

A magazine of innovation, new ideas and how the world is changing

Face value

Long

Face value

Researchers recently claimed to be able to use AI to infer the inherent 'criminality' of faces. It's an idea with a long, dubious and dangerous history

26th January 2017

By Rhodri Marsden

Illustration by Luke Best

When the BBC broadcast the recent documentary by [Louis Theroux](#) that looked back at the time he spent in the company of Jimmy Savile, there was disbelief across social media that no one had stepped in to stop Savile from committing his crimes. Some blamed the BBC, some blamed those in Savile's immediate circle, but others blamed a simple error of human judgment.

"He literally couldn't look more like a paedophile," read one post – one of many to state a supposedly incontrovertible truth: that Savile's criminal tendencies could have been detected from the shape of his features, his eyes, his hair. Moreover, this has nothing to do with the benefit of hindsight and should have been picked up at the time. His looks, they suggested, were a moral indicator, with a wealth of compelling visual

evidence to support the claim.

We know that paedophiles, murderers and other violent criminals come in many shapes and sizes. If we knew nothing about their criminal history, some of their photos might even appear attractive. But the idea that someone's features betray their character is something rooted deep within us; it's the reason why certain photos perform well on dating apps, or why trustworthy-looking politicians might rack up votes. But how wrong are our hunches of perceived criminality?

A recent [paper](#), published by Xiaolin Wu and Xi Zhang of Shanghai's Jiao Tong University, claims to be the first to use machine learning and neural networks to attempt a fully automated inference of criminality from facial images, removing prejudice from the equation and testing the validity of our gut feelings. "What facial features influence the average Joe's impulsive and yet consensual judgments on social attributes?" they ask. Through a study of 1,856 images ("controlled for race, gender, age and facial expression") they claim to have established the validity of "automated, face-induced inference on criminality, despite the historical controversy surrounding this line of enquiry."

In other words, they believe that they've found a relationship between looking like a criminal and actually being one.

The idea that someone's features betray their character is something rooted deep within us

It's a claim that's been made many times over the years. Physiognomy, the

'science' of judging people by their appearance, was first theorised by the ancient Greeks in around the 5th century BC. Aristotle's pronouncement that "it is possible to infer character from features" led to a number of works relating to 'Physiognomica', a word derived from *physis* (nature), *nomos* (law) and (or) *gnomon* (judge or interpreter).

All of Greek society, it was claimed, could benefit from this skill: it could assist with choosing an employee, a slave or a spouse, while its inherent vagueness made it intriguing to philosophers and useful for scientists who bent the theories to support their own beliefs. It became a recognised science in the Islamic world, and was used and taught in Europe throughout late antiquity and the early Middle Ages, despite nagging doubts among thinkers and physicians of the day. In the early 16th century, Leonardo da Vinci claimed not to "concern myself with false physiognomy, because these chimeras have no scientific foundation."

Theories of physiognomy, however, would persist beyond the Renaissance. In 1586, Italian scholar Giambattista della Porta published a book, **De humana physiognomonia libri IIII**, which established him as the 'father of Physiognomy'. Della Porta's thinking was based on the 'doctrine of signatures'; the idea that the appearance of plants and animals offers clues to their nature. For example, as one writer of the time suggested, walnuts are good for curing headaches because they're shaped a bit like a human head. The theories in della Porta's book were supported by dozens of detailed illustrations which, by comparing human faces to those of animals, suggested that they must surely share similar character traits.

In the 17th century, Swiss poet Johann Caspar Lavater took della Porta's methodology and ran with it, commissioning artists to illustrate his popular **Essays On Physiognomy** – which, to the chagrin of his contemporary, the writer Hannah More, sold for "fifteen guineas a set... while in vain we boast that philosophy [has] broken down all the

strongholds of prejudice, ignorance, and superstition."

Lavater's work was criticised for being ridden with bias (black faces rarely emerged well from his analyses) but he was right in one respect: "Whether they are or are not sensible of it," he wrote, "all men are daily influenced by physiognomy."

Many studies have been done into our psychological response to faces, and it's clear that a so-called halo effect will inevitably work its magic.

"Attractive people are regarded as better at everything," says Professor Peter Hancock, lecturer in Psychology at Stirling University. "And we can't shake that off because there's some truth to it. Good genes produce intelligent people, attractive faces, fit bodies, and we imagine that they're going to be good at everything else, too. We don't have good insight into our own behaviour. We tend to think we understand what we're doing, but we don't."

Hancock describes attending a conference where one speaker showed a series of black faces and white faces to students (who were mostly white) and asked them what they thought the experiment was about. "They knew that he was trying to assess whether they would rate the black ones as more criminal," says Hancock. "But then they did!"

We attribute social characteristics based on opinions we already hold about certain kinds of faces: whether they look unusual in some way, whether they resemble a partner, a family member or even ourselves, or perhaps have some other cultural association. Physiognomy ultimately stems from what Alexander Todorov, professor of psychology at Princeton University, calls an 'overgeneralisation hypothesis'. "People," he wrote, "use easily accessible facial information (eg an expression such as a smile, cues to gender and ethnic group) to make social attributions congruent with this information (eg a nice person)."

Studies show that people with

stereotypically 'untrustworthy' faces tend to receive harsher treatment

In a social media age, the pictures we choose to represent ourselves online are a form of self-presentation driven by those social attributions and the knowledge that our pictures are being judged.

Experiments at Princeton found that we take less than one tenth of a second to form an opinion of strangers from their pictures, and those opinions tend to stand firm even if we're exposed to those pictures for a longer period of time. That tendency to judge instantly gives rise to a number of selfie tropes that are deemed to elicit positive responses, particularly when it comes to photos on dating profiles: certain angles, particular expressions, minute adjustments of eyebrows and lips that might appear to be about narcissism and vanity, but are more about a fear of being incorrectly assessed. After all, false suppositions based on people's faces are hugely influential within society, and in extreme cases they can have a huge impact on people's lives.


When retired teacher Christopher Jefferies was held by police in connection with the murder of Joanna Yeates in Bristol back in 2010, more than half a dozen newspapers gave his unusual appearance particular scrutiny and made assumptions accordingly, which in turn influenced public opinion. This culminated in [substantial damages for defamation](#), two convictions for contempt of court and a [painful ordeal for Jefferies](#), who was entirely innocent.

This kind of deep-seated bias looms large throughout physiognomic works of the 19th and 20th centuries, from absurdities such as Vaught's [Practical Character Reader](#) of 1902 (handy if you want to find out what a "deceitful

chin" looks like) to more inherently troubling volumes such as Cesare Lombroso's **Criminal Man**.

After performing a number of autopsies on criminals, the Italian physician claimed to have discovered a number of common characteristics, and it's worth listing them if only to establish the supposed criminality of pretty much everyone you know:

Unusually short or tall height; small head, but large face; fleshy lips, but thin upper lip; protuberances on head and around ear; wrinkles on forehead and face; large sinus cavities or bumpy face; tattoos; receding hairline; large incisors; bushy eyebrows, tending to meet across nose; large eye sockets but deep-set eyes; beaked or flat nose; strong jaw line; small and sloping forehead; small or weak chin; thin neck; sloping shoulders but large chest; large, protruding ears; long arms; high cheek bones; pointy or snubbed fingers or toes.



Some of the face samples used in the controversial article [Automated Inference on Criminality using Face Images](#), which investigates whether AI can rate the inherent 'criminality' of faces. The paper raises questions about ethics in AI

In a woeful misreading of Darwinian theory, Lombroso unwittingly founded the field of **anthropological criminology**, and more specifically the idea of the born criminal: a hereditary quality that posed a danger to society and must be rooted out. His theories became discredited during the 20th century, but the kind of bias displayed by Lombroso can still be found in legal systems across the world; studies show that people with stereotypically 'untrustworthy' faces tend to receive harsher treatment than

those who don't. There's evidently some consensus over people's attitudes toward certain faces, but it doesn't follow that the consensus is correct.

The only attributes that we're reasonably good at detecting, according to research done at the University of Michigan in the 1960s and later tested at the University of Stirling in 2007, are extroversion and conscientiousness. For other traits there's insufficient evidence that our hunches are correct, with anomalies explained by our evolved aversion to 'ugliness', established links between broader faces and powerful physiques, or cultural associations with certain demographics which are reinforced with nagging regularity by newspapers, books, television and film.

Data-driven studies, based upon huge quantities of facial data, would seem to offer the final word on this. Since 2005, computational models have used various techniques to test for links between social attributes and facial features, resulting in suggestions that our faces can betray, for example, political leanings, sexual orientation and criminality. One [BBC Future article from 2015](#) even describes the 'discipline' of physiognomy as 'gaining credibility'. But Todorov details many problems with these studies, pointing out the challenging nature of doing such experiments with sufficient rigour – not least because different images of the same people can prompt wildly differing results.

The aforementioned study at Shanghai's Jiao Tong University, with its enthusiastic, data-driven analyses of such questions as "What features of a human face betray its owner's propensity for crimes?" prompted a wave of press coverage.

Different images of the same people can prompt wildly differing results

The vision outlined in these articles is of an unethical dystopia where neural networks can assess our faces and establish a likely score for criminality – but Todorov is scathing about this paper, too. "The main problem is the sampling of the images," he says. "There is not enough information about the [nature of] the images of the people who were convicted. Second, clearly, there are huge differences between the two samples [of convicts and non-convicts] [in terms of] education and socio-economic status."

In other words, your appearance is affected by the kind of life you've led, so the classifiers within the computer program are simply distinguishing between different demographics rather than detecting a propensity for criminal behaviour.

Todorov is also wary of these classifiers misidentifying more 'innocent' people than identifying actual criminals, and accuracy is a concern shared by Peter Hancock. "Networks don't assess faces in the same way that we do," he says. "One of our systems, which is a deep network, has a recognition engine which generates an ordered list of how similar various faces are. And sometimes you get good matches – but other times you look at them and say, well, it's the wrong race! To humans they look completely different. And that underlines the fact that the networks are working in a different sort of way, and actually you don't really know how they're working. They're the ultimate black box."

Popular

[Why the world wide web will break up](#)

[Is there such thing as a 'criminal' face?](#)

[What lettuces tell us about deregulating Britain](#)

This isn't to say that the use of big data, and particularly the use of composite imagery (digitally blending together certain types of faces)

doesn't give us useful information and fascinating correlations. "You can, for example, take a given face and use computer software to make it look more or less trustworthy," says Hancock. "I remember a colleague playing with this and he made a less trustworthy version of George W Bush – and how shifty did he look! I'm surprised that they're not using these techniques in political advertising, because you couldn't tell that anything had been done [to the picture], but when you look at it you think 'I wouldn't trust him'."

The revitalisation of the theory of physiognomy by the Shanghai students is, according to Todorov, deeply problematic on a theoretical level. "Are we back to Lombroso's theory," he asks, "that criminals were anomalous creatures, evolutionary degenerates? How does one become criminal, and what role do various life forces play into this? There are people making claims that you just need to look at the face to predict personality and behaviour, but many of these people have not given much thought to their underlying assumptions."

While it's true that we judge books by their covers, covers are more than just faces; we piece together all kinds of cues from people to form our impressions of them. Jimmy Savile's appearance was unusual by any standards, but we absorbed a great deal of information about him over the years that will have influenced our opinions – not least from the [original Louis Theroux programme](#) from 2000 that was reexamined in that recent BBC documentary. Savile's vague resemblance to the Child Catcher from the film Chitty Chitty Bang Bang is convenient but ultimately misleading, and the way it reinforces the idea of what a paedophile might 'look like' is unfortunate; not least because it helps to sustain a low-level belief in the 'science' of physiognomy, despite its tendency to crumble under the slightest cross examination.

Share this article