ANOMALOUS GENDER PRESENTATION DETECTED
DECONSTRUCTING GENDER CLASSIFICATION WITH DIGITAL MEDIA

by

SAMUEL A. POSNER

**SUBMITTED TO POMONA COLLEGE IN PARTIAL FULFILLMENT
OF THE DEGREE OF BACHELOR OF ARTS**

PROFESSOR TRAN

DECEMBER 11, 2015

Computer vision is a key weapon in the technological arsenal of contemporary surveillance and biopolitics, and facial recognition algorithms are among the latest in a long history of technologies that have worked to functionalize the concept of human identity. Digital technologies of biometric identification reproduce the ideology of their nineteenth-century forbearers; namely, that identity can be read unproblematically from the human body. This principle is the founding myth of a military and academic industrial complex that produces algorithms to divine identifying information – race, gender, sexual orientation, etc. – from images of human faces. Contrary to the positivist rhetoric of their promoters, these algorithms encode the biases of their human designers and, inevitably, erase non-normative bodies. *Anomalous Gender Presentation Detected* (*AGPD*) scrutinizes a subset of facial recognition technologies that classify faces based on gender. An interactive, web-based experience, *AGPD* draws upon Foucault's biopolitics and Žižek's theory of human subjectivity to deconstruct the human bias that underlies supposedly objective algorithms and to counter the dehumanizing effects of opaque surveillance technology.

Contemporary biometric technologies can trace their origins to criminal identification systems developed in the 1870s. The results of what Shoshana Magnet calls a "nineteenth-century desire to force the body to speak the truth of its identity,"[1] these early technologies were designed to support new European laws that distinguished between first-time offenders and recidivist criminals[2]. In "Identity without the Person," Giorgio Agamben argues that nineteenth-century technologies such as Alphonse Bertillon's body-measurement (Bertillonage) and Francis

[1] Magnet, *When Biometrics Fail: Gender, Race, and the Technology of Identity*, Introduction
[2] Agamben, "Identity without the Person," p. 49

Galton's fingerprint identification were sufficient to rearticulate the concept of identity in terms of biological phenomena:

> What now defines my identity and recognizability are the senseless arabesques that my inked-up thumb leaves on a card in some police station. This is something with which I have absolutely nothing to do, something with which and by which I cannot in any way identify myself or take distance from: naked life, a purely biological datum."[3]

With digital technology, biometrics became viable means of general surveillance. This proliferation of biometric identification technologies "beyond the police stations and immigration offices to penetrate the sphere of everyday life"[4] fits Foucault's model of the historical shift from the repressive power of sovereigns to the biopolitical power of modern states. As developed by Foucault, biopolitics signifies the ways in which power regulates human life at a biological level. Whereas classical regimes had the power to end individual lives, modern biopolitical power is normalizing – it "operates at the level not of the management of the individual but of the administration of aggregate populations."[5]

As mediated records of human biology, facial images lie at the intersection of biopolitics and the politics of representation. As such, facial recognition technologies do not merely encode biased, normative models of the human body – their failure to recognize non-normative bodies works to sever those bodies' semiotic grasp on "human," rendering their owners non-human a priori. The effacement of such bodies from the symbolic order by what Judith Butler calls "normative schemes of intelligibility"[6] is an invisible act of violence that serves to regulate human life at the biological level. As Magnet argues in *When Biometrics Fail: Gender, Race,*

---

[3] *Ibid*., p. 50
[4] *Ibid*., p. 51
[5] Steinmann, "Apparatus, Capture, Trace: Photography and Biopolitics"
[6] Butler, "Precarious Life"

and the Technology of Identity, the "failure" of biometric models to classify some individuals is inevitable – a condition of their possibility; a feature, not a bug.

The undecidable nature of identity is ultimately a feature of human subjectivity, not of computation. By way of Levinas, Butler argues that representation can signify the human only by failing to do so, and then making visible its own failure. In this view, "[f]or representation to convey the human […] representation must not only fail, but it must *show* its failure."[7] By locating the Real in a "gap" between the human as such and the human as represented, Butler mirrors Slavoj Žižek's ontological reading of Lacanian subjectivity. In *Less Than Nothing*, Žižek argues that the Real subject – the subject as such, or subject as object – is comprised of a "self-relating negativity,"[8] a being-less Void that is literally less than nothing. The subject can only apprehend objects in the world while this Void – its own "objectal correlate" – is withdrawn.[9] This fundamental lack on the part of the subject is what Lacan would call desire. For Lacan, the subject represses the Real cause of desire in order to forestall trauma, substituting for it an object cause; the *objet petit a*. In Žižek's interpretation, everything the subject apprehends – in Kantian terms, things as they appear – serve as objects cause of a fundamental metaphysical lack. A representation of such a subject would follow the method of psychoanalysis, bringing to light the subject's repressed desire – in this case, its lack of being. To do so, representation must lay bare its own failure to capture the Real, exposing the *relation* between subject and Real.

A psychoanalysis of human-computer interaction (HCI) would reveal that it too is structured by a fundamental gap: the ontological distinction between human and machine. New media can take advantage of this condition to comment on subjectivity more generally, and the

---

[7] *Ibid*., p. 144
[8] Žižek, "Conversations with Zizek," p. 61 in Hickman, "Zizek's Agon: The Failure of Things"
[9] Žižek, *Less Than Nothing*, Interlude 5: Correlationism and Its Discontents

digital face – a key symptom of the repressed desire for technological unity – offers a unique artistic opportunity. While facial recognition technologies only entered the mainstream in the late 2000s, faces themselves – from Apple's smiling icons to Microsoft's anthropomorphic "personal assistants" – have been staples of consumer electronics design since the 1990s. Anna Munster calls this trend in HCI "facialization," an "attempt by interface designers to eliminate the interface as a space of differentiation between human and computer."[10] In "Interfaciality: From the Friendly Face of Computing to the Alien Terrain of Informatic Bodies," Munster argues that by representing the digital as human *and* the human as digital, new media can dismantle the interface once and for all. Digital representations of the human body can promote a state of *inter*faciality that "swings psychotically between two poles: the anthropomorphization of the machine on the one hand, and the technological makeover of the organic by digital technologies on the other."[11]

The "psychosis" of interfaciality can be understood in the psychoanalytic sense: as a breakdown of the symbolic order. For Lacan, the Symbolic is the domain of the "big Other," a law-giving construct (traditionally a Father figure) that prohibits the fulfillment of desire. In HCI, this role is played by the interface – subjects imagine that, but for the crude instruments that mediate between the digital and physical worlds, they could become one with the machine. As this desire is impossible inherently, its prohibition is merely symbolic. In Žižek's words, the symbolic order serves to obfuscate "inherent impossibility in order to sustain the illusion that, were it not for the externally imposed prohibition, the full ('incestuous') gratification would be possible."[12] Without this pretense, the subject's fundamental lack becomes visible. Munster

[10] Munster, "Interfaciality: From the Friendly Face of Computing to the Alien Terrain of Informatic Bodies," p. 127
[11] *Ibid*., p. 132
[12] Žižek, "Cyberspace, Or, How to Traverse the Fantasy in the Age of the Retreat of the Big Other," p. 491

argues that in interfaciality, the digitization of the human and the humanization of the digital makes the interface superfluous, opening up the "possibility of 'the gap,' the 'unrepresentable,' the disjunctive interval between 'the body' and embodiment."[13] In Butler's language, digital representations of the human body constitute failures of representation. By making them visible, new media are able to convey the human after all.

Anomalous Gender Presentation Detected aims to humanize both the agents and subjects of facial recognition by making visible its failure to capture human gender. In the narrative of this interactive, web-based experience, a digitized version of the author classifies the genders of visitors, inviting them to "prove [him] wrong" by manipulating their faces in real-time. The arrogant demeanor of the anthropomorphic host serves to ridicule the veneer of algorithmic objectivity that surrounds computer vision – humanizing the technological agents of facial recognition. As participants modify their features (by widening their eyes or squaring their jaws, for example), the algorithm's indecisiveness[14] makes the representational failures of its underlying model obvious. If forced to oscillate rapidly between "male" and "female," the algorithm will eventually "give up" and declare that the participant's gender presentation is anomalous. By making transparent the limits of facial models as a means of gender classification, participants reclaim for themselves some of the humanity that surveillance-scale facial recognition systems silently invalidate. Or, in Žižek's terms – by apprehending the failure of their own representation, individuals encounter the fundamental lack that structures their subjectivities.

[13] Munster, "Interfaciality: From the Friendly Face of Computing to the Alien Terrain of Informatic Bodies," p. 139
[14] On-screen text flickers dramatically between "man" and "woman" as the output of the algorithm changes in real-time

AGPD is indebted to a number of artistic traditions within new media – it shares formal and rhetorical techniques with online interventions, but has more in common conceptually with traditional new media installations. Like *CV Dazzle* (Harvey, 2010 – present), *AGPD* is a web-based project that draws its rhetorical strength from engagement with actual facial recognition technology. Where *AGPD* encourages participants to discover the limits of gender classification themselves, *CV Dazzle* provides practical, DIY fashion advice for avoiding recognition in the first place – it presents makeup patterns and hairstyles that subvert normative models of the human face. Closer to *AGPD* in form is *Gay-Check Online* (NETRO, 2014), a web app that purports to analyze sexual orientation via webcam. However, it does not actually perform facial recognition, and the result is always the same: "CONGRATULATIONS, YOU ARE GAY."

Unlike in *Gay-Check Online*, in *AGPD* the physical performance of the participant is central – as a result, it has more in common conceptually with physical installations. For instance, the *Einstein's Brain Project* (Dunning and Woodrow, 1996 – 2007) consists of a series of installations that feature interactive models of human biology. In *The Crucible* (2001), participants encounter a model of the human head. As described by the project website, "the head contains pressure sensitive pads and light and sound sensors, corresponding to the 37 organs of mental and moral faculties identified by phrenologists."[15] Triggering the sensors causes images to be projected on the walls, but, as Munster observes, "the relation of sensor activation to image call-up is randomly programmed."[16] Like *AGPD*'s shifting gender classifications, these arbitrary connections between physical sensors and digital images serve to

[15] Dunning and Woodrow, "The Crucible"
[16] Munster, "Interfaciality: From the Friendly Face of Computing to the Alien Terrain of Informatic Bodies," p. 143

undermine the authority of the model at stake – in the case of *The Crucible*, that of phrenologist pseudoscience.

The production of *AGPD* was constrained by two major decisions: first, that the project should employ actual facial recognition technology, and second, that it must run as a web application. The logical choice for projects that require any sort of real-time computer vision is OpenCV, the Open Source Computer Vision library maintained by Russian research and development firm Itseez. OpenCV is a performance-critical C and C++ library that primarily targets desktop platforms. As an HTML5 web application written in JavaScript, *AGPD* cannot call C++ functions directly. Instead, it uses Google's Native Client technology to interface with a modified version of OpenCV running in a sandbox. For every frame in the webcam video stream, the JavaScript captures an image and sends it to this C++ module. Using a cascade of Haar facial features included with OpenCV[17], the module locates the participant's face within the frame. Finally, it classifies the gender of this sub-image according to a custom-trained Fisherface model and reports the result back to the web application.

Proposed in 1997 by Belhumeur, Hespanha, and Kriegman[18], Fisherfaces is a facial recognition algorithm that uses Linear Discriminant Analysis (LDA) to classify human faces. It is a modified version of Eigenfaces, a more general algorithm based on Principal Component Analysis (PCA). Both algorithms compare samples against a model generated from a set of training images, which are treated as elements of a high-dimensional vector space. Eigenfaces reduces the dimensionality of this space using PCA, creating a subspace that consists of principle components of the training data. These components – eigenvectors of the matrix of covariance in

[17] haarcascade_frontalface_alt.xml, included in the OpenCV 2.4 repository
[18] Belhumeur, Hespanha, and Kriegman, "Eigenfaces vs. Fisherfaces: Recognition Using Class Specific Linear Projection"

the training space, or "eigenfaces" – are linear combinations of the features that account for the data's variability. In other words, Eigenfaces models the features that make individual faces unique. In traditional facial recognition, the model is trained with images of targeted individuals. To "recognize" a given face – to determine if a subject is one of the targets – an image of the subject's face is projected onto the PCA subspace and compared to projections of the target faces. If the subject's projection is sufficiently similar to one of the target projections, then a match is reported.

In Fisherfaces, each training image is labeled as a member of a pre-defined class. The algorithm uses LDA to create a subspace of features that account for variability between members of these classes. The "fisherfaces" that make up this subspace model the features that distinguish between two or more classes of individuals. In the case of binary gender classification, the model is trained with images labeled as male or female. The resulting subspace has only one element – a fisherface that represents the features that determine gender. To classify the gender of a face, a sample image is projected onto this fisherface and compared to projections of the training images. The subject is determined to share the gender of the projection with which it is most similar.

OpenCV's implementation of Fisherfaces was written by Philipp Wagner, a semi-anonymous German programmer. As far as I am aware, it was Wagner who suggested using Fisherfaces for gender classification in the first place[19]. I followed his guide to the OpenCV interface[20] and wrote a Linux command line program to train a Fisherfaces model given a set of labeled training images, using it to generate a gender classifier from the "faces94" dataset

[19] Wagner, "Gender Classification with the Fisherfaces Algorithm"
[20] Wagner, "Gender Classification with OpenCV"

published by University of Essex computer scientist Libor Spacek[21]. I wrote another program to load the model and classify faces, and quickly finished a desktop prototype of my project. The next step was to convert this prototype into a Native Client module that could run in the browser.

Most client-side web applications are written in JavaScript, a high-level language that is interpreted by the browser. This process protects the user from malicious code, but causes web applications to run more slowly than programs compiled to the CPU's native assembly language. Currently only available on Google Chrome, Native Client allows web applications to run native code in a "sandbox," restricting certain operations but allowing developers to call C++ libraries like OpenCV. Native Client was designed for corporate developers with stringent security requirements – it is poorly documented and complex. The process of compiling and running my Fisherfaces code as a Native Client module was the most time-consuming and frustrating aspect of *AGPD*'s production. A demo written by Matt McDonnell illustrated running various OpenCV algorithms in the browser[22], but sidestepped an issue crucial to my project – loading a model from a data file. McDonnell loaded the Haar cascade for his face detection demo by hardcoding the data in a single massive C++ string. A two megabyte YAML file, my Fisherfaces model was far too large to load in this manner. I had to use Native Client's obscure input/ output API to map the file on my server to a local filesystem and then load the model from a separate thread of execution.

AGPD requires well-lit images to work as intended, and does not work at all on browsers other than Chrome. After receiving feedback suggesting that these technical requirements were unclear, I added a "landing page" to provide context. This page was designed to parody the

[21] Spacek, "Collection of Facial Images: Faces94"
[22] McDonnell, "OpenCV in Chrome Using NaCl"

marketing rhetoric used by surveillance technology companies – presented as the website of a fictional "Facelook LLC," the obnoxiously "Web 3.0" site sells participants on an "Advanced Gender Perception Device." From this page, participants are thanked for their interest in the technology and encouraged to try a demo. The ruse may have been too effective – after presenting the final product I realized that, but for my direction, participants may not have realized that *AGPD* is a deconstruction. Despite the haughty invitation to "prove [it] wrong," some participants were frustrated with their initial classification and did not engage the application without explicit external guidance. Additionally, some individuals were unable to influence the algorithm at all. In these cases, *AGPD* failed to make visible the failure of its internal representation, assuming (at best) the appearance of a poorly-written program and at worst reproducing the very violence it was intended to critique. While the majority of participants seemed to enjoy subverting the algorithm and were able to achieve an "anomalous gender presentation," *AGPD* has made me more conscious of my work's (in)accessibility. It also forced me to reflect on my own relationship with technology. As the self-congratulatory tone of the section on Native Client may indicate, I take a somewhat perverse pleasure in the brute mechanics of programming. I will continue to balance this intrinsic love of creation with my theoretical and political goals as I develop as an artist.

Bibliography

Agamben, Giorgio. “Identity without the Person,” *Nudities*. Trans. David Kishik and Stefan Pedatella. Stanford, CA: Stanford UP, 2010.

Belhumeur, Peter N., Joao P. Hespanha, and David J. Kriegman. "Eigenfaces vs. Fisherfaces: Recognition Using Class Specific Linear Projection." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 19.7 (1997): 711-20. Web.

Butler, Judith. “Precarious Life,” *Precarious Life: The Powers of Mourning and Violence*. London: Verso, 2004.

Dunning, Alan, and Paul Woodrow. "The Crucible." *The Einstein's Brain Project*. University of Calgary, n.d. Web. 19 Nov. 2015. <http://people.ucalgary.ca/~einbrain/new/main.html>.

Hickman, S.C. "Zizek's Agon: The Failure of Things." *Dark Ecologies*. N.p., 01 Jan. 2013. Web. 11 May 2015. <http://darkecologies.com/2013/01/01/zizeks-agon-the-failure-of-things/>.

Magnet, Shoshana. *When Biometrics Fail: Gender, Race, and the Technology of Identity*. Durham: Duke UP, 2011. Kindle.

McDonnell, Matt. "OpenCV in Chrome Using NaCl." *Matt's Site*. Web. 28 Nov. 2015. <https://www.matt-mcdonnell.com/code/NaCl/ImageProc/index.html>.

Munster, Anna. “Interfaciality: From the Friendly Face of Computing to the Alien Terrain of Informatic Bodies.” *Materializing New Media: Embodiment in Information Aesthetic*. Dartmouth: Dartmouth College Press, 2006.

Spacek, Libor. "Collection of Facial Images: Faces94." *Computer Vision Research Projects*. University of Essex. Web. 28 Nov. 2015. <http://cswww.essex.ac.uk/mv/allfaces/faces94.html>.

Steinmann, Kate. "Apparatus, Capture, Trace: Photography and Biopolitics." *Fillip*. Fillip 15, Fall 2011. Web. 29 Oct. 2015.

Wagner, Philipp. "Gender Classification with OpenCV." *OpenCV 2.4.120 Documentation*. Web. 28 Nov. 2015. <http://docs.opencv.org/2.4/modules/contrib/doc/facerec/tutorial/facerec_gender_classification.html>.

Wagner, Philipp. "Gender Classification with the Fisherfaces Algorithm." *Bytefish*. 17 Aug. 2011. Web. 28 Nov. 2015. <http://bytefish.de/blog/gender_classification/>.

Zizek, Slavoj. "Cyberspace, Or, How to Traverse the Fantasy in the Age of the Retreat of the Big Other." *Public Culture* 10.3 (1998): 483-513.

Žižek, Slavoj. *Less Than Nothing: Hegel and the Shadow of Dialectical Materialism*. New York: Verso, 2012. Kindle.

Note on citations

All the resources that I used in writing this paper are digital. Kindle books lack absolute page numbers, so I give chapter and section headings.