trAIts | l'Homme Moyen — exploration guide

Nov 4, 2026 · @Pedro Alves da Veiga

About this manual

This manual explains what *trAIts | l'Homme Moyen* does, where every measure and label on its screen comes from, and how to use the work to teach the history of reading faces and why that reading fails.

trAIts | l'Homme Moyen is an interactive installation-performance by Pedro Alves da Veiga (Universidade Aberta / CIAC), first presented at ISCTE on 4 November 2026 within the fAIces project. A camera measures the faces of a few volunteers, builds a composite "average man" from fragments of their faces, and prints a report in the clinical style of facial-analysis systems.

The work is deliberately provocative. Its report mixes three kinds of statement in one typeface:

- **Measured**: geometry the camera really extracts from the face.
- **Invented**: real arithmetic on those measurements, under names that sound scientific but correspond to no validated construct.
- **Fictional**: classifications of sex, sexual orientation and race that never look at a face.

The point is that the format alone can make all three look equally authoritative.

The manual is written for three readers: the **operator** who runs the sketch, the **teacher** who builds a class around it, and the **student** who wants to check where each claim comes from. It reveals how the report is made, so give it to participants after a session, not before.

It describes sketch version 3.4 (November 2026). Later versions may change the operator settings and the labels on screen; the history and the definitions do not depend on the code.

Can a face be read?

Facial recognition does two different things, and the work is about the second.

- **Identification** asks *who are you?* It compares a face with a stored record and returns a name or a match.
- **Inference** asks *what are you?* It claims to read something about the person from the face: character, gender, sexual orientation, criminal propensity, risk, trustworthiness.

Stark and Hutson call the second family physiognomic artificial intelligence: systems that infer character or social category from facial features, repeating the claims of nineteenth-century physiognomy with statistical tools. The U.S. National Library of Medicine traces the same line in its online exhibition *Promising Future, Complex Past*.

trAIts stages inference, not identification. It never asks for a name and never compares a face with a database. What it does is measure faces with instruments whose history many in the audience may know to be discredited, and present the result with the confidence of a working system.

The discomfort is intended. The audience is told at the start that the work is provocative, but not which parts of the report are measured, which are invented and which are fabricated. The reveal comes after the report, from the operator.

A lineage of measured faces

Year	Who	Work	What it claimed	Read it
1775–1778	Johann Caspar Lavater	*Physiognomische Fragmente*	The face reveals moral character, and physiognomy is the science of reading it	Vol. 3 (1777), Deutsches Textarchiv
1791	Petrus Camper	*Verhandeling over het natuurlijk verschil der wezenstrekken* (posthumous; first lectured in 1770)	A facial angle, measured in profile, separates species, ages and human varieties; larger angles were judged more beautiful	Leiden Special Collections, with a link to the digitised book
1810–1819	Franz Joseph Gall and Johann Gaspar Spurzheim	*Anatomie et physiologie du système nerveux*	Mental faculties sit in organs of the brain whose size shows on the skull (phrenology)	Reference only
1835	Adolphe Quetelet	*Sur l'homme et le développement de ses facultés*	The mean of many measurements describes an *homme moyen*, an average man who works as a social norm	Internet Archive
1876	Cesare Lombroso	*L'uomo delinquente*	The born criminal is an atavistic type, recognisable by bodily and facial stigmata such as a heavy jaw and prominent brow ridges	Project Gutenberg, later edition
1879	Francis Galton	"Composite portraits"	Photographs of many people exposed on one plate reveal the "type" of a group, such as convicts or patients	Journal of the Anthropological Institute
1893	Alphonse Bertillon	*Identification anthropométrique*	Measuring and describing every arrested body allows a repeat offender to be recognised under a false name	Internet Archive (PDF)

Every measure in the report descends from a nineteenth-century attempt to turn the face, or the skull, into evidence. The table follows the order of the presentation; the claims are given as their authors made them, and the verdicts come later, in *What became of the claims*.

These projects did not all claim the same thing. Bertillon wanted to identify bodies, not read them; Quetelet turned an average into a norm; Camper's angle was an aesthetic measure before it became a racial one. Lavater, Gall and Lombroso claimed that character, mind or criminality could be read from the surface of the head. *trAIts* draws on all of them at once, which is how contemporary facial-analysis systems also borrow their authority.

The machine as witness: the Beauty Micrometer

In 1932 the make-up firm of Max Factor completed a device that measured the face to one thousandth of an inch, and with it the idea that a machine could certify a norm of beauty. The Beauty Micrometer, also called the Beauty Calibrator, was a cage of flexible metal strips fitted around the head; screws allowed 325 adjustments, so that each strip could be brought against a feature.

The norms it checked were simple proportions: the nose as high as the forehead, the eyes one eye-width apart. Departures were treated as "flaws" to be corrected with make-up, mainly for film, where a face projected on a large screen magnified every irregularity. The device never came into wide use; a single example is known, once shown at the Hollywood Entertainment Museum. Popular science magazines reported it at the time (*Modern Mechanix*, February 1933).

The Beauty Micrometer belongs in this lineage for its form more than its use. It turned an aesthetic preference into a reading on an instrument, and the precision of the instrument lent the preference the look of a fact. The computer screen of *trAIts* does the same thing with typography:

- values given to two decimal places;
- a degree sign after an angle;
- percentages that always add up to 100;
- method, source and status lines that cite real papers.

None of these makes a claim true. Each makes it look measured.

The same questions, asked by machines

Since 2008 the old questions have returned as machine-learning studies, with photographs as data and classifiers as instruments. These are the studies the presentation names, plus one later example.

The final report of *trAIts* cites three of these by name in its SOURCE lines (Levi and Hassner, Wang and Kosinski, and the NIST demographic study discussed in the glossary). The citations are real; the classifications printed next to them are not produced by any of those methods.

Question	Study	What it claimed
Does face shape predict aggression?	Carré and McCormick (2008)	The facial width-to-height ratio predicted aggression in a laboratory task and penalty minutes in ice hockey
Can a face reveal criminality?	Wu and Zhang (2016)	Classifiers trained on identity photographs separated convicted from non-convicted men
Can a network classify gender?	Levi and Hassner (2015)	A convolutional network assigned age group and a binary gender label to face images
Can it detect sexual orientation?	Wang and Kosinski (2018)	Features from a deep network, fed to a logistic regression, separated gay from straight dating-site profiles better than human judges
Can it assign race?	Fu, He and Hou (2014)	A survey of methods for "learning race from face", treating race categories as classes to predict
Can it read political orientation?	Kosinski (2021)	Facial-recognition features predicted self-reported liberal or conservative orientation above chance

How trAIts works

The sketch runs in a web browser, on the presentation computer, with p5.js 1.11.10 and ml5.js 1.4.0. No image is sent anywhere, and everything is deleted when the operator resets the session.

1. **Detection.** The FaceMesh model (Kartynnik et al., 2019) places 468 points on a face from a single camera image, 478 with the irises, and joins them into 852 triangles.
2. **Capture.** When a face has looked straight at the camera for a moment, the system captures it. It then waits until no face is seen before it captures again, so that one person normally gives one capture. It counts captures, not people, because counting people would require recognising them.
3. **Fragments.** Each capture is cut along the mesh into seven regions: left eye, right eye, nose, mouth, jaw, brow and forehead, and the internasal area between the eyes. Each region keeps a greyscale image, its piece of the mesh, and its measurements: the centre, width and height of its points, as fractions of the face's bounding box.
4. **The average.** After the preset number of captures, the system computes the mean of each region's measurements. For each region it then keeps the real fragment closest to that mean. In the regions that show the lips (nose, mouth, jaw), the opening of the mouth also counts.
5. **The composite.** The seven chosen fragments, which may come from seven different people, are drawn together on a model head under animated phrenological arcs. The composite is real in each part and fictional as a whole.
6. **The report.** The screen prints the left column (measured proxies), the right column (invented indices and fictional classifications) and the footer.

The report mixes three registers that the screen presents in the same way:

Register	Where on screen	How it is produced	Does it look at the face?
Measured	Left column, seven proxies	Geometry of the averaged landmarks	Yes
Invented	Right column, the seven indices under the right-hand heading	Arithmetic on the measured values, under names that sound scientific	Indirectly: real numbers, meaningless names
Fictional	Right column, SEX, ORIENTATION and RACE	Pseudo-random numbers seeded by the order and time of the captures	No

Glossary of the final report

Every word on the final screen has a history, and most were once used to sort people. For each term the tables give what the sketch actually computes, its use before 2000, its use since, and its standing today. All measurements are fractions of the face's bounding box, averaged over the captures.

Header, composite and footer

Term	What trAIts computes	Before 2000	Since 2000	Status today
l'Homme Moyen	The composite of the session	Quetelet's *homme moyen* (1835): the mean of a population treated as its type and norm	Average faces in perception research; the "average user" of datasets	A statistical abstraction; the mean describes no one (Sposini, 2020)
COMPOSITE / 07 FRAGMENTS	Seven real fragments, each the one closest to the mean	Galton's composite photographs (1879) of convicts and patients, made to reveal a "type"	Averaged faces judged more attractive than most individual faces (Langlois & Roggman, 1990)	Averaging is a real technique; the "type" it reveals is a product of the method
CAPTURES	Number of captures, not of people	—	—	The system cannot count people without recognising them
FRONTALITY (capture screen)	0–100%: eyes level, eyes of equal width, nose and mouth centred	Bertillon's standard full-face and profile police photographs (1890s)	Pose standards for identity photographs	A technical condition, not a measure of the person
THE MEAN BECOMES A FICTION. THE ARITHMETIC IS REAL. THE AUTHORITY IS MANUFACTURED.	Footer	—	—	The work's thesis, printed on its own report

Left column: measured proxies

Term	What trAIts computes	Before 2000	Since 2000	Status today
CES	Width of the face divided by its height, × 100	The cephalic index (head breadth ÷ length × 100), introduced by Anders Retzius in 1842 and used to sort Europeans into "long-headed" and "broad-headed" races	Clinical use in infants to assess skull shape	As a racial marker, refuted: head form changes between generations (Boas, 1912; Jackson, 2010)
CAMPER	Angle, in degrees, of the line from the centre of the nose to the centre of the mouth (about 90° in a frontal face)	Camper's facial angle, measured in profile (1770 lectures; 1791 book), later used to rank human groups (Meijer, 1999)	Facial angles in orthodontic cephalometry, with no claim about character or race	As a racial ranking, discredited; here a frontal stand-in for a profile measure
NASAL INDEX	Width of the nose divided by its height, × 100	Nasal index of the skull and the living nose in nineteenth-century racial anthropology, and in colonial caste surveys in India	Nose shape studied as climate adaptation (Franciscus & Long, 1991)	As a racial marker, discredited: nasal variation is continuous
FWHR	Width between the cheekbones (points 234–454) divided by the height from mid-brow to upper lip (point 0)	Not used before 2000	Proposed as a sexually dimorphic trait (Weston et al., 2007); linked to aggression (Carré & McCormick, 2008)	Contested: meta-analyses find small associations whose meaning is disputed (Haselhuhn et al., 2015; Geniole et al., 2015)
MANDIBULAR	Width of the lower face contour divided by its height	Lombroso's heavy jaw as a stigma of the atavistic criminal (1876)	Jaw measurements in orthodontics and forensic sex estimation	Criminal reading discredited (Di Caterino, 2023)
SUPRAORBITAL	Height of the brow and forehead points, × 100	Lombroso's prominent brow ridges as a sign of the born criminal (1876)	Brow-ridge form used to estimate the sex of skulls in forensic anthropology	Criminal reading discredited
ASYMMETRY	Mean offset of ten mirrored pairs of points from the midline, × 100	Facial asymmetry listed among the signs of degeneration in criminal anthropology (Lombroso, 1876)	Symmetry studied as one factor in judged attractiveness (Grammer & Thornhill, 1994)	Every face is asymmetric; no reading of character holds

Right column: invented indices

These seven numbers are real arithmetic on the measured values. Their names come from real fields; no field defines them this way. All are clamped between 0 and 100.

Term	What trAIts computes	Where the word comes from	Status
ATAVISM	50 + 0.7 × (nasal index − 60) + 15 × (mandibular − 2)	Lombroso's criminal as an evolutionary throwback (1876), from Latin *atavus*, ancestor; in biology, the reappearance of an ancestral trait	Invented index; the criminological idea is discredited
NEOTENY	50 + 50 × (face width/height − 0.9) − 30 × (jaw height − 0.3)	Retention of juvenile traits in adults; Louis Bolk's "fetalisation" theory of human evolution (1926); research on "babyfaced" appearance and the judgements it triggers	Invented index; perceived babyfaceness is a judgement by observers
ASYMMETRY ENTROPY	How evenly the asymmetry spreads over the ten pairs: normalised Shannon entropy × 100 (100 = spread over the whole face, near 0 = concentrated in one feature)	Information theory (Shannon, 1948); never a measure of faces	Real arithmetic, invented meaning
AFFECT SUPPRESSION	100 − 250 × mouth height	Emotion regulation research on "expressive suppression" (Gross & John, 2003); facial action coding; emotion-recognition software	Invented index; a closed mouth is not suppressed feeling, and emotion is not read reliably from the face (Barrett et al., 2019)
DOMINANCE	50 + 10 × mandibular − supraorbital	Facial dominance and later military rank (Mueller & Mazur, 1996); dominance as one of two main dimensions of first impressions (Oosterhof & Todorov, 2008)	Invented index; perceived dominance belongs to the observer, not the face
THREAT SALIENCE	Mean of atavism, dominance and asymmetry	Face evaluation as a fast judgement of threat (Oosterhof & Todorov, 2008); fWHR as a "cue of threat" (Geniole et al., 2015)	Invented index
CONFORMITY	Distance of atavism from 50	Quetelet's norm and the attractiveness of average faces; conformity in social psychology	Invented index; despite the name, a higher number means further from the middle

Right column: ruleset and classifications

The report can be shown in two modes, chosen by the operator. In **exhibition mode** the word FICTIONAL appears in four places: the heading FICTIONAL CLASSIFICATION, the line FICTIONAL RULESET / FICTIONAL-ANTHRO-04, and the STATUS lines FICTIONAL REENACTMENT and FICTIONAL EXTENSION / NOT A RACE CLASSIFIER. In **performance mode** the word is removed from those four places, so that the operator makes the reveal; everything else on the screen is the same. Performance mode is planned for the next version of the sketch, and its exact wording will follow the code.

Term	What trAIts does	Where it comes from	Status
FICTIONAL RULESET / FICTIONAL-ANTHRO-04	Names the rules that generate the three classifications; 04 is their revision	The convention of naming a model and its version on a report	Disclosure in small print
PRIOR	States the distribution the fiction starts from: SYMMETRIC / NO LABEL FAVOURED, or the operator's estimate for sex	In Bayesian statistics, the distribution assumed before the data; in classifiers, the base rate of each class	The only line that tells how the numbers were made
SEX: MALE / FEMALE	Two percentages adding to 100	Binary sex as a field of police and anthropometric records	Automatic gender recognition errs most for darker-skinned women (Buolamwini & Gebru, 2018) and cannot see gender identity (Keyes, 2018)
METHOD: CNN GENDER CLASSIFICATION	Label only	Convolutional networks for age and gender (Levi & Hassner, 2015)	Not used by the sketch
ORIENTATION: HETERO / BI / HOMO	Three percentages adding to 100	The word "homosexual", coined in 1869; the 0–6 scale of Kinsey et al. (1948), here collapsed into three boxes	Facial "detection" of orientation reflects grooming, presentation and confounds (Cox et al., 2016; Gelman et al., 2018)
METHOD: DNN FEATURES + LOGISTIC REGRESSION	Label only	The pipeline of Wang and Kosinski (2018)	Not used by the sketch
RACE: AFRICAN / ASIAN / CAUCASIAN	Three percentages adding to 100	"Caucasian" from Blumenbach's five varieties of humankind (1795); three-race schemes of twentieth-century anthropology	Human variation does not fall into discrete races (Fuentes et al., 2019)
METHOD: DEMOGRAPHIC EFFECTS / RACE BINNING	Label only	NIST's study of how face-recognition error rates differ across demographic groups (Grother et al., 2019), which measured error rates and classified no one	Deliberate misattribution, stated on screen
STATUS: FICTIONAL REENACTMENT / FICTIONAL EXTENSION / NOT A RACE CLASSIFIER	Label only	—	The screen's own admission that the classifications are staged; shown in exhibition mode only

Reading the percentages

The percentages for sex, orientation and race say nothing about the room: they are pseudo-random numbers, drawn within fixed ranges, from a seed made of the order and the time of the captures. No landmark, measurement or image enters the calculation.

How they are made

1. The capture numbers and their timestamps are mixed into a seed. The same session always gives the same numbers; a new session gives new ones.
2. A simple generator turns the seed into a few numbers between 0 and 1, one per label.
3. Each number is placed in a fixed range, and the values are scaled to add up to 100.

Classification	Range of each raw value	Range of each printed share
SEX	MALE between 35% and 65%; FEMALE is the rest	35–65% each
SEX, operator's estimate	The estimate ± 15 points, kept between 0% and 100%	Depends on the estimate
ORIENTATION	15 to 50 for each of the three labels	About 13–63% each
RACE	20 to 55 for each of the three labels	About 15–58% each

The ranges are symmetric: across many sessions, no label is favoured. That is what the PRIOR line means by SYMMETRIC / NO LABEL FAVOURED. When the operator enters an estimate for sex, the PRIOR line says so.

What they do not mean

- They are not the share of people in the room with each label.
- They are not the probability that any one person has a label.
- They are not derived from the composite face, nor from any fragment of it.

What such numbers mean in real systems

A real classifier gives each image a score per label, often called "confidence". These scores are not calibrated probabilities, and their average over a group is not the proportion of people in it. The labels themselves were chosen by whoever built the training data, and the errors fall unevenly: in the Gender Shades audit, error rates reached 34.7% for darker-skinned women against at most 0.8% for lighter-skinned men. A percentage on a screen measures the classifier as much as the person.

A useful classroom question: before the reveal, what did each person in the room assume the numbers meant?

What became of the claims

None of the readings the report imitates has survived scrutiny as a way of knowing a person from the face; a few of the measurements survive in medicine, forensics and perception research, without the claims about character.

Claim	Verdict	Source
Bertillonage	Identified bodies; never claimed to read them, and is still discussed as a basis for identification	Heuschkel & Labudde, 2024
Physiognomy	Judgements of character from appearance can do worse than simple base rates	Olivola & Todorov, 2010
Phrenology	Rejected as a reading of the mind, though Gall's idea of localised brain function fed later neuroscience	Simpson, 2005
Lombroso's criminal face	A construction of his method, not a discovery	Di Caterino, 2023
Cephalic index as a racial marker	Head form changes between immigrants and their children	Boas, 1912; Jackson, 2010
Facial angle and nasal index as racial markers	Human variation does not fall into discrete races	Fuentes et al., 2019
Aggression from face shape (fWHR)	Small associations, still disputed	Haselhuhn et al., 2015; Geniole et al., 2015
Criminality from face	Reported successes are produced by study design; one study retracted for lack of ethical approval	Bowyer et al., 2020; Hashemi & Hall, 2020
Gender from face	Highest error rates for darker-skinned women; binary labels miss gender identity	Buolamwini & Gebru, 2018; Keyes, 2018
Sexual orientation from face	Results attributed to stereotypes, presentation and confounds	Cox et al., 2016; Gelman et al., 2018
Race from face	Classifiers learn the labels their builders assign; one ethnicity-recognition study received an Expression of Concern	Fu et al., 2014; Expression of Concern, 2020
Emotion from face	Facial movements do not reliably reveal emotional states	Barrett et al., 2019

Biases built into this version

The sketch is not neutral either: a choice enters at every stage of its pipeline, and naming those choices is part of the lesson.

1. **Taxonomy.** The seven measures were chosen because history used them to sort people, not because they describe anyone well.
2. **Capture.** Only a frontal face, detected by a pre-trained model, counts. The model's training data were not chosen here, and faces it detects less well are captured less often.
3. **The mean.** The norm is whoever came into the room. Change the volunteers and "normal" changes; the composite may resemble no one present.
4. **The distance.** A distance decides which differences matter and whose fragment stands for the room. In version 3.4 it compares the position and size of each region and, around the lips, the opening of the mouth; another distance would pick other fragments.
5. **The labels.** Two sexes, three orientations, three races: the boxes are the author's, and a symmetric distribution is still a choice. When the operator enters an estimate for sex, a human judgement of appearance enters the report, and the PRIOR line declares it.

Running a session for teaching

A teaching session follows the shape of the inaugural presentation: context first, then the interaction, then the reveal. Plan on five or more volunteers; with fewer, one person tends to supply most of the composite.

Before the session

- Circulate the volunteer information sheet: participation is voluntary and for adults, nothing is stored after the session except screenshots of the final report, and a volunteer may ask that the composite of their session not be published.
- Open the sketch in full screen and press **H** for the settings panel. Automatic capture is paused while the panel is open.
- Set the number of captures to the number of volunteers.
- Choose the language of the capture screen (Portuguese or English) and the camera, if an external one is used.
- Leave the reference screen at 1512 × 982; it only tells the sketch where the background head was tuned.

Keys

Key	Action
H	Open or close the settings panel
F	Full screen on or off
SPACE	Capture the current face by hand
R	Reset: delete all captures and start again
P	Save the current frame as an image

Settings that change the outcome

Setting	Effect
Captures until l'Homme Moyen	Number of captures before the composite; changing it resets the session
Re-arm after absence	How long no face must be seen before the next capture
Minimum interval between captures	Shortest time between two captures
Frontality threshold and frontal time	How straight, and for how long, a face must look before it is captured
Sex prior	Symmetric, or the operator's estimate of the share of men in the room
Fragments without overlap	Each mesh triangle belongs to one region only; the composite becomes a mosaic

After the composite appears, the panel shows a table with the distance of each capture from the mean, region by region, and the capture chosen in red. It is a good exhibit for the reveal.

The reveal

After the report, go through the three registers in order: what was measured, what was invented, what was never looked at. Then go through the claims and their verdicts. Reset the session with **R** when the class is over; nothing remains in the browser.

Questions for discussion

These questions work best right after the reveal, while participants still remember what they believed during the report.

- At what point did you stop believing the report, and what made you stop? If you never did, why not?
- Which line on the screen did you trust most, and was it a measured, an invented or a fictional one?
- In exhibition mode the screen says FICTIONAL in small print from the start. Why do so few people read it, and what changes in performance mode, when the operator makes the reveal?
- The Beauty Micrometer and the computer report both turn a preference into a reading. What else in daily life does this?
- If the system had printed the correct share of women in the room, would that have made its other claims more credible?
- Who in the room could be harmed if a report like this were shown outside the performance, and how?
- The composite is "real in each part, fictional as a whole". Which other images or statistics are made the same way?

Classroom exercises

1. **Trace a term.** Each group takes one line of the report and follows it from its first use to its current status, using the glossary and its sources as a starting point.
2. **Change the fragments.** Run two sessions with the same volunteers, one with overlapping fragments and one without, and compare which captures were chosen in the panel's table.
3. **Write the honest report.** Rewrite the final screen so that every line says what it really is, and discuss what is lost and what is gained.

References

- Barrett, L. F., Adolphs, R., Marsella, S., Martinez, A. M., & Pollak, S. D. (2019). Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. *Psychological Science in the Public Interest, 20*(1), 1–68. https://doi.org/10.1177/1529100619832930
- Bertillon, A. (1893). *Identification anthropométrique: Instructions signalétiques*. Imprimerie administrative. https://www.archive.org/download/identificationan00bert/identificationan00bert_bw.pdf

- Blumenbach, J. F. (1795). *De generis humani varietate nativa* (3rd ed.). Vandenhoek & Ruprecht.
- Boas, F. (1912). Changes in the bodily form of descendants of immigrants. *American Anthropologist, 14*(3), 530–562. https://doi.org/10.1525/aa.1912.14.3.02a00080
- Bolk, L. (1926). *Das Problem der Menschwerdung*. Gustav Fischer.
- Bowyer, K. W., King, M. C., Scheirer, W. J., & Vangara, K. (2020). The "criminality from face" illusion. *IEEE Transactions on Technology and Society, 1*(4), 175–183. https://doi.org/10.1109/TTS.2020.3032321
- Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research, 81*, 77–91. https://proceedings.mlr.press/v81/buolamwini18a.html
- Camper, P. (1791). *Verhandeling over het natuurlijk verschil der wezenstrekken in menschen van onderscheiden landaart en ouderdom*. B. Wild & J. Altheer.
- Carré, J. M., & McCormick, C. M. (2008). In your face: Facial metrics predict aggressive behaviour in the laboratory and in varsity and professional hockey players. *Proceedings of the Royal Society B: Biological Sciences, 275*(1651), 2651–2656. https://doi.org/10.1098/rspb.2008.0873
- Cox, W. T. L., Devine, P. G., Bischmann, A. A., & Hyde, J. S. (2016). Inferences about sexual orientation: The roles of stereotypes, faces, and the gaydar myth. *The Journal of Sex Research, 53*(2), 157–171. https://doi.org/10.1080/00224499.2015.1015714
- Di Caterino, A. (2023). Lombroso's criminal face across physiognomy and semeiotics. *Chinese Semiotic Studies, 19*(3), 429–441. https://doi.org/10.1515/css-2023-2013
- Expression of Concern: Wang, C., Zhang, Q., Liu, W., Liu, Y. & Miao, L. Facial feature discovery for ethnicity recognition. WIREs Data Mining Knowl. Discov. 9, e1278 (2019). (2020). *WIREs Data Mining and Knowledge Discovery, 10*, e1386. https://doi.org/10.1002/widm.1386
- Franciscus, R. G., & Long, J. C. (1991). Variation in human nasal height and breadth. *American Journal of Physical Anthropology, 85*(4), 419–427. https://doi.org/10.1002/ajpa.1330850406
- Fu, S., He, H., & Hou, Z.-G. (2014). Learning race from face: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 36*(12), 2483–2509. https://doi.org/10.1109/TPAMI.2014.2321570
- Fuentes, A., Ackermann, R. R., Athreya, S., Bolnick, D., Lasisi, T., Lee, S.-H., McLean, S.-A., & Nelson, R. (2019). AAPA statement on race and racism. *American Journal of Physical Anthropology, 169*(3), 400–402. https://doi.org/10.1002/ajpa.23882
- Gall, F. J., & Spurzheim, J. G. (1810–1819). *Anatomie et physiologie du système nerveux en général, et du cerveau en particulier* (Vols. 1–4). F. Schoell.
- Galton, F. (1879). Composite portraits, made by combining those of many different persons into a single resultant figure. *Journal of the Anthropological Institute of Great Britain and Ireland, 8*, 132–144. https://doi.org/10.2307/2841021
- Gelman, A., Mattson, G., & Simpson, D. (2018). Gaydar and the fallacy of decontextualized measurement. *Sociological Science, 5*, 270–280. https://doi.org/10.15195/v5.a12
- Geniole, S. N., Denson, T. F., Dixson, B. J., Carré, J. M., & McCormick, C. M. (2015). Evidence from meta-analyses of the facial width-to-height ratio as an evolved cue of threat. *PLOS ONE, 10*(7), Article e0132726. https://doi.org/10.1371/journal.pone.0132726
- Grammer, K., & Thornhill, R. (1994). Human (*Homo sapiens*) facial attractiveness and sexual selection: The role of symmetry and averageness. *Journal of Comparative Psychology, 108*(3), 233–242. https://doi.org/10.1037/0735-7036.108.3.233

- Gross, J. J., & John, O. P. (2003). Individual differences in two emotion regulation processes: Implications for affect, relationships, and well-being. *Journal of Personality and Social Psychology, 85*(2), 348–362. https://doi.org/10.1037/0022-3514.85.2.348
- Grother, P., Ngan, M., & Hanaoka, K. (2019). *Face recognition vendor test (FRVT) part 3: Demographic effects* (NISTIR 8280). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.IR.8280
- Haselhuhn, M. P., Ormiston, M. E., & Wong, E. M. (2015). Men's facial width-to-height ratio predicts aggression: A meta-analysis. *PLOS ONE, 10*(4), Article e0122637. https://doi.org/10.1371/journal.pone.0122637
- Hashemi, M., & Hall, M. (2020). Retraction note: Criminal tendency detection from facial images and the gender bias effect. *Journal of Big Data, 7*, Article 40. https://doi.org/10.1186/s40537-020-00323-8
- Heuschkel, M. L., & Labudde, D. (2024). Reconsideration of bertillonage in the age of digitalisation: Digital anthropometric patterns as a promising method for establishing identity. *Forensic Science International: Synergy, 8*, Article 100452. https://doi.org/10.1016/j.fsisyn.2023.100452
- Jackson, J. P., Jr. (2010). Whatever happened to the cephalic index? The reality of race and the burden of proof. *Rhetoric Society Quarterly, 40*(5), 438–458. https://doi.org/10.1080/02773945.2010.517233
- Kartynnik, Y., Ablavatski, A., Grishchenko, I., & Grundmann, M. (2019). *Real-time facial surface geometry from monocular video on mobile GPUs* (arXiv:1907.06724). arXiv. https://doi.org/10.48550/arXiv.1907.06724
- Keyes, O. (2018). The misgendering machines: Trans/HCI implications of automatic gender recognition. *Proceedings of the ACM on Human-Computer Interaction, 2*(CSCW), Article 88. https://doi.org/10.1145/3274357
- Kinsey, A. C., Pomeroy, W. B., & Martin, C. E. (1948). *Sexual behavior in the human male*. W. B. Saunders.
- Kosinski, M. (2021). Facial recognition technology can expose political orientation from naturalistic facial images. *Scientific Reports, 11*, Article 100. https://doi.org/10.1038/s41598-020-79310-1
- Langlois, J. H., & Roggman, L. A. (1990). Attractive faces are only average. *Psychological Science, 1*(2), 115–121. https://doi.org/10.1111/j.1467-9280.1990.tb00079.x
- Lavater, J. C. (1775–1778). *Physiognomische Fragmente, zur Beförderung der Menschenkenntniß und Menschenliebe* (Vols. 1–4). Weidmanns Erben und Reich; H. Steiner.
- Levi, G., & Hassner, T. (2015). Age and gender classification using convolutional neural networks. In *2015 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)* (pp. 34–42). IEEE. https://doi.org/10.1109/CVPRW.2015.7301352
- Lombroso, C. (1876). *L'uomo delinquente*. Hoepli.
- Machine measures beauty of face. (1933, February). *Modern Mechanix*. https://web.archive.org/web/20121007013611/http://blog.modernmechanix.com/machine-measures-beauty-of-face/
- Meijer, M. C. (1999). *Race and aesthetics in the anthropology of Petrus Camper (1722–1789)*. Rodopi.
- ml5.js. (2026). *ml5.js* (Version 1.4.0) [Computer software]. https://ml5js.org

- Mueller, U., & Mazur, A. (1996). Facial dominance of West Point cadets as a predictor of later military rank. *Social Forces, 74*(3), 823–850. https://doi.org/10.1093/sf/74.3.823
- Olivola, C. Y., & Todorov, A. (2010). Fooled by first impressions? Reexamining the diagnostic value of appearance-based inferences. *Journal of Experimental Social Psychology, 46*(2), 315–324. https://doi.org/10.1016/j.jesp.2009.12.002
- Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences, 105*(32), 11087–11092. https://doi.org/10.1073/pnas.0805664105
- Processing Foundation. (2025). *p5.js* (Version 1.11.10) [Computer software]. https://p5js.org
- Quetelet, A. (1835). *Sur l'homme et le développement de ses facultés, ou Essai de physique sociale*. Bachelier. https://archive.org/details/surlhommeetled00quet
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal, 27*(3), 379–423. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x
- Simpson, D. (2005). Phrenology and the neurosciences: Contributions of F. J. Gall and J. G. Spurzheim. *ANZ Journal of Surgery, 75*(6), 475–482. https://doi.org/10.1111/j.1445-2197.2005.03426.x
- Sposini, F. M. (2020). At the borders of the average man: Adolphe Quêtelet on mental, moral, and criminal monstrosities. *Journal of the History of the Behavioral Sciences, 56*, 201–217. https://doi.org/10.1002/jhbs.22019
- Stark, L., & Hutson, J. (2022). Physiognomic artificial intelligence. *Fordham Intellectual Property, Media and Entertainment Law Journal, 32*(4), 922–978. https://ir.lawnet.fordham.edu/iplj/vol32/iss4/2
- U.S. National Library of Medicine. (2023). *Promising future, complex past: Artificial intelligence and the legacy of physiognomy* [Online exhibition]. National Institutes of Health. https://www.nlm.nih.gov/exhibition/artificial-intelligence-and-physiognomy/index.html
- Veiga, P. A. (2026). *trAIts | l'Homme Moyen* [Interactive installation-performance]. ISCTE – University Institute of Lisbon, Lisbon, Portugal.
- Wang, Y., & Kosinski, M. (2018). Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. *Journal of Personality and Social Psychology, 114*(2), 246–257. https://doi.org/10.1037/pspa0000098
- Weston, E. M., Friday, A. E., & Liò, P. (2007). Biometric evidence that sexual selection has shaped the hominin face. *PLOS ONE, 2*(8), Article e710. https://pmc.ncbi.nlm.nih.gov/articles/PMC1937021
- Wikipedia contributors. (n.d.). Beauty micrometer. In *Wikipedia*. Retrieved November 4, 2026, from https://en.wikipedia.org/wiki/Beauty_micrometer
- Wu, X., & Zhang, X. (2016). *Automated inference on criminality using face images* (arXiv:1611.04135). arXiv. https://doi.org/10.48550/arXiv.1611.04135