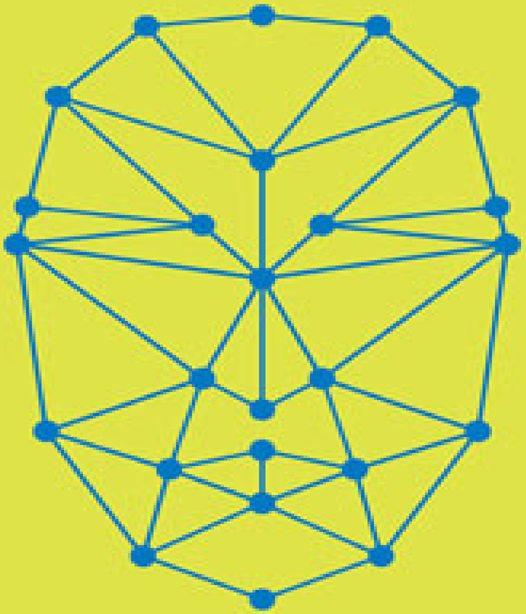

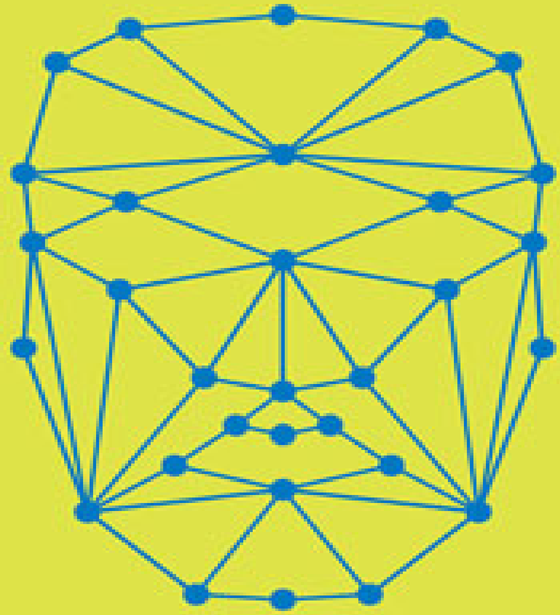

FACIAL RECOGNITION

Mark Andrejevic
Neil Selwyn

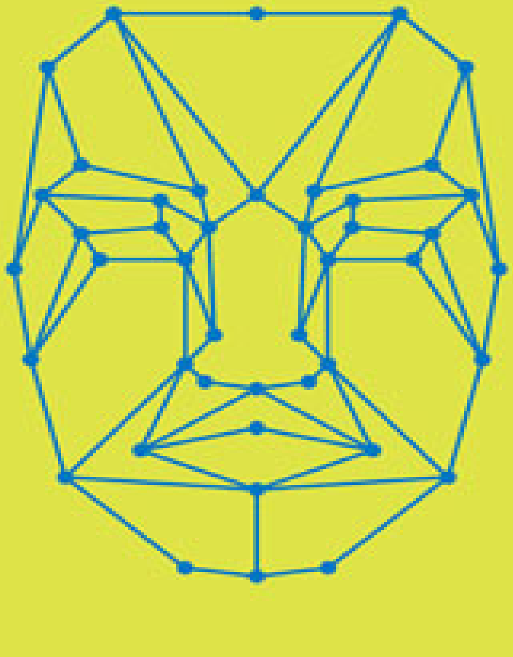

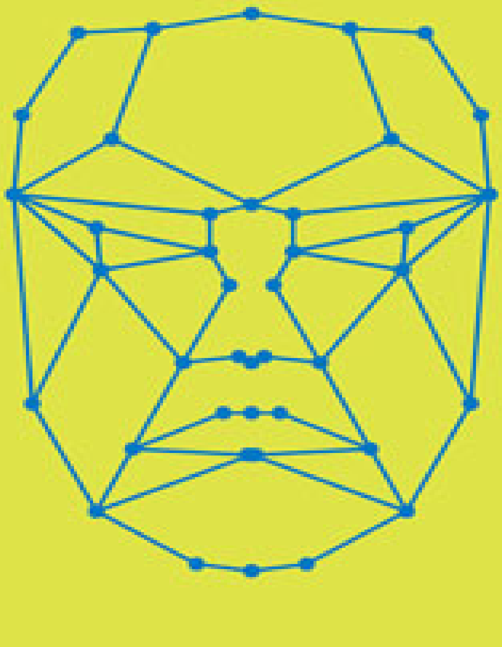

Table of Contents

List of Illustrations

Chapter 1

Facial Recognition

Mark Andrejevic and Neil Selwyn

polity

Copyright Page

First published in 2022 by Polity Press

Polity Press

65 Bridge Street

Cambridge CB2 1UR, UK

Polity Press

111 River Street

Hoboken, NJ 07030, USA

ISBN-13: 978-1-5095-4732-6

ISBN-13: 978-1-5095-4733-3(pb)

A catalogue record for this book is available from the British Library.

Library of Congress Control Number: 2022930699

by Fakenham Prepress Solutions, Fakenham, Norfolk NR21 8NL

The publisher has used its best endeavours to ensure that the URLs for external websites referred to in this book are correct and active at the time of going to press. However, the publisher has no responsibility for the websites and can make no guarantee that a site will remain live or that the content is or will remain appropriate.

Every effort has been made to trace all copyright holders, but if any have been overlooked the publisher will be pleased to include any necessary credits in any subsequent reprint or edition.

For further information on Polity, visit our website: politybooks.com

Acknowledgements

The issues we address in the following pages have come to public attention through the efforts of many people whose work we draw upon to write this book – including Ruha Benjamin, Joy Buolamwini, Kelly Gates, Timnit Gebru, Chris Gilliard, Woodrow Hartzog, Os Keyes, Safiya Noble, Deborah Raji, Evan Selinger, Luke Stark, Simone Browne and many others. We look forward to the books that these (and other) critics are writing on this topic. Our book offers an initial perspective on a topic that is being worked on by hundreds of critically minded and concerned people.

The writing of this book arises as part of a research project funded by the Australian Research Council (DP200100189). We would especially like to thank our co-researchers Chris O’Neill, Gavin Smith and Xin Gu for all their ongoing conversations around these issues, as well as other research colleagues including Liz Campbell, Robbie Fordyce, Zala Volcic, Miguel Vatter and Brett Hutchins.

We are also grateful to the industry, regulatory and computer science experts who have spoken with us as part of the ARC research project, taking time to explain the intricacies of the facial recognition industry and how this technology is being implemented in everyday society.

Preface

The human visual system is remarkably good at recognizing faces. It is estimated that most adults are capable of recognizing around 5,000 different faces, taking an average of around 0.2 seconds to work out who someone is. A very small number of individuals are exceptionally good at remembering the faces of different people they come across. These ‘super-recognizers’ can recall up to 80 per cent of the faces they see (as compared to the 20 per cent level that most people are capable of) and are highly sought after by police, intelligence agencies, banks and casinos to provide specialist face identification. Nevertheless, even the most capable human is unable to maintain this level of recognition over prolonged periods of time, nor to expand their recognitive abilities to the scale of entire national populations. The attempt to recognize individuals at this scale has been an ongoing project that – like so much else – has recently been adopted by the developers of automated digital technology. For the first time in human history, we are confronting the prospect of systems that would have the capacity to recognize individuals by automatically comparing them to databases on the scale of national populations – and beyond.

Computer-based facial recognition carries with it the promise of the infallible and all-insightful eyewitness account. This is the compelling power of ‘seeing’ something happen and knowing exactly who is involved, even if we require a camera to do the seeing and a computer to do the recognizing. Because this technology relies on the unblinking eye of the machine, it also ushers in the prospect of always-on, ubiquitous identification at-a-distance – the ability to augment the world around us with a ‘recognitive’ overlay. In this regard, facial recognition technology marks a decisive shift in monitoring capability, promising the advent of spaces that recognize us wherever we go, even if the people around us do not. We might describe this as the definitive end of a certain kind of privacy. Perhaps, once upon a time, it was reasonable to presume that we retained some anonymity in public spaces, with the majority of our actions and activities going unremarked and unrecorded. If so, then the widespread use of facial recognition-enabled cameras marks the end of such a time.

In many ways, it might seem surprising that such technology would gain widespread purchase, given the demonstrable threat to privacy and concerns about accuracy and bias. However, the experience of the past couple of decades has been shaped by the widespread implementation of increasingly comprehensive and granular forms of monitoring in exchange for the convenience and affordances of various data-driven digital technologies. The ever-growing forms of surveillance and ‘dataveillance’ associated with the use of smartphones and social media have been well documented and discussed. In light of these general conditions, it would be surprising if we collectively ended up deciding to finally and definitively draw the line at facial recognition technology. ‘Go ahead and track wherever I go, all of my communications and online interactions, everything I look at, write, do or share online . . . just don’t scan my face.’

Drawing on the lessons of the recent past and, in particular, the rise of the online surveillance economy, this book starts from the premise that various forms of facial recognition technology will likely become steadily embedded in the minutiae of daily life. In a practical sense, numerous facial recognition applications already exist to allow us to unlock our phones without having to key in any numbers or to provide cardless, touchless access to mass transit, ATM machines, offices, homes and cars. In a more abstract sense, this is technology that promises to give us the satisfaction (at least to some) of recognition – a feeling that the places through which we move know us and respond to us personally, as is increasingly the case in online spaces. In some ways, then, this is technology that offers us the comforting sense that we make an impression. Moreover, as with other recognition-based interactive technologies, the lure of facial recognition technology comes with the promise of convenience, efficiency *and* comfort. This is technology that is framed by its proponents as the solution to inconvenience, inefficiency and risk in a variety of contexts. In the friction-free world envisioned by proselytizers of the technology, for example, face recognition opens up a world of convenience: cars that open at a glance and automatically adjust the seat, mirrors and soundtrack to suit the driver; shops without checkout queues; speedy transit through borders and checkpoints of all kinds, from international customs to secure office spaces (no more key cards!).

For those on the less privileged side of this technology, however, recognition is not necessarily a benefit. This is technology that can be used to track those who fall under suspicion, to bar access, to scrutinize and to sort. This is also technology that the IT industry is currently attempting to develop in forms that can assess someone’s employability, potential threat and/or creditworthiness. In this respect, facial recognition technology may usher in yet another dimension of the digital divide – significant disparities between those who use and control the technology as opposed to those who are subjected to it. As we shall see throughout this book, there are numerous instances where such disparities are already beginning to emerge – from casino managers using facial recognition to decide which patrons get filtered into fast-track VIP queues through to authoritarian states using facial recognition technology to identify and track ethnic minorities.

Yet proponents of FRT are quick to reason that we should not lose sight of more positive uses of the technology, such as deployment of the technology to reunite missing children with their families. In 2020, facial recognition was used to reunite Mao Yin from Xi’an with the birth parents from whom he had been kidnapped 32 years earlier. Authorities reportedly used a photo of the child to create a simulated image of what he would look like as an adult and then searched for matches using facial recognition technology. Elsewhere, researchers have explored the use of facial recognition technology to identify families separated by natural disasters, even if they have sustained facial injuries. These examples are deployed by advocates of the technology to highlight its real pro-social uses.

Those in favour of FRT would also point to uses that many support for increased security: the ability to identify criminals who might otherwise have escaped the consequences of their actions; or helping to prevent fraud that costs people tens of billions of dollars a year. These forms of facial recognition technology speak to the feeling of frustration that comes with the uncertainty that can plague law enforcement: the strong suspicion that someone must be guilty while being unable to find them, and/or the need to know for sure whether a guilty verdict is fully justified. Whether or not facial recognition technology can follow through on the promise to provide certainty is, however, a different matter. As we shall see, there is no such thing as completely certain identification when it comes to automated face matching. Moreover, this is technology that has long been plagued by issues of misrecognition and bias, rendering it less reliable for certain populations. While some experts might argue that such flaws have all been eliminated, others suggest that it is, even in theory, impossible to eradicate entirely.

So, all told, this is technology that appears in many guises: from authoritarian control to personal assistant. As in the case of technology more generally, its future will depend on how we choose to use it, and who the ‘we’ end up being. One thing seems clear: this is technology that is likely to become more widespread in the near future, which means the time to anticipate and assess its social consequences is now.

1
Facial Recognition: An Introduction

Introduction

The beginning of the 2020s marked a point when facial recognition technology was thrust to the forefront of political and mainstream concern. This was a time when some of the most high-profile tech companies vowed to suspend their development of the technology, or else cease operations completely. In the midst of the 2020 Black Lives Matter protests in the United States against police brutality, CEO Arvind Krishna sent an open letter to Congress pledging that ‘IBM no longer offers general purpose IBM facial recognition or analysis software’ (IBM 2020). Soon after, Amazon joined Microsoft in also announcing that it was ‘implementing a one-year moratorium on police use of Amazon’s facial recognition technology’ (Amazon 2020). Twelve months later, these corporate stances were reaffirmed, and various US city, state and federal bans on government use of facial recognition technology (FRT) were pursued. Then, towards the end of 2021 Facebook proclaimed that it was switching off the FRT auto-tagging feature on the social media platform, and deleting its vast dataset of over one billion facial scans in response to what it framed as ‘many concerns about the place of facial recognition technology in society’ (Hill and Mac 2021).

Perhaps the most damning indictment came from within the academic computer science community. The ACM (Association for Computing Machinery), while hardly the most politically motivated or publicity-seeking organization, issued a sternly worded ‘statement on principles and prerequisites for the development, evaluation and use of unbiased facial recognition technologies'. This June 2020 statement pulled no punches:

> when rigorously evaluated, the technology too often produces results demonstrating clear bias based on ethnic, racial, gender, and other human characteristics recognizable by computer systems. . . . Such bias and its effects are scientifically and socially unacceptable . . . [the committee] urges an immediate suspension of the current and future private and governmental use of facial recognition technologies in all circumstances known or reasonably foreseeable to be prejudicial to established human and legal rights. (ACM 2020)

These statements and actions were welcomed at the time by activists and critics as marking the end of what had come in some circles to represent a fundamentally oppressive, discriminatory and retrograde aspect of digital technology development. Nevertheless, as with any multi-billion-dollar industry based on cutting-edge technology development, the story is not so simple.

Soon after their announcements, many of the Big Tech companies began to face challenges over the details of their apparent changes of heart. What exactly did IBM mean by 'general purpose' facial recognition, and what other specific purposes did that leave open? Amazon only offered to suspend police use of its Rekognition technology, while proudly boasting that it would continue with its humanitarian uses by other organizations. Most of the US city bans on facial recognition were directed at specific uses by municipal services. Even the ACM statement was actually a call for improving the accuracy, transparency and accountability of what it acknowledged was 'powerful' technology that was 'likely to improve in the future' with 'potential to help meet significant societal needs'. By placing concerns around bias at the heart of the push-back against the technology, some bans served as an incitement to further develop the technology and provided it with a potential alibi: if this technology could be developed without systematic forms of bias, then perhaps it might be acceptable.

Similarly, the high-profile Facebook promise to delete its vast dataset of facial recognition scans deflected attention from the fact that the company was retaining the DeepFace model that it had developed from this dataset, and plans to incorporate FRT into future products including its plans for the 'metaverse' and for smart glasses – stressing that 'every new technology brings with it potential for both benefit and concern, and we want to find the right balance' (Hill and Mac 2021). Proclamations such as this are best seen as Big Tech companies 'responding to the controversy by pivoting rather than pulling back' (Bass and Bergen 2021). Investment in facial recognition ventures continues to rise, US companies continue to sell facial recognition products to police and security forces overseas, as well as develop less controversial domestic markets.

So, while industry acknowledgement of the need to reconsider the use of facial recognition technology was welcome, this seemingly significant turning-point has subsequently proved to make little difference to the ongoing rise of this technology around the world. A well-crafted proposal to Congress – the Facial Recognition and Biometric Technology Moratorium Act – for federal legislation that sought to place an indefinite ban on police use of the technology received no industry support, and eventually sank without trace. At the same time, the development and implementation of FRT continued apace. As it turned out, the opening years of the 2020s proved to be a time when FRT was introduced into everything from $400 home security camera systems through to 'pay-by-face’ kiosks in burger chains. A couple of months after the Amazon and ACM announcements, it was estimated that the market for ‘facial biometrics’ would exceed US$15billion by 2027 (Burt 2020).

Against this background, this book attempts to make sense of one of the most far-reaching – and controversial – technologies of recent times. How did we reach a point where the *New York Times* shifted from enthusing in 2008 that ‘Facial recognition software holds great promise’ (*New York Times* 2008) to presenting ‘A Case for Banning Facial Recognition’ (*New York Times* 2020a) just over a decade later? How do controversies over the use of this technology for law enforcement compare with more banal applications in shopping malls and sports stadia, and with humanitarian uses such as for finding missing children? We take these seemingly disparate applications of the technology to be held together by what Kelly Gates describes as a shared logic of control based on increasingly comprehensive monitoring and individualization. As she puts it in her groundbreaking work on facial recognition technology:

> the possibility of digital biometric identification should not be understood in narrow terms as the natural outgrowth of technical advancements in identification systems. Instead, these technologies are being envisioned and designed to fulfil certain perceived social necessities and political–economic demands of large-scale, late capitalist societies . . . The expansion of computer networks has created new problems of communication-without-bodies, necessitating new techniques for identifying people, verifying their legitimate identities, and otherwise gaining knowledge about who they are. (Gates 2011: 16)

In this opening chapter, we consider how this once niche ‘biometric’ technology grew to be a relatively inexpensive and easy ‘plug-in’ to even the most innocuous everyday devices and applications. First, we look back over the history of computers being given the task of matching faces with people – an application that computer scientists have been grappling with since the 1960s. This history is an important part of making sense of the present and future uses of this technology, with a number of logics already set in motion well before this technology came to mainstream public attention. We then set the scene for the rest of the book by considering the social implications of how this technology is currently being used, as well as the future applications that lie ahead. More significantly, perhaps, we begin to consider the promised benefits and overarching imperatives that shape the current deployment of the technology. Despite highly publicized concerns and backlash, FRT continues to develop apace and is poised to transform the surveillance landscape. Why is it that despite recent controversies, the technology nonetheless continues to work its way into so many aspects of contemporary technology and its anticipated future uses?

A history of computers and facial recognition

1960s to 1990s: establishing a technical proof of concept

The association between computers and facial recognition is usually traced back to the work of US researcher Woodrow (Woody) Wilson Bledsoe and his collaborators Helen Chan Wolf and Charles Bisson. These three spent much of the 1960s working under the guise of the Panoramic Research company based in Palo Alto. Previously, Bledsoe had been involved in developing the ‘n-tuple’ approach to automated pattern recognition – a technique that divides images of shapes onto a grid of cells, assigning values of 1 (full) and 0 (empty) to each pixel, and then computing a unique score that could be later matched to other close-scoring patterns. Extending this logic, the Panoramic team saw faces as a computationally challenging (in Bledsoe’s words, ‘noisy’) pattern to match.

That said, Bledsoe’s interest in face matching was driven by broader ambitions. In a detailed investigative report on Panoramic Research over fifty years later, journalist Shaun Raviv (2020) recounted Bledsoe’s innocuous ambitions for expanding this technology beyond merely recognizing patterns, and instead building a mechanic ‘computer friend’: ‘I could see it, or a part of it, in a small camera that would fit on my glasses, with an attached earplug that would whisper into my ear the names of my friends and acquaintances as I met them on the street . . . For you see, my computer friend had the ability to recognize faces.’

In this formulation, Bledsoe touches on a characteristic theme of facial recognition technology that highlights both its appeal and its ‘creep factor’. The idea that the computerized automated interfaces that increasingly populate our interactive world might come to recognize us certainly makes the technology feel less alienating – especially the idea that this is technology capable of getting to ‘know’ us. In a world where we spend more and more time ‘at the interface’, many people feel an implicit desire to humanize the technologies that surround us – partly in an attempt to make our digitally mediated lives a bit less isolating. At the same time, of course, there is something uncanny, alarming, and creepy about machines that seem to recognize us and divulge our identities – in part because this endows them with an opaque power. This is technology that may know our details while what goes on behind the interface remains obscure to any onlooker. What does the machine know, what information is it collecting and whom is it sharing it? There is always something suspect lurking in the promise of machinic recognition: a pastiche of human interaction.

Panoramic, however, took a more pragmatic approach than that outlined by Bledsoe, pitching the technology's capability for a variety of military intelligence and law enforcement applications. For the remainder of the 1960s, nearly all the company's work took the form of classified projects funded by a succession of unspecified US intelligence agencies. Bledsoe's initial proposal to conduct 'a study to determine the feasibility of a simplified facial recognition machine' was necessarily small-scale – seeking to program a computer to recognize ten different faces. When it became apparent that this feat could not be achieved by the computer alone, the Panoramic team adopted a 'man-machine' approach with human operators using electronic tablets to manually mark coordinates of various facial features including the eyes, nose, hairline and mouth. This data could then be transformed with the n-tuple methods. While progress was slow, by 1967 Panoramic researchers had successfully developed a system that could match police mugshots within a photographic database of '400 adult male Caucasians'. The initial proof of concept for facial recognition had been achieved.

Bedsloe's work established many of the basic principles that facial recognition developers continue to develop fifty years later, while also encountering many of the field's enduring problems. For example, the Panoramic team established the idea of digitizing images and using pointillistic methods, as well as making early attempts to rotate images to account for head tilt, lean and rotation. More fundamentally, even the most sophisticated facial recognition systems continue Bledsoe's basic approach of creating scores for images and comparing similarities. The

Panoramic team also encountered now-familiar challenges to facial recognition developers – such as dealing with variations in facial expression, hair growth and the effects of aging, as well as how photographs are lit and composed.

Perhaps more significantly, Bedsloe also unwittingly pre-empted much of the ethical and moral complexities of this new branch of computer science. Panoramic took regular funding from various state agencies for highly classified applications of facial recognition. Although Panoramic pitched a Defence Department project in 1965 to use facial recognition to identify people's racial background, little interest tended to be shown in training their systems to recognize diverse sets of images. As Raviv (2020) notes, 'I did not see images of women or people of colour, or references to them, in any of Woody's facial-recognition studies.' Bledsoe's dream of an all-seeing 'computer friend' was already mired in the realities of 1960s US society and politics.

These initial efforts were then followed by a succession of US research and development projects over the next thirty years. A paper in May 1971 by Jay Goldstein, Leon Harman and Ann Lesk (all engineers at Bell Telephone Laboratories) outlined a refined method for manual facial recognition relying on 22 key 'markers'. These included features such as ear protrusion and eyebrow separation – with the team estimating that being able to match as few as seven of these markers could result in a match for unique identification. As with Bledsoe's work, this was a 'man-machine' system relying on researchers manually marking photographs. Again, all the examples given (including their composite 'Mr. Average' image) were white middle-aged men.

These techniques continued to be refined over the 1970s and 1980s, with researchers using larger training datasets and becoming less reliant on manual coding. Using an expanded set of 850 digitized photographs, the Japanese researcher Takeo Kanade was soon able to program the automated extraction of key facial features such as eyes, mouth and nose. Further refinements continued up until the end of the 1980s when Lawrence Sirovich and Michael Kirby – mathematicians from Brown University – made what is now considered to be a significant breakthrough. This involved the use of a linear algebra technique called ‘eigenvectors’ to produce a relatively small ‘basis set’ of low-dimensional face images (‘eigenfaces’). These eigenface images were composite images of hundreds of actual faces, resulting in blurry low-resolution patterns, which often look little like recognizable faces.

Sirovich and Kirby reasoned that any human face could be uniquely identified by its variation from a baseline ‘average face’ eigenface and the extent to which its features are present in other eigenface images. Significantly, Sirovich and Kirby found that most people could be identified on the basis of how their face matched with fewer than one hundred of these eigenface images. The advantage of this approach was that it could recognize and store an individual’s face as a series of values corresponding to each of the ‘basis set’ eigenfaces being used in the system. Relying on numbers rather than digital photographs meant that massive amounts of facial information could be collated and stored.

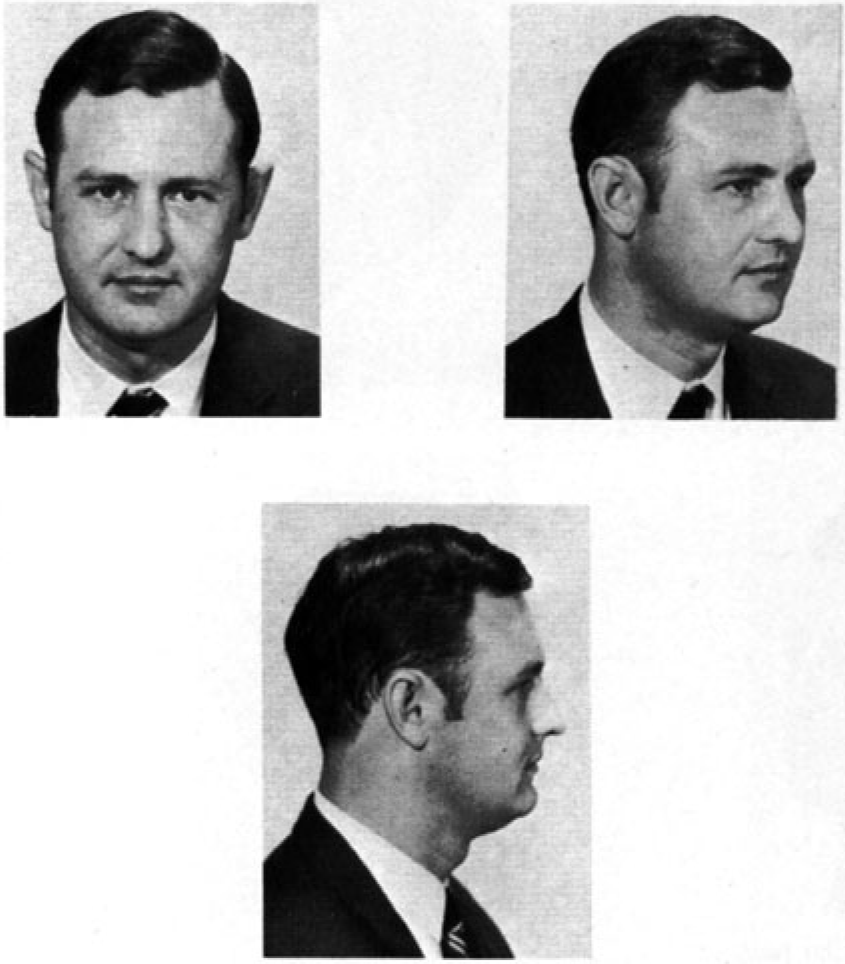

Figure 1.1 Goldstein, Harmon and Lesk’s (1971) profile photographs of an ‘average’ face. © IEEE

A few years later, these efficiencies allowed Alex Pentland and Matthew Turk from MIT’s Media Lab to refine the eigenface technique to extract images of human faces from their background environments, and then make quick matches. Three decades on from Bledsoe’s original plans, a working system for real-time automatic facial detection had been realized. As Turk and Pentland (1991) put it, their development of the eigenface approach ‘was motivated by information theory, leading to the idea of basing face recognition on a small set of image features that best approximate the set of known face images, without requiring that they correspond to our intuitive notions of facial parts and features'. This move away from having to recognize a 'face' per se was acknowledged to be 'a practical solution that is well fitted to the problem of face recognition. It is fast, relatively simple, and has been shown to work well in a somewhat constrained environment.'

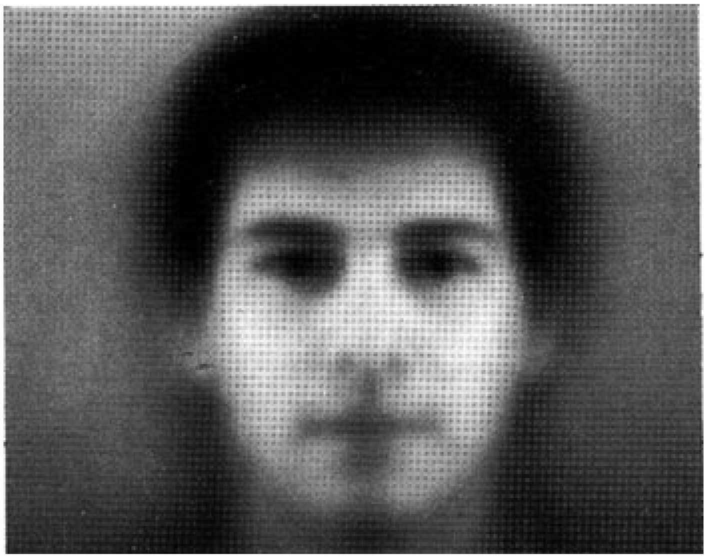

Figure 1.2 'Average face' based on ensemble of 115 faces (Sirovich and Kirby 1987). Reprinted with permission from © The Optical Society.

1990s to 2010s: establishing commercial opportunities

With these technical precedents having been set, attention then shifted to developing a commercial market for facial recognition technology. Central to these efforts in the United States were the government agencies DARPA (the Defence Advanced Research Projects Agency) and NIST (the National Institute of Standards and Technology). These agencies collaborated during the 1990s to run the FERET (Face Recognition Technology) research program, with the goal of collating a large database of high-quality images for commercial developers to use. NIST subsequently launched a program to provide independent government evaluations of commercial facial recognition systems from the time they were being developed or through to their entry into the commercial market. These Face Recognition Vendor Tests were held four times during the 2000s, alongside a couple of Face Recognition Grand Challenges – all designed to encourage the development of increasingly accurate and adaptable systems.

Despite this extensive government support, facial recognition remained a frustratingly inconsistent and unreliable technology throughout the 2000s. The computational challenges inherent in the process of extracting and processing large numbers of potential facial images meant that facial recognition systems remained difficult to deploy at scale. One notable setback was the 2001 covert deployment of facial recognition at Super Bowl XXXV in Tampa. Here federal authorities and state police monitored the stadium and surrounding bars and clubs as a test case for the FaceFinder large-scale surveillance system. While claiming success in detecting 19 known ‘petty criminals’ from around 100,000 attendees, no arrests were made, and there was a general sense that the system had proven incapable of coping with large crowds. Public reaction to what news media soon dubbed the ‘Snooper Bowl’ remained muted.

Yet, after forty years of incremental progress, the technological development that eventually tipped facial recognition over into being reliable and powerful enough to be deployed at scale was not directly related to computer vision or pattern matching per se. Instead, one of the most significant shifts in facial recognition was the rise of social media platforms such as Myspace and the appropriately named Facebook. In particular, one unforeseen consequence of people’s insatiable appetite for sharing images of themselves on social media (at all angles, and all