The Profiling Potential of Computer Vision and the Challenge of Computational Empiricism

Jake Goldenfein*

Cornell Tech

New York, New York

jg2323@cornell.edu

ABSTRACT

Computer vision and other biometrics data science applications have commenced a new project of profiling people. Rather than using ‘transaction generated information’, these systems measure the ‘real world’ and produce an assessment of the ‘world state’ - in this case an assessment of some individual trait. Instead of using proxies or scores to evaluate people, they increasingly deploy a logic of revealing the truth about reality and the people within it. While these profiling knowledge claims are sometimes tentative, they increasingly suggest that only through computation can these excesses of reality be captured and understood. This article explores the bases of those claims in the systems of measurement, representation, and classification deployed in computer vision. It asks if there is something new in this type of knowledge claim, sketches an account of a new form of computational empiricism being operationalised, and questions what kind of human subject is being constructed by these technological systems and practices. Finally, the article explores legal mechanisms for contesting the emergence of computational empiricism as the dominant knowledge platform for understanding the world and the people within it.

CCS CONCEPTS

• **Social and professional topics** → **Computing / technology policy**; *Surveillance*; • **Computing methodologies** → **Computer vision**; Machine learning; • **Applied computing** → *Law*;

KEYWORDS

Computer Vision, Data Science, Biometrics, Law and Policy, Computational Empiricism

ACM Reference Format:

Jake Goldenfein. 2019. The Profiling Potential of Computer Vision and the Challenge of Computational Empiricism. In *FAT* '19: Conference on Fairness, Accountability, and Transparency (FAT* '19), January 29–31, 2019, Atlanta, GA, USA*. ACM, New York, NY, USA, Article 4, 10 pages. <https://doi.org/10.1145/3287560.3287568>

*Dr Jake Goldenfein is a Postdoctoral Research Fellow at the Digital Life Initiative, Cornell Tech, Cornell University and a Lecturer in Law at Swinburne University of Technology.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

FAT '19, January 29–31, 2019, Atlanta, GA, USA*

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6125-5/19/01...\$15.00

<https://doi.org/10.1145/3287560.3287568>

1 INTRODUCTION

“Profiling” means different things in different traditions. The proliferation of machine learning and data science in profiling however, make the European General Data Protection Regulation (GDPR) definition especially meaningful [19, 24, 58]. Under that regime, “profiling” includes the automated evaluation of certain traits about a person [2], reminding us of the legal significance of “classification” even without subsequent discrimination. This definition also applies to an emerging class of automated image-based profiling technologies that evaluate “traits” about individuals. ‘Computer vision’ techniques ‘make sense’ of image or video data by measuring, using machine learning to transform those measurements into “representations”, and subsequently into decisions and classifications. The ‘online profiling’ discussed over the past several decades, typically concerned knowledge production from ‘transaction generated information’ created by interacting with informational environments [41], framing individuals in terms of proxies and scores [32, 45]. Computer vision (and other biometric) techniques however, use extremely granular measurements of “the real world” and increasingly deploy a logic of *exposing* or *revealing* the truth of reality and people within it. This article explores the nature of those knowledge claims and what might be the appropriate register of legal intervention to address data science’s dominance as a paradigm for knowing people.

Computer vision generates classifications and knowledge about scenes, objects, events and people. Computer vision is not the only domain of data science articulating this new epistemological stance. Similar ideas are evident in numerous data science applications, often those involving biometrics. Computer vision does however, provides a useful entry point into what is qualitatively new about the empirical and epistemological position being developed. When applied to people, computer vision has clear legal and social significance. Facial recognition for instance, one such notorious application, links a person’s portrait to some form of institutional identity like a driver’s licence [15, 53, 73]. By connecting an individual to their behaviour in physical space, facial recognition couples identity across physical and computational registers, enabling automated decisions to be articulated into physical environments. (Another example is vehicle registration and automated number plate recognition [12].) Both the intrinsic limitations and problematic applications of these technologies are well documented [49, 62], with calls for regulation not far behind [77]. However, there is already much more that computer vision can do when ‘looking at people’.¹ Beyond “identification”, computer vision can also “classify” people [41]. As well as linking your image and spatial location to

¹The term ‘looking at people’ is an umbrella for research on computer vision applied to persons, see e.g. [23].

an institutional identity, computer vision can make decisions about non-visual attributes such as what type of person you are, what you are doing, what you are feeling, and how you are likely to act in the future. These decisions may be about identity, gender, emotional state or future behaviours. But more controversial classifiers have also addressed questions of sexuality, wangkosinski criminal propensity, political position, IQ, workplace suitability, and paedophilic tendencies. That type of ‘computational physiognomy’ is one element of an emerging field of ‘personality computation’, often described as Apparent Personality Analysis (APA) or Apparent Personality Recognition (APR) [42].

Personality computation uses faces, postures, movements, actions, gestures, interactions, and emotions (as well as whether those emotions and expressions are real or fake) to infer personality traits [42]. Different analytic techniques operate on different data inputs. Some use dynamic information such as what a person *does*, or the way a physical morphology changes.² Some use static information such as how a person *looks*. Some use multiple modalities like combinations of visual and audio data in different configurations. Whatever the particular decision or classification, the premise of these techniques is generating knowledge by applying data science to images in a way that transforms *measurements* of the world in the form of pixel vectors into classifications.

The scientific merit of computer vision profiling, especially when applied to non-visual characteristics of persons, has been rigorously challenged [10, 65]. But the scientific validity of these projects is not the focus of this article. As Ian Hacking reminds us, there is a long history of producing social knowledge through new systems of measurement, both at the level of groups and individuals, and that ‘new classifications and new enumerations are inseparable’ [31]. The purpose of this article is therefore to examine those systems of enumeration and classification, try and understand the basis on which they produce knowledge about people, and explain why this presents unique challenges. The profiling exercises outlined here are thus presented as examples of the shift towards *computational empiricism* or *positivism* as a dominant knowledge system likely to have ongoing effects in juridical and political practices. While the knowledge claims of computer vision profiling are often modest or tentative, in certain papers authors have begun to argue that *only through computation* are we able to access all the information available in human faces and truly know people. In other words, only through computation can the excesses of the physical world and the people within it be perceived, processed and understood.

While the trope that technology facilitates greater knowledge of nature is a familiar one, it still requires close analysis in the specific context of computer vision. That analysis has become more imperative in the context of Trevor Paglen’s identification that ‘Human visual culture has become a special case of vision, and exception to the rule’ in a situation where actually ‘The overwhelming majority of images are now made by machines for other machines’ [63]. While computer vision “images” may be increasingly “invisible”, they still operate within the dynamics of representation and knowledge, and they still interact with human visuality. Western

thought has spent a good deal of time reflecting on ‘the connection between actuality and photo’ [9], as well as the consequences of living in a world coded by undecipherable ‘technical images’ [27]. This material offers a useful touchstone for understanding the computer vision project of ‘looking at people’, and its political and epistemological consequences. While there has been a great deal of exceptional work on the problematic application of computer vision technologies, the goal of this article is thus less about exposing bias or lack of fairness, and more about understanding the “programs” of the computer vision “apparatus”.³ Rather than challenge computer vision at the level of problematic applications, this article instead argues that, to avoid becoming a vector of data science’s politics of completion, legal narratives should highlight the fundamental contingency of computational empiricism and its knowledge claims.

This article begins with a historical and technological contextualisation of personality computation. It looks at similar exercises from the past, and tracks the history of combining photographic portraiture with statistical analyses. That section demonstrates the ongoing ideologies behind this type of profiling, and how they are serviced by ‘big data epistemology’. However, this is identified as only one of the political problems of this type of classification – another being the reverence for the forms of empiricism operationalised. Through the historical analysis, the idea is to demonstrate not how statistics and mathematics are becoming instruments of real or symbolic violence, but rather the nature of the change from visual to statistical forms of knowing. The article then seeks to sketch an outline of this emerging computational empiricism by connecting the knowledge claims of computer vision profiling to antecedent media technologies. The goal is to understand the ongoing role of “representation” within the systems of measurement at work in computer vision, and the politics of their dissimulation. Finally, the article explores possibilities for legal and technical intervention. In that section, it is argued that the narratives animating existing legal protections are becoming deficient in the face of technological practices that see humans in a new way, primarily as patterns of information. Accordingly, the article suggests that when dealing with profiling, privacy and data protection now need to focus as much on “metaphysics” as they do on “metadata” [57]. Rather than the elimination of bias or unfairness, or the total prohibition of certain applications (both of which have their place within legal armatures), the argument here is for legal mechanisms that marginalize purely computational ways of knowing. The goal is not to ‘put the genie back in the bottle’ but rather highlight and contest the ontological power of computer vision when classifying people [47].

2 COMPUTER VISION PROFILING

Computer vision is a family of technologies and practices by which computational systems come to understand and make decisions about the physical world and the people within it. When ‘looking

²See e.g. [38] where they discuss the problems of Paul Ekman’s ‘Facial Action Coding System’, which measured the way people’s faces change when they express emotions, suggesting the use of dynamic information does not necessarily solve any problems.

³These terms are intended to follow the definitions provided by Vilém Flusser in [27], the book includes a Glossary: At (83) ‘Apparatus (pi. -es): a plaything or game that simulates thought [trans. An overarching term for a non-human agency, e.g. the camera, the computer and the ‘apparatus’ of the State or of the market]; organization or system that enables something to function.’ At (84) ‘Program: a combination game with clear and distinct elements [trans. A term whose associations include computer programs, hence the US spelling].’

at people', applications range from relatively benign to highly problematic. The more problematic applications do, however, typically reveal more in terms of embedded technical, ideological, and epistemological assumptions. Forms of computational analytics have been used in personality analysis since the 1990s. But the history of combining photography with statistics to explore correlations between external appearance and internal personality goes back much further. Kate Crawford reminds us that, 'If algorithms present us with a new knowledge logic, then it is important to consider the contours of that logic and by which histories and philosophies it is most strongly shaped.' [17] Connecting computer vision projects to previous photographic ones thus demonstrates what lurks behind these systems of human classification, and how previous claims about the 'mechanical objectivity' of the camera have been displaced into the computational objectivity of statistics. But rather than demonstrate how data science can be a vector for ideology, the goal is to demonstrate what else is at stake in the computer vision project when 'looking at people' - in particular the proliferation of computational empiricism, the difference between visual and statistical knowledge, and the idea that data science can know people better than they know themselves.

2.1 Classification and Ideology

Criminal portraits were one of the first datasets systematically produced after the invention of the daguerreotype photographic process [43]. Founder of British Eugenics, bio-metrician, and Charles Darwin's cousin, Sir Francis Galton used these portraits to create "composite" images intended to expose the "mean" appearance of criminality [28]. Both phrenology and physiognomy had already been scientifically discredited by the time of Galton's experiments in 1877. However, belief in photography's 'mechanical objectivity' gave Galton a new (though crude) statistical tool to investigate his hypothesis of biological degeneration.⁴ Galton's experiments were ultimately failures. He found that the *visual* similarities of the criminal classes disappeared through the composite process which instead revealed the 'common humanity in all' [28]. Despite this failure however, and despite the stigma associated with eugenics and 'Social Darwinism', the experiment has been repeated using new photographic techniques, statistical methodologies, and biological theories. One important differentiation however, is the movement in knowledge paradigm from qualitative visual searches for similarity amongst groups towards purely quantitative statistical analysis.

In parallel with ongoing physiognomic experimentation, there has been a growing psychological literature exploring how 'first impressions' are generated from looking at faces (including impressions of criminality) [5, 81]. That material demonstrates how faces are a form of non-verbal communication and part of the intrinsic heuristics we use to navigate daily social interactions. However, it has also led to the claim that 'research on appearance-based [personality assessment] might seem more credible if society were

"not so enamoured of the idea that because a person's appearance ought not to make a difference, it does not," blaming the dearth of this research on the "naturalistic fallacy" (that is, confusing how things are with how they ought to be).' [76] In rejecting the 'naturalistic fallacy', the theoretical justifications for this work have at least moved away from the biological or genetic determinisms of Social Darwinism, and is justified with other ideas.⁵ Some explicitly avoid the politics of investigating correspondences between personality and appearance by claiming their work interrogates the pseudoscience of physiognomy in order to understand human subjective judgments (i.e. stereotyping). Others deploy Darwin's finding of adaptive significance in the capacity to make quick social judgments - what is called 'environment of evolutionary adaptation studies' [?] - to provide an "ecological" rather than "biological" theory. On that basis, researchers like Valla, Ceci and Williams argue that 'faces, even in the absence of dynamic behaviour, may be informative' [76]. In their study of the accuracy of physiognomic interpretation, they cite multiple studies showing that accurate physiognomic judgments can be made by humans.⁶ Whether it is genetic determinism, biased social information processing, or "ecological" accounts, there has however, never emerged a stable and accepted physiognomic theory. More than the success of 'evolutionary adaption' theories, it is the absence of an accepted or unified theoretical justification that has become the mechanism by which these physiognomic experiments continue. Although the relationship between causation and correlation in scientific method is complex, these projects seem content, in Kuhn's terms, to travel the road from scientific law to measurement in reverse, and focus on correlations entirely unpinned from theoretical accountability [51]. To that end, explorations of facial morphology through statistical methods have proliferated and gained disciplinary legitimacy, with entire fields and competitions (sponsored by industry actors like Microsoft, Facebook, Amazon and even Disney) for personality and social interaction typing through image and video data. With its capacity to surpass the visual with the quantitative, computer vision represents the technological boon that Galton was missing.

2.2 Computational Personality Analysis

The success of machine learning and particularly convolutional neural networks in computer vision mean APA and APR have become significant sub-fields in computer vision research. Whatever the specific technique, this research assumes a stable statistical relationships between a stimuli and social perception of personality (in APA), or true personality characteristics (for APR) [42]. Whereas

⁴Lorraine Daston and Peter Galison explain the term 'mechanical objectivity' in their text *Objectivity* [18] as 'the insistent drive to repress the wilful intervention of the artist-author, and to put in its stead a set of procedures that would, as it were, move nature to the page through a strict protocol, if not automatically'. They note how the camera could 'quiet the observer so nature could be heard', and "[l]et nature speak for itself".

⁵See e.g. [4, 75] where Alley neatly categorises common approaches in physiognomy including: (a) seeking resemblances in appearance between humans and animals, assuming similarities in appearance indicate shared psychological qualities (e.g., a person whose face is reminiscent of a fox is believed to be sly and cunning); (b) an inductive process whereby others sharing distinctive facial characteristics with a known person (or group) are believed to share some of the familiar person's (group's) psychological characteristics; (c) an approach based on functional analogies whereby inferences are made according to the size, shape, or presence of facial features and their common function(s); and (d) an approach based on the facial expressions engendered by various emotional and cognitive states that looks for traces of these expressive facial postures as clues to the common emotional and cognitive states of an individual.

⁶One example is a 1962 study on inferences of criminality that claimed success in replicating Galton's process. Kozeny's 'Experimental investigation of physiognomy utilizing a photographic-statistical method' used 730 criminal portraits to create composites for 16 crime categories, which were then physiognomically classified (from the visual information) in Kozeny [48].

APA avoids interrogating the accuracy of its claims, instead focusing on personality traits attributed by others (often Amazon Turk workers), APR on the other hand, seeks to identify the true personality of the individual. Although these disciplines explore multiple stimuli including dynamic facial information, handwriting, and speech, they also include what could be called ‘computational physiognomy’. Early work on computational physiognomy through computer vision was based on very simplistic models using Euclidean distances (the straight-line measurement between 2 points in a geometric plane) and geometric angles to classify facial features corresponding with the pioneering physiognomer Lavator’s division of the face into 32 classes [25, 30, 83]. Initially, these experiments were framed in terms of investigating whether computers could replicate the trait evaluation performed by humans (i.e. first impressions analysis). There were no complex machine learning methods, or assessments of accuracy. It was simply a translation of the task of physiognomic measurement into a computer vision system. From that point however, this type of personality computation became far more sophisticated.⁷ It was not long until convolutional neural networks were predicting intelligence and other personality characteristics based on images [84].⁸ Competitions for computational personality analysis (both apparent and true) began in 2012, with ‘deep learning’ dramatically improving the early unimpressive results in classifier accuracy [42]. Deep learning enables “measurement” at a radically different scale than earlier classifiers that used engineered low-level image “features” (like facial landmarks) or intermediate level “attributes” (like nose, glasses, moustache, hair, and forehead height).

How neural networks change the basis of knowledge production as compared to visual systems is demonstrated very clearly in several computer vision and data science projects, for instance Wu and Zhang’s 2016 paper ‘Automated Inference on Criminality Using Face Images’ [82]. That paper used 1856 government ID images, about half of which included individuals with a criminal conviction, as a dataset for a criminal-propensity classifier.⁹ The theoretical grounding of this research is similar to other APR exercises, ostensibly testing the social inference hypothesis [82]. But the research is more revealing in terms of how computer vision systems claim to generate knowledge about individuals. Most telling is the authors’ acknowledgment that the variance between criminal and non-criminal populations is not evident from visual assessments or simple Euclidean measurements [82]. In other words, visual information, in the sense of information that is visually perceivable and interpretable, what Galton had unsuccessfully relied on, was insufficient for physiognomic purposes. Only through the higher dimensional computational analysis of numerical quantitative measurements was this statistical separation discernible. This finding represents the premise and utility of neural networks for this application [52], and is highly illustrative of the shift from qualitative visual to quantitative statistical ways of knowing people. This change in the nature of knowledge production is made even

more explicit in Wang and Kosinski’s controversial physiognomic work ‘Deep Neural Networks are More Accurate than Humans at Detecting Sexual Orientation from Facial Images’ [79].

That project used 14438 facial images extracted from an existing database, of which 6076 were identifiably gay according to correlation with a Facebook Audience Insights platform, and self-declared sexual interest on Facebook. A deep neural network called VGG-Face (originally deployed for facial recognition) that transforms facial images into 4096 particular scores, was used for pattern analysis and classification.¹⁰ While the authors begin from the proposition that ‘DNNs are increasingly outperforming humans in visual tasks such as image classification, facial recognition, or diagnosing skin cancer’, they ultimately conclude that ‘our faces contain more information about sexual orientation than can be perceived or interpreted by the human brain’ [79]. This finding suggests that if the data sets can be constructed, and the results are to be accepted, there are few limitations to what characteristics might be inferred from images.

There is something historically significant in physiognomy being a platform for the advancement of computational empiricism. Criminologist Nicole Rafter has brilliantly argued with respect to the earlier (analogue) practices of physiognomy and phrenology that their scientific invalidity was insignificant when compared to the epistemological paradigm that they ushered in [68]. Those pseudosciences shifted how we understand people and their behaviour away from metaphysics and theology and towards ‘analytical empiricism’. Rafter shows how phrenology, the discredited science of ‘the correspondences between the external and internal man, the visible superficies and the invisible contents’ but based on reading character traits from skull morphology, therefore ‘produced one of the most radical reorientations in ideas about crime and punishment ever proposed in the Western world.’ [68] It was radical for its providing a measurable explanation of criminal behaviour and other social phenomena on the basis of positivist reasoning with empirical methods, and it transformed criminology, penology, and jurisprudence by changing how we generate knowledge about people. It is argued that the profiling techniques described above should be thought of as data points in the movement towards computational empiricism as a dominant knowledge system. Beyond “laundering” various ideological, even eugenicist, politics through technological neutrality [10], the claim here is that the foundational knowledge claims of these technical systems exhibit their own political, even philosophical, position.

Computational physiognomy, like its analogue predecessor, is not the only application of this empirical paradigm. Rather, it is best understood as a harbinger of an evolving epistemological environment. Other biometric data science applications operate on a similar knowledge-logic. For instance, automated analysis of voice and speech are being used to assess ethnic origins for migration assessments [7], as well as making customs assessments of migrants and tourists [35]. There is also a growing health-tech industry, where recording of particular bodily functions like sleeping and eating, speaking, or physical interactions with a computer are being computed into assessments of physical and psychological health

⁷These advances were typically based on research into object detection using convolutional neural networks in [50].

⁸Work of this type often uses the OCEAN model for personality description.

⁹These authors built a criminality classifier not on the basis of a statistical mean as sought by Galton, but by measuring variance from a non-criminal normal (i.e. based on the idea that there is a greater resemblance in non-criminal faces).

¹⁰The results were post-hoc justified with a prenatal hormone theory (that over or under exposure of androgens during gestation affects both facial appearance and sexual orientation), which gave a biological foundation to the statistical findings.

[3, 55]. The accuracy of these systems is not going to determine their epistemological validity. Projects are being funded, and the political will is there to deploy them. The empirical sciences accordingly follow, building new ways to classify and know persons. Data science, and its particular systems of measurement and classification, are thus becoming the prime-movers in building this new epistemological terrain.

3 COMPUTATIONAL EMPIRICISM

The claim that nature can only be understood through computation articulates data science as an attempt to access the hidden mathematical sub-structures of reality. In many respects, this claim is neither unique nor controversial, but it remains significant in the context of a technological capacity that is proliferating so rapidly and widely. It is also far from universally accepted. Some computer scientists sensibly argue that ‘deep learning can’t extract information that isn’t there, and we should be suspicious of claims that it can reliably extract hidden meaning from images that eludes human judges.’¹¹ But such admissions seem peculiar when, really, deriving meaning from measurements too granular for non-computational analysis is the very premise of machine learning. Further, as argued, the accuracy of those applications is unlikely to be the primary determinant of whether their outputs are interpreted as “true”. To that end, some have begun to argue that claiming access to the hidden structures of reality has now become the organizing principle of data science [57].

Many technologies of representation claim to reveal previously unseen information. Photography, for instance, was understood as a way to ‘see into nature’s cabinet’ [56]. It challenged the ‘optical unconscious’. Walter Benjamin describes this access to nature as opening up ‘in this material the physiognomic aspects of the world of images, which reside in the smallest details, clear and yet hidden enough to have found shelter in daydreams.’ [9] Other optical technologies offer similar narratives. The telescope gave access to celestial knowledge - the microscope to cellular knowledge. X-ray imaging and radiographic measurements of material density afforded a ‘New Sight’ [1] that could reveal ‘hidden existence’ [54]. Roberta McGrath describes public attitudes around x-rays, noting ‘The body itself is thus perceived as literally ghost-like, immaterial, only flesh, as being merely a thin veneer, literally a skin, covering a hidden, deeper reality which will, like truth, be uncovered, revealed.’ [56] Artists and intellectuals, like Umberto Boccioni, similarly responded, ‘Who can still believe in the opacity of bodies, since our sharpened and multiplied sensibilities have already penetrated the obscure manifestation of mediums?’ [21] The idea of the body as readable medium and material manifestation of a deeper informational truth is also highlighted in Mark Andrejevic’s analysis of “neuromarketing”, where he describes how: ‘The notion that bodies are, for marketing purposes, more truthful than the words they utter is emerging as a recurring theme in the promotion of neuromarketing, which promises to render obsolete the allegedly quaint and outdated techniques of surveys and focus groups... They can,

thanks to new technologies, cut through directly to the underlying truths revealed by the brain.’ [6]

Andrejevic places this claim within ‘The appeal of techniques for bypassing discursive forms of representation (by cutting ‘straight to the brain’)’ against ‘the popularised and reflexive mediated critiques of discursive forms of representation for their potentially deceptive, indeterminate, and constructed character.’ These approaches can thus be genealogically situated amongst the behavioural sciences, wherein the premise of stable relationship between stimuli and character afforded mechanisms for personality and behaviour computation such as ‘digital phenotyping’ [?]. They also represent an historical trajectory of understanding human subjectivity as an information pattern [22, 34], bound within the material container of the body that becomes its expression. As William Gibson said, ‘data made flesh’ [29]. Computer vision and biometric data science add another dimension to these narratives however. Photographs, telescopic observations, radiographic photograms, and even neuromarketing enabled a form of discovery that still requires human interpretation. In computer vision, the measurement, encoding and decoding, and knowledge discovery are increasingly automated. Through machine learning and neural networks, even the selection of representations - that is, the mechanisms by which the world is presented for analysis - is also automated. The elements of images from which meaning is derived, often called “features”, are selected through automated learning processes instead of the laborious manual process of ‘feature engineering’. Through these processes, the clinical is replaced with the empirical, image comprehension is replaced by image computation, observation replaced with measurement, with the truth - the deeper hidden reality - only becoming available in the high-dimensionality that neural networks can process.¹² This is not the combination of inputs from ‘multiple large data sources to generate new hypotheses about the way the world works and prescriptions for how to act upon that knowledge’ typically associated with big data and predictive analytics [40]. It involves a deeper belief in how the world can be understood, premised on new systems of measurement, new systems of representation, new understandings of human subjectivity, and new forms of statistical analysis.

3.1 Measurement and Representation

Accessing the hidden mathematical substructure of reality requires numbers. Reducing the real world to numbers requires measurement. Accordingly, computer vision systems do not see - they measure. They measure visual data ($\mathbf{x}$) in order to determine a ‘world state’ ($\mathbf{w}$) [66]. As Sheila Jasanoff notes, ‘Any form of data collection involves, to begin with, an act of seeing and recording something that was previously hidden and possibly nameless’ [40]. In the same way that photography bypassed the optical unconscious, computer vision profiling is about noticing, measuring, and analysing that which was previously not available to human perception and cognition. The process is complex however because the relationship between $\mathbf{x}$ and $\mathbf{w}$ is not one-to-one. Photographic images are reductions of the three-dimensional world into a two-dimensional set of

¹¹[10] although those authors do then give the conflicting argument that ‘On a scientific level, machine learning can give us an unprecedented window into nature and human behaviour, allowing us to introspect and systematically analyze patterns that used to be in the domain of intuition or folk wisdom.’

¹²See Kichin in [46] where he says, ‘Big Data analytics enables an entirely new epistemological approach for making sense of the world; rather than testing a theory by analysing relevant data, new data analytics seeks to gain insights “born from the data”.

measurements. There are however multiple configurations of the three-dimensional world that might result in any particular two-dimensional output. In other words, ‘there may be many real-world configurations that are compatible with the same measurements.’ [66] Thus the chance that any possible world state is present or true is described using probability. The probabilistic techniques of statistical clustering used to generate meaning about the world state typically use machine learning. But machine learning is computationally demanding, and it cannot operate on absolutely all the data that could be derived from an image. Instead, it is necessary to marshal selections of data into “representations”.

When computer vision emerged as a discipline in the 1960s, it became apparent that those systems could not respond to the totality of recorded signal (an image). Image data typically comes as RGB values per pixel inscribed by the translation (or transduction) of light energy to voltage in a sensor. The computational resources required to perform statistical analysis on that amount of data is immense. Therefore, instead of ‘image level computation’, computational economy requires transformations into symbolic representations [13]. The inability of computer vision systems to respond directly to the totality of registered signal has been acknowledged by some as a failure [66], and others as a fundamental limitation [20]. Nevertheless, it remains the primary mechanism by which computer vision translates measurements of the world into the computational register. That means “features” of visual data have to be selected for analysis, but as noted above that selection is increasingly automated. In the last few years significant research has gone into ways to avoid manual assessments of images to identify what might be relevant for building classifiers (i.e. feature engineering) [8] in favour of a statistical pattern recognition that can identify a relevant feature on the basis of its probabilistic relationship to other features in the environment [67]. This is sometimes called ‘representation learning’, where a system is fed with “raw data” ‘to automatically discover the representations needed for detection or classification.’ [52] ‘Deep learning’ means multiple layers of representation, automatically generated, where each level of representation is more complex and abstract than the previous one. As LeCun, Bengio and Hinton explain, ‘The key aspect of deep learning is that these layers of features are not designed by human engineers: they are learned from data using a general-purpose learning procedure.’ [52]

The inherent paradox of this epistemological platform is the use of correlation to create invisible or inaccessible automated representations in order to surpass the use of manual representation. In other words, overcoming the limits of representation with more representations.¹³ Framing it this way prompts us to reconsider imagistic representation within data scientific contexts. Much has been written about how knowledge is modulated through natural representations like the index and icon [33]. But now similar investigations are necessary for mathematical imagistic representations like ‘functional representations’ - where a function is fit to the discrete and finite set of measurements that constitute an image such as the pixel coordinates and values; ‘linear representations’ - where

images are unwound into a vector matrices; ‘spatial frequency representations’ - that measure the speed at which a particular quality changes across an image surface; ‘relational representations’ - where images are represented with graphs; and ‘probabilistic representations’ - where mathematical tools are used to estimate the best version of a particular image given a measurement of a corrupted or noisy image [67]. Representation in computer vision has changed into a project of finding the target thing (i.e. the pattern) from all the information within an image. With real consequences for how we derive meaning from the visual world.

3.2 Computational Empiricism as a Dominant Epistemology

As society becomes more statistical [31], it makes sense that both representation and knowledge take less deterministic forms. However, from the above, we can begin to draw an outline of new forms of computational empiricism (or positivism) as applied to understanding people.¹⁴ First, it operates on the basis that external measurement or observation is a more reliable pathway to knowledge than the symbolic output of a subject. This is, of course, not unique to these practices or technologies, it is simply one element of the schema. It is also not a uniquely visual phenomenon. In addition to the techniques described above, the technology of the stethoscope and the practice of “auscultation” - listening to the body at a physical distance - offers another useful example of technological mediation making the body an object of knowledge. Histories of diagnostic practices with stethoscopes have been invoked to demonstrate the movement from theoretical to perceptual ways of knowing the body, achieved through the combination of rationality and empiricism [72]. This line of thinking describes the ascendance of empiricism in this (medical) context as entwined with the construction of a new subject, and productive of a new medical epistemology of pathological anatomy. The updating of remote sensing in medical applications with new sensors and data science is similarly oriented towards constructing epistemologies of pathological behaviours [3, 78]. However, empirical data science engages with a different human subject from those older techniques, one that is simultaneously interpreted and constructed an informational pattern rather than embodied puzzle.

A second element to the schematic is a specific type of computational intervention in the relationship between measurement and classification. On one hand, computational systems differ from other representational technologies because of their capacity to both measure and process quantities of data that are too large for human tabulation, too discrete for human perception, and too complex for human cognitive analysis. This is also not unique to computer vision or data science. Where these techniques differ however, is the degree of automation in the selection of elements (or features or dimensions) of a measurement deemed to be meaningful. In

¹³This is also discussed by Andrejevic in [6].

¹⁴Note the term ‘computational empiricism’ was used by Paul Humphries in [39], where he discussed the idea in the context of changing scientific methods associated with the adoption of powerful instrumentation and powerful computational devices. His point is to suggest notions of ‘logical empiricism’ require updating by computationally oriented methods. The term ‘computational positivism’ is best described in Narasimha [60], where he uses the term to describe methodologies in exact sciences and mathematics that focus on matching algorithms to observations rather than drawing conclusions from axioms and models.

other words, decisions about what elements of a measurement inform each layer of representation are increasingly displaced into automated learning systems. Those are highly political decisions about how the real world is translated into the symbolic register of computation, and are handed off to automated systems. Of course, humans participate in selecting input data (i.e. what the sensor captures) and defining the accuracy of outputs, which each participate in tuning the selection of representations and their parameters (i.e. weightings), but the process of producing representations is dissimulated into the architecture of the system.

A third element is a belief that this process is working towards exposing the fundamental substructures of reality. The belief that increasingly granular measurement and high dimensionality analysis has the capacity to reveal hidden truths is visible to a greater or lesser degree in research projects using high resolution sensors and learning applications. This computational variant of realism has been described by Dan McQuillan as ‘machinic neoplatanism’ [57]. That is, a metaphysical commitment to a world of truth, form, and idea existing behind, and only imperfectly imprinting on, the world of the humanly sensible, accessible only through mathematics. Rather than being a specific tool or method, for McQuillan, data science thus represents an automated and applied philosophy, maintaining an epistemological reverence for ‘[a] hidden layer of reality which is ontologically superior, expressed mathematically and apprehended by going against direct experience.’ [57] None of these observations are unique to technologies of computer vision, the practices of machine learning, or the disciplines of data science. However, when schematised and directed at understanding people, they represent an idiosyncratic system of knowledge production that challenges the way legal narratives (as well as many others) have been deployed to protect individuals in the context of online profiling.

4 LAW IN THE WORLD STATE

A classic critique of technology insists that technological mediation inhibits access to ‘the real’ or ‘the event’. In the case of photography for instance, we are reminded that the images we create, while supposed to be windows or maps for understanding the world around us, actually operate more like “screens”. Rather than expose the truth of the world, our images saturate the world, producing a veneer under which ‘the real’ slowly decays. A form of this critique is often levelled at digital profiling, wherein data produced through interactions with information environments are used as proxies for defining characteristics about us. That is the world of actuarial risk assessment and the ‘scored society’. Critiques of those technologies describe how such “scores” inadequately capture or represent individuals. Proxies result in reduction, distortion and error [26, 61]. Justin Clemens and Adam Nash offer a useful account of this process wherein information ‘must first be digitised to data, then modulated between storage and display in an endless protocol-based negotiation that both severs any link to the data’s semantic source and creates an ever-growing excess of data weirdly related to, but ontologically distinct from, its originating data source.’ [16] In other words, the mathematical codes and conventions used to analyse and parse already hyper-mediated digital information omit or marginalise their “natural” starting point. Clemens and Nash

thus claim that only through modulation into a display register does digital information obtain meaning. This intervention frames the harm to persons from profiling in terms of loss. It also grounds, for instance, data protection’s animating principle of “transparency”. Legal narratives around transparency, particularly in the context of data protection, emerged to deal with that potential for misrepresentation, inaccuracy, or distortion in cases where individuals are transcribed into computational systems and categorised (i.e. profiled) on the basis of pattern analysis. Data subject rights of access and rectification are thus deployed to help individuals maintain the ‘borderlines of meaning’ about themselves as data is disclosed and processed across contexts [14]. However, as technical narratives change, they also challenge this legal mechanism.

The scores and proxies used in, for example, computer vision profiling are of a different class and scale than those used in actuarial analysis and other forms of data mining. The scores of ‘behavioural computation’ and ‘personality computation’ take the shape of a ledger, through which behavioural characteristics can be collated [69]. Through its new epistemological program, data science then starts to move from a logic of “approximation” to a logic of “revelation”. From the computational empiricist position in which the quantitative is the path to knowledge, the world of infinite possibility now exists in the data rather than in the human spirit. The “natural” starting point is no longer omitted, but rather measured in a dimensionality that humans can neither access nor interpret. Profiling systems claiming deeper knowledge about persons thus challenge the utility of transparency as a legal narrative and dismiss the role of the captured subject in forming their own identity. Thinking through how law or any other system of governance might address these technologies accordingly means attending to this re-arrangement of technological practice and how it constructs persons.

Legal thinking has so far responded to these technologies and techniques in different ways. One register seeks improvement of automated systems’ encoding of the world. The goal is to achieve a “fairer” computational translation of the real world by exposing and limiting bias and prejudice. Unfair discrimination can find its way into automated systems in multiple ways [11], and improving the outcomes of automated decision systems remains critically important. But scholars like Frank Pasquale have begun to ask whether this form of accountability adequately considers the question of ‘accountability to whom’ [64]. Without proper attention to that question, this form of accountability risks becoming part of the feedback mechanism that continually improves and thus proliferates automated decision making through legal optimisation. Yarden Katz similarly comments that ‘If AI runs society, then grievances with society’s institutions can get reframed as questions of “algorithmic accountability”’. This move paves the way for AI experts and entrepreneurs to present themselves as the architects of society.’ [44] Fairness and accountability through decision system optimisation thus fail to address the problem that computational empiricism should be understood as simply one way of knowing, especially when applied to persons. Rather than challenge the dominance of computational empiricism, “fairness” projects focused exclusively on system improvement risk law becoming enrolled in data science’s politics of completion. As Joshua Scannell notes, ‘it is not at all clear that biometric accountability, accuracy, and

fairness are mechanisms for achieving justice, or even a baseline common humanity. That is not what biometric systems do. Instead, they reflect the institutional demands of the entities that employ them - demands that are always intended to sort people into those with access or without; secured by or made insecure by the state; capacitated or incapacitated by a credit score.' [70]

Another legal register challenges specific applications of data science on ethical or political grounds. This legal mode identifies where data science applications have too pernicious an effect on society according to liberal democratic values [71]. This is typically a more traditional "privacy" mechanism defending the fundamental dignity and opacity of persons against overreaching applications. These projects similarly deploy their force at the level of application rather than at the epistemological or political bases of computational empiricism. Beyond abandoning or optimising the technologies however, there remains an important space for demonstrating the contingency of computational empiricism's knowledge claims and challenging data science's move from "metadata" to "metaphysics" [57]. This is what Mireille Hildebrandt, for instance, means when she talks of 'speaking law to the power of statistics' [37] - a program of re-inscribing uncertainty into automated knowledge production.

4.1 Law and Computer Vision Profiling

There are already important legal limitations on automated decision making and automated profiling in, for instance, the GDPR. Alongside well explored 'access and rectification rights',¹⁵ and the fundamental principles of processing in Article 5,¹⁶ Article 22 gives data subjects rights to not be subject to automated profiling decisions based solely on automated processing where those decisions produce legal or similar effects. There are numerous limitations to that latter provision, including explicit consent, enabling legislation, or satisfaction of a contract.¹⁷ What constitutes legal or similar effects, as well as purely automated processing, are also unclear. That a human decision maker anywhere in the process might remove an automated system from the purview of the Article seems a problematic limitation, especially considering the growing evidence that human decision makers rarely contest the outcomes of decision support systems. Further, guidance by the Court of Justice of the European Union will be essential to a more complete understanding of these provisions. A right to "explanation" has also been read into Art 15(1)(h), producing rigorous debate over what the Article truly affords data subjects, and how useful it may be. While a similar (and arguably broader) provision has been in place since 1995, a broad reading may give data subjects useful mechanisms for not only understanding, but properly contesting and challenging automated decisions. However, it still places a substantial impetus on the data subject to protect its own rights, and only indirectly challenges the knowledge logic of computational empiricism. Explainability may even be harmful if entrenching automated decision making and narrowing the types of reasons to be given for decisions and thus the grounds for contestation. Being subjected to automated decisions

without understanding how or why that decision was made may be problematic. But receiving inadequate automated explanations without recourse might be worse. In that format, "explanation" similarly becomes a legal vector for proliferating decision automation. However, other approaches appear to more directly challenge data science's epistemological dominance.

Hildebrandt, for instance, outlines several concepts that challenge data science's program of quantification - what she describes as the 'overcomplete datafication of anything and everything based on the idea that the mathematics that grounds all these machines reveals the ultimate layer of a hidden reality.' [36] She offers a right to human non-computability built on the philosophical principle of indeterminate identity. This is not necessarily a return to older opacity paradigms of privacy, a relocation of the black box to the level of the individual, but rather a mechanism for limiting certain classes of knowledge claims. It is achieved in her vision through "agonistic" machine learning systems, capable of 'demanding that companies or governments that base decisions on machine learning must explore and enable alternative ways of datafying and modelling the same event, person or action.' Agonistic systems would demonstrate how each act of computation relies on a specific system of measurements, representations, and analytics. Rather than challenging human computability wholesale, the project highlights how humans can be computed in multiple ways, in order to 'ward off monopolistic claims about the "true" or the "real" representation of human beings'. This is arguably also a form of explanation and accountability, but one targeted more appropriately at the decision making process itself rather than the relationship between input and outcome (i.e. counter-factual analysis). In other words, an explanation mechanism informing how to transcend the solution space of any particular decision. This type of mechanism is appealing because it more directly addresses the construction of the human subject as information pattern. It also recognises the significance of representations in a manner similar to technical approaches like 'disentangling the factors of variation' in order to produce 'disentangled representations' [74]. What this manner of legal intervention would actually look like requires more legal, conceptual, and technical thinking however. But it at least offers a pathway to envisioning new projects built for a technical age that sees humans in a very different way than our existing legal systems insist on seeing them.

5 CONCLUSION

This article has attempted to outline some emerging challenges produced by a profound new technical capacity. Lawmakers now have to contend with a radically extrapolated enlightenment philosophy insisting that measuring everything and drawing knowledge from those measurements is the path to truth. Permitting such measurement and classification without limitation risks generating a fair, transparent, and non-prejudicial totalitarianism. This has been described as the 'insanity proper to logic', whereby measuring everything, logic conceives 'a world in which all things are relative, makes itself absolute, and denying the whole of nature, establishes its own artifices.' [59] Addressing this reality means clarifying how computer vision systems and data science applications are merely apparatuses, combining and computing symbols that encode the

¹⁵See e.g. Arts 13, 14, 15 (access), 16 (rectification), 17 (erasure), 18 (restriction), 19 (notification), 20 (data portability).

¹⁶These include lawfulness, fairness, transparency, purpose limitation, data minimisation, accuracy, accountability, storage limitation, integrity and confidentiality.

¹⁷Art 22(2)(a), (b), and (c).

world a particular way, according to particular programs. Without intervention at that level, digital systems risk simultaneously constituting both the world and the dominant understanding of it [80]. As Flusser reminds us, if humans cease to decode technical images and instead project them unencoded back onto the world 'out there', the world 'itself becomes like an image - a context of scenes, of states of things.' [27] Critiquing data science therefore means exposing the difference between world and the 'world state', and challenging the idea that such systems access a 'hidden reality' instead of producing and operationalise a para-reality built from para-visual representations.

Computer vision has been presented as one vector of that knowledge logic, and one target of a new kind of legal thinking. However, this is certainly not the case for computer vision alone. It is relevant for any data science application translating the real world into the symbolic register of computation. The goal is to open up space for legal ideas that introduce contingency and automated un-decidability into those translations that are building the environments ordering social experience. Such projects seem more and more imperative, especially in the face of a new form of computational empiricism designed for 'looking at people'.

ACKNOWLEDGMENTS

The author would like to thank Dan Hunter, Helen Nissenbaum, James Parker, Joel Stern, Rose Parfitt, Luis Eslava, Pete Chambers, Erica Du, Lauren van Haaften-Schick, the Digital Life Initiative at Cornell Tech, and Swinburne Law School, for supporting this research, and for invaluable discussions and tech-fixes. This work was presented as a work in progress at the 'Fascism and the International' Workshop, Melbourne Australia, as part of the 'Eavesdropping' project by Liquid Architecture and Melbourne Law School, and at the DLI Seminar at Cornell Tech.

The author would also like to thank the anonymous referees for their valuable comments and helpful suggestions. The work was supported by the National Science Foundation under Grant:1650589.

REFERENCES

- [1] 1898. The New Sight and the New Photography.
- [2] 2016-05-04. REGULATION (EU) No 2016/679 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *OJ L* 119 (2016-05-04), 1–88.
- [3] Saeed Abdullah, Mark Matthews, Ellen Frank, Gavin Doherty, Geri Gay, and Tanzeem Choudhury. 2016. Automatic detection of social rhythms in bipolar disorder. *Journal of the American Medical Informatics Association* 23, 3 (2016), 538–543. <https://doi.org/10.1093/jamia/ocv200>
- [4] Thomas Alley. 1988. Physiognomy and Social Perception. In *Social and Applied Aspects of Perceiving Faces*, Thomas Alley (Ed.). Laurence Erlbaum Publishers.
- [5] Nalini Ambady and John Skowronski (eds). [n. d.]. *First Impressions, year = 2008*, isbn = 978-1593857165, publisher = The Guilford Press.
- [6] Mark Andrejevic. 2012. Brain Whisperers: Cutting Through the Clutter with Neuromarketing. *Somatechnics* 2(2) (2012), 198–215.
- [7] Emily Apter. 2016. Shibboleth: Policing by Ear and Forensic Listening in Projects by Lawrence Abu Hamdan. *October* 156 (2016), 100–115.
- [8] Yoshua Bengio, Aaron C. Courville, and Pascal Vincent. 2012. Unsupervised Feature Learning and Deep Learning: A Review and New Perspectives. *CoRR* abs/1206.5538 (2012). [arXiv:1206.5538](http://arxiv.org/abs/1206.5538) <http://arxiv.org/abs/1206.5538>
- [9] Walter Benjamin. 1980. A Short History of Photography 1934. In *Classic Essays on Photography*, Alan Trachtenberg (Ed.). Leete Island Books.
- [10] Margaret Mitchell Blaise Agüera y Arcas and Alexander Todorov. 2017. Physiognomy's new Clothes. <https://medium.com/@blaisea/physiognomys-new-clothes-f2d4b59fdd6a>
- [11] Solon Borocas and Andrew Selbst. 2016. Big Data's Disparate Impact. *California Law Review* 104(3) (2016), 671–732.
- [12] James Bridle. 18 December 2013. How Britain Exported Next-Generation Surveillance. <https://medium.com/matter/how-britain-exported-next-generation-surveillance-d15b5801b79e>
- [13] C. M. Brown. 1984. Computer Vision and Natural Constraints. 224, 4655 (1984), 1299–1305. <https://doi.org/10.1126/224.4655.1299> [arXiv:http://mag.org/content/224/4655/1299.full.pdf](http://mag.org/content/224/4655/1299.full.pdf)
- [14] Lee Bygrave. 2010. Privacy and Data Protection in an international perspective. *Scandinavian Studies in Law* 56(8) (2010), 165–200.
- [15] Alvaro Bedoya Clare Garvie and Jonathan Frankle. 18 October 2018. The Perpetual Lineup: Unregulated Police Face Recognition in America. <https://www.perpetuallineup.org/>
- [16] Justin Clemens and Adam Nash. 2015. Being and Media: Digital Ontology after the Event of the End of Media. *The Fibreculture Journal* 24 (2015), 6–32.
- [17] Kate Crawford. 2016. Can an Algorithm be Agonistic? Ten Scenes from Life in Calculated Publics. *Technology & Human Values* 41(1) (2016), 77–92.
- [18] Lorraine Daston and Peter Galison. 2007. *Objectivity*. Zone Books.
- [19] Christopher Millard Dimitra Kamarinou and Jatinder Singh. 2016. *Machine Learning with Personal Data*. Technical Report 247. Queen Mary University of London School of Law Legal Studies Research Paper.
- [20] Herbert Dreyfus. 2007. Why Heideggerian AI failed and how fixing it would require making it more Heideggerian. *Artificial Intelligence* 171(18) (2007), 1137–1160.
- [21] Didier Ottinger (ed). 2008. *Futurism*. Centre Pompidou Five Continent Editions.
- [22] Paul Edwards. 1996. *The Closed World: Computers and the Politics of Discourse in Cold War America*. MIT.
- [23] Sergio Escalera, Xavier Baró, Hugo Jair Escalante, and Isabelle Guyon. 2017. ChaLearn Looking at People: Events and Resources. *CoRR* abs/1701.02664 (2017). [arXiv:1701.02664](http://arxiv.org/abs/1701.02664) <http://arxiv.org/abs/1701.02664>
- [24] Fred H. Cate et al. 2017. *Machine Learning with Personal Data: Is Data Protection Law Smart Enough to Meet the Challenge?* Technical Report 2634. Articles by Maurer Faculty.
- [25] V Dominique et al. 1997. What Represents a Face? A Computational Approach for the Integration of Physiological and Psychological Data. *Perception* 26 (1997), 1271–1288.
- [26] Virginia Eubanks. 2018. *Automating Inequality*. St Martins Press.
- [27] Vilém Flusser. 2000. *Towards a Philosophy of Photography*. Anthony Mathews trans in Reaktion Books.
- [28] Francis Galton. 1879. Composite Portraits, Made by Combining Those of Many Different Persons into a Single Resultant Figure. *Journal of the Anthropological Institute of Great Britain* 3 (1879), 132–144.
- [29] William Gibson. 1984. *Neuromancer*. ACE Books.
- [30] John Graham. 1961. Lavater's "Physiognomy": A Checklist. *The Papers of the Bibliographical Society of America* 55(4) (1961), 297–308.
- [31] Ian Hacking. 1990. *The Taming of Chance*. Cambridge University Press.
- [32] Bernard Harcourt. 2015. *Exposed: Desire and Disobedience in the Digital Age*. Harvard University Press.
- [33] Charles Hartshorne and Paul Weiss (eds). 1932. *Collected Papers of Charles Sanders Peirce*. Harvard University Press.
- [34] N Katherine Hayles. 1999. *How We Became Post-Human: Virtual Bodies in Cybernetics, Literature and Infomatics*. University of Chicago Press.
- [35] Douglas Heaven. 31 October 2018. An AI lie detector will interrogate travellers at some EU borders. <https://www.newscientist.com/article/mg24032023-400-an-ai-lie-detector-will-interrogate-travellers-at-some-eu-borders/>
- [36] Mireille Hildebrandt. [n. d.]. Privacy as Protection of the Incomputable Self: From Agnostic to Agonistic Machine Learning. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3081776
- [37] Mireille Hildebrandt. 2018. Law as Computation in the era of artificial legal intelligence: Speaking Law to the Power of Statistics. *University of Toronto Law Journal* 68, 1 (Jan. 2018), 12–35. <https://doi.org/10.3138/utlj.2017-0044>
- [38] Anna Hoffman and Luke Stark. 11 September 2015. Are Face Recognition Systems Accurate? Depends on Your Race. <https://www.technologyreview.com/s/601786/are-face-recognition-systems-accurate-depends-on-your-race>
- [39] Paul Humphries. 1996. Computational Empiricism. In *Topics in the Foundation of Statistics*, Bas C. van Fraassen (Ed.). Springer.
- [40] Sheila Jasanoff. 2017. Virtual, visible, and actionable: Data assemblages and the sightlines of justice. *Big Data & Society* 4, 2 (2017), 2053951717724477. <https://doi.org/10.1177/2053951717724477> [arXiv:http://arxiv.org/abs/10.1177/2053951717724477](http://arxiv.org/abs/10.1177/2053951717724477)
- [41] Oscar H Gandy Jr. 2012. Statistical Surveillance: Remote Sensing in the Digital Age. In *Routledge Handbook of Surveillance Studies*, Kevin D Haggerty Kristie Ball and David Lyon (Eds.). Routledge.
- [42] Júlio C. S. Jacques Júnior, Yagmur Güçlütürk, Marc Pérez, Umut Güçlü, Carlos Andújar, Xavier Baró, Hugo Jair Escalante, Isabelle Guyon, Marcel A. J. van Gerven, Rob van Lier, and Sergio Escalera. 2018. First Impressions: A Survey on Computer Vision-Based Apparent Personality Trait Analysis. *CoRR* abs/1804.08046 (2018). [arXiv:1804.08046](http://arxiv.org/abs/1804.08046) <http://arxiv.org/abs/1804.08046>

- [43] Jens Jäger. 2001. Photography: a means of surveillance? *Judicial Photography* 1850 – 1900. *Crime, History & Societies* 5(1) (2001), 27–51.
- [44] Yarden Katz. [n. d.]. Manufacturing an Artificial Intelligence Revolution. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3078224
- [45] Danielle Keats Citron and Frank Pasquale. 2014. The scored society: Due process for automated predictions. *Washington Law Review* 89 (2014), 1–33.
- [46] Rob Kichin. 2014. Big data, New Epistemologies and Paradigm Shifts. *Big Data and Society* 1(1) (2014), (online).
- [47] Rob Kichin and Martin Dodge. 2011. *Code/Space: Software and Everyday Life*. MIT.
- [48] E Kozeny. 1962. Experimentelle Untersuchungen zur Ausdruckskunde mittel photographisch-statistischer Methode. *Archive für die gesamte psychologie* 114 (1962), 55–71.
- [49] Robin Kramer and Kay Ritchie. 14 December 2016. The Trouble with Facial Recognition technology (in the real world). <https://theconversation.com/the-trouble-with-facial-recognition-technology-in-the-real-world-69685>
- [50] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems* 25, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 1097–1105. <http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf>
- [51] Thomas Kuhn. 1961. The Function of Measurement in Modern Physical Science. *Isis* 52(2) (1961), 161–193.
- [52] Yann LeCun, Yoshua Bengio, and Geoffrey E. Hinton. 2015. Deep learning. *Nature* 521, 7553 (2015), 436–444. <https://doi.org/10.1038/nature14539>
- [53] Monique Mann and Marcus Smith. 2017. Automated facial recognition technology: Recent developments and approaches to oversight. *University of New South Wales Law Journal* 40, 1 (2017), 121–145.
- [54] Mary Warner Marien. 2015. *Photography: A Cultural History (4th ed)*. Laurence King Publishing.
- [55] Mark Matthews, Stephen Volda, Saeed Abdullah, Gavin Doherty, Tanzeem Choudhury, Sangha Im, and Geri Gay. 2015. In Situ Design for Mental Illness: Considering the Pathology of Bipolar Disorder in mHealth Design. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '15)*. ACM, New York, NY, USA, 86–97. <https://doi.org/10.1145/2785830.2785866>
- [56] Roberta McGrath. 2002. *Seeing her sex: Medical archives and the female body*. Manchester University Press.
- [57] Dan McQuillan. 2017. Data Science as Machinic Neoplatanism. *Philosophy & Technology* 31(2) (2017), 253–273.
- [58] Isak Mendoza and Lee Bygrave. 2017. A Short History of Photography 1934. In *EU Internet Law: Regulation and Enforcement*, Taitani Synodinou et al (Ed.). Springer.
- [59] Jacques-Alain Miller and Richard Miller. 1987. Jeremy Bentham's Panoptic Device. *October* 41 (1987), 3–29. <http://www.jstor.org/stable/778327>
- [60] Roddam Narasimha. 2001. Axiomatism and Computational Positivism: Two Mathematical Cultures in Pursuit of Exact Sciences. *Journal of Medical Humanities* 22(2) (2001).
- [61] Cathy O'Neil. 2016. *Weapons of Math Destruction*. Crown Books.
- [62] Mike Orcutt. 6 July 2016. Are Face Recognition Systems Accurate? Depends on Your Race. <https://www.technologyreview.com/s/601786/are-face-recognition-systems-accurate-depends-on-your-race>
- [63] Trevor Paglen. 8 December 2016. Invisible Images (Your Pictures are Looking at You). <https://thenewinquiry.com/invisible-images-your-pictures-are-looking-at-you/>
- [64] Frank Pasquale. 20 Aug 2018. Odd Numbers. <http://reallifemag.com/odd-numbers/>
- [65] Frank Pasquale. 2018. When Machine Learning is Facially Invalid. *Commun. ACM* 61, 9 (Aug. 2018), 25–27. <https://doi.org/10.1145/3241367>
- [66] Simon JD Prince. 2012. *Computer Vision: Models, Learning and Inference*. Cambridge University Press.
- [67] Hairong Qi and Wesley E. Snyder. 2004. *Machine Vision*. Cambridge University Press.
- [68] Nicole Rafter. 2005. The Murderous Dutch Fiddler: Criminology, History and the Problem of Phrenology. *Theoretical Criminology* 9(1) (2005), 65–96.
- [69] Vlad Savov. 17 May 2018. Google's Selfish Ledger is an Unsettling Vision of Silicon Valley Social Engineering. <https://www.theverge.com/2018/5/17/17344250/google-x-selfish-ledger-video-data-privacy>
- [70] Joshua Scannell. 17 September 2018. Controlled Measures. <http://reallifemag.com/controlled-measures/>
- [71] Evan Selinger and Woodrow Hartzog. 22 June 2015. Opinion: It's time for an about-face on facial recognition. <https://www.csmonitor.com/World/Passcode/Passcode-Voices/2015/0622/Opinion-It-s-time-for-an-about-face-on-facial-recognition>
- [72] Jonathan Sterne. 2001. Mediate Auscultation, the Stethoscope, and the "Autopsy of the Living": Medicine's Acoustic Culture. *Journal of Medical Humanities* 22, 2 (01 Jun 2001), 115–136. <https://doi.org/10.1023/A:1009067628620>
- [73] Jordan G Teicher. 18 July 2018. What Do Facial Recognition Technologies Mean for our Privacy? <https://www.nytimes.com/2018/07/18/lens/what-do-facial-recognition-technologies-mean-for-our-privacy.html>
- [74] Valentin Thomas, Jules Pondard, Emmanuel Bengio, Marc Sarfaty, Philippe Beaudoin, Marie-Jean Meurs, Joelle Pineau, Doina Precup, and Yoshua Bengio. 2017. Independently Controllable Factors. *CoRR* abs/1708.01289 (2017). [arXiv:1708.01289](http://arxiv.org/abs/1708.01289)
- [75] Alexander Todorov. 2017. *Face Value: The Irresistible Influence of First Impressions*. Princeton University Press.
- [76] Jeffrey M Valla, Stephen J Ceci, and Wendy M Williams. 2011. The Accuracy of Inferences about Criminality Based on Facial Appearance. *Journal of Social, Evolutionary, and Cultural Psychology* 5(2) (2011), 66–91.
- [77] Jonathan Vanian. 13 July 2018. Microsoft President: Facial Recognition Technology Needs Government Regulation. <http://fortune.com/2018/07/13/microsoft-facial-recognition-regulation/>
- [78] Rui Wang, Weichen Wang, Min S. H. Aung, Dror Ben-Zeev, Rachel Brian, Andrew T. Campbell, Tanzeem Choudhury, Marta Hauser, John Kane, Emily A. Scherer, and Megan Walsh. 2017. Predicting Symptom Trajectories of Schizophrenia Using Mobile Sensing. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 3, Article 110 (Sept. 2017), 24 pages. <https://doi.org/10.1145/3130976>
- [79] Yilun Wang and Michal Kosinski. 2018. Deep Neural Networks are More Accurate than Humans at Detecting Sexual Orientation from Facial Images. *Innovations in Social Psychology* 114(2) (2018), 246–257. <http://arxiv.org/abs/1701.02664>
- [80] Leif Weatherby. 2018. Digital Metaphysics: The Cybernetic Idealism of Warren McCulloch. https://iasc-culture.org/THR/THR_article_2018_Spring_Weatherby.php
- [81] J Willis and A Todorov. 2006. First Impressions: Making up your mind after 100ms exposure to a face. *Psychological Science* 17(7) (2006), 592–598.
- [82] Xiaolin Wu and Xi Zhang. 2016. Automated Inference on Criminality using Face Images. *CoRR* abs/1611.04135 (2016). [arXiv:1611.04135](http://arxiv.org/abs/1611.04135)
- [83] Hee-Deok Yang and Seong-Whan Lee. 2006. Automatic Physiognomic Analysis by Classifying Facial Component Feature. , 1212–1215 pages. <https://doi.org/10.1109/ICPR.2006.1196>
- [84] Ting Zhang, Ri-Zhen Qin, Qiu-Lei Dong, Wei Gao, Hua-Rong Xu, and Zhan-Yi Hu. 2017. Physiognomy: Personality traits prediction by learning. *International Journal of Automation and Computing* 14, 4 (01 Aug 2017), 386–395. <https://doi.org/10.1007/s11633-017-1085-8>