a degree of intractability in all artistic production; photographers and photo subjects may have their own agendas that exceed the technical specifications laid out by the state. These traces of creative expression will linger, even after the images are repurposed to new ends.<sup>26</sup>

In Halim’s photo, my skin—normally the color of a lightly fried plantain—is lustrous and pink. The face in the photo is not the one I see in the mirror; at best, it is my poreless and sunburnt twin. But it is the photo that now appears in my passport, which I renewed in February 2020. Presumably, this photo also sits in some government database, quiescent and waiting for the day it will become useful once again.

### Function Creep

Over the years, NIST turned to other databases already held by the US government as sources of photos to be included in its Facial Recognition Vendor Test.<sup>27</sup> In 2010, the FRVT began testing on a second, larger dataset of non-immigrant visa images (6,249,392 images of 5,738,141 subjects), this time from around the world, whereas the previous visa dataset had included only Mexicans.<sup>28</sup> A large dataset of FBI mugshot images (2,407,768 images of 1,802,874 subjects) was also harvested for use in the 2010 test.<sup>29</sup>



***Figure 5. Border detainee photos included in the final report for the 2013 Facial Recognition Vendor Test. The population of this dataset is listed simply as "Central America, adult." Images originally appeared in Patrick Grother and Mei Ngan, Face Recognition Vendor Test (2014),<sup>29</sup> a US government report in the public domain.***

Later, in 2013, an additional dataset of “border detainee booking photos” was added to the test. These booking photos were taken from law enforcement records; the people in them were apprehended near the US-Mexico border and were most likely deported shortly after their photos were taken. It is difficult to look at the images and not feel as if they are, in some way, haunted.<sup>30</sup>

The border detainee photos are of much poorer quality than the visa photos. They were taken using an “inexpensive webcam” and “are in considerable violation of most quality-related clauses of all

face recognition standards” due to deficiencies such as “non-frontal pose (associated with the rotational degrees of freedom of the camera mount), low contrast (due to varying and intense background lights), and poor spatial resolution (due to inexpensive camera optics).”<sup>31</sup> In other words, the images are grainy, the lighting is poor, and many of the subjects are not looking directly at the camera.

If the images were of such low quality, why even include them in the FRVT dataset? By 2013, facial recognition algorithms had improved significantly since the first evaluation was held nearly ten years before. As the algorithms got better, the test needed to evolve as well; additional milestones were required to see how the best algorithms would perform under more constrained conditions. The FRVT organizers explained that the border detainee images allowed them to “show how recognition accuracy degrades in non-ideal poorly-controlled situations,” allowing them to estimate whether the technology was good enough to identify people’s faces from grainy surveillance footage.<sup>32</sup> In order to meet their new objective—assessing facial recognition algorithms’ ability to identify people based on poor quality images—NIST yet again turned to a pre-existing government database, this time containing images of Central Americans who were detained at the southern US border.

In recent years, critics have been alarmed by research demonstrating that facial recognition models return less accurate results for people with darker skin. In 2019, NIST released a “Demographic Effects” report which confirmed that the majority of contemporary facial recognition models are less accurate for certain people, finding that “false positive rates are highest in West and East African and East Asian people.”<sup>33</sup> The NIST report expanded upon prior research by Joy Buolamwini, Timnit Gebru, and Deborah Raji, and, by testing a larger set of facial recognition models and using a larger database of images, provided further empirical evidence that these discriminatory outcomes exist.<sup>34</sup>

One of the challenges of conducting audits for facial recognition models has been the dearth of diverse training and testing examples; it is difficult to measure whether facial recognition models are biased because there are not enough databases containing high-quality images of non-white faces. To perform her audit of commercial facial analysis algorithms, Joy Buolamwini created an original dataset that featured headshots of members of parliament in Finland, Iceland, and Sweden (to represent light-skinned faces) and Rwanda, Senegal, and South Africa (to represent darker-skinned faces). Buolamwini built the dataset from images of parliamentarians because “they are public figures with known identities and photos available under non-restrictive licenses posted on government websites.”<sup>35</sup> The final parliamentarian dataset contained images of 1,270 faces. Compare this to the FRVT, which in 2019 analyzed “a total of 18.27 million images of 8.49 million people.”<sup>36</sup> Considering the lack of diverse datasets, the US government databases of immigrant photos have become especially valuable, a scarce resource in an industry that is increasingly motivated to improve facial recognition performance across racial categories.<sup>37</sup>

The diversity of the visa photo database can be attributed to larger geopolitical trends, such as immigration patterns of who comes to the United States, as well as political conditions that stipulate which countries’ citizens are required to apply for a visa to enter the US. Citizens of only thirty-nine countries, the majority of them in Europe or East Asia, are currently permitted to enter the US without a visa. This leaves citizens from all other countries at risk of having their photo appear in a facial recognition evaluation dataset. Indeed, the 2019 evaluation included photos from “more than 100 countries involved in immigrant and non-immigrant application processes.”<sup>38</sup> The

relationship between immigration trends and dataset composition was especially noticeable in the 2002 and 2013 FRVT evaluations, which featured datasets that contained only Mexican and Central American subjects.

As facial recognition technology advanced, new sources of data were needed to challenge the best-in-class models and improve accuracy across the field. Over time, the NIST evaluators made a habit of culling data from pre-existing government databases to keep up with the pace of development of facial recognition technology. Looking back on the history of the FRVT, we see a recurring pattern in which more and more immigrant faces—including those of children and infants—were added to the evaluation datasets. The logic is one of accumulation, as if adding more data will overcome the issues of inaccuracy and bias that plague facial recognition. By 2019, the “Demographic Effects” evaluation was performed in the following way:

We used these algorithms with four large datasets of photographs collected in U.S. governmental applications that are currently in operation:

- Domestic mugshots collected in the United States.
- Application photographs from a global population of applicants for immigration benefits.
- Visa photographs submitted in support of visa applicants.
- Border crossing photographs of travelers entering the United States. (Patrick Grother, Mei Ngan, and Kayee Hanaoka, *Face Recognition Vendor Test Part 3: Demographic Effects* [Gaithersburg, MD: National Institute of Standards and Technology, December 2019], 1)

Three of these four categories are intrinsically tied to immigration and borders. The immediate privacy concern is that the facial photos of millions of immigrants have become raw material for the refinement of facial recognition software without their knowledge or consent. Although these individuals may have willingly submitted their photos and personal information as part of their application for visas or immigration benefits, these photos have been repurposed for a use that goes far beyond their original intent. This exemplifies a phenomenon known as *function creep* or *surveillance creep*, in which systems “originally intended to perform narrowly specified functions are expanded ... thereby sidestepping or pushing the limits of legal frameworks meant to protect issues of privacy and data protection.”<sup>39</sup> Applicants who are not granted a US visa may never enter the country, yet their faces and personal information will exist in perpetuity in US government databases. There is a certain irony for those whose applications are denied; the data body crosses borders, while the biological body must remain in place. This transnational flow of data is particularly perverse when it involves photographs of people who were immobilized at or by the border, such as applicants who were denied a US visa or Central Americans who were deported from the US.

In the introduction, I shared promotional materials from leading facial recognition vendors, like IDEMIA and NEC, in which they boast about their contracts with U.S. Customs and Border Protection, the Transportation Security Administration (TSA), and major international airports. This is because one of the primary applications for facial recognition is the enforcement of national borders. In this regard, the visa photo has come full circle. However, those same facial recognition algorithms are used for a variety of purposes beyond border checkpoints. In an immediate sense, this presents a technical problem: models that have been optimized to perform well on visa

application photos and photos taken at border crossings will likely be less accurate when applied to different use-cases—whether it’s scanning your face to unlock your phone or using facial recognition to prove your identity to access government services.<sup>40</sup> On a deeper level, however, what does it mean that these images are used as the benchmark for testing facial recognition model performance? What are the implications when algorithms designed and built to perform well on photos of immigrant faces are later applied to a variety of commercial purposes?

## PART TWO: THE AFTERLIFE

### The Afterlife of Images

In the section above, I traced the life cycle of a visa photo from the point when a person applies for a visa to the United States. Their photograph is taken, their face captured, their identity flattened into a standardized frame. The photo is submitted with their visa application, which may be accepted or denied. The result is unimportant. Indeed, the person themselves becomes unimportant—their dreams, their desperations, their reasons for applying for a US visa.<sup>41</sup> Regardless of the outcome of this process, their visa photo will enter a US government database and may one day be used to train computer vision models as a data point for the Facial Recognition Vendor Test.<sup>42</sup> The person in the photo will never know about the role they have played in developing facial recognition technology.

The conversion of a personal photo to the raw data used to train computer vision models is an all-too-common path.<sup>43</sup> Many of the popular image datasets used to train facial recognition algorithms are made up of millions of pictures obtained covertly from social media or photo sharing websites like Flickr and Yahoo.<sup>44</sup> Tech developers refer to these images as “faces in the wild,” since the photos exhibit “natural variations” in dimensions like facial expression, lighting, and pose.<sup>45</sup> In 2019, *New York Times* reporters Kashmir Hill and Aaron Krolik tracked down some of the people whose photos had been scraped from their Flickr albums and included in databases of images used to train facial recognition algorithms.<sup>46</sup> “It’s gross and uncomfortable,” said a teenager whose photo had been uploaded by their parents when they were still a minor. “I wish they would have asked me first if I wanted to be part of it. I think artificial intelligence is cool and I want it to be smarter, but generally you ask people to participate in research. I learned that in high school biology.”<sup>47</sup>

For more than a decade, the public was generally unaware of the scope of biometric databases and the incredible number of photos that were scraped from the internet. Within the past five years, academics and journalists have begun to take note, converging into an emerging field of research that Nanna Thylstrup calls “critical dataset studies.” By providing a critical history of the Facial Recognition Vendor Test, this article contributes to ongoing critical conversations about how the images in computer vision datasets are sourced.<sup>48</sup> However, whereas previous work on this topic has discussed the over-/under-representation of certain demographics, the reductive labels attached to image datasets, and the ethics of scraping photos off the internet, this article addresses a somewhat different concern. The images of immigrants included in the Facial Recognition Vendor Test are not “faces in the wild,” they are faces captured for specific purposes in institutional settings, which raises a different set of questions about ethics and technology. The inherent coerciveness of the US border regime means that legal conventions about privacy and consent are insufficient for protecting the biometric data of visa applicants and travelers.

Photojournalist Martina Bacigalupo engages in a different type of photo recycling in her piece *Gulu Art Studio*, a series of found photographs collected from a photo studio in northern Uganda. Bacigalupo has explained that, because of the equipment constraints of this particular studio, when customers wanted an ID photo, it was cheaper and easier to print a full-size portrait and punch out the relevant section around the face. The leftovers were simply thrown away—until Bacigalupo came to collect them.



***Figure 4. Martina Bacigalupo, Gulu Real Art Studio, 2011–12. © The artist. Photo courtesy of the artist and The Walther Collection.***

In *Gulu Art Studio*, we see rows of headless figures with empty white squares where the faces should be, like a poem in which the spaces between words are what speaks the loudest. Upon further inspection, certain details become apparent: army fatigues, a purse clutched in a woman’s lap, an oversized blue blazer puddling around a man’s wrists. Though their faces are absent, the figures are still, in some way, identifiable. And while the standardized format of the identity photo produces a sense of interchangeability, *Gulu Art Studio* foregrounds the subjects’ uniqueness and individuality. Through the serial arrangement of the discarded photos, Bacigalupo tells us that we can learn something about a person from what is left out of the frame of the identity photo. She recounts, “During the editing process I went over hundreds of leftovers, where small details—a hidden sign, a comic posture, a bitter aspect—were revealing different stories.”<sup>49</sup> In particular, she recalls one man’s muddy boots, or “the purple napkin between the knees of a man who obviously walked a long way to the studio and made sure his face was clean before the photo was taken.”<sup>50</sup> By turning our eye to the clothing and objects worn by the people posing in the photos, Bacigalupo reminds us of how ID photos are aspirational, imbued with hopes of a different life.<sup>51</sup> In rescuing the “leftovers” of these identity photos, Bacigalupo endows them with an afterlife.

In Indonesia, where Halim took my passport photo, there is a historical tradition of repurposing identity photos (*pasfoto*) as personal mementos. In the early days of the post-colonial Indonesian nation-state, the *pasfoto* was often the first (and perhaps only) photo that people had of themselves. Visual anthropologist Karen Strassler has examined how *pasfoto* took on new meaning as they were exchanged among friends, pasted into photo albums, and integrated into funeral ceremonies. In this way, “identity photographs enter into spheres of social relations and identification that challenge the state’s claim to be the sole agent of recognition within the nation.”<sup>52</sup> This reappropriation of passport photos is a personal reclaiming of a certain representation of the self, the very opposite of web-scraping without consent.

As mobile objects, photos move in and out of context as they are repurposed toward different ends. Each in their own way, both Strassler and Bacigalupo call attention to the many lives of the passport photo and its plasticity as an image that simultaneously facilitates border crossing, acts as a mode of self-representation, and enables the calibration of facial recognition algorithms.<sup>53</sup>

## The Afterlife of Data

The afterlife of the visa photo is its transformation into raw data. The photo is uprooted from its original context, and the identity of the person in it is rendered unimportant. In its afterlife, the image is rendered into pure information—impersonal and atemporal. This metamorphosis of an identity photograph into raw data is the first afterlife. There is a second, murkier afterlife that follows, when the data become immortalized in a facial recognition model.

Computer models are byproducts of the specific data used to create them. In the words of the Stanford Institute for Human-Centered Artificial Intelligence, “user data does not only exist in its raw form in a database, it is also implicitly contained in models trained on that data.”<sup>54</sup> Computer models learn a certain representation of the world based upon the data they have been fed.<sup>55</sup> If you were to train a facial recognition model using 10,000 pictures of white men, its entire conception of what a face looks like would be skewed because it has seen only the faces of white men.<sup>56</sup> If you were to input different training images, you would generate a fundamentally different model. For this reason, a researcher attempting to create a less biased facial recognition model would likely train the model on a more diverse dataset of faces. However, even if you were to train your facial recognition model on 10,001 pictures of white men, this would *still* result in a different model from the one trained on 10,000 images. Each data point leaves its mark, which means that adding or deleting a single entry would alter the entire model.

As a further illustration of this point, consider the following thought experiment offered by legal scholar Tiffany C. Li:

A serial killer and conceptual artist harvests 100 human ears and creates a sculpture in which he places each of the ears on a surface of wet clay, which then dries into a terrible blob of clay that is covered in human ears. One of the serial killer’s victims miraculously survived and now requests his ear back. Detectives and/or art gallery personnel are able to remove the ear from the surface of the sculpture, but the imprint of the ear persists. While the victim may have the ear back, to use at his disposal, he can never erase the imprint of his ear on the resulting sculpture. The victim still suffers harm—both the harm of the physical violence as well as the psychological and emotional harm of seeing the imprint of his ear on the sculpture.

The ear-less man is analogous to the data subject of a privacy violation, whose data (ear) has been stolen and misused by a privacy violator (serial killer) to create a machine learning model (clay and ear sculpture). While the data subject can request that their data be deleted (that the ear be pulled from the sculpture), the subject cannot remove the persistent algorithmic shadow, the imprint of their data on the resulting machine learning model.<sup>57</sup>

Li's example is both grisly and evocative. The imagery of severed ears underscores how some consider the seizure of personal data to be an exercise of power rooted in violence. And through the imaginary surface of wet clay, Li emphasizes the persistence of a single data point, which continues to leave a mark on an algorithm, even after the original data have been deleted.

This leads us to an open research problem in the field of computer science, the “data deletion conundrum.” Researchers are faced with the following dilemma: how do you delete data from the training set without drastically affecting the model? As James Zou, a leading expert in artificial intelligence, explains, “[when] training our machine learning models, bits and pieces of data can get embedded in the model in complicated ways. That makes it hard for us to guarantee a user has truly been forgotten without altering our models substantially.”<sup>58</sup>

This growing sub-field of computer science research underscores the difficulty of disentangling the training data from the model. This research agenda has been spurred by recent legal regulations that can require technology developers to delete data from their models, especially if the data were obtained unscrupulously. For instance, “right to be forgotten” laws in places like the European Union and Argentina allow people to request that their personal data be deleted from the databases of research institutions and private companies. In the United States, the Federal Trade Commission (FTC) has ordered some tech companies to delete data, as well as models produced from those data, if the data were determined to be collected without users’ consent.

A 2021 decision by the FTC found that Everalbum, a California-based company, had deceived its customers about its data storage and use policies. Rohit Chopra, the FTC Commissioner at the time of the ruling, issued a public statement that said Everalbum had “enhanced their facial recognition technology by allegedly baiting consumers into using Ever, a “free” app that allowed users to store and modify photos.”<sup>59</sup> While the public face of Everalbum was a (seemingly) innocuous photo editing app, the company appropriated users’ photos to improve its facial recognition technology and then sold that technology to clients like the US Department of Defense and the Air Force through a separate arm of the company named Paravision. As a penalty for misleading customers about how their photos were being used, Everalbum was ordered by the FTC to “to forfeit the fruits of its deception” and delete the ill-gotten images from its servers.<sup>60</sup>

Such legal requirements create a problem for tech companies. Even when they are legally obligated to delete user data, companies do not want to sacrifice their entire facial recognition model, which would be prohibitively expensive to rebuild from scratch. Model developers must therefore figure out a way to delete training data with minimal harm to their models.

Some companies have even chosen to preemptively delete controversial datasets, even before they were required to do so by law, as a response to public outcry and negative media coverage. However, while companies like Microsoft<sup>61</sup> and Meta<sup>62</sup> have pledged to delete large databases of people’s faces, the original images live on in facial recognition models. Because facial recognition

models are inextricably entangled with the photos used in their development, deleting objectionable images may not be sufficient to redress past harms.

The data that contribute to computer models have a long afterlife. Traces linger, long after models have been trained, tested, refined, and deployed in the real world. This fact has important consequences. Even if the Facial Recognition Vendor Test were to stop using immigrant photos in future evaluations, that would not be enough to remove the imprint of those photos on the field of computer vision.

Recent research has demonstrated that it is possible to reverse-engineer a model to reveal the data that were used to train the model.<sup>63</sup> Essentially, this means that a person who only has access to a model could work backwards and reconstruct the training dataset using clues that have been built into the model. Computer scientist Jonathan Brophy explains that because “machine learning models have a ‘memory’ of the data they are trained on, they can then leak information about that data. Even if this data has been deleted ... the data continues to exist in their learned models.”<sup>64</sup>

Computer models have memories, Brophy tells us. How far back do they remember? In this essay, I have traced the life cycle of the visa photo in order to illuminate both its prehistory and its afterlife. I have described the downstream impact of visa photographs that have played an important—but mostly invisible—role in the development and commercialization of facial recognition technology. In addition, I have reflected on the origin of these images, whose very existence is the result of long histories of imperialism and colonialism that determine which countries’ citizens are required to apply for visas in the first place. I have argued that the images themselves are imbued with politics due to the distinct aesthetic of visa photos, which are taken by photographers following explicit guidelines about composition and pose to conform to government standards. By structuring my investigation in this way, I suggest that attending to the afterlives of data also requires a return to their origins.

## ACKNOWLEDGEMENTS

I would like to thank Margaux Fitoussi, Myriam Amri and Valentina Ramia for their constant reassurance and brilliance. Without their support, I would still be staring at an empty page. I am also grateful to the participants of the STS-MIGTEC Workshop in Athens—particularly Silvan Pollozek, Margie Cheesman, Salah El-Khalil, Mirko Forti, Romain Lanneau, and Wouter van Rossem—for their feedback on an early draft and for reinvigorating my enthusiasm for the subject.

## ENDNOTES

<sup>1</sup> NEC, “NEC Face Recognition Technology Ranks First in NIST Accuracy Testing,” *NEC News Room*, October 3, 2019, [https://www.nec.com/en/press/201910/global\\_20191003\\_01.html](https://www.nec.com/en/press/201910/global_20191003_01.html).

<sup>2</sup> Ibid.

<sup>3</sup> Dave Gershgorn, “NEC Is the Most Important Facial Recognition Company You’ve Never Heard Of,” *OneZero*, February 21, 2020, <https://onezero.medium.com/nec-is-the-most-important-facial-recognition-company-youve-never-heard-of-12381d530510>.

<sup>4</sup> In 2020, Customs and Border Protection officials testified to Congress that they had used facial recognition on over 43.7 million travelers crossing US borders. In 2021, the agency reported that it processed an additional 23 million travelers in fiscal year 2020. See Kyle Wiggers, “U.S. Homeland Security Has Used Facial Recognition on over 43.7 Million People,” *VentureBeat*, February 6, 2020, <https://venturebeat.com/2020/02/06/u-s-homeland-security-has-used-facial-recognition-on-over-43-7-million-people/>; US Customs and Border Protection. “CBP Trade and Travel Report Fiscal Year 2020,” February 2021, <https://www.cbp.gov/sites/default/files/assets/documents/2021-Feb/CBP-FY2020-Trade-and-Travel-Report.pdf>.

<sup>5</sup> Ayang Macdonald, “Idemia’s Facial Recognition Accuracy Confirmed as Solution Maintains Top Ranking in NIST FRVT,” *Biometric Update*, August 3, 2021, <https://www.biometricupdate.com/202108/idemias-facial-recognition-accuracy-confirmed-as-solution-maintains-top-ranking-in-nist-frvt>.

<sup>6</sup> “CyberLink FaceMe® AI Facial Recognition Earns 99.7% Accuracy Rate, High Ranking on Industry-Leading NIST Leaderboard,” *BusinessWire*, April 6, 2020, <https://www.businesswire.com/news/home/20200406005229/en/CyberLink-FaceMe%C2%AE-AI-Facial-Recognition-Earns-99.7-Accuracy-Rate-High-Ranking-on-Industry-Leading-NIST-Leaderboard>.

<sup>7</sup> Paravision, “Paravision Extends Face Recognition Accuracy, NIST FRVT Report,” March 2021, <https://www.paravision.ai/news/paravision-accuracy-nist-frvt-feb-2021/>.

<sup>8</sup> Innovatrics, “Innovatrics: Top Performer in Every NIST FRVT Category,” November 2020. <https://www.innovatrics.com/awards/frvt/>.

<sup>9</sup> The full history of facial recognition technology is outside the scope of this article, which narrows its focus to Facial Recognition Vendor Test. While the FRVT was a major turning point in the development and normalization of facial recognition technology, by the time the test was established, research on facial recognition had already been underway since the 1960s, with more concerted activity throughout the 1990s. See Kelly Gates’ *Our Biometric Future* (New York: New York University Press, 2011) for an extended history.

<sup>10</sup> P. Jonathon Phillips, P. J. Rauss, and Sandor Z. Der, *FERET (Face Recognition Technology) Recognition Algorithm Development and Test Results* (Adelphi, MD: Army Research Laboratory, October 1996), 2.

<sup>11</sup> Inioluwa Deborah Raji, and Genevieve Fried, “About Face: A Survey of Facial Recognition Evaluation,” *ArXiv:2102.00813 [Cs]*, February 1, 2021, <http://arxiv.org/abs/2102.00813>.

<sup>12</sup> Patrick J. Rauss et al., *FERET (Face Recognition Technology) Program* (Boston: Proc. SPIE 2935, Surveillance and Assessment Technologies for Law Enforcement 1997), 2–3.

<sup>13</sup> *Ibid.*, 5.

<sup>14</sup> P. Jonathon Phillips et al., *Face Recognition Vendor Test 2002 Evaluation Report* (Arlington, VA: NIST 2003), 10.

<sup>15</sup> Duane M. Blackburn, Mike Bone, and P. Jonathon Phillips, *Facial Recognition Vendor Test 2000 Evaluation Report* (Dahlgren, VA: DoD Counterdrug Technology Development Program Office 2001), ii.

<sup>16</sup> *FRVT 1:1 Verification* (NIST 2022) <https://pages.nist.gov/frvt/html/frvt11.html>.

<sup>17</sup> Given the increased enthusiasm for biometric information after 9/11, an alternative pathway might have been to place the FRVT under the purview of the newly-formed Department of Homeland Security rather than NIST — a choice that would have had far different implications. See Toft Djanegara, Nina. “How 9/11 Birthed America’s Biometrics Security Empire.”

*FastCompany*, September 10, 2021. <https://www.fastcompany.com/90674661/how-9-11-sparked-the-rise-of-americas-biometrics-security-empire>

<sup>18</sup> Phillips et al., *Face Recognition Vendor Test 2002 Evaluation Report*, 4.

<sup>19</sup> The US government makes ample use of technology in border enforcement, from policing to basic bureaucratic administration; see Iván Chaar-López, “Alien Data: Immigration and Regimes of Connectivity in the United States,” *Critical Ethnic Studies* 6, no. 2 (2020); Camilla Fojas, *Border Optics: Surveillance Cultures on the US-Mexico Frontier* (New York: New York University Press, 2021); Ana Muñoz, *Borderland Circuitry: Immigration Surveillance in the United States and Beyond* (Oakland: University of California Press, 2022). The FRVT, however, represents something different: instead of technology as a tool for border enforcement or an experimental site for testing new technologies, the US border regime was regarded as a *source* of data.

<sup>20</sup> Kate Crawford, *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence* (New Haven: Yale University Press, 2021), chap. 3.

<sup>21</sup> Patrick J. Rauss et al. “FERET (Face Recognition Technology) Program,” in *Surveillance and Assessment Technologies for Law Enforcement*, eds. A. Trent DePersia, Suzan Yeager, and Steve M. Ortiz (Boston: 2-11, 1997), 4.

<sup>22</sup> See Lisa Gitelman’s “*Raw Data*” *Is an Oxymoron* (Cambridge, MA: MIT Press, 2013) for discussion of both the history and foolhardiness of this notion.

<sup>23</sup> In this essay, I conflate visa and passport photographs because the technical standards are identical for both categories of images. The terms visa photo, passport photo, and identity photo are all used interchangeably.

<sup>24</sup> This name is a pseudonym.

<sup>25</sup> “8 FAM 402.1 PASSPORT PHOTOGRAPHS,” Foreign Affairs Manual, Office of Origin: CA/PPT/S/A: U.S. Department of State, September 7, 2021. <https://fam.state.gov/fam/08fam/08fam040201.html>.

<sup>26</sup> My point here is not to suggest that edited passport photos might compromise the accuracy of facial recognition algorithms or offer a means of camouflage for those whose photos have been manipulated—though the effect of makeup and beautification techniques is an active area of research for computer scientists, as well as artists and anti-surveillance activists (see Antonio Cerella, “Dressing for a Machine-Readable World: An Interview with Adam Harvey,” *Security Dialogue Blog*, SAGE Publications, July 25, 2019, <https://blogs.prio.org/SecurityDialogue/2019/07/dressing-for-a-machine-readable-world-an-interview-with-adam-harvey/>; Christian Rathgeb, Pawel Drozdowski, and Christoph Busch, “Makeup Presentation Attacks: Review and Detection Performance Benchmark,” *IEEE Access* 8 (2020): 224958–73; Esther Shein, “Using Makeup to Block Surveillance,” *Communications of the ACM* 65, no. 7 (June 21, 2022): 21–3). Instead, I’m interested in the tension between rigid technical systems and the everyday practices of people working in and around those systems. My analysis is also inspired by historians who have examined the racial histories of identity documents and immigration control in the United States and Canada during the late 19th and early 20th centuries. In their work on the use of photography to document the identities of Chinese migrants, scholars like Anna Pegler-Gordon, Lily Cho, and Kitty Calavita have illustrated how Chinese migrants played a role in shaping their own self-representation, working within the confines of immigration law and the strict genre of identity photography. For instance, Calavita argues that some Chinese migrants strategically used photographs to make a visual argument that they deserved to remain in the United States by making themselves appear to be upper-class merchants, bedecked in silken finery. Similarly, Lily Cho asserts that by “shaping how they would be identified by the state,” Chinese migrants were contesting their racial

marginalization (Lily Cho, “Anticipating Citizenship: Chinese Head Tax Photographs,” in *Feeling Photography*, eds. Elspeth H. Brown and Thy Phu [Durham: Duke University Press, 2014], 170). While the balance of power is tilted overwhelmingly in favor of the state, practices like photo manipulation are a reminder of the state’s inability to fully observe and control its subjects, as well as the agency of people who are immobilized and exploited by repressive border regimes.

<sup>27</sup> See also the “Perpetual Lineup” report by Georgetown Law, which details how the FBI drew upon existing databases to create a massive facial recognition database that contains profiles for over 117 million Americans, more than half of the country’s adult population (Claire Garvie and Alvaro M. Bedoya, “The Perpetual Line-Up” [Center on Privacy & Technology at Georgetown Law, 2016], <https://www.perpetuallineup.org/>.)

<sup>28</sup> Patrick Grother, George W. Quinn, and P. Jonathon Phillips, *Report on the Evaluation of 2D Still-Image Face Recognition Algorithms* (Gaithersburg, MD: NIST, 2011), 13.

<sup>29</sup> *Ibid.*, 13.

<sup>30</sup> When considering the history of surveillance, we can see how the state’s gaze is often aimed at certain groups and populations who are often at the peripheries of society. For instance, through her historical survey of the surveillance of Blackness, sociologist Simone Browne has argued that surveillance locates certain people as more threatening, more deviant, and more necessary to watch. She offers the term “racializing surveillance” to describe how surveillance is structured by and reifies preexisting racial categories. When viewed in this light, it is unsurprising that border detainees and non-citizens have had their images captured without consent, placed in US government databases, and used as experimental tools. Simone Browne, *Dark Matters: On the Surveillance of Blackness* (Durham: Duke University Press, 2015), 8. See also Khiara Bridges, *The Poverty of Privacy Rights* (Stanford: Stanford Law Books, an imprint of Stanford University Press, 2017); Rachel E. Dubrofsky and Shoshana Magnet, *Feminist Surveillance Studies* (Durham: Duke University Press, 2015); and Virginia Eubanks, *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor* (New York: St. Martin’s Press, 2018). An alternative historical chronology might place identity photos in the same genealogy as colonial anthropological portraits and “racial type” photography, which share similar aesthetic conventions, as well as a sense of ownership over the subject. See Deborah Poole, *Vision, Race, and Modernity: A Visual Economy of the Andean Image World* (Princeton: Princeton University Press, 1997) and Anne Maxwell, “Modern Anthropology and the Problem of the Racial Type: The Photographs of Franz Boas,” *Visual Communication* 12, no. 1 (February 2013): 123–42.

<sup>31</sup> Patrick Grother and Mei Ngan, *Face Recognition Vendor Test (FRVT)* (Gaithersburg, MD: NIST, 2014), 12.

<sup>32</sup> The NIST evaluators were pleasantly surprised by how the algorithms performed on the low-quality border detainee dataset: “Exceptionally, however, the most accurate algorithm fails in only 11.3% of searches.” This was seen as a promising result, since the findings suggested that facial recognition would now be feasible for other situations where only grainy images are available, such as surveillance footage or bank ATM cameras. Grother and Ngan, *Face Recognition Vendor Test*, 2–3.

<sup>33</sup> Patrick Grother, Mei Ngan, and Kayee Hanaoka, *Face Recognition Vendor Test Part 3: Demographic Effects* (Gaithersburg, MD: NIST, 2019), 5.

<sup>34</sup> Joy Buolamwini and Timnit Gebru, “Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification,” in *Proceedings of the Conference on Fairness, Accountability and Transparency* (2018), 77–91; see also Inioluwa Deborah Raji and Joy Buolamwini. “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*, Honolulu HI (New York: ACM, 2019), 429–35.

<sup>35</sup> Buolamwini and Gebru, “Gender Shades,” 5.

<sup>36</sup> Grother, Ngan, and Hanaoka, *Face Recognition Vendor Test Part 3: Demographic Effects*, 1.

<sup>37</sup> For public statements by IBM, Microsoft, and Amazon, see Michele Merler et al., “Diversity in Faces,” *ArXiv:1901.10436 [Cs]*, April 8, 2019; John Roach, “Microsoft Improves Facial Recognition to Perform Well across All Skin Tones,” *The AI Blog*, June 26, 2018, <https://blogs.microsoft.com/ai/gender-skin-tone-facial-recognition-improvement/>; Matt Wood, “Thoughts on Recent Research Paper and Associated Article on Amazon Rekognition,” *Amazon Web Services*, January 26, 2019, <https://aws.amazon.com/blogs/machine-learning/thoughts-on-recent-research-paper-and-associated-article-on-amazon-rekognition/>.

<sup>38</sup> Grother, Ngan, and Hanaoka, *Face Recognition Vendor Test Part 3: Demographic Effects*, Annex 3.

<sup>39</sup> D. Broeders, “The New Digital Borders of Europe: EU Databases and the Surveillance of Irregular Migrants,” *International Sociology* 22, no. 1 (January 1, 2007): 81.

<sup>40</sup> As computer scientist Inioluwa Deborah Raji explains, “The images [used in the FRVT] are well lit and in a constrained environment with a consistent background. If the model does not do well on this test, that’s significant, but doing well on NIST does not at all mean the model can handle real-world conditions,” quoted in Dave Gershgorin, “NEC Is the Most Important Facial Recognition Company You’ve Never Heard Of.” See also Erik Learned-Miller et al., “Facial Recognition Technologies in the Wild: A Call for a Federal Office” (Cambridge, MA: Algorithmic Justice League, May, 2020), <https://www.ajl.org/federal-office-call>; Daniel Ho, Emily Black, Maneesh Agrawala, and Fei Fei Li. “Evaluating Facial Recognition Technology: A Protocol for Performance Assessment in New Domains” (Stanford: Institute for Human-Centered Artificial Intelligence, Stanford University, November 2020).

<sup>41</sup> Kate Crawford makes a similar point in her discussion of the ethical implications of NIST’s use of mugshot photos. Says Crawford, “A person standing in front of a camera in an orange jumpsuit, then, is dehumanized as just more data. The history of these images, how they were acquired, and their institutional, personal, and political contexts are not considered relevant.” Kate Crawford, *Atlas of AI*, 98.

<sup>42</sup> Puck Lo, “The Databases That Feed the System,” Community Justice Exchange, 2022, <https://abolishdatacrim.org/en/bestiary>.

<sup>43</sup> The most notorious example of non-consensual repurposing of images for facial recognition is probably that of Clearview AI, a company that has seized billions of people’s photographs, using a process called “web scraping” that trawls the internet for photos and amasses them into a giant database. As of April 2022, Clearview’s database contains approximately 20 billion photographs, with aims of reaching 100 billion photographs by 2023. The company’s ultimate goal is to make “almost everyone in the world ... identifiable,” according to a Clearview presentation for investors obtained by *The Washington Post*. The controversial company has been taken to court by tech giants like Facebook and Google, who claim that Clearview violated their terms of service by scraping images from their websites; meanwhile, national governments in Australia, Canada, France, and Italy have banned or sanctioned Clearview for violating data privacy laws. See Matt O’Brien and Tali Arbel, “Face Scanner Clearview AI Aims to Branch out beyond Police,” *ABC News*, April 1, 2022, <https://abcnews.go.com/Business/wireStory/face-scanner-clearview-ai-aims-branch-police-83819522>; Drew Harwell, “Facial Recognition Firm Clearview AI Tells Investors It’s Seeking Massive Expansion beyond Law Enforcement,” *The Washington Post*, February 16, 2022, <https://www.washingtonpost.com/technology/2022/02/16/clearview-expansion-facial-recognition/>.

<sup>44</sup> Adam Harvey and Jules LaPlace, “Transnational Flows of Face Recognition Image Training Data,” July 7, 2019a, [https://exposing.ai/research/munich\\_security\\_conference](https://exposing.ai/research/munich_security_conference).

<sup>45</sup> Gary B. Huang et al., “Labeled Faces in the Wild.”

<sup>46</sup> Kashmir Hill and Aaron Krolik, “How Photos of Your Kids Are Powering Surveillance Technology,” *The New York Times*, October 11, 2019, <https://www.nytimes.com/interactive/2019/10/11/technology/flickr-facial-recognition.html>.

<sup>47</sup> Ibid.

<sup>48</sup> See, for example, Kate Crawford and Trevor Paglen. “Excavating AI: The Politics of Images in Machine Learning Training Sets,” September 19, 2019, <https://www.excavating.ai>; Amandalynne Paullada et al., “Data and Its (Dis)Contents: A Survey of Dataset Development and Use in Machine Learning Research,” *ArXiv:2012.05345 [Cs]*, December 9, 2020; Vinay Uday Prabhu and Abeba Birhane, “Large Image Datasets: A Pyrrhic Win for Computer Vision?” *ArXiv:2006.16923 [Cs, Stat]*, July 23, 2020; Morgan Klaus Scheuerman et al., “How We’ve Taught Algorithms to See Identity: Constructing Race and Gender in Image Databases for Facial Analysis,” *Proceedings of the ACM on Human-Computer Interaction* 4, no. CSCW1 (May 28, 2020): 058:1-058:35; Emily Denton et al. “On the Genealogy of Machine Learning Datasets: A Critical History of ImageNet,” *Big Data & Society* 8, no. 2 (2021).

<sup>49</sup> Martina Bacigalupo and the Walther Collection, eds., *Gulu Real Art Studio* (Göttingen: Steidl, 2013), 17.

<sup>50</sup> Ibid., 17.

<sup>51</sup> See Tina Campt’s *Listening to Images* (Durham: Duke University Press, 2017) for further analysis of Bacigalupo’s work, including discussion of how passport photos are imbued with certain aspirations about mobility—of both the social and geographical sorts. My engagement with the NIST databases throughout this article aligns with Campt’s methodology of “listening” to photos. For Campt, this entails looking past the instrumental purpose or format of images and attending instead to the “unspoken relations that structure them.” Campt’s methodology reminds us that even when carefully posed subjects appear to be “mute supplicants of governmentality,” there still may exist hidden stories of resistance and refusal. Campt, *Listening to Images*, 8–9.

<sup>52</sup> Strassler, Karen, *Refracted Visions: Popular Photography and National Modernity in Java* (Durham: Duke University Press, 2010), 21.

<sup>53</sup> This framing is inspired by anthropologist Jennifer Bajorek’s work on the circulation of photographs and their “unruliness and mutability” as they move across contexts and spaces. My understanding of identity photographs also bears a resemblance to Star and Griesemer’s concept of boundary objects, namely objects that sit in the interstices between communities of practice, enabling linkages as divergent actors come together to achieve common goals. Jennifer Bajorek, “Of Jumbled Valises and Civil Society: Photography and Political Imagination in Senegal,” *History and Anthropology* 21, no. 4 (December 1, 2010): 434; Susan Leigh Star and James R. Griesemer, “Institutional Ecology, ‘Translations’ and Boundary Objects,” *Social Studies of Science* 19, no. 3 (1989): 387–420.

<sup>54</sup> Andrew Myers, “A New Approach to the Data-Deletion Conundrum,” Stanford HAI, September 24, 2021, <https://hai.stanford.edu/news/new-approach-data-deletion-conundrum>.

<sup>55</sup> There is an important distinction between training and test data. Training data are used to teach a model, while test data are used to evaluate how well a model has learned its task. In other words, training data is associated with open-source benchmark datasets like ImageNet, while test data is associated with closed evaluations like the Facial Recognition Vendor Test. Although the algorithms submitted to the FRVT evaluation are not trained on the immigrant datasets, developers are likely to select for models that perform well on these benchmark datasets. As

Bernard Koch et al. have noted, “Because advancement on established benchmarks is viewed as an indicator of progress, researchers are encouraged to make design choices that maximize performance on benchmarks, as this increases the legitimacy of their work.” Bernard Koch et al., “Reduced, Reused and Recycled: The Life of a Dataset in Machine Learning Research,” *Proceedings of NeurIPS 2021* (2021), 1.

<sup>56</sup> The thought experiment I offer here is an oversimplification of the problem of biased facial recognition. For a more in-depth technical explanation of this phenomenon, see Joy Buolamwini and Timnit Gebru, “Gender Shades,” 77–91; K. S. Krishnapriya et al., “Issues Related to Face Recognition Accuracy Varying Based on Race and Skin Tone,” *IEEE Transactions on Technology and Society* 1, no. 1 (March 2020): 8–20. My point here is to illustrate the effect of a single data point on a computer model, not to offer recommendations about how to resolve demographic disparities in facial recognition.

<sup>57</sup> Tiffany C. Li, “Algorithmic Destruction,” *SMU Law Review* (forthcoming), 11–12, <https://ssrn.com/abstract=4066845> or <http://dx.doi.org/10.2139/ssrn.4066845>.

<sup>58</sup> Zou, quoted in Andrew Myers, “A New Approach to the Data-Deletion Conundrum.”

<sup>59</sup> Rohit Chopra, Statement of Commissioner Rohit Chopra in the Matter of Everalbum and Paravision, Commission File No. 1923172 § (2021).

<sup>60</sup> *Ibid.*

<sup>61</sup> In 2019, Microsoft made the decision to delete MS-Celeb-1M, a popular dataset of one million celebrity photos that has been widely used to train facial recognition models. This move came shortly after researchers Adam Harvey and Jules LaPlace called attention to the fact that the MS-Celeb-1M dataset employed a dubiously broad definition of what kind of person counted as a celebrity; faces in the dataset included those of anti-privacy activists and journalists. As Harvey and LaPlace note, “Many of the names in the MS Celeb face recognition dataset are merely people who must maintain an online presence for their professional lives.” Though MS-Celeb-1M has now been removed from Microsoft’s website, the facial recognition models built from this dataset continue to be in circulation; research teams from Facebook and IBM, as well as Chinese tech giants like Baidu and Sensetime are among those developers who made use of MS-Celeb-1M before its deletion. See Adam Harvey and Jules LaPlace “Exposing AI: MS-CELEB-1M,” 2019b, <https://exposing.ai/msceleb/>; Karen Hao, “Deleting Unethical Data Sets Isn’t Good Enough,” *MIT Technology Review*, August 13, 2021, <https://www.technologyreview.com/2021/08/13/1031836/ai-ethics-responsible-data-stewardship/>.

<sup>62</sup> Similarly, in November 2021, Meta, the parent company of Facebook, announced that it planned to delete more than one billion faceprints—digital templates of people’s unique facial features that the company derived from photos that users uploaded to Facebook. Meta framed this as an important step in its moral reckoning over privacy, personal data, and the societal applications of facial recognition technology. However, it is important to note that Meta will continue to retain its proprietary algorithm called DeepFace, which was built on the back of users’ facial biometric data. See Jerome Pesenti, “An Update on Our Use of Face Recognition.,” *Meta* (blog), November 2, 2021, <https://about.fb.com/news/2021/11/update-on-use-of-face-recognition/>.

<sup>63</sup> The process of inferring training data using General Adversarial Networks (GAN) is experimental and complex; this is a nascent field of research, though early results suggest that model inversion and data reconstruction are possible. Researchers in this field are generally concerned by the privacy implications of reconstructed data, as well as by the need to prevent adversarial attacks. See Mai Guangcan et al., “On the Reconstruction of Face Images from Deep Face Templates,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, no. 5 (May 1, 2019): 1188–1202; Yuheng Zhang et al., “The Secret Revealer: Generative Model-

Inversion Attacks Against Deep Neural Networks,” in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 250–8 (Seattle: IEEE, 2020). For discussion of the legal implications of data reconstruction, see Michael Veale, Reuben Binns, and Lilian Edwards, “Algorithms That Remember: Model Inversion Attacks and Data Protection Law,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376, no. 2133 (November 28, 2018).

<sup>64</sup> Jonathan Brophy, “Exit Through the Training Data: A Look into Instance-Attribution Explanations and Efficient Data Deletion in Machine Learning,” (Eugene, OR: University of Oregon, February 14, 2020), 9. See also Zachary Izzo et al., “Approximate Data Deletion from Machine Learning Models,” in *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS)* (San Diego, 2021).

## REFERENCES

- “8 FAM 402.1 PASSPORT PHOTOGRAPHS.” Foreign Affairs Manual. Office of Origin: CA/PPT/S/A: U.S. Department of State, September 7, 2021. <https://fam.state.gov/fam/08fam/08fam040201.html>.
- Bacigalupo, Martina, and Walther Collection, eds. *Gulu Real Art Studio*. 1st ed. Göttingen: Steidl, 2013.
- Bajorek, Jennifer. “Of Jumbled Valises and Civil Society: Photography and Political Imagination in Senegal.” *History and Anthropology* 21, no. 4 (December 1, 2010): 431–52.
- Blackburn, Duane M., Mike Bone, and P. Jonathon Phillips. “Facial Recognition Vendor Test 2000 Evaluation Report.” Dahlgren, VA: DoD Counterdrug Technology Development Program Office, February 16, 2001.
- Bridges, Khiara M. *The Poverty of Privacy Rights*. Stanford: Stanford Law Books, an imprint of Stanford University Press, 2017.
- Broeders, D. “The New Digital Borders of Europe: EU Databases and the Surveillance of Irregular Migrants.” *International Sociology* 22, no. 1 (January 1, 2007): 71–92.
- Brophy, Jonathan. “Exit Through the Training Data: A Look into Instance-Attribution Explanations and Efficient Data Deletion in Machine Learning.” University of Oregon, February 14, 2020. <https://www.cs.uoregon.edu/Reports/AREA-202009-Brophy.pdf>.
- Browne, Simone. *Dark Matters: On the Surveillance of Blackness*. Durham: Duke University Press, 2015.
- Buolamwini, Joy, and Timnit Gebru. “Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification.” In *Proceedings of the Conference on Fairness, Accountability and Transparency*, 77–91. New York: ACM, 2018.
- BusinessWire. “CyberLink FaceMe® AI Facial Recognition Earns 99.7% Accuracy Rate, High Ranking on Industry-Leading NIST Leaderboard.” BusinessWire, April 6, 2020. <https://www.businesswire.com/news/home/20200406005229/en/CyberLink-FaceMe%C2%AE-AI-Facial-Recognition-Earns-99.7-Accuracy-Rate-High-Ranking-on-Industry-Leading-NIST-Leaderboard>.
- Calavita, Kitty. “The Paradoxes of Race, Class, Identity, and ‘Passing’: Enforcing the Chinese Exclusion Acts, 1882–1910.” *Law & Social Inquiry* 25, no. 1 (2000): 1–40.
- Campt, Tina. *Listening to Images*. Durham: Duke University Press, 2017.
- Cerella, Antonio. “Dressing for a Machine-Readable World: An Interview with Adam Harvey.” *Security Dialogue Blog*. SAGE Publications, July 25, 2019.

- <https://blogs.prio.org/SecurityDialogue/2019/07/dressing-for-a-machine-readable-world-an-interview-with-adam-harvey/>.
- Chaar-López, Iván. "Alien Data: Immigration and Regimes of Connectivity in the United States." *Critical Ethnic Studies* 6, no. 2 (2020).
- Cho, Lily. "Anticipating Citizenship: Chinese Head Tax Photographs." In *Feeling Photography*, edited by Elspeth H. Brown and Thy Phu, 158–80. Durham: Duke University Press, 2014.
- Chopra, Rohit. Statement of Commissioner Rohit Chopra in the Matter of Everalbum and Paravision, Commission File No. 1923172 § (2021).
- Crawford, Kate. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven: Yale University Press, 2021.
- Crawford, Kate, and Trevor Paglen. "Excavating AI: The Politics of Images in Machine Learning Training Sets." September 19, 2019. <https://www.excavating.ai>.
- Denton, Emily, Alex Hanna, Razvan Amironesei, Andrew Smart, and Hilary Nicole. "On the Genealogy of Machine Learning Datasets: A Critical History of ImageNet." *Big Data & Society* 8, no. 2, 2021.
- Dubrofsky, Rachel E., and Shoshana Magnet, eds. *Feminist Surveillance Studies*. Durham: Duke University Press, 2015.
- Eubanks, Virginia. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. 1st ed. New York: St. Martin's Press, 2018.
- Fojas, Camilla. *Border Optics: Surveillance Cultures on the US-Mexico Frontier*. New York: New York University Press, 2021.
- Garvie, Claire, and Alvaro M. Bedoya. "The Perpetual Line-Up." Center on Privacy & Technology at Georgetown Law, 2016. <https://www.perpetuallineup.org/>.
- Gates, Kelly. *Our Biometric Future: Facial Recognition Technology and the Culture of Surveillance*. New York: New York University Press, 2011.
- Gershgorn, Dave. "NEC Is the Most Important Facial Recognition Company You've Never Heard Of." *OneZero*, February 21, 2020. <https://onezero.medium.com/nec-is-the-most-important-facial-recognition-company-youve-never-heard-of-12381d530510>.
- Gitelman, Lisa, ed. "*Raw Data*" Is an Oxymoron. Infrastructures Series. Cambridge, MA: The MIT Press, 2013.
- Grother, Patrick, and Mei Ngan. "Face Recognition Vendor Test (FRVT)." Gaithersburg, MD: National Institute of Standards and Technology, 2014.
- Grother, Patrick, Mei Ngan, and Kayee Hanaoka. "Face Recognition Vendor Test Part 3: Demographic Effects." Gaithersburg, MD: National Institute of Standards and Technology, December, 2019.
- Grother, Patrick, George W. Quinn, and P. Jonathon Phillips. "Report on the Evaluation of 2D Still-Image Face Recognition Algorithms." Gaithersburg, MD: National Institute of Standards and Technology, 2011.
- Guangan, Mai, Kai Cao, Pong C. Yuen, and Anil K. Jain. "On the Reconstruction of Face Images from Deep Face Templates." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, no. 5 (May 1, 2019): 1188–1202.
- Hao, Karen. "Deleting Unethical Data Sets Isn't Good Enough." *MIT Technology Review*, August 13, 2021. <https://www.technologyreview.com/2021/08/13/1031836/ai-ethics-responsible-data-stewardship/>.
- Harvey, Adam, and Jules LaPlace. "Transnational Flows of Face Recognition Image Training Data." July 7, 2019a. [https://exposing.ai/research/munich\\_security\\_conference](https://exposing.ai/research/munich_security_conference).
- . "Exposing AI: MS-CELEB-1M." 2019b. <https://exposing.ai/msceleb/>.

- Hill, Kashmir, and Aaron Krolik. “How Photos of Your Kids Are Powering Surveillance Technology.” *The New York Times*, October 11, 2019. <https://www.nytimes.com/interactive/2019/10/11/technology/flickr-facial-recognition.html>.
- Ho, Daniel, Emily Black, Maneesh Agrawala, and Fei Fei Li. “Evaluating Facial Recognition Technology: A Protocol for Performance Assessment in New Domains.” Institute for Human-Centered Artificial Intelligence, Stanford University, November 2020.
- Huang, Gary B., Marwan Mattar, Tamara Berg, and Eric Learned-Miller. “Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments,” 2008. <http://vis-www.cs.umass.edu/lfw/lfw.pdf>.
- Innovatrics. “Innovatrics: Top Performer in Every NIST FRVT Category,” November 2020. <https://www.innovatrics.com/awards/frvt/>.
- Izzo, Zachary, Mary Anne Smart, Kamalika Chaudhuri, and Zou James. “Approximate Data Deletion from Machine Learning Models.” In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS)*, San Diego, CA. n.p.: MLResearchPress, 2021.
- Koch, Bernard, Emily Denton, Alex Hanna, and Jacob G Foster. “Reduced, Reused and Recycled: The Life of a Dataset in Machine Learning Research.” In *Proceedings of NeurIPS 2021*, 2021.
- Krishnapriya, K. S., Vitor Albiero, Kushal Vangara, Michael C. King, and Kevin W. Bowyer. “Issues Related to Face Recognition Accuracy Varying Based on Race and Skin Tone.” *IEEE Transactions on Technology and Society* 1, no. 1 (March 2020): 8–20.
- Learned-Miller, Erik, Vicente Ordóñez, Jaimie Morgenstern, and Joy Buolamwini. “Facial Recognition Technologies in the Wild: A Call for a Federal Office.” Cambridge, MA: Algorithmic Justice League, May, 2020. <https://www.ajl.org/federal-office-call>
- Li, Tiffany C. “Algorithmic Destruction,” *SMU Law Review* (forthcoming), 11–12. <https://ssrn.com/abstract=4066845> or <http://dx.doi.org/10.2139/ssrn.4066845>.
- Lo, Puck. “The Databases That Feed the System.” Community Justice Exchange, 2022. <https://abolishdatacrim.org/en/bestiary>.
- Macdonald, Ayang. “Idemia’s Facial Recognition Accuracy Confirmed as Solution Maintains Top Ranking in NIST FRVT.” Biometrics Research Group, August 3, 2021. <https://www.biometricupdate.com/202108/idemias-facial-recognition-accuracy-confirmed-as-solution-maintains-top-ranking-in-nist-frvt>.
- Mai, Guangcan, Kai Cao, Pong C. Yuen, and Anil K. Jain. “On the Reconstruction of Face Images from Deep Face Templates.” *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, no. 5 (May 1, 2019): 1188–202.
- Maxwell, Anne. “Modern Anthropology and the Problem of the Racial Type: The Photographs of Franz Boas.” *Visual Communication* 12, no. 1 (February 2013): 123–42.
- Merler, Michele, Nalini Ratha, Rogerio S. Feris, and John R. Smith. “Diversity in Faces.” *ArXiv:1901.10436 [Cs]*, April 8, 2019.
- Muñiz, Ana. *Borderland Circuitry: Immigration Surveillance in the United States and Beyond*. Oakland: University of California Press, 2022.
- Myers, Andrew. “A New Approach to the Data-Deletion Conundrum.” Stanford HAI, September 24, 2021. <https://hai.stanford.edu/news/new-approach-data-deletion-conundrum>.
- NEC. “NEC Face Recognition Technology Ranks First in NIST Accuracy Testing.” NEC News Room, October 3, 2019. [https://www.nec.com/en/press/201910/global\\_20191003\\_01.html](https://www.nec.com/en/press/201910/global_20191003_01.html).
- National Institute of Standards and Technology (NIST). “FRVT 1:1 Verification,” March 18, 2022. <https://pages.nist.gov/frvt/html/frvt11.html>.

- Paravision. "Paravision Extends Face Recognition Accuracy, NIST FRVT Report," March 2021. <https://www.paravision.ai/news/paravision-accuracy-nist-frvt-feb-2021/>.
- Paullada, Amandalynne, Inioluwa Deborah Raji, Emily M. Bender, Emily Denton, and Alex Hanna. "Data and Its (Dis)Contents: A Survey of Dataset Development and Use in Machine Learning Research." *ArXiv:2012.05345 [Cs]*, December 9, 2020.
- Pegler-Gordon, Anna. *In Sight of America: Photography and the Development of U.S. Immigration Policy*. Berkeley: University of California Press, 2009.
- Pesenti, Jerome. "An Update on Our Use of Face Recognition." *Meta* (blog). November 2, 2021. <https://about.fb.com/news/2021/11/update-on-use-of-face-recognition/>.
- Phillips, P. Jonathon, Patrick Grother, Ross J. Micheals, Duane M. Blackburn, Elham Tabassi, and Mike Bone. "Face Recognition Vendor Test 2002 Evaluation Report." Arlington, VA: National Institute of Standards and Technology, March 2003.
- Phillips, P. Jonathon, P. J. Rauss, and Sandor Z. Der. "FERET (Face Recognition Technology) Recognition Algorithm Development and Test Results." Adelphi, MD: Army Research Laboratory, October 1996.
- Poole, Deborah. *Vision, Race, and Modernity: A Visual Economy of the Andean Image World*. Princeton Studies in Culture/Power/History. Princeton: Princeton University Press, 1997.
- Prabhu, Vinay Uday, and Abeba Birhane. "Large Image Datasets: A Pyrrhic Win for Computer Vision?" *ArXiv:2006.16923 [Cs, Stat]*, July 23, 2020.
- Raji, Inioluwa Deborah, and Joy Buolamwini. "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*, 429–35. Honolulu: ACM, 2019.
- Raji, Inioluwa Deborah, and Genevieve Fried. "About Face: A Survey of Facial Recognition Evaluation." *ArXiv:2102.00813 [Cs]*, February 1, 2021.
- Rathgeb, Christian, Pawel Drozdowski, and Christoph Busch. "Makeup Presentation Attacks: Review and Detection Performance Benchmark." *IEEE Access* 8 (2020): 224958–73.
- Rauss, Patrick J., Jonathan Phillips, Hyeonjoon Moon, Syed A. Rizvi, Mark K. Hamilton, and A. Trent DePersia. "FERET (Face Recognition Technology) Program." In *Surveillance and Assessment Technologies for Law Enforcement*, edited by A. Trent DePersia, Suzan Yeager, and Steve M. Ortiz, 2–11. Boston: 1-12, 1997.
- Roach, John. "Microsoft Improves Facial Recognition to Perform Well across All Skin Tones." *The AI Blog*, June 26, 2018. <https://blogs.microsoft.com/ai/gender-skin-tone-facial-recognition-improvement/>.
- Scheuerman, Morgan Klaus, Kandra Wade, Caitlin Lustig, and Jed R. Brubaker. "How We've Taught Algorithms to See Identity: Constructing Race and Gender in Image Databases for Facial Analysis." *Proceedings of the ACM on Human-Computer Interaction* 4, no. CSCW1 (May 28, 2020): 058:1–058:35.
- Shein, Esther. "Using Makeup to Block Surveillance." *Communications of the ACM* 65, no. 7 (June 21, 2022): 21–3.
- Star, Susan Leigh, and James R. Griesemer. "Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907–39." *Social Studies of Science* 19, no. 3 (1989): 387–420.
- Strassler, Karen. *Refracted Visions: Popular Photography and National Modernity in Java*. Durham: Duke University Press, 2010.
- Thylstrup, Nanna Bonde. "The Ethics and Politics of Data Sets in the Age of Machine Learning: Deleting Traces and Encountering Remains." *Media, Culture & Society*, April 28, 2022.

- Toft Djanegara, Nina. "How 9/11 Birthed America's Biometrics Security Empire." *FastCompany*, September 10, 2021. <https://www.fastcompany.com/90674661/how-9-11-sparked-the-rise-of-americas-biometrics-security-empire>.
- U.S. Customs and Border Protection. "CBP Trade and Travel Report Fiscal Year 2020." February 2021. <https://www.cbp.gov/sites/default/files/assets/documents/2021-Feb/CBP-FY2020-Trade-and-Travel-Report.pdf>.
- Veale, Michael, Reuben Binns, and Lilian Edwards. "Algorithms That Remember: Model Inversion Attacks and Data Protection Law." *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376, no. 2133 (November 28, 2018): 20180083.
- Wiggers, Kyle. "U.S. Homeland Security Has Used Facial Recognition on over 43.7 Million People." *VentureBeat*, February 6, 2020. <https://venturebeat.com/2020/02/06/u-s-homeland-security-has-used-facial-recognition-on-over-43-7-million-people/>.
- Wood, Matt. "Thoughts on Recent Research Paper and Associated Article on Amazon Rekognition." Amazon Web Services, January 26, 2019. <https://aws.amazon.com/blogs/machine-learning/thoughts-on-recent-research-paper-and-associated-article-on-amazon-rekognition/>.
- Zhang, Yuheng, Ruoxi Jia, Hengzhi Pei, Wenxiao Wang, Bo Li, and Dawn Song. "The Secret Revealer: Generative Model-Inversion Attacks Against Deep Neural Networks." In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 250–58. Seattle: IEEE, 2020.

## AUTHOR BIO

Nina Dewi Toft Djanegara is a Ph.D. candidate in Anthropology at Stanford University and the Associate Director of the Technology & Racial Equity Initiative at Stanford's Center for Comparative Studies in Race & Ethnicity (CCSRE). Her research uses ethnographic and archival methods to explore how computer vision is applied to "solve" political problems. Her dissertation, titled *Taken at Face Value: Facial Recognition Technology and the Cultural Politics of Identity*, investigates the use of facial recognition for US border enforcement.