

Chapter 2: Eigenface

Background: Eigenface in the Development of AFR Technology

Cultural historian Wolfgang Schivelbusch reminds us that, “as technological methods recede in importance, they reappear as an object of study.”¹ It has been almost thirty years since the AFR method known as eigenface was first developed. Since its introduction, more sophisticated methods of AFR have been developed and have become more widely used. Yet the eigenface method remains important to consider when investigating the visuality of AFR. The eigenface method served as a basis for the development of other procedures that went on to expand the possibilities of AFR. Its success validated the use of facial recognition algorithms and spurred on the development of facial recognition technology. For example, the development of the eigenface method made possible the introduction of fisherfaces, a more refined method allowing for a more precise recognition process. The methods of facial recognition more often used today rely on deep neural networks and other feature-based methods. The Viola-Jones algorithm (2001), for instance, uses a method of Haar Cascade to detect objects through superimposition, training an algorithm to differentiate true positives from increasing variations of false negatives. The methods that have since developed do not produce an image or deploy an ab-

1 Wolfgang Schivelbusch, *The Railway Journey: The Industrialization of Time and Space in the 19th Century* (Berkeley, CA: The University of California Press, 1986), xiii.

stract notion of “seeing” through an algorithmic process. Programmers argue that today’s AFR algorithms do not remotely resemble the human capacity of vision.

The eigenface method was introduced after what has been termed the “AI Winter” of the 1980s, when pessimism surrounded the technology development sector and funding for new technology declined. Nevertheless, this decade also saw the first use of the term “biometrics” in the media and in public forums to describe automated recognition systems. This was the result of a growing awareness in the field of automated recognition technologies² that led to a “Coming of Age” of technological development in the 1970s. In the 1990s, the decade in which eigenface was introduced, the first “Biometric Consortium” was held, organized by the research division of the US National Security Agency (NSA).³

The eigenface method was first developed as a fully automated biometric facial recognition system by two MIT scientists, Matthew Turk and Alex Pentland, who developed it in conjunction with Arbitron,⁴ a television ratings company, for the purpose of monitoring ratings. Working within this consumer-marketing context, their goal was, as they have stated, to “develop a computational model of face recognition that is fast, reasonably simple, and accurate in constrained environments such as an office or a household.”⁵ It was designed to be used in TV sets to determine which individuals within a household were watching TV at which times, feeding this information into consumer ratings for specific television programs. Essentially, eigenface was designed to be integrated into a kind of TV that watches you as you watch it. Given this context, the ideal situation for the algorithm to operate in is a real-time situation in which people are sedentary. Ideally, the faces to be recognized would all be positioned conveniently, that is, squarely in front of the TV screen, in a neutral, forward-facing pose. This pose is familiar from identification photographs, so it was not a great leap to imag-

2 Wayman, “Scientific Development of Biometrics,” 266.

3 *Ibid.*, 269.

4 Turk, “Over Twenty Years of Eigenface,” 2.

5 Turk and Pentland, “Eigenfaces for Recognition,” 71.

ine that eigenface might potentially have uses outside of the context of TV ratings, particularly in areas involving the use of identification documents. The multi-faceted application of facial recognition methods can be understood as a bleeding through from the sphere of consumer interests to the socio-political arenas of risk and control.

At the time of its introduction in 1991, the eigenface method was considered one of the first facial recognition methods successfully to perform face detection and recognition in real time. Before eigenface, automated recognition methods had focused on “feature extraction,” that is, the recognition of isolated features (also termed “landmarks”) of the face, such as the eyes, nose and mouth, and the measurement of the distances between these features. The eigenface method departs from this earlier approach by relying on a representational mechanism that takes into account a holistic representation of the face rather than its isolated features. In doing so, the eigenface algorithm had a built-in capacity to detect faces, as well as to locate, track and classify a subject’s face.

Bledsoe: “The Model Method in Facial Recognition”

The introduction of the eigenface method marked a shift in the development of AFR technology not only because its seemingly simple technique was successful but also because it performed recognition differently from the AFR systems that came before it (and from those that would come after). The eigenface method shifted the approach of AFR methods towards a holistic representation of the human face. To understand why eigenface was considered successful, it helps to understand the original problems and challenges to which this technology was developed as a response. The first attempts to codify and automate facial recognition in an operable process were documented in reports authored by the computer scientist Woodrow Wilson Bledsoe, considered one of the founders of artificial intelligence. Two of these reports were only recently made publicly available (in 2014). They had previously been classified, while references to these reports described them as being commis-

sioned by an “unnamed intelligence agency.”⁶ Alongside the recent discovery and publication of these reports, it has also emerged that they were funded by a CIA front organization, the King-Hurley Research Group.⁷

Figure 1: “Examples of photograph pairs used in the study,”
Woodrow W. Bledsoe, 1964.



6 Michael Ballantyne, Robert S. Boyer, and Larry Hines, “Woody Bledsoe: His Life and Legacy,” *AI Magazine* 17, no. 1 (Spring 1996): 7–20, <https://doi.org/10.1609/aimag.v17i1.1207> and Wayman, “Scientific Development of Biometrics,” 264.

7 The two Bledsoe reports were made publicly available thanks to the efforts of researcher Justin Lange, who, in 2013, as a master’s student at the Interactive Telecommunications Program at New York University, was able to successfully retrieve them from the Dolph Briscoe Center for American History at the University of Texas. On the basis of his research into Bledsoe’s reports, Lange concludes that the King Hurley Research Group, whose name is included in the title of one of these reports, was a front organization for the CIA. Lange’s claim is corroborated by Christopher Robbins’s book *The Invisible Air Force: The Story of the CIA’s Secret Airline* (London: Macmillan, 1981). Thank you to artist Kyle McDonald for information about this backstory and for connecting me with Justin Lange.

Thus, AFR technology seems to have originated as a mechanism for surveillance and intelligence accumulation in the context of national security operations, specifically under auspices of the research arm of the CIA.

Figure 2: “Double exposure shows that the two different subjects are surprisingly similar on a point by point basis,” Woodrow W. Bledsoe, 1964.





Figure 3: “Double exposure shows that the two poses of the same subject have very little in common when considered on a point by point basis,” Woodrow W. Bledsoe, 1964.



There are many fascinating details in these original reports of the pioneering and systematic attempt to automate the recognition of the face, but the most relevant to this discussion is Bledsoe's account of the problems and challenges of the task and how he chooses to visualize these. In Bledsoe's report, “The Model Method in Facial Recognition,” Bledsoe notes that “one of the most challenging areas of pattern recognition is the identification of human

photographs by machine.”⁸ He goes on to describe the difficulties of recognition on the basis of facial photographs, given the variations in age, expression, the angle and rotation of the face and in the direction and intensity of light hitting the face. This original account of the challenges faced by automated recognition is still mentioned in scientific papers published on AFR technology today. As part of this report, Bledsoe included a collection of training images, a dataset that he drew upon (figure 1). It is still not known where these original training images originated from, but they are all of white males of various ages.⁹

Bledsoe uses the method of double exposure to visualize some of these challenges confronting automated recognition. He pairs two portraits together that differ in the direction of facial rotation and superimposing one over the other to make clear the differences between a reference image and a capture image. Figure 2 is an example of this. It presents the facial faces of two men of different ages, yet these images exhibit the same direction of lighting and the same head rotation. Once superimposed, the faces merge. According to Bledsoe’s caption, “two different subjects are surprisingly similar on a point-by-point basis.”¹⁰ In a contrasting series of images, Bledsoe presents two images of the same man, now with different lighting direction and head rotation (figure 3). Here, he presents a double-exposure image depicting a jumble of ears, eyes and hair. Bledsoe describes how the two images of the same subject have very little in common. Bledsoe’s exercise in superimposed photographic depiction conveys the central problem of similarity and difference in AFR technology, and thus the potential for false positives and false negatives. Bledsoe’s early visual experiments with the overlaying of facial images foreshadow the eventual solution found in the eigenface method. The method of statistical pattern recognition

8 Woodrow Wilson Bledsoe, “The Model Method in Facial Recognition,” Technical Report PRI 15 (Palo Alto, CA: Panoramic Research, Inc., 1964), 2.

9 Based on the demographics, Lange speculates that these portraits are from a criminal database. Bledsoe himself cites the work of Alphonse Bertillon and Cesare Lombroso on the “criminal man.”

10 Bledsoe “The Model Method,” 7.

used in eigenface essentially utilizes these differences and similarities and encodes them. This early production of a binary composite reveals the kinds of challenges faced in the initial stages of the development of AFR technology. Yet Bledsoe's composites also anticipate the eigenface approach, which utilizes the representation of similarity and difference through superimposition and transforms this into a mechanism of successful recognition.


Representational Mechanisms and the Machinic Observer

The inspiration for the eigenface algorithm was the thought that it might be possible for algorithmic processes to mimic the processes of human recognition. This was expressed through a shift in the algorithmic modeling of the face from a concentration on its isolated features to a holistic representation of the face. Turk and Pentland were working at the intersection between physiology, information theory and the psychology of face recognition. They argued that human face recognition does not occur through the perception of individual facial features and the relationships between them, as previous algorithmic research had suggested. They stated: "individual features and their immediate relationships comprise an insufficient representation to account for the performance of adult human face identification."¹¹ Instead, Turk and Pentland set out to build an algorithmic facial recognition technique that could produce a holistic representation of the face. This shift is important to consider. Algorithms depend on some form of representation and reproduction of the face. The form and design of this representation in the algorithm provides a foundation that allows the AFR technology to learn to recognize a human face. Marcin Miłkowski has defined the representational mechanism in computational learning processes as having three capacities: "it can refer to a target of representation, it can identify information about the target that is relevant for its own interests and goals, and it can evaluate the value of the infor-

11 Turk and Pentland, "Eigenfaces for Recognition," 72.

mation based on environmental feedback.”¹² Here I use the term “representational mechanism” to refer to a part of the AFR process that is responsible for identifying information relating to a face that is relevant to a process of recognition. Eigenface relies on the use of a statistical method of pattern recognition as a representational mechanism. It is this representational mechanism which is depicted in the eigenface image.

*Figure 4: “Sample face on top and its caricature below it.”
Sirovich, Kirby, 1987.*



In order to produce the primary representational mechanism for eigenface, Turk and Pentland drew on the work of two scholars of applied mathematics at Brown University, Lawrence Sirovich and

12 Marcin Milkowski, *Explaining the Computational Mind* (Cambridge, MA: MIT Press, 2013), 156.

Michael Kirby. In 1987, Sirovich and Kirby published a paper on the use of a statistical method of pattern recognition called Principal Component Analysis (PCA), which they applied to facial images in order to produce what they called “eigenpictures.”¹³ In their paper, Sirovich and Kirby speculate about how humans perform the function of recognizing faces and note how adept humans are at this complex task: humans can recognize an almost infinite number of different faces. They propose that humans are able to recognize so many faces because we engage in a process of deduction in relation to facial characteristics; that is, we recognize the characteristics of a face that depart from a kind of characteristic mean. They additionally propose a mathematical translation of this process, applied to multiple facial images, as a possible model for how humans recognize faces.¹⁴ For example, at one point they refer in the text to photographs of two faces (figure 4), with one photograph having had this mathematical translation of the deduction process (referred to as the Fourier method) applied to it, a photo they refer to as a “caricature,” and the other without having had this reduction applied to it. The photographs, they say, appear “virtually the same to us.”¹⁵ The original (on top) is a still from a video, whereas the other, having undergone a mathematical reduction of pattern recognition, is its transformation into a computational image. When looking at these two facial images, I see a significant difference between the “caricature” and the original; the caricature looks as though it has undergone what is referred to in photography as a solarization process, whereby a tonal inversion occurs through the developing process of a photographic image. In the caricature, the mid-tones and shadow areas become darker while any highlighted areas, in contrast, become brighter. It is as if the tonal spectrum of the image has been compressed and the differences between tones are made more extreme. Sirovich and Kirby, however, conclude that, since the images

13 Sirovich and Kirby, “Low-dimensional Procedure,” 521. Although Sirovich and Kirby were the first to apply PCA to facial images, Turk and Pentland were the first to design an automated recognition system utilizing PCA.

14 *Ibid.*, 519.

15 *Ibid.*, 523.

appear the same to them, they provide evidence that “our own visual apparatus does a similar subtraction.”¹⁶ In other words, our own perceptual processes of recognition most likely involve a process of some form of deduction.


In their recorded observations and conclusions it is possible to detect an interesting and altogether separate confrontation that occurs in these initial eigenpictures, and that is a confrontation between human and machine perception. It is exemplified by Sirovich and Kirby’s reference to “our own visual apparatus.” For the mathematical caricature, the video still and our own processes of perception all represent intersections of perceptual relationships between multiple apparatuses – photographic, video, computational and bare human perception – that inform an exercise in recognition. With their speculation about the human ability to perceive faces and the possibility of expressing this capacity in these abstract and mathematically based caricatures, Sirovich and Kirby stray into the realm of what John Johnston has described as “machinic vision,” that is, “not only an environment of interacting machine and human-machine systems but a field of decoded perceptions that, whether or not produced by or issuing from these machines, assume their full intelligibility only in relation to them.”¹⁷ Two key movements that Johnston outlines in relation to machinic vision are a deterritorialization and a reterritorialization of vision.¹⁸ The former occurs when visual perception is freed from the person that is doing the seeing, and the latter occurs when that seeing is “recoded,” that is, recontextualized and expressed in a new form so as to produce new meaning. Sirovich and Kirby’s speculation may be understood along these lines: the caricature is an expression of vision, in a process of recognition, freed from human cognition.

16 Ibid.

17 Johnston, “Machinic Vision,” 27.

18 Ibid., 28. Here Johnston is drawing on Gilles Deleuze and Félix Guattari’s use of the term “machinic.”

Figure 5: AT&T Laboratories, Eigenfaces of faces from the ORL face database.



Sirovich and Kirby pay close attention to their own visual perception of these caricatures. They believe that the conditions of recognition can be revealed through them. In this way, the caricatures also reveal the conditions of human recognition. What is also important to consider is that, for Sirovich and Kirby, the caricature makes visible the human recognition process; in other words, it makes it possible to see how we see. In this way, it reveals an “observer,” in the sense of the term Jonathan Crary explains in his study of the historical construction of vision: “one who sees within a prescribed set of possibilities, one who is embedded in a system of conventions and limitations.”¹⁹ Crary argues that it is only through the observer that vision, in history, is able to “materialise, to become itself visible.”²⁰ These early experiments with eigenpictures express an early assemblage between human cognitive processes of recognition and a machinic translation of that recognition. In the caricature, a kind

¹⁹ Crary, *Techniques of the Observer*, 6.

²⁰ *Ibid.*, 5.

of “machinic observer” is revealed, and we can begin to see how vision can be codified within a specific operation and within the conventions of an operation of recognition. In other words, Sirovich and Kirby’s early eigenpictures and caricatures give a visual form to the conditions of recognition.

Three Aspects of Eigenface

The German prefix *eigen-* means “*inherent, own, individual, peculiar, specific, and characteristic*.”²¹ As this suggests, the facial recognition method is supposed to be an algorithmic ability to distinguish what is characteristic of an individual’s face in order for the algorithm to determine the individual’s identity. The eigenface algorithm is designed to do just this. Eigenface is based on the premise that the most relevant information about an individual face has to do with the ways it is different from another. Eigenface has been successful in demonstrating an ability both to detect faces and to encode the differences between multiple faces. As Turk and Pentland state, “A simple approach to extracting the information [...] is to somehow capture the variation in a collection of face images [...] and use this information to encode and compare individual face images.”²² Indeed, the primary difficulty in developing a successful AFR system is, as Bledsoe had earlier realized, that human faces vary endlessly in appearance. The eigenface method takes this difficulty and transforms it into a recognitive capacity through a tool of differentiation. In this way, variation is utilized in encoding an individual face. Yet, far from distinguishing particular characteristics, the eigenface image depicts a very different process: an erasure of all individual facial particularities (figure 5). All that is specific and particular to a human face seems to dissolve in a blur. This paradox between the method and the image harbors a contradiction in the modalities of recognition between algorithm and human. To further elaborate on

21 “Eigen,” *Wiktionary*, last modified April 7, 2019, <https://en.wiktionary.org/wiki/eigen#German>.

22 Turk and Pentland, “Eigenfaces for Recognition,” 73.

the mode of recognition by eigenface, I will describe briefly here three key technical aspects of the eigenface process that produce this image and that constitute its representational mechanism and operation of recognition. These three aspects are Principal Component Analysis, the eigenvector and the face space.

Principal Component Analysis is a statistical procedure that has primarily been used as a classification tool and as way of producing predictive models based on a statistical method of mean centering. In mathematical terms, PCA treats each facial image as a point or a vector on a grid with a high-dimensional space allowing for high degrees of variation. This high-dimensional coordinate space can be understood as Cartesian space gone digital. Each collected facial image in the training set is translated into a unit of measurement, or a weight, within this virtual space. Averages are calculated from the different weights of facial images. Each average takes into account all the possible variations of each weight. The averaged or mean face is described as “the center of gravity for all the faces combined.”²³ This averaged face delimits the highest degrees of variation, that is, the farthest directions of deviation from the average that exist between the collected facial images. Turk and Pentland explain that “any collection of face images can be approximately reconstructed by storing a small collection of weights for each face.”²⁴ The PCA procedure calculates a mean by averaging the value of each pixel across the face images. PCA is able to extract the principal components, or the primary differences, between multiple faces and encode this variation. Eigenface programmers describe this as revealing the internal structure of the data. Sirovich and Kirby state: “It seems reasonable to assume that an efficient procedure for recognizing and storing pictures concentrates on departures from the mean. With this in mind, the deviation or departure from the mean is defined.”²⁵ What they describe is a way of defining the char-

23 Jeremy Kun, “Eigenfaces, for Facial Recognition,” *Math Programming* (blog), July 27, 2011, <https://jeremykun.com/2011/07/27/eigenfaces/>.


24 Turk and Pentland, “Eigenfaces for Recognition,” 73.

25 Sirovich and Kirby, “Low-dimensional Procedure,” 519.

acteristics of faces through the differences between faces. The *mean*, therefore, becomes a base from which to designate a difference.

As in many facial recognition algorithms, the application of PCA begins with a training set produced from multiple facial images. In 1986, Sirovich and Kirby were the first to experiment with building a training set of faces on which to apply the method of PCA. They report taking still analogue shots of video recordings.²⁶ Their first test group of faces came from their immediate environment, the relatively homogenous population of “the undergraduate male population” that dominated the mathematics department at Brown University, a group they describe as made up of “smooth-skinned caucasian males.” They recorded the faces for the training set using video, on top of which they overlaid a cross-hair aligned vertically with the midline of the face and horizontally with the pupils. They were able to adjust the depth of field of each video still so that the width of each face could be equalized. These images were then digitized and turned into gray-scale pictures through an image processor. The result of PCA is a mean of all the faces of the training set and is depicted in an image in Sirovich and Kirby’s paper (figure 6).

Figure 6: Sirovich, Kirby, “Average face based on an ensemble of 115 faces.” 1987.



²⁶ *Ibid.*, 522.

These initial experiments with training sets and the application of PCA are interesting to examine because the average male face was constructed from a training set collected from the exclusively white, male population of the applied mathematics department at Brown University. In choosing faces whose characteristics were closer in similarity, it raised the threshold of success for recognition. Sirovich and Kirby describe how they purposefully chose faces that were similar to each other in order to produce the best outcome.²⁷ Higher rates of positive recognition are thus dependent on a smaller degree of difference between the faces included within the modeling of the average face.²⁸ Similarities between the physical characteristics of faces is part of the logic of recognition. Even at this stage, the building a viable training set, there is already a kind of reduction being applied, a normative categorization of faces according to characteristics of gender and race.

There is another technical reduction process that occurs at this initial stage. The facial images that make up the training sets are converted to gray-scale (if originally in color) and to low pixel resolution. This reduction is indicative of perceptual conditions that contrast with those of human processes of recognition. The conversion to low resolution (usually at 200 x 180 pixels) and to gray-scale, that is, to values of light and dark, reduces the amount of visual information available to the human eye. Yet, for the algorithm, this reduction provides clarity by way of “simplifying” the images – to use the vernacular of digital post-production terminology – meaning that it leaves only the information necessary for the operation and gets rid of the rest. Reduction by way of gray-scale and low resolution eliminates the obscurity or *extra* information that detracts from the ingredients the algorithm deems important, that is, what can be measured and calculated through pixel values.

The reduction of the information in the image provides a better “palette” for pattern recognition to take place. These initial processes of reduction, both in the format of the images and in the choice

27 Ibid., 524.

28 This method made possible improvements to the eigenface algorithm and the development of the method of “fisherfaces.”

of demographic from which the faces are drawn, are part of constructing the representational mechanism of the algorithm; that is, they help shape what is deemed salient and relevant to an operation of recognition. In applying statistical processes of reduction, PCA performs a kind of transformation of facial variations into a “working object.” As Lorraine Daston and Peter Galison explain, a “working object” as, “any manageable, communal representative of the sector of nature under investigation” such as atlas images and type specimens, that at times replace the natural specimens they represent.²⁹ Organic forms produce endless variation and are unable to “cooperate” in generalizations and comparisons. In contrast a “working object” provides a common object inquiry. In describing scientific atlases as a “working object,” Daston and Galison explain how it served to “teach how to see the essential and overlook the incidental, which objects are typical and which are anomalous, what the range and limits of variability in nature are.”³⁰ The principle of the “working object” was based on allowing for collective scientific inquiry to occur through the standardization of natural forms. As Daston and Galison show, the creation of working objects was central to the work of scientific inquiry and the classification of natural phenomena. PCA performs the task of transforming faces and facial variations into a manageable and workable object by defining the range of facial variations in an operation of recognition. The statistical method of PCA is able to merge multiple natural forms of faces into a single conglomerate. As with the working object, PCA is able to refine the facial variations and envelop them within a readable (at least by a machine) working object of the averaged face, transforming the variations into a manageable form. In this way, the statistical process of PCA functions like a scientific atlas of the algorithm to train it to see the essential aspects of the human face, found in the average but also to see the “incidental” or rather the deviations from the average as a method of recognition.

29 Lorraine Daston and Peter Galison, *Objectivity* (Cambridge: Zone Books, 2007), 19.

30 *Ibid.*, 26.

Training sets are the data banks and, as such, the source of knowledge for an algorithm. They provide knowledge for the algorithm on what it is allowed to see and recognize. In this way, for an algorithm, the training set is the link between knowing and seeing. The training set is what enables the algorithm to know certain faces and thus to recognize them. The early examples of images of average faces produced by Sirovich and Kirby's training sets reveal a bias toward white men as a primary and normative category of recognition. Although the training sets used in their research were part of an experiment with eigenface and not examples of the actual application of AFR technology, the demographic of the people in the images corresponds with that of the people in the training sets in the initial experiments with AFR technology conducted by Bledsoe. As such, there is a history of the training sets used in the development of automated processes of recognition predominantly involving the figure of the white male. In their analysis of facial recognition systems, Lucas D. Introna and David Wood remark that such reductions are where bias can be located in the algorithmic process.³¹ Although they do not scrutinize specific AFR methods in great detail, based on their examination of training databases they speculate that, through the reductive process of both image-based and feature-based algorithms, minorities are most likely to deviate from statistical averages that result from facial recognition systems. In analyzing the problematics of reduction, they conclude that minorities of Asian and African-American descent are the easiest to recognize in virtue of this deviation from the mean.³² This, they argue, contradicts the claims of suppliers and security analysts in the biometric industry about the neutrality of AFR systems.³³ In light of Introna and Wood's findings and the "average faces" presented in the papers on the development of eigenface, we may conclude that representation in facial recognition systems has been dominated by the white male, presenting all other demographics as deviations from the norm.

31 Introna and Wood, "Picturing Algorithmic Surveillance," 186.

32 *Ibid.*, 190.

33 *Ibid.*, 191.

The application of PCA to training set images creates “eigenvectors,” that is, mathematical objects that display the degree of variability or deviation between facial characteristics and an average. Each eigenvector represents the greatest degree by which the facial images may vary – i.e., the highest eigenvalues. Multiple eigenvectors result from applying PCA to training sets, creating a mechanism to classify unknown faces based on their deviation from these eigenvectors. The eigenvector is a virtual model of “known” faces and serves as a reference point for the classification of unknown faces. Like a map, an eigenvector stands as a kind of idealized model; it acts as a primary referent on the basis of which the algorithm is able to measure the distance or variation between it and an unknown face. As a virtual model, the eigenvector is a form of representation that transforms the pictorial, individual representations of known faces into a geometrical space defined by facial measurements. In this way the eigenvector comes to represent faces based solely on their relationships to other faces. The eigenvector is a representation of the differences and similarities between faces and in this way functions as a unit of facial measurement.

It is only when an eigenvector is displayed to meat eyes, that is, to human vision, that it is referred to as an eigenface. The greater the variation of an eigenvector, the more blurred the eigenface appears. Programmers have referred to eigenvectors as “ghost faces”³⁴ because of their phantom-like appearance. These programmers are describing the form of these eigenvectors, which is characterized by the multiplicity that is inherent in the statistical process. The programmers’ reference to ghost faces evokes a notion of imagery put forward by W. J. T. Mitchell, who describes a type of imagery that is perceptual and occupying, “a kind of border region where physiologists, neurologists, psychologists, art historians, and students of optics find themselves collaborating with philosophers

34 Müge Çarıkçı and Figen Özen, “A Face Recognition System Based on Eigenfaces Method,” *Procedia Technology* 1 (2012): 122, <https://doi.org/10.1016/j.protcy.2012.02.023>. [118-123]

and literary critics.”³⁵ He further describes this imagery as playing the role of fantasmata and as existing as “revived versions of those impressions called up by the imagination in the absence of the objects that originally stimulated them.” The eigenface image can be understood as this type of fantasmatic image. Indeed, one of the primary aesthetic features of the eigenface image is the absence of the individual face, which disappears in the midst of its conglomerate form. Instead, the facial appearances in the eigenface image function symbolically to construct the virtual facial model, which acts as the central referent in the recognition process.

The collection of eigenvectors create a subspace referred to as the “face space.” Eigenface developers have described the face space as a virtual subspace that is defined and framed by the measured distances between a collection of eigenvectors. The concept of a face space derived from the field of psychology in interpreting how faces are processed by human recognition. The face space is originally defined as a, “multidimensional psychological space, in which faces can be represented according to their perceived properties” and with the, “assumption...that faces (or concepts) could be represented as a collection of interchangeable parts.”³⁶ The face space in the eigenface algorithm is defined by the range of variability of these “interchangeable parts.” Conceptually, the face space spans all possible variations of faces. Turk explains that any kind of:

image deviations (whether due to image noise or other factors, such as illumination, pose, expression, occlusions, etc.) push an image away from the space, and the *distance from face space* can be used to determine how likely an image is to be a face in the first place, thus providing a built-in mechanism for face detection.³⁷

35 W. J. T. Mitchell, *Iconology: Image, Text, Ideology*. Chicago: University of Chicago Press, 1986, 10.

36 Tim Valentine, Michael B. Lewis, and Peter J. Hills, “Face-Space: A Unifying Concept in Face Recognition Research,” *Quarterly Journal of Experimental Psychology* 69, no. 10 (2016): 1996-2019, <https://doi.org/10.1080/17470218.2014.990392>.

37 Turk, “Over Twenty Years of Eigenface,” 2-3 (italics in original).

Turk's description thus identifies the face space as the source of facial detection in the algorithm. The face space operates as a virtual data bank, storing an algorithm's knowledge, enveloping all the variations of possible faces that the eigenface algorithm can conceivably recognize. The collection of eigenvectors with highest eigenvalues, that is, highest measured variations, is used as a referential source for algorithmic knowledge. In this way, the face space is like a virtual filing cabinet – the primary bureaucratic mechanism of identification in criminology. But instead of a filing cabinet storing individual records, the face space collects statistical averages or eigenvectors to serve in the process of recognition and verification.

The actual recognition process in eigenface involves projecting the captured image of the individual who needs to be identified on to the face space. The image is compared with the face space by calculating the Euclidian distances between the eigenvectors and the captured image. If there is a small distance between the capture and the eigenface, there may be a match. The distances between them are then expressed in numerical values and a data set is created. This data set then represents a person's identity and is entered into a database. An individual is classified within a biometric database not through their image but rather through numerical code. If there are large distances between the capture and the eigenvectors within the face space, then there is no match. If there is no match, the captured face can be incorporated into the algorithm, adding a new variation within the eigenvector. Turk and Pentland expanded on Sirovich and Kirby's use of the PCA method by incorporating this machine-learning technique within the eigenface method. They state: "The concept of face space allows the ability to learn and subsequently recognize new faces in an unsupervised manner."³⁸ Turk and Pentland describe the process by which the detection of unrecognized faces builds new patterns in the algorithm:


When an image is insufficiently close to face space but is not classified as one of the familiar faces, it is initially labeled as "unknown."

38 Turk and Pentland, "Eigenfaces for Recognition," 79.

The computer stores the pattern vector and the corresponding unknown image. If a collection of “unknown” pattern vectors cluster in the pattern spaces the presence of a new but unidentified face is postulated.³⁹

The statistical average or eigenvector defines the recognition procedure and is the form of measurement in relation to which a face is or is not recognized. In machine learning, the algorithm is designed to continue developing over time as facial variations are continuously added and learned.

Figure 7: Vincent Scheib, “Eigenfaces of UNC”.



Face spaces provide visual sources of information about the eigenface process. As collections of eigenvectors, face spaces depict a wide spectrum of measured distances of similarity and difference. I present three face spaces here in order to illustrate the kinds of variation they exhibit, as well as the issues and problems that arise in the construction of these spaces. By presenting eigenvectors side by side, face spaces can illustrate specific aspects of variation in these eigenvectors. The face space in figure 7 is typical of the face space created by programmers working with eigenface. One thing that we can see clearly in this face space is the variation between the dif-

³⁹ Ibid.

ferent eigenvectors.⁴⁰ The eigenvector at the top left has the highest eigenvalues, representing the average of all the 201 faces within the training group. The eigenvector immediately to the right captures the overall brightness of the face in the picture. The two following eigenvectors to the right capture the direction from which the face is illuminated. The eigenvectors on the bottom row capture variations in face shape. As the programmer states, the rest of the eigenvectors, of which there were 201, capture more subtle details.

Figure 8: Alexandra Feldman, “Face Recognition: Final Project CS 129, Spring 2011,” Computer Science at Brown University




The greater the variation of an eigenvector, the more blurred it appears. For example, figure 8 is a face space created by a computer science student at Brown University.⁴¹ This face space stood out to

40 Vincent Scheib, “Eigenfaces of UNC,” accessed April 25, 2015, <http://www.scheib.net/school/uncfaces/index.html>.

41 Alexandra Feldman, “Face Recognition: Final Project CS 129, Spring 2011,” Computer Science at Brown University, accessed April 25, 2015, <http://cs.brown.edu/courses/csci1290/2011/results/final/amfi/>.

me from other face spaces because each individual eigenvector is constructed through extreme contrasts of light and shadow, creating a grouping that is aesthetically characterized by a painterly effect, where each conglomerate facial feature resembles a brush stroke. Rather than the usual blur, the facial features are defined through blunt markings of light and shadow. The eyes of the figure in the upper-middle eigenvector are blacked out, like two large ink spots, while the eigenvector in the lower left sports a glowing beard. From a programming perspective the student who constructed this face space was considered to have made a gross error (and was required to fix the contrast levels in her algorithm),⁴² its tonal extremities reveal an aspect of the process by which face spaces are built. Increasing the contrast levels so dramatically causes the feature similarities within each eigenvector to become more distinct – so much so that these distinctions begin to merge into each other and construct other (facial) forms out of the composited similarities. The eigenvector in the upper-left corner is rendered almost completely black. This eigenvector has the highest eigenvalues in the face space. The higher value, which equates to a larger amount of information for machine perception, amounts, paradoxically, to the least visually coherent image for human vision.

Figure 9: Wonju Lee, Minkyu Cheon, Chang-ho Hyun, and Mignon Park, “Best Basis Selection Method Using Learning Weights for Face Recognition,” *Sensors (Basel)* 13, no. 10 (October 2013)



42 Ibid.

Figure 9 is a face space included in an article written by four engineers who were experimenting with an alternative method for selecting eigenvectors for face spaces. In this method, the eigenvectors were selected not on the basis of their eigenvalues but rather through groupings of similar, closely related faces.⁴³ The face space thus created highlights the extreme differential values produced by misaligned faces and is expressed through varying tones of illumination. The silhouette of differently positioned faces creates a halo effect. These “halos” consist of light and shadow that correspond to high eigenvalues. As in the previous face spaces, the eigenvalues descend in order, starting in the top-left corner. The students who created this face space argue that the eigenvectors in group (a), the three images on the top row, are “unimportant eigenfaces” because they do not “have discriminant information”⁴⁴ that would allow the technology to perform an act of recognition. They argue that the faces in the training set should be cropped even closer so as not to include the illuminated edges within the eigenvalues of an eigenvector. Their rationale here is that the variation in illumination was found to have weakened the overall recognition process, and their aim was to find ways to more tightly define facial variation in order to allow for more precise techniques of recognition.

The Eigenface Image

I have outlined these three technical aspects, the PCA, the eigenvector and the face space, because they structure the ways in which the algorithm performs recognition. They also form the empirical basis for my own analysis of how the algorithm perceives. These three aspects provide routes into an understanding of the conditions of recognition in an AFR system. They explain how a face comes to be

43 Wonju Lee, Minkyu Cheon, Chang-ho Hyun, and Mignon Park, “Best Basis Selection Method Using Learning Weights for Face Recognition,” *Sensors (Basel)* 13, no. 10 (October 2013): 12830–51, <https://doi.org/10.3390/s131012830>.

44 *Ibid.*, 12834.

known and how information is produced and archived by the algorithm through successful operations of recognition. In order to understand these processes within the cultural and socio-political contexts in which they take place, I plan not to examine their more technical elements but instead to place them into a dialogue with ideas from the field of visual culture theory and to recognize them as enculturating an algorithmic way of seeing and medium of thought. What I mean by “medium of thought” here is that eigenface is not a static technical process; it is rather designed to learn new faces continuously and to incorporate them within its face space. The representational mechanism of PCA and the production of eigenvectors are the eigenface algorithm’s means of knowledge accrual. The eventual operation of recognition produces knowledge of an individual’s identity. In this way, the algorithmic way of seeing is intertwined with modes of contemporary knowledge production. These technical aspects thus reveal the ways in which the algorithm sees and knows. John Berger opens his seminal book *Ways of Seeing* by outlining the intimate relationship between seeing and knowing: “Seeing comes before words [...] The relation between what we see and what we know is never settled.”⁴⁵ Berger describes a productive and “always-present gap” between knowing and seeing that is fundamental to the experience of visual perception.⁴⁶ In contrast, seeing by way of recognition through an automated algorithmic process reverses the order Berger describes. Knowledge comes before seeing. In this section, I examine this reversal by asking how it affects the subject being seen and by tracing the ways of seeing present in the aspects of the eigenface algorithm discussed above.

The production of the eigenface image marks a kind of *pictorial turn*, albeit a brief one, in the development of automated facial recognition technology. In this turn, images come to supplement equations in the operations of the algorithmic process. The eigenface image is a visualization of a statistical process. It depicts a *statistical way of seeing* in which the eigenvector, or the aggregate face, becomes the source of knowledge for the algorithm. Statistics, as a

45 Berger, *Ways of Seeing*, 7.

46 Ibid.

way of seeing, is not new. In fact, statistics has been described as a tool of visualization, as a way of seeing on a scale beyond the capacity of human senses. In his critical genealogy of the term “information,” media theorist John Durham Peters describes how statistics merged with an understanding of information in contemporary society. Peters notes that the etymological origins of the word lie in the German *statistik*, meaning the comparative (and competitive) study of states. Statistics arose as a tool of politics and state governance. Peters states, “The scale of the modern state presents its managers and citizens with a problem: it is out of sight and out of grasp. It must be made visible.”⁴⁷ He goes on: “statistics arose as the study of something too large to be perceptible – states and their climates, their rates of birth, marriage, death, crime [...] [Statistics is] a set of techniques for making those processes visible and interpretable.”⁴⁸ As such, statistics is a visual tool that is inextricably bound up with its original context of use: state governance and control.

Peters highlights an analogy between the acquisition of knowledge and the body, describing statistics as the “empiricism of the state,” whereby “the state becomes the knower, bureaucracy its senses, and statistics its information.”⁴⁹ When information comes to be understood in the form of statistics, the process of knowledge accumulation relocates from the site of the body to the site of the governing institution. Peters characterizes information produced through statistics simply as “knowledge with the human body taken out of it.”⁵⁰ Considering the use of the statistical method of PCA as the representational mechanism of eigenface, this suggests an ontological paradox. The eigenface method, and other AFR methods, are wholly reliant on statistical pattern recognition processes to produce information, but this information is constituted by and issues from the body itself. This presents us with a situation in which the physical phenomenon deemed invisible and ungraspable is none other than the body itself.

47 Peters, “Information,” 14.

48 Ibid.

49 Ibid.

50 Ibid., 15.

The foregoing discussion about statistics as a tool of vision raises the issue of the treatment of the body as information. Btihaj Ajana, a scholar working within the field of digital cultures, sets out an important approach to the understanding of the body as information with regard to biometric practices. Drawing on Eugene Thacker's concept of "biomedia," she argues for an understanding of the use of biometrics as "less as a tool and more as a *process*, less as an instrument and more as an *act*."⁵¹ Along these lines, Ajana argues that the result of these biometric *processes* and *acts* is that the body is rendered as "both the 'medium' (the means by which 'measurement' is performed) and the 'mediated' (the 'object' of measurement)."⁵² Through the differential calculations of the eigenvector, the recognition process in the eigenface method realizes the convergence between these two roles that the body plays. Irma van der Ploeg claims that the treatment of the body as information in biometric practices introduces a new body ontology. She describes a process of the "informatization of the body," in which bodies are represented in digital code, which "construe[s] the body in terms of flows of information and communication patterns."⁵³ She describes the body as a historical construction (much as Peters describes information) that is "implicated in a process of co-evolution with technology – information technologies, but also surgical, chemical and genetic and visualization techniques, and combinations of these."⁵⁴ She describes how biometrics produces new forms of knowledge that transform our understanding not only of technology but of the body itself.

The eigenface image, as a visual artifact of the algorithmic process, allows us to investigate the visuality of the informatized body of AFR systems. The image presents us with the way faces are depicted, read and treated as information. The visualization of its patterns and form is a result of a statistical representational mechanism, PCA. Algorithms are often invisible, operating in a "black box" and leaving no trace of their processes of computation behind.

51 Ajana, *Governing through Biometrics*, 23 (italics in original).

52 Ibid.

53 Van der Ploeg, "Biometrics and the Body," 64.

54 Ibid.

The algorithm does not need to produce pictures to understand its own process of recognition. The picture is for human eyes. The eigenface image serves as a training image, not for the algorithm but rather for us; it allows human eyes to see like algorithmic eyes. As a window into this process, the image answers one of the sociological criticisms of AFR systems, namely, that they lack reciprocity – its systems identify people without *identifying with* people. The eigenface image provides a form of visual reciprocity, depicting what and how the algorithm “sees.” But observation of the eigenface image provides the very opposite of clarity concerning a person’s identity. The eigenface image depicts a moment of stasis between the multiple inputs of data from the training set and the singular output of recognition. The images compiled in the eigenface training sets offer a multitude of possibilities and a wealth of variation, while the operation of recognition reduces the output to one possible match. Positioned at this in-between phase of the algorithm, the eigenface image not only presents a statistical process but also preserves a moment at which multiple possibilities remain open.

Turk has explained that he and Pentland designed the production of the eigenface image as part of the algorithm because they “wanted to keep clear of the ‘black box’ approach of [...] neural networks [...] in order to have a better ability to understand and debug the method.” The eigenface image is part of the recognition operation; it is really information and not an image in the traditional sense. It is an example of what the artist Harun Farocki has termed an “operational image,” “the aesthetics of which were not intended [...] instead of representing the objects in the world, these images are doing things in the world, they are part of a process [...] they are information and not really images.”⁵⁵ Operational images are produced by a machine and are self-reflexive in the sense that they depict both the conditions of observation and what is observed by the machine. In this way, the eigenface image is, in a sense, pure information. Yet, for the human observer, it gives a sense of a modality of machinic recognition, complete with an inherent aesthetic, as is suggested by its description as a “ghost face.” In developing the

55 Farocki, *War at a Distance*.

concept of the operational image further, specifically with regard to surgical imaging, Aud Sissel Hoel and Frank Lindseth say that they are “generative,” that they “differentially intervene,” distilling characteristic patterns that would not be seen in other ways.⁵⁶ The eigenface image can also be read in this way, as depicting a process of the statistical pattern recognition of multiple faces through a process of differentiation that results in an operation of recognition.

Sirovich and Kirby claim at one point in their paper that the “eigenpicture” is a way of “making matters more concrete.” This claim about the function of the eigenface image relates to its role as a technical image, as Vilém Flusser terms images that structure information in contemporary society and replace other forms of communication. Flusser differentiates the technical image from what he calls “traditional images.”⁵⁷ Technical images, he says, operate at the “intervals” of understanding. They “translate particles” and “bits of information” that could otherwise not be seen into something that is “graspable, conceivable, tangible.”⁵⁸ Flusser says that one of the functions of the technical image is specifically to make graspable information that has become abstract through processes of mathematical calculation. Importantly, Flusser describes the technical image as having an ability “to turn from extreme abstraction back into the imaginable.”⁵⁹ The concepts of the operational image and the technical image offer ways of understanding the role of the eigenface image within the context of its production, as a part of the operation of recognition. But they also offer a way out: that is, they draw our attention to how the eigenface functions outside of the programming context and within a wider ecology of images that break from traditional notions of representation. In particular, two features of images are useful in

56 Aud Sissel Hoel and Frank Lindseth, “Differential Interventions: Images as Operative Tools,” *Photomediations: A Reader*, ed. Kamila Kuc and Joanna Zylinska (London: Open Humanities Press, 2016), 181.

57 Vilém Flusser, *Into the Universe of Technical Images*, trans. Nancy Ann Roth (Minneapolis, MN: University of Minnesota Press, 2011), 10.

58 *Ibid.*, 16.

59 *Ibid.*, 21.

approaching the eigenface method within the context of notions of recognition and identity: their “generative” character and their ability to render things “imaginable.”