

This is a repository copy of *Regimes of recognition on algorithmic media*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/203591/>

Version: Published Version

Article:

Jacobsen, Benjamin orcid.org/0000-0002-6656-8892 (2023) Regimes of recognition on algorithmic media. *New media & society*. pp. 3641-3656. ISSN: 1461-7315

<https://doi.org/10.1177/14614448211053555>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



Article

Regimes of recognition on algorithmic media

new media & society

1–16

© The Author(s) 2021



Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/14614448211053555

journals.sagepub.com/home/nms



Benjamin N Jacobsen 

Durham University, UK

Abstract

This article examines ways in which people are seen, recognised, and made to matter by social media platforms. Drawing on Louise Amoore's notion of 'regimes of recognition', I argue that social media platforms can be conceptualised as increasingly powerful arbiters of recognisability, determining the conditions of possibility of how people are seen and come to matter. Through an analysis of Twitter's saliency detection algorithm, which automatically crops images uploaded to the platform, the article highlights how social media platforms participate in producing novel modes of recognisability, that is, conditions by which people are rendered visible and invisible within or by the platform. Moreover, the article highlights how regimes of recognition on algorithmic media shape people's parameters of attention and perception more generally through what I call the automatic production of 'consistent' lines of sight. Ultimately, the article seeks to highlight how the notion of recognition is increasingly arbitrated in and through algorithmic media and how this is fraught with political issues and tension. As such, there is an ongoing need to critically examine the power of social media to render people visible and invisible.

Keywords

Algorithms, attention, automation, bias, data, machine learning, perception, social media

Introduction

In 2014, researchers at Facebook developed a deep learning algorithm called DeepFace. The facial recognition model was trained on over 4 million images of Facebook users,

Corresponding author:

Benjamin N Jacobsen, Department of Sociology, Wentworth College, University of York, Heslington, York YO10 5DD, UK.

Email: benjamin.jacobsen@durham.ac.uk

and it was reported to be able to recognise people featuring in any images uploaded to the social media platform at a level of accuracy virtually equal to humans (Lowensohn, 2014). Although platforms like Facebook have since developed much larger and more sophisticated algorithmic models, the central idea remains the same: social media are in the business of recognition. Platforms such as Facebook, Instagram, and Twitter, what Taina Bucher (2018) calls 'algorithmic media', have increasingly become image-sharing and image-processing platforms. They have become spaces of 'ubiquitous photography' (Hand, 2012), where images are increasingly algorithmically mediated (Uricchio, 2011). Yet, rather than merely facilitating more image-sharing among people, social media platforms are increasingly seeking to understand the content of images on a more granular level. That is, to algorithmically identify and recognise the persons and objects within these images at the level of the pixel.

In this article, I argue that social media platforms have become powerful arbiters of recognisability. They determine not only what is made visible to users or what is ignored, but also constitute the conditions for what can be perceived and recognised by the platforms themselves. I suggest that in an age of algorithms, the capacity to arbitrate recognisability on social media, and the social implications of this, are increasingly important to critically examine, especially since recognition is crucial for how we act and interact with others in society (Crandall, 2005; Honneth, 1992). The article thus examines the politics of the capacity to arbitrate recognisability on algorithmic media and the ways in which it is inextricably linked to what comes to matter within social media images.

Conceptually, the article draws on Louise Amoore's (2020) notion of 'regimes of recognition' to examine how digital images are algorithmically processed, analysed and shaped by social media platforms. In the context of this article, regimes of recognition refer to the techniques and tools used by social media to arbitrate recognition. The focus on the notion of recognisability foregrounds the politics of how algorithmic systems 'see' the world and make it matter in certain ways, shaping people's parameters of attention and perception. In this context, regimes of recognition refer specifically to the algorithmic practices, tools, techniques and technologies used by social media platforms to process, analyse, tweak and circulate images. As this article points out, regimes of recognition on algorithmic media can shape what features are brought to the fore within an image and which features are not. This has implications for how the social world is seen and enacted, for as Jonathan Crary (1990) suggests, 'an observer is more importantly one who sees within a prescribed set of possibilities, one who is embedded in a system of conventions and limitations' (p. 6). I argue that regimes of recognition on algorithmic media circumscribe how people see the social world; they shape the prescribed set of possibilities of what can be seen and what is rendered invisible.

More specifically, the article explores this algorithmic power through an examination of Twitter's saliency detection neural network algorithm. As will be expanded upon further in the article, the saliency algorithm seeks to determine what aspects of an image count as 'salient', that is, worthy of attention, and what aspects should be automatically cropped out. In September 2020, Twitter came under media scrutiny when accusations surfaced that the algorithm favoured white faces over black faces in images that had been uploaded to the platform, to the point where black faces were repeatedly being cropped out of images (Hern, 2020). Drawing on a critical analysis of company reports published

by Twitter as well as the computer science literature on saliency detection algorithms, this article explores the power of platforms to arbitrate recognisability as well as the ethicopolitical implications of this.

In order to explore how social media platforms can be understood as powerful arbiters of recognisability, this article is divided into the following four main parts: First, I discuss the ways in which machine learning algorithms learn to ‘see’ the world as well as Amoore’s (2020) notion of ‘regimes of recognition’. Second, drawing on both company reports and computer science literature, I provide a detailed description of the functionality of Twitter’s saliency detection algorithm and the incident of September 2020. Here, I foreground the socio-technical assumptions and logic that underlie its operations, showcasing how saliency detection algorithms fundamentally divide the world into salient regions and ‘redundant’ ones, regions deemed worthy of attention and those that are not. The following section then explores the implications of this underlying logic. I argue that social media’s capacity to arbitrate recognisability constitutes the power to render (certain) people visible and invisible within platforms. The final section, I argue that the capacity to arbitrate recognisability is also the power to shape people’s perception through the algorithmic construction of what I call ‘consistent’ lines of sight. That is, social media platforms seek to circumscribe people’s parameters of attention, and in turn, create subjects that are productive, predictable and consistent. Ultimately, I argue that through an analysis of so-called regimes of recognition on algorithmic media, we may further unpack the power and politics of algorithms in everyday life and how they shape the conditions of what becomes visible to people on social media platforms.

Regimes of recognition: how algorithms ‘see’ the world

Much research has already been conducted into the social power and politics of algorithms (Beer, 2009, 2017; Cheney-Lippold, 2016; Gillespie, 2014; Kotliar, 2020; Willson, 2017). From a computational perspective, algorithms are fundamentally calculative procedures employed in computer software to process input data in order to generate target outputs (Kitchin, 2017). Yet, they are not merely operations performed on data, but are also powerful social actors, shaping the social world in various ways. Algorithms are often used to mine, sort and order the social world, which has led some scholars to define the social power of algorithms as a ‘soft biopower’ (Cheney-Lippold, 2011). As such, algorithms should not merely be conceived in mathematical or computational terms. Rather, they are ‘ethicopolitical arrangements of values, assumptions, and propositions about the world’ (Amoore, 2020: 6).

The social power of machine learning algorithms has also been critically examined in relation to social media platforms (Bucher, 2018; Flyverbom, 2019; Gillespie, 2010; Hallinan and Striplas, 2016). For Taina Bucher (2012b), for instance, social media platforms generate new ‘modalities of visibility’. She draws on the example of Facebook’s EdgeRank algorithm, which calculates the content posted on the platform and filters it in a way that creates a hierarchy of visibility – where some posts attain more attention than others. For Bucher, this generates a ‘threat of invisibility’ (Bucher, 2012b: 1164) among users. As a result, one of the crucial aspects of the social power and politics of algorithms

resides in the various ways in which they help to shape the parameters of visibility and how people can see the world.

This raises the question of how algorithmic systems learn to ‘see’ and ‘recognise’ aspects of the social world. This is a question which is becoming increasingly pressing given the proliferation of tools such as facial recognition in society (Andrejevic and Selwyn, 2019; Bueno, 2019; Crawford and Paglen, 2019). Indeed, the intersections of algorithmic systems and issues of visibility and perception are foregrounded through a proliferation of visual concepts presented in the scholarly literature. Notions such as ‘data gaze’ (Beer, 2019), ‘algorithmic gaze’ (Kotliar, 2020), ‘platform seeing’ (Mackenzie and Munster, 2019) and ‘optical unconscious of big data’ (Agostinho, 2019) have all been proposed to make sense of the heterogeneous ways in which data companies, social media platforms and algorithmic systems see people and how they render the social world visible in different ways – all of which have ethicopolitical implications for how people come to see the world.

In her book *Cloud Ethics*, Louise Amoore (2020) explores the processes by which machine learning algorithms, and neural networks in particular, learn to see the world. That is, how they learn to recognise features within images and videos such as faces, animals, objects, events or scenery. Amoore (2020) points out that convolutional neural networks are trained on big datasets of images, where they learn to recognise particular patterns and features at the pixel level of the image. The input data are then assigned a series of weightings or parameters that determine its significance within the model. As a result, the algorithm learns over time to weight some patterns or clusters in the pixel values more than others. For instance, learning to recognise particular breeds of dog in an image or learning to link the image of a face to a concrete individual. Moreover, this means that neural networks possess a certain level of autonomy. As Amoore states, ‘when deep neural network algorithms learn, then, they adjust themselves in relation to the features of their environment’, that is, they learn how ‘to afford weight or value to one pixelated part of an image over others’ (p. 74).

Although this capacity for the algorithm to self-adjust and adapt to data inputs may raise concerns regarding issues of transparency, opacity and accountability (Ananny and Crawford, 2018; Burrell, 2016), it also highlights a more fundamental question: who or what counts as recognisable to the algorithm? Amoore (2020: 69) argues,

Machine learning algorithms do not merely recognize people and things in the sense of identifying – faces, threats, vehicles, animals, languages – they actively generate recognizability as such, so that they decide what or who is recognizable as a target of interest in an occluded landscape.

Through adjusting its weights or parameters over time, weighting some data elements in an image more than others, the algorithm does not simply learn to see the world. Rather, it must actively generate what can be recognised – whether that is a particular face, a particular feature of a face, an animal or a potential criminal. As such, algorithmic systems produce modes of recognition when they iteratively learn over time to weight and recognise certain features in images rather than others. Arbitrating recognisability therefore becomes a highly political process, for what the algorithm learns to recognise is, in fact,

‘what is normal and anomalous at each parse of the data’ (Amoore, 2020: 68). Amoore (2020) therefore argues that regimes of recognition are political actors in terms of ‘both arbitrating recognizability and outputting a desired target that is actionable’ (p. 72).

However, Amoore’s notion of regimes of recognition should not be understood as purely technical or computational constructs. Rather, they are socio-technical achievements, generated by assemblages of human and algorithmic actors. Amoore (2020) states that there are multiple moments where both can be seen to be implicated in the production of regimes of recognition: ‘The selection of training data; the detection of edges; the decisions on hidden layers; the assigning of probability weightings; and the setting of threshold values’ (p. 71). A good example of this is how algorithms adjust their own parameters and how they are tweaked by developers, because it has been shown that a small adjustment in the weightings and parameters of a neural network can generate a fundamentally different set of outputs (Su et al., 2019). In other words, an image of a car may come to be seen as a cat by the algorithm. Regimes of recognition thus constitute ‘intimate entanglements’ (Latimer and Gómez, 2019) between developer and algorithm.

Crucially, therefore, the politics of algorithms resides not only in what they output, but also the ways in which they are adjusted or arranged. As Amoore (2020) suggests, ‘the shifting of the thresholds for that recognition embodies all the valuations, associations, prejudices, and accommodations involved in determining what is “useful” or “good enough”’ (p. 68). For this reason, Amoore (2020) argues that ‘the processes and arrangements of weights, values, bias, and thresholds in neural nets . . . must be presented as questions and political claims in the world’ (p. 81). Algorithmic regimes of recognition are produced through these various processes and mechanisms. They shape who or what can be recognised, becoming the condition of recognisability on social media platforms.

In this article, the concept of regimes of recognition refers to the algorithmic tools, techniques and practices that social media platform use to arbitrate recognisability within their platforms. That is, how they make people and objects visible, determining what comes to matter on the platform. This topic needs to be critically examined, as social media platforms not only incorporate sophisticated computational systems into the way in which they operate. They also often figure at the forefront of cutting-edge research on algorithms such as machine vision, object and facial recognition, image and video segmentation, and natural language processing. This is evident through the many projects that are made public in forums such as Facebook Engineering, Facebook AI Research, Google AI, Instagram Engineering and Twitter Engineering. As a result, it is crucial that social media platforms are critically examined in terms of how they arbitrate recognisability and generate the conditions by which people are rendered visible and knowable to the algorithm. The arbitration of recognisability participates in shaping what becomes visible on the platform and the parameters of what people can or should not see. In what follows, I examine Twitter’s saliency detection neural network algorithm, how it learns to ‘recognise’ people in social media images, as well as its potential social ramifications.

Neural networks and the value of being salient on Twitter

In September 2020, Twitter came under media scrutiny for their use of an image-cropping algorithm. Users complained that the algorithm could be seen to favour white faces

over black faces within certain images that were uploaded to the platform. One widely shared example showed that the algorithm consistently cropped out an image of former president Barack Obama while presenting a 'normally' cropped image of US senator Mitch McConnell. This bias was repeatedly highlighted in various other experiments such as with stock photo models, cartoon characters and even white and black dogs (Hern, 2020). Twitter's chief design officer, Dantley Davis, acknowledged the racial bias in the algorithm, stating 'We tested for bias before shipping the model and didn't find evidence of racial or gender bias in our testing', but also adding 'it's clear that we've got more analysis to do' (Murdock, 2020). However, Davis suggested that the algorithm is not explicitly racist since it does not make its decision based on particular faces but rather on the contrasts that are calculated from the pixel values of the image. Davis finally stated, 'It's 100% our fault. No one should say otherwise. Now the next step is fixing it' (Murdock, 2020).

Twitter's image-cropping algorithm is a deep convolutional neural network. It is trained to recognise the most relevant aspects of an image and automatically crop out the 'redundant' parts. In a blog titled 'Speedy Neural Networks for Smart Auto-Cropping of Images', Lucas Theis, a machine learning researcher, and Zehan Wang, a software engineer, outline how the neural network used by Twitter operates and the core underlying aspects of its functionality (Theis and Wang, 2018). They begin by stating that since photo sharing is 'an integral part of the Twitter experience' the platform automatically crops images that have been uploaded to the timeline to 'improve consistency and to allow you to see more Tweets at a glance' (Theis and Wang, 2018). In other words, the image-cropping algorithm is deployed to supposedly provide a more frictionless experience of the platform, which in turn is aimed at optimising user engagement and activity. Having previously used face detection software, Twitter abandoned this approach since, of course, not all images contain faces or people. As such, Twitter implemented what Theis and Wang (2018) call 'deep saliency prediction networks' which would allow the platform to auto-crop images based on the idea of 'saliency,' that is, 'the most interesting part of the image' (Theis and Wang, 2018). As a result, the result of Twitter's regime of recognition is the automatic production of a particular kind of ready-made images, algorithmically processed and cropped.

In order to critically examine the implications of this, it is crucial to gain a deeper understanding of the technical aspect of Twitter's algorithm. This can be achieved through a reading of some of the computer science literature on saliency detection in deep neural networks. According to Ardizzone et al. (2013), the aim of visual saliency detection methods used by platforms such as Twitter is 'to build a saliency map that replicates the human visual system behaviour in the visual attention process' (pp. 1–2). Extracting a saliency map from an image is based on analyses of the way in which so-called 'interest points' or fixation points are distributed in the image. In other words, images are presumed to be constitutive of salient regions which are surrounded by 'unnecessary background areas' (Ardizzone et al., 2013: 2). A region is considered to have high saliency when a person looking at an image is more likely to focus on that particular area rather than another. As such, a saliency map can be conceptualised as a 'normalized probability distribution over fixation locations' (Theis et al., 2018: 2). In short, what areas or regions of an image are people likely to fixate on.

Moreover, as Li and Yu (2015) point out, saliency is derived from ‘visual contrast as it intuitively characterizes certain parts of an image that appear to stand out relative to their neighboring regions or the rest of the image’ (p. 5455). The deep neural network used for image-cropping is trained using regions from a set of already labelled saliency maps, which allow the algorithmic model to evaluate the pixel contrasts between different image regions. This, in turn, enables the neural network to predict what aspects of an image can be considered salient or irrelevant to the user. Moreover, through training, Li and Yu (2015) claim that the neural networks become capable of ‘inferring the saliency score of every image region from the multiscale CNN features extracted from nested windows surrounding the image region’ (p. 5456). Through exposure to training data, in other words, neural networks iteratively adjust their weights or parameters to better be able to recognise salient regions in images, which in turn, enable them to infer a saliency score from these regions.

As the literature on saliency detection algorithm shows, there are some powerful assumptions at play in Twitter’s algorithm, assumptions which have political implications. First, there is an assumption that some things are worthy of attention and some things are not – and that this can be inferred predominantly from the pixel values of an image. The algorithm processes social media images in terms of a binary relation; determining salient or redundant regions. As a result, the extent to which an image is automatically cropped, and the extent to which certain features are made relevant in an image, is predicated on this binary or Boolean selection process. As Theis and Wang (2018) suggest, the Twitter neural network algorithm uses these inferences ‘to center a crop around the most interesting region’ of a given image, predicting what users are most likely wanting to see. In other words, saliency detection allows Twitter’s deep convolutional neural network to learn a saliency model, identify what it thinks is the most important regions or areas of an image, cropping and resizing it to fit with the platform it is uploaded to while discarding areas of the image considered irrelevant.

Second, there is an assumption that Twitter should take an active role in algorithmically cropping people’s images, rendering certain elements of images redundant. The underlying attitude seems to be: since we have an algorithmic model for it, we *should*. Through these following two assumptions, we see a particular regime of recognition emerge, one based on the supposed capacity of neural networks to identify, recognise and infer the saliency score of particular regions of social media images and to automatically crop them accordingly. It is a regime predicated on both the algorithmic capacity to crop images and the proactive role assumed by social media platforms in ‘cleaning’ people’s uploaded images. In the following section, I want to argue that the social power to arbitrate recognisability produces not only novel modes of recognition but also nonrecognition, that is, ways in which people are rendered both visible and invisible by algorithmic media. Afterwards, I want to suggest that the regime of recognition on Twitter can be seen to condition the parameters of users’ perceptual experience within the platform, determining what is included in images and what is discarded as redundant.

Regimes of nonrecognition: when algorithms ‘look through’

As the example from Twitter shows, algorithms have the capacity to render people recognisable, which, in turn, determines who comes to matter within the platform. Yet, it

also demonstrates that for this to work they must also arbitrate nonrecognisability, shaping the conditions of who is not made visible through saliency detection. In terms of Twitter's neural network, the notion saliency detection is grounded in the fundamental assumption that social media images are composed of salient regions that are, according to Ardizzone et al. (2013) 'surrounded by unnecessary background areas' (p. 2). Of course, in terms of human vision, Jonathan Crary (1999) points out that attention is a process of selection: 'an activity of *exclusion*, of rendering parts of a perceptual field unperceived' (pp. 24–25, original emphasis). For Crary, the annulment of some sensory data is integral to human perception. Yet, regimes of recognition establish the conditions by which perception as an activity of exclusion is delimited, circumscribed and enacted within social platforms. In other words, what or who must count as an 'unnecessary background area' is a highly political question, because it impacts how neural networks are weighted and what kinds of outputs or auto-crops they produce.

On one level, this activity of exclusion is predicated on the dataset used to train the deep learning algorithm. As Amoores (2020) writes, 'whether someone or something can be recognized depends on what the algorithm has been exposed to in the world' (p. 72). As such, this brings to the fore issues of bias and prejudice and how they may be baked into datasets (Angwin et al., 2016; Benjamin, 2019; Eubanks, 2018; Noble, 2018; O'Neil, 2016). Through a study of facial recognition systems, Buolamwini and Gebru (2018), for instance, argue that these repeatedly miscategorised darker-skinned woman at a much higher rate than any other group. In other words, these scholars argue that algorithms encode patterns of racial oppression, while also reflecting entrenched socio-cultural prejudices and inequalities in society.

Twitter's saliency detection algorithm highlights what Crawford and Paglen (2019) call 'the politics of images in machine learning training sets'. In their exhibition *Excavating AI*, Crawford and Paglen (2019) present the various datasets, composed of images scrapped from places such as Instagram, used to train facial recognition systems. Even though these images display humans, Crawford and Paglen argue that these are mostly looked at by algorithms. They are used to train facial recognition how to see humans. As Paglen (2014) has argued elsewhere, such images are predominantly 'operational', that is, they do not represent the particular contours of humans but are rather part of larger computational processes and operations. These datasets comprised of images, Crawford and Paglen suggest, can therefore be used to not only learn algorithms to recognise the presence of individuals but also to detect and infer emotional states, which has been highly critiqued in recent scholarship (McStay, 2018). They can also be heavily skewed in terms of who they represent and, more importantly, who they do not represent. The question of who is excluded from a training dataset (e.g. minority populations) is not only an issue of inclusion and diversity; rather, I argue that it is a fundamentally ethico-political question for it delimits the scope of what an algorithm is capable of recognising. As a result, algorithmic systems constitute both regimes of recognition and nonrecognition, determining who or what counts or comes to matter and who does not.

Yet, it is also crucial to look at the underlying processes by which algorithms learn to recognise certain features and not others. As Louise Amoores (2020) states, 'As the algorithm adjusts itself according to the specific object features in the training data as well as the weights of its own calculations, it is becoming the contemporary condition of

recognizability as such' (p. 73). In the case of Twitter's saliency detection algorithm, the politics of recognition is not merely a question of utilising biased datasets to train the algorithm. It is also, importantly, a question of the ways in which neural network algorithms are weighted, what the parameters are, how pixel values are calculated within individual images, as well as what is considered the threshold of saliency. These questions determine what visual features within an image an algorithm is able to see and recognise. Thus, what Taina Bucher (2012b) calls 'the threat of invisibility' on social media platforms is seen here to be predicated not only on users' levels of participation (or lack thereof), but also on a certain arrangement of probability weightings and parameters within the algorithm that produces a certain output. The bias is not merely located in the historical data used to train algorithms. In fact, they are also located on a more granular level, in the parameters, weightings and thresholds of neural networks. These processes also determine what counts as salient within images and what is automatically cropped out.

As algorithmic media such as Twitter arbitrate recognisability within their platforms, they increasingly shape the boundaries of what can be recognised and what is rendered nonrecognisable, and hence, the boundaries of what users are able to see. This brings to the fore questions regarding the relationship between social media platforms and opportunities of political recognition in society. For instance, Jordan Crandall (2005) suggests that

Being-seen is an ontological necessity; we strive to be accounted for within the dominant representational matrices of our time. We are not only talking about a gaze that is intrusive and controlling. We are talking about a gaze that provides the condition for action – the gaze for which one acts. (p. 20)

In other words, to be seen is not always a matter of surveillance. Rather, it is also a fundamental human need. Of course, fights for equal recognition, especially for minorities, has a long history in Western society (Honneth, 1992; Taylor, 1994). Yet, in the age of algorithms, struggles for recognition increasingly take place at the level of the pixel. They take place at a level of granularity that is imperceptible to the human eye, below the thresholds of human perception. It takes place at the level where the weights and parameters of neural network algorithms are tweaked and optimised. To be seen on social media is not merely a question of how one participates within it. To be seen on social media is to be seen and recognised by the algorithm at the level of the pixel.

As the example of Twitter's saliency detection algorithm also shows, to be *looked through* – to be rendered nonrecognisable and invisible – is a crucial implication of the capacity to arbitrate recognisability. Here, paradoxically, the weights and parameters of Twitter's saliency detection algorithm have the capacity to render certain people as weightless, as invisible, as un-outputtable, as without matter. As the algorithm disregards users' embodied, socio-cultural contexts in its processing of images, it is unsurprising that it is seen to depict a disembodied image, an image where the person was not recognised as such at the level of the pixel. In short, the social power of regimes of recognition on algorithmic media consists not only in how they see people, objects and events, but also in ways in which they look through them. Social media platforms produce not only modes of recognition but also modes of nonrecognition.

These processes, through the ever-pervading integration of recognition technologies, are normalised on social media platforms. Rather than seeing incidents like the one on Twitter as an ‘accident’ or as a deliberate act of racial prejudice, there is a need to see it as an inevitable outcome of the logic of regimes of recognition. Rather than being an ‘accident’, the incident on Twitter highlights the power of regimes of recognition. The politics of looking through is baked into the fundamental logic of the algorithm. In order for some aspects of images to be salient and recognisable, others have to be rendered unimportant and nonrecognisable. As a result, regimes of recognition are always simultaneously regimes of nonrecognition, and their operations on algorithmic media demonstrate the extent to which the social world can and should be rendered visible and thus actionable, which in this case has implications for struggles for political recognition.

The automatic production of ‘consistent’ lines of sight

Regimes of recognition on algorithmic media not only have the capacity to render certain people nonrecognisable; they also have the power to circumscribe people’s parameters of attention and perception more generally. The idea of attention is generally understood as ‘the cognitive processes of selecting and focusing upon certain aspects of information while ignoring others’ (Fazi, 2019: 3). In the context of Twitter’s saliency detection algorithm, more specifically, I argue that the capacity to arbitrate recognisability is the power to shape people’s attention through the algorithmic construction of what I call ‘consistent’ lines of sight.

This becomes evident when revisiting the blog written by Theis and Wang (2018). As they state, the saliency detection algorithm used by Twitter to automatically crop uploaded images forms part of the platform’s efforts to ‘improve consistency and to allow you to see more Tweets at a glance’ (Theis and Wang, 2018). The idea here is that the algorithm is supposed to enable a more frictionless, smooth and consistent user experience of the platform. Yet, what is meant by ‘consistent’ here? The word is by no means politically neutral. Rather than merely signifying a better user experience, I suggest that the notion of ‘consistency’ in this context is indicative of the sort of work that is currently being done by regimes of recognition on algorithmic media. For instance, the production of perceptual consistency is fundamentally a political question, because it is predicated on the assumption that some features within images can and should be rendered visible and others can and should be automatically cropped. As such, Twitter’s regime of recognition operates as what Taina Bucher (2012a) calls a ‘technicity of attention’, participating in shaping, organising and automating users’ parameters of attention and perception – creating a subject that is productive, predictable and ‘consistent’.

How are these so-called ‘consistent’ lines of sight generated? In what follows, I argue that they are the product of two interweaving processes: separating and binding. In terms of the first process, separating, Twitter’s regime of recognition can be conceptualised as a mode of ‘algorithmic attentiveness’ (Amoore, 2009). By this, Amoore means a particular attentiveness to the world that breaks up the visual field of an image solely in terms of pixel values. This algorithmic attentiveness is ‘a means of apportioning, segregating, singling out for our collective attentions’ (Amoore, 2009: 19). Similarly, Twitter’s regime of recognition functions by separating the images one encounters on social media from

the social world and people's lived experiences. By processing, analysing and automatically cropping images before they enter people's field of vision, the saliency algorithm is a means by which social media platforms relocate vision 'to a plane severed from a human observer' (Crary, 1990: 1).

This means that social media images can be understood increasingly as 'disrelated images' (Denson, 2020), referring less to socio-cultural realities and embodied subjects and more to way these images are 'seen' and processed by algorithms as well as what regions of images algorithms are able to recognise. Echoing Jonathan Crary (1999), regimes of recognition on social media platforms are 'not primarily concerned with looking at images but rather with the construction of conditions that individuate, immobilize, and separate subjects, even within a world in which mobility and circulation are ubiquitous' (p. 74). In other words, what parts of a social media image are unrecognised by the algorithm is less a result of an embodied process of attention and more a direct by-product of the ways in which social media platforms operate as arbiters of recognisability. Here, the visual is further 'abstracted' (Crary, 1990) by algorithmic technologies, severed from embodiment and the socio-cultural contexts from which it arises. As such, I argue that the politics of regimes of recognition on social media can be said to constitute a continual effort to separate and unmoor users' perceptual experience from their messy, socio-cultural contexts and ways of seeing the world, and instead algorithmically produce 'consistent' and frictionless lines of sight. Consistency, as proposed by Theis and Wang (2018), therefore, constitutes a set of algorithmic 'dividing practices that render ways of life economic, make them amenable to management, trading, or exchange' (Amoore, 2009: 19).

Furthermore, I argue that the production of consistent lines of sight also involves processes of binding. What is meant by this? Twitter's neural network algorithm is predicated on predefined notions of relevance and saliency. That is, pixel values within a social media image are processed in relation to underlying assumptions about what constitutes meaningful regions in an image, which in turn determines their relevance to the overall saliency score of the image. The notion of binding therefore foregrounds another important way in which consistent lines of sight are generated. The regime of recognition on Twitter not only unmoors people's perceptual experiences of social media images from their lived experiences and socio-cultural contexts; it also *rebinds* them in relation to what can be seen and recognised by the algorithm. Algorithmic modes of attentiveness operate at a level that is imperceptible to the human eye, namely the pixel. Users' perceptual experiences are, therefore, organised and fixed in relation to this logic. As such, Twitter's regime of recognition can also be understood as 'operations of fixing, of fastening' (Crary, 1999: 332). Yet, this operation of fixing is by no means neutral for what the algorithm perceives as recognisable and thus 'normal' (redundant regions become equivalent to deviations from this norm and are therefore automatically cropped). As such, Twitter's regime of recognition engenders a new valuation of visual experience; one which aligns itself with the logic of consistency, frictionlessness and what meanings algorithms can extract from pixel values.

However, the production of consistent lines of sight by regimes of recognition on algorithmic media should not be reduced to a set of technologically deterministic processes and effects. This is because, first, it is important to see social media platforms

comprising an intricate ‘attention ecology’ (Citton, 2017), where attention is distributed among both human users and algorithmic systems. In this view, regimes of recognition constitute an aspect of this attention ecology. For instance, Fazi (2019) argues that ‘we think alongside machines that are, in a sense, already thinking. Similarly, we pay attention alongside machines that are, in a sense, already paying attention’ (p. 85). The result of this plethora of attentive agents within present ecologies, Fazi argues, is a restructuring of the conditions for the human capacity to pay attention (Fazi, 2019: 99).

Social media platforms can thus be said to constitute ecologies of diverse human and nonhuman modalities of attention, intimately entangled. For instance, where humans would quite easily see a person in an image, the neural network algorithm recognises them based on the way it has been trained and how connections within the network are weighted. For the algorithm, a person is equivalent to a cluster of salient regions recognised at the level of the pixel. In other words, the particular regimes of recognition of algorithmic media such as Twitter and Facebook further introduce and perpetuate, as Mark B. Hansen (2015: 4) put it, ‘levels of operationality that impact our experience without yielding any perceptual correlate’.

Second, the production of consistent lines of sight should not be understood in deterministic terms given the tension that these overlapping agencies and diverse modalities of attention may produce on social media. Indeed, I suggest that algorithmic media can be understood as *spaces of perceptual contestation*. As Kate Crawford (2016) argues, ‘algorithmic decision making is always a contest’ (p. 87). That is, ‘the spaces of intersection between humans and algorithms can be competitive and rivalrous, rather than being purely dictated by algorithms that are divorced from their human creators’ (Crawford, 2016: 82). Here, Twitter’s regime of recognition can be seen to engender a similar space of contestation in terms of what can and should be seen in social media images. In arbitrating recognisability, or failing to do so, it produces instances of tensions, mishaps and incidents of racism. A way to critically examine the politics of regimes of recognition on algorithmic is therefore to pay attention to such instances where arbitrating recognisability on social media seems to fail. Yet, they are not simply instances of algorithmic aberrations. As Louise Amoore (2020) argues, they are instead moments in which algorithms ‘give account of themselves’, that is, their politics and their intrinsic logic. The incident with Twitter’s saliency detection algorithm can therefore be understood as highlighting the social power of regimes of recognition on algorithmic media and how they seek to produce consistent lines of sight.

Conclusion

In this article, I have sought to showcase how algorithms may shape how we see and experience the world. This article has sought to add to our understanding of the power and politics of algorithms through the way in which they shape people’s perceptual experiences on social media through the algorithmic reconfiguration of images. Drawing on Louise Amoore’s (2020) idea of ‘regimes of recognition’, I have argued that social media platforms can be understood as powerful ‘arbiters of recognisability’, shaping the conditions of what can be seen and what is rendered invisible. Drawing on an analysis of industry reports as well as computer science literature, this article provided an in-depth

study of Twitter's so-called deep saliency algorithm. It attempts to detect salient regions within an image and to automatically crop out what are considered redundant regions. As I pointed out, the algorithm is a mechanism to optimise and maximise user engagement with platform images, ensuring a continuous and frictionless shareability within the social media platform.

More specifically, I argued that social media's capacity to arbitrate recognisability is fundamentally political in the following two main ways: first, it constitutes the power to render people visible and invisible within platforms. Here, I show that what is shown in images uploaded to Twitter is partly a result of algorithmic processes of highlighting salient regions in images and automatically cropping out the redundant parts. This was seen to raise issues of racial bias, as the Twitter algorithm had been shown to crop out black faces because it did not recognise their being there. Rather than an accident, however, this case foregrounds the logic of the algorithm and the power of the capacity of arbitrating recognisability. That is, in order for some aspects of images to be salient and recognisable, others have to be rendered unimportant and nonrecognisable. According to this algorithmic logic, for some people to be rendered visible others have to be rendered redundant.

Second, I argued that the capacity to arbitrate recognisability is also the power to shape perception through the algorithmic construction of so-called 'consistent' lines of sight. The notion of consistency is not only Twitter's marketing talk but also highlights the underlying desire of social media platforms to algorithmically shape and organise people's parameters of attention, and in turn, create subjects that are productive, predictable and perceptually 'consistent'. I theorised that this attempt to create consistent lines of sight is the product of two interweaving processes, namely separating and binding. In other words, people's perception is separated from their socio-cultural realities and bound to the conditions by which algorithms see and recognise the world, namely pixel values.

As Jordan Crandall (2005) suggested, 'being-seen is an ontological necessity' (p. 20). It is a human need to be recognised by others, a need which 'provides the condition for action' (Crandall, 2005: 20). To be seen and recognised is crucial for how we act in the world. Yet, what I have attempted in this article is to argue that this ontological necessity for recognition is increasingly being arbitrated by algorithmic media. As our lives become ever more entangled with social media platforms, they increasingly determine the conditions of possibility for how we are rendered visible and/or invisible to others. What is seen and recognised by the algorithm is valued. In this view, Twitter's saliency algorithm is not only a point of interest because of its racist 'mistake'. Rather, I argue that the algorithm signifies a broader development in society, whereby how we see and are seen by others is increasingly arbitrated and shaped by algorithmic regimes of recognition. That is, social media platforms make the world appear in particular ways, shaping the conditions of how people can be recognised and seen. I argue that to be seen on social media is to be algorithmically recognised at the pixel level.

However, it is also crucial that this notion of 'regimes of recognition' not be seen in fixed terms. As Roland Barthes (1974: 11) put it, we need to be 'attentive to the plural' inherent in the concept, acknowledging the multiplicity of interlinking ways, practices, tools, techniques, algorithms and systems used by social media platforms to arbitrate

recognisability, that is, how they render people recognisable or nonrecognisable within their platforms. There is therefore a need to further explore ways in which users are made recognisable within platforms and what the social implications of this may be. For example, how are users seen and recognised by platforms such as Facebook, Instagram, YouTube or TikTok? What particular algorithmic techniques are used to know the user more intimately and render them visible or invisible on the platform? These questions remain unexplored in this article, but they remain fruitful avenues for future research. As people perceive within ‘a prescribed set of possibilities’ (Crary, 1990: 6), there is an ongoing need to critically investigate the algorithmic regimes used to circumscribe this set of possibilities.

Author Note

Benjamin N Jacobsen is now affiliated to Department of Geography, Durham University, Lower Mountjoy, Durham, UK.

Funding

The author(s) received no financial support for the research, authorship and/or publication of this article.

ORCID iD

Benjamin N Jacobsen  <https://orcid.org/0000-0002-6656-8892>

References

- Agostinho D (2019) The optical unconscious of Big Data: datafication of vision and care for unknown futures. *Big Data & Society*. Epub ahead of print 5 February. DOI: 10.1177/2053951719826859.
- Amoore L (2009) Lines of sight: on the visualization of unknown futures. *Citizenship Studies* 13(1): 17–30.
- Amoore L (2020) *Cloud Ethics: Algorithms and the Attributes of Ourselves and Others*. Durham, NC: Duke University Press.
- Ananny M and Crawford K (2018) Seeing without knowing: limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society* 20(3): 973–989.
- Andrejevic M and Selwyn N (2019) Facial recognition technology in schools: critical questions and concerns. *Learning, Media and Technology* 45(2): 115–128.
- Angwin J, Larson J, Mattu S, et al. (2016) Machine bias. *ProPublica*, 23 May. Available at: <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Ardizzone E, Bruno A and Mazzola G (2013) Saliency based image cropping. In: Petrosino A (ed.) *Image Analysis and Processing – ICIAP*. Berlin: Springer-Verlag, pp. 773–782.
- Barthes R (1974) *S/Z* (trans. R Miller). New York: Hill and Wang.
- Beer D (2009) Power through the algorithm? Participatory web cultures and the technological unconscious. *New Media & Society* 11(6): 985–1002.
- Beer D (2017) The social power of algorithms. *Information, Communication & Society* 20(1): 1–13.
- Beer D (2019) *The Data Gaze: Capitalism, Power and Perception*. London: SAGE.
- Benjamin R (2019) *Race after Technology: Abolitionist Tools for the New Jim Code*. New York: Wiley.

- Bucher T (2012a) A technicity of attention: how software ‘makes sense’. *Culture Machine* 13: 1–23.
- Bucher T (2012b) Want to be on top? Algorithmic power and the threat of invisibility on Facebook. *New Media & Society* 14(7): 1164–1180.
- Bucher T (2018) *If. . . Then: Algorithmic Power and Politics*. Oxford: Oxford University Press.
- Bueno CC (2019) The face revisited: using Deleuze and Guattari to explore the politics of algorithmic face recognition. *Theory, Culture & Society* 37(1): 73–91.
- Buolamwini J and Gebru T (2018) Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research* 81: 1–15.
- Burrell J (2016) How the machine ‘thinks’: understanding opacity in machine learning algorithms. *Big Data & Society*. Epub ahead of print 6 January. DOI: 10.1177/2053951715622512.
- Cheney-Lippold J (2011) A new algorithmic identity: soft biopolitics and the modulation of control. *Theory, Culture & Society* 28(6): 164–181.
- Cheney-Lippold J (2016) Jus algoritmi: how the national security agency remade citizenship. *International Journal of Communication* 10: 1721–1742.
- Citton Y (2017) *The Ecology of Attention*. Cambridge: Polity Press.
- Crandall J (2005) Envisioning the homefront: militarization, tracking and security culture. *Journal of Visual Culture* 4(1): 17–38.
- Crary J (1990) *Techniques of the Observer: On Vision and Modernity in the Nineteenth Century*. Cambridge, MA and London: MIT Press.
- Crary J (1999) *Suspensions of Perception: Attention, Spectacle, and Modern Culture*. Cambridge, MA: MIT Press.
- Crawford K (2016) Can an algorithm be agonistic? Ten scenes from life in calculated publics. *Science, Technology & Human Values* 41(1): 77–92.
- Crawford K and Paglen T (2019) Excavating AI: the politics of images in machine learning training sets. Available at: <https://www.excavating.ai/>
- Denson S (2020) *Discorrelated Images*. Durham, NC and London: Duke University Press.
- Eubanks V (2018) *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. New York: St. Martin’s Press.
- Fazi MB (2019) Distraction machines? Augmentation, automation and attention in a computational age. *New Formations: A Journal of Culture, Theory, Politics* 98: 85–100.
- Flyverbom M (2019) *The Digital Prism: Transparency and Managed Visibilities in a Datafied World*. Cambridge: Cambridge University Press.
- Gillespie T (2010) The politics of ‘platforms’. *New Media & Society* 12(3): 347–364.
- Gillespie T (2014) The relevance of algorithms. In: Gillespie T, Boczkowski PJ and Foot KA (eds) *Media Technologies: Essays on Communication, Materiality, and Sociality*. Cambridge, MA and London: MIT Press, pp. 167–193.
- Hallinan B and Striplas T (2016) Recommended for you: the Netflix prize and the production of algorithmic culture. *New Media & Society* 18(1): 117–137.
- Hand M (2012) *Ubiquitous Photography*. Cambridge: Polity Press.
- Hansen MB (2015) *Feed-Forward: On the Future of Twenty-First-Century Media*. Chicago, IL and London: The University of Chicago Press.
- Hern A (2020) Twitter apologises for ‘racist’ image-cropping algorithm. *The Guardian*, 21 September. Available at: <https://www.theguardian.com/technology/2020/sep/21/twitter-apologises-for-racist-image-cropping-algorithm>
- Honneth A (1992) *Struggle for Recognition: The Moral Grammar of Social Conflicts*. Cambridge: Polity Press.
- Kitchin R (2017) Thinking critically about and researching algorithms. *Information, Communication & Society* 20(1): 14–29.

- Kotliar DM (2020) Data orientalism: on the algorithmic construction of the non-Western other. *Theory and Society* 49: 919–939.
- Latimer J and Gómez LD (2019) Intimate entanglements: affects, more-than-human intimacies and the politics of relations in science and technology. *The Sociological Review Monograph* 67(2): 247–263.
- Li G and Yu Y (2015). Visual saliency based on multiscale deep features. In: *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, Boston, MA, 7–12 June 2015, pp. 5455–5463. Piscataway, NJ: IEEE.
- Lowensohn J (2014) Facebook's working on facial verification that's 'nearing human-level performance'. *The Verge*, 17 March. Available at: <https://www.theverge.com/2014/3/17/5518808/facebook-working-on-facial-verification-thats-nearing-human-levels>
- Mackenzie A and Munster A (2019) Platform seeing: image ensembles and their invisualities. *Theory, Culture & Society* 36: 3–22.
- McStay A (2018) *Emotional AI: The Rise of Empathetic Media*. London: Sage.
- Murdock J (2020) Twitter CDO acknowledges 'racist' photo algorithm: 'it's 100 percent our fault'. *Newsweek*, 21 September. Available at: <https://www.newsweek.com/twitter-image-preview-neural-network-racial-bias-dantley-davis-1533195>
- Noble S (2018) *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press.
- O'Neil C (2016) *Weapons of Math Destruction*. New York: Broadway Books.
- Paglen T (2014) Operational images. *E-Flux* 59: 1–3.
- Su J, Vargas DV and Sakurai K (2019) One pixel attack for fooling deep neural networks. *IEEE Transactions on Evolutionary Computation* 23(5): 828–841.
- Taylor C (1994) The politics of recognition. In: Gutmann A (ed.) *Multiculturalism*. Princeton, NJ: Princeton University Press, pp. 25–75.
- Theis L and Wang Z (2018) Speedy neural networks for smart auto-cropping of images. *Twitter Engineering*, 24 January. Available at: https://blog.twitter.com/engineering/en_us/topics/infrastructure/2018/Smart-Auto-Cropping-of-Images.html
- Theis L, Korshunova I, Tejani A, et al. (2018) Faster gaze prediction with dense networks and fisher pruning. Available at: <https://arxiv.org/abs/1801.05787>
- Uricchio W (2011) The algorithmic turn: photosynth, augmented reality and the changing implications of the image. *Visual Studies* 26(1): 25–35.
- Willson M (2017) Algorithms (and the) everyday. *Information, Communication & Society* 20(1): 137–150.

Author biography

Dr Benjamin N Jacobsen is a Postdoctoral Research Associate on the Algorithmic Societies project at Durham University. His book *Social Media and the Automatic Production of Memory* (co-authored with David Beer) was published in April 2021.