

CLASSIFYING HUMANS: THE INDIRECT REVERSE OPERATIVITY OF MACHINE VISION

Linda Kronman

To cite this article: Linda Kronman (2023) CLASSIFYING HUMANS: THE INDIRECT REVERSE OPERATIVITY OF MACHINE VISION, *photographies*, 16:2, 263-289, DOI: 10.1080/17540763.2023.2189160

To link to this article: <https://doi.org/10.1080/17540763.2023.2189160>



© 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



Published online: 25 May 2023.



Submit your article to this journal [↗](#)



Article views: 174



View related articles [↗](#)



View Crossmark data [↗](#)

ARTICLE

Linda Kronman 

CLASSIFYING HUMANS: THE INDIRECT REVERSE OPERATIVITY OF MACHINE VISION

Classifying is human. Classifying is also what machine vision technologies do. This article analyses the cybernetic loop between human and machine classification by examining artworks that depict instances of bias when machine vision is classifying humans and when humans classify visual datasets for machines. I propose the term ‘indirect reverse operativity’ – a concept built upon Ingrid Hoelzl’s and Remi Marie’s notion of ‘reverse operativity’ – to describe how classifying humans and machine classifiers operate in cybernetic information loops. Indirect reverse operativity is illustrated through two projects I have co-created: the Database of Machine Vision in Art, Games and Narrative and the artwork Suspicious Behavior. Through ‘artistic audits’ of selected artworks, a data analysis of how classification is represented in 500 creative works, and a reflection on my own artistic research in the Suspicious Behavior project, this article confronts and complicates assumptions of when and how bias is introduced into and propagates through machine vision classifiers. By examining cultural conceptions of machine vision bias which exemplify how humans operate machines and how machines operate humans through images, this article contributes fresh perspectives to the emerging field of critical dataset studies.

Introduction

‘To classify is human.’¹ Classifying is also what machine vision is trained to do. When Adrian Mackenzie writes about ‘machine learners,’ he turns machine learning into

CONTACT Linda Kronman linda.kronman@uib.no Department of Linguistic, Literary and Aesthetic Studies, University of Bergen, Sydnesplassen 7,5007 Bergen, Norway

photographies, 2023

Vol. 16, No. 2, 263–288, <https://doi.org/10.1080/17540763.2023.2189160>

© 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

a concept which refers both to humans and machines. After all, as Mackenzie asserts: ‘Machine learners are often simply called “classifiers.”’² His statement highlights how ‘learning’ in machine learning often corresponds to training the machine to classify. In machine learning for computer vision, classifiers learn from visual data. And like all data, visual data ‘must be classified in some way to be put to use.’³ In the context of machine vision, data becomes useful when images are collected and classified into datasets. There are many ways to assemble visual datasets. Typically, vast amounts of images or videos are downloaded from the internet. Then classifying humans — data curators and crowd-sourced on-demand annotators — translate those images into machine readable data. Classification is thus a core practice of developing artificially intelligent machine vision.⁴ Hence, in this article ‘classifying humans,’ like Mackenzie’s ‘machine learners,’ refers to both human and machine classifiers.

Classifying humans are in the centre of attention as this article investigates cultural conceptions of machine vision bias. Bias is an overloaded term which is defined differently in different disciplines. It comes with multiple meanings even when the context is narrowed to machine learning. In this article bias is understood through an alignment of operations on images. For example, machine vision bias arises when a machine classifier misgenders images of faces. This is often due to a biased or non-representative training dataset underrepresenting the misgendered population. Dataset bias in turn aligns with humans performing bias when classifying and collecting images into datasets. For example, Trevor Paglen and Kate Crawford have demonstrated that ImageNet, which many models are trained on, includes highly problematic classification of images of people.⁵ Particularly image recognition technologies deployed to classify humans tend to lead to harmful bias showing prejudice against individuals or groups of people.⁶ Acknowledging that different types of biases arise throughout a machine learning models lifecycle, this article is limited to addressing bias in the dataset. Biased datasets can be caused by lack of representation, as when facial recognition algorithms are trained primarily on white men. The biased representation can be a historical artefact. For instance, ImageNet uses WordNet’s classification system from the 1990s, which includes stereotypes that would not be acceptable today.⁷ We could also imagine how even an unlabelled training dataset from the 1950s would visually portray gender roles as quite different from today. Such dataset bias can cause representational harm such as reinforcing stereotypes. On the other hand, in the near end of this article in my discussion of *Suspicious Behavior* I address instances in which classifying humans perform bias when annotating visual datasets. This entails both how humans interpretate and are instructed to interpretate images. It involves annotation interfaces which facilitate perception at scale and speed. And it is about annotators learning to ‘perceive on the bias’ meaning that they embody ‘a temporal and rhythmical apparatus’ to interpretate images.⁸

Artworks and artistic research have been influential in communicating to a wider audience that visual datasets are fundamental for machine perception.⁹ Artists have audited, exposed, excavated, and exhaustively watched publicly available visual datasets and brought to attention the lack of diversity in datasets, revealing prejudiced and racist taxonomies and privacy concerns in dataset assembly.¹⁰ Thus, artworks provide a rich context to examine the following questions: Why does dataset bias lead to harmful machine vision? When is bias introduced into machine vision? And why does machine vision bias propagate? I engage with these questions of machine vision

bias both as a researcher and an artist through two projects I have been involved in. The first project I will discuss is the ‘Database of Machine Vision in Art, Games and Narratives’ (henceforth, Machine Vision Database).¹¹ The second is *Suspicious Behavior* (2020), an artwork I co-created with Andreas Zingerle (as the artist duo KairUs).¹²

I start by analysing four artworks that I propose viewing as ‘artistic audits.’ I have chosen four ‘artistic audits’: Max Dovey’s performance *How to be More or Less Human* (2015), Joy Buolamwini’s video work *AI, Ain’t I a Woman?* (2018), Nouf Aljowaysir’s artwork *Salaf* (2020), and Nakeema Stefflbauer and Nushin Isabelle Yazdani’s artistic video essay *Future Tense: AI from the margins* (2020). These four artworks illustrate why and how representation and historical bias leads to harmful discrimination when machine vision is deployed. For a broader understanding of cultural conceptions of classifying humans, I turn to the 500 creative works in the Machine Vision Database, and use data visualisations to explore who is classifying and who gets classified in what we call ‘Machine Vision Situations.’ This data analysis allows me to explore patterns that emerge out of interactions between technologies and other agents implying how bias is conceptualized in a larger corpus of creative works.

Finally, I turn to my second project, *Suspicious Behavior*, a speculative annotation tutorial, which contributes the fresh perspective of artistic research to the emerging field critical dataset studies. Artworks and artistic research have mainly focused on scrutinizing the content (images, labels, and taxonomies) of visual datasets. Less has been done to understand how and when labels get attached to images. My reflections upon the design of an annotation interface for *Suspicious Behavior* complicates general assumptions of how and when bias is introduced into datasets. In the sections analysing the four artistic audits and the data from the Machine Vision database, my focus is on bias that arises when machine vision operates on humans through images. Then, in my artistic research, I change perspective and bring to attention humans who perform bias by operating machine classifiers through images.

Both when machines classify images of humans and when humans classify images for datasets, images are not primarily interpretive or affective artifacts, but *operational*: they are means to an end. Based on Harun Farocki’s frequently cited definition of ‘operative images’ as ‘images that do not represent an object, but rather are part of an operation,’¹³ images have been understood to hold different degrees of operativity.¹⁴ Images discussed in this article are operationalized as visual data in the context of supervised or semi-supervised machine learning. They thus fit the narrowest definitions of operativity since they function as elements in technical processes. Expanding on Farocki’s ‘operative image,’ Ingrid Hoelzl and Rémi Marie introduce the concept of ‘reverse operativity.’ Using Google Street View (GSV) as an example, they bring attention to cybernetic information loops in which humans operate upon images which in reverse operate back on humans:

In the case of GSV, operation which is not restricted to user navigation, but which is part of a larger circular operation of data exchange, with the users’ trajectories feeding back into the database. This ‘reverse operativity’ reveals the more problematic side of the algorithmic turn: For if we are operating GSV images, they are at the same time operating us.¹⁵

Hoelzl and Marie use the term ‘reverse operativity’ to describe how images are ‘aimed at us.’¹⁶ For instance, in commercial context images are aimed at us when they are used to trace mobility patterns of customers. Each time we operate GVS images, data is collected. Machine classifiers make predictions based on the data, and aim it back at us, e.g. as personalized advertisement. In other words, reverse operativity describes instances in which ‘humans operate machines and machines operate humans through images’¹⁷

Hoelzl’s and Marie’s example of reverse operativity implies that Google’s images instantaneously operate directly back on the person operating an image. However, when humans operate on images and make them functional for machine learning, then reverse operativity acts differently. The images are not instantaneous nor directly operating back on the person operating the image. Instead, images operate back on humans in indirect ways. Thus, building upon Hoelzl’s and Marie’s example of reverse operativity I suggest the term ‘indirect reverse operativity.’ The term arises from how classifying humans (data curators and annotators) operate on images and those images indirectly operate back on humans when machine vision technologies are deployed to classify us.

The main differences between ‘reverse operativity’ and ‘indirect reverse operativity’ is in the distribution, temporality, and the direction of operations. Classifying images for machine learning is a distributed operation involving both human and machine interpretation. I return to this point in the section on *Suspicious Behavior*. The reversal of an operation — the image aiming back at us — is not instant. There can be long gaps between classifying images for a dataset and the deployment of a machine vision product trained on that dataset. When machine vision is finally deployed, it is not aimed at those who classified the image. Instead, reverse operativity is directed towards humans in general. Due to this indirectness of reverse operativity harmful machine vision bias is often discovered only after machine vision is deployed. Hence, I start by analysing artistic audits which test and evaluate machine vision bias and discuss why machine vision bias leads to harmful discrimination particularly when humans are classified. After pinpointing harms that arise when machines operate through images on humans, I change point of view to reflect upon how classifying humans operate on images.

What emerges when analysing these artworks as instances of indirect reverse operativity is that bias propagates in cybernetic loops. Although this article is focused on a subcategory of AI — namely, machine vision — classification as production of knowledge and a mechanism of knowledge construction is of concern for the whole field of AI. Hence, the concept of indirect reverse operativity could be further expanded to think more generally about the ways we operate on data and how that data operates back on us. By reverse engineering cybernetic loops of machine vision bias this article demonstrates that classifying humans are still central in defining how machines perceive and classify the world. Critical dataset studies aim to ‘bring people back in’ to our understandings of datasets, as Nanna Bonde Thylstrup has argued it ‘requires us to think through how to visibilise humans in machine cultures, but also raises ethical questions about how to encounter these humans with empathy and care.’¹⁸

The concept of indirect reverse operativity enables this, supporting nuanced analysis of datasets that makes human actions visible. The concept is particularly useful as data curation and annotation are increasingly automated, making human actions less obvious.

Artistic audits — when machine vision bias leads to harmful discrimination

Artistic audits, I propose, are artworks which bring attention to the constructed ideological borders that human classifiers have encoded into classification products. Artistic audits demonstrate how certain types of automated bias lead to discriminating machine vision behaviour. Here I outline a few examples of artistic audits to articulate why we need to pay attention to indirect reverse operativity. The chosen artistic audits show how human decisions to classify images play out when they operate back on us. For example, in his performance *How to be More or Less Human*, Max Dovey plays with how gender stereotypes are encoded into machine vision. ‘Max Dovey’¹⁹ is dressed in a suit and is posing in front of a web camera trying to be identified as a man. If ‘Dovey’ is dressed, the image recognition software installed in the camera classifies him ‘correctly’ as a man. However, when he starts undressing, an ambiguous threshold is crossed, and his body is classified as a woman. What I want to point out with this example is how a worldview ingrained in a dataset, in other words dataset bias, plays out when a machine learning model is trained on images that indirectly reverse operates when machine vision is deployed to classify humans.²⁰ In this case the machine vision application is biased to recognize nudity as a female attribute.

Because datasets are situated artifacts there can be many sources for such bias. A visual dataset is a product of a particular time and place. Thus, it matters when and from where images are collected. Historical bias becomes a source of harmful discrimination when a model reinforces gender or racial stereotypes for example in image search results often leading to representational harm.²¹ Representation bias, also called sample bias, occurs when types of images occur in a class more often than others. For example, ImageNet is contains a higher female-to-men ratio in classes containing nudity (e.g. bra, bikini, and maillot and pornographic images).²² The bias towards classifying nudity as a female attribute in *How to be More or Less Human* could also be caused by similar overrepresentation of females in images containing nudity. Bias can also arise from how data is classified. In the case of automated gender classification, applications typically recognize gender as binary male/female and all other gender expressions are bluntly ignored.²³ Classification requires choices; thus, some features are given more importance and consequently others less when a classifier predicts between chosen classes.

In her artwork *AI, Ain't I a Woman?* Joy Buolamwini depicts how iconic Black women are repeatedly misgendered by popular gender classification products. In Buolamwini’s video we see screenshots of how Google, Amazon, IBM, Microsoft, and Face++ label iconic Black women like Serena Williams, Michelle Obama and

Sojourner Truth as ‘male,’ ‘men,’ or ‘gentlemen’ when predicting externalized gendered appearance. Buoloamwini is also a computer scientist, and in an intersectional auditing of gender classification products Buolamwini further demonstrates racial bias in predicting gender. Tested systems performed well on light-skinned men, yet poorly on darker-skinned women.²⁴ Such bias is caused by historical overrepresentation of white males in facial recognition datasets. Lack of diversity in datasets turns into harmful bias when all other demographics are rendered as deviations from the white male norm when images operate back on us.

Artistic audits do not only depict gender and racial bias in machine vision classifiers. Nouf Aljowaysir’s artwork *Salaf* demonstrates how machine vision classifiers exhibit Americentric and Eurocentric bias.²⁵ When Aljowaysir was working with photographs from her ancestral archive she realized how image recognition software repeatedly misclassified subjects wearing traditional Saudi and Iraqi clothes. The algorithm attached labels like ‘military,’ ‘army,’ or ‘soldiers’ to the photographs, demonstrating how the Western colonial gaze produces stereotypes about the Arab world by encoding certain bodies as threats. As an artistic audit, *Salaf* and the other works demonstrate that if we ignore how images indirectly reverse operate on humans, then machine vision, like other technological systems, will continue to reproduce ‘the values of white supremacist heteropatriarchy, capitalism, ableism, and settler colonialism.’²⁶

Nakeema Stefflbauer and Nushin Isabelle Yazdani’s artwork *Future Tense: AI from the margins* articulates how machine vision bias leads to harmful discrimination. In this video work ‘Nakeema Stefflbauer’ explains why people like she, ‘not privileged,’ ‘not white,’ experience AI technologies as intrusive and unfair. ‘Stefflbauer’s’ monologue describes how her racialized, gendered body is ‘hypervisible’ as there are no ‘unoccupied,’ ‘unmonitored’ spaces to escape the gendered and racially biased technologies monitoring every aspect of her life. Yet, at the same time, she is ignored and invisible by design. Numerous examples of harmful bias are taken up in the video: machine vision regularly failing to correctly detect the faces of dark-skinned females, search engines classify images of Black women as apes or gorillas, smart cameras perceiving “‘Asian” eyes as closed,’ and emotion detection interpreting Black basketball players as angrier than white players.²⁷ What Stefflbauer describes are instances in which machine vision technologies are employed with the assumption of being more objective or neutral than prior technologies, yet they reproduce existing inequalities. This technological encoding of discrimination is what Ruha Benjamin in *Race After Technology* calls the ‘New Jim Code.’²⁸ In other words, what *Future Tense: AI from the margins* lists is one example of ‘New Jim Code’ after another. The examples of bias revealed in this artistic audit are not isolated cases of poorly designed machine vision classifiers. Machine vision systematically discriminates those in the margins. Stefflbauer and Yazdani’s almost eighteen-minute-long video ends with a feminist manifesto written into an editing tool. Referencing the ‘matrix of domination,’ a concept developed by Patricia Hill Collins in *Black Feminist Thought*,²⁹ the manifesto breaks down how the given examples reflect dominant power relations in ‘our capitalist, patriarchal, white supremacist society.’³⁰ In a way, the artwork collects evidence of how machine vision reproduces existing inequalities. It demonstrates that intersections of race, gender, and other identities shape how people

experience these technologies.³¹ Through ‘Stefflbauer’s’ persona the artwork articulates what is also highlighted in Virginia Eubanks’s book *Automating Inequality*: those oppressed by AI already belong to systematically marginalized populations.³²

These four artistic audits demonstrate how images indirectly operate back on humans with differing accuracy, and how this can lead to experiences of discrimination. Negative attitudes towards biased machine classifiers are not only apparent in these artworks, but also present in a larger sample of creative works including artworks, games and narratives in the Machine Vision Database which I turn to in the next section.

Data analysis – who is classifying and who gets classified in machine vision situations?

The Machine Vision database is a collection of 500 ‘Creative Works’ (190 artworks, 77 games and 233 narratives), and thus presents an opportunity to explore cultural conceptions of machine vision classification in a larger sample of works. I take this opportunity here to examine more broadly who is classifying or being classified in what we call machine vision situations, and what does the interactions between classifying agent reveal about indirect reverse operativity?

However, before getting to the data analysis, I need to make a short detour and explain a few things about the database. As a core member of the ‘Machine Vision in Everyday Life’ project I was involved in building this database, contributing mainly by collecting and analysing digital artworks and tagging them with interpretative metadata. I want to start by stressing that the corpus of works in the Machine Vision database cannot be claimed as fully representative. In addition, the following data analysis and visualisations are by necessity reductive and cannot express the full range of expression in each work collected into the database. My main unit of analysis is what we call ‘Machine Vision Situations’ (hereafter referred to in short form as ‘Situations’).³³ For a detailed description of our backgrounds, data collection practices, the database architecture, exported datasets, definitions, ethical concerns, and dataset biases see our data paper ‘Representations of Machine Vision Technologies in Artworks, Games and Narratives: A Dataset.’³⁴

The unit ‘Situations’ was developed to ‘capture granular details of what humans, technologies, and other agents are doing in specific interactions with machine vision technologies.’³⁵ The research team identified and registered instances of interaction with machine vision and entered structured information about agents based on their interpretation of the situation. There can be one or more ‘Situation’ in each ‘Creative Work’ in the database. [Figure 1](#) presents a diagram of the database structure and shows how ‘Verbs’ are used to describe the actions of machine vision technologies and other agents in this ‘Situation.’ In the following data visualisations, I focus on the verbs ‘Classifying’ and ‘Classified’ which are among the most commonly used actions used to describe interactions in ‘Situations’ (‘Classifying’ is used 155 times and ‘Classified’ 95 times)

Creative Work (500)

Suspicious Behavior

Scan to see the database page

Fields: Year, Creator, Country, Publication Type, Technologies referred or used, Topics, Sentiment, Characters, Description

Character(778)

Annotator (Suspicious Behavior)

Age: Adult
Species: Human
Gender: Unknown
Race/Ethnicity: Unknown
Individual or group?: Individual
Customizable?: Not customizable

Description: The user of the artwork is given the role of a image annotator (crowdsworker) who is labeling images for machine learning purposes.
Works the character appears in: Art, Online art

Publication Type: Art, Online art
Year: 2020

Fields: Age, Species, Gender, Race/ethnicity, Individual or group?, Customizable?

Machine Vision Situation (874)

Suspicious Behavior (labeling images)

Brief description
In this artwork the user becomes part of the AI pipeline taking the role of a underpaid crowdworker annotating images for machine learning purposes. The task is to decide if a situation is suspicious or not. The client (corporation) is giving certain guidelines of what authorities define as suspicious based on this the annotator needs to do quick interpretation relying rather on intuition than reflection. The annotator is told that the images will be used to train machines to detect suspicious behavior.

Work that the situation appears in

Title	Publication Type	Year	Creator
Suspicious Behavior	Art, Online art	2020	KairUs

Screenshots or images

Who does what?
This character is **Annotator (Suspicious Behavior)** is **Labeling** **Classifying** **Interpreting** **Deciding** **Guided**
This technology is **Machine learning** **Surveillance cameras** is **Trained** **Detecting** **Analysing** **Developed** **Classifying**
This entity is/are **Corporation** is/are **Categorising** **Guiding** **Outsourcing** **Controlling** **Training**
This entity is/are **Image** is/are **Labeled** **Interpreted** **Classified**
This entity is/are **Humans in general** is/are **Classified** **Analysed**

Technology (26)
Fixed vocabulary: 3D scans, AI, Augmented reality, Biometrics, Body scans, Camera, Cameraphone, Deepfake, Drones, Emotion recognition, Facial recognition, Filtering, Hologram, Image generation, Interactive panoramas, Machine learning, Microscope/Telescope, Motion tracking, Non-visible spectrum, Object recognition, Ocular implant, Satellite images, Surveillance camera, UGV, Virtual reality, Webcams.

Entity (10)
Fixed vocabulary: Corporation, Creator, Environment, Government, Humans in general, Image, Law enforcement, Military, Object, User.

Verbs
Open vocabulary 1104 unique verbs

Fig. 1. Diagram of the Machine Vision Database with the artwork *Suspicious Behavior* used as an example of a 'Creative Work.' The diagram depicts how each 'Creative Work' link to at least one 'Situation' and what type of data was logged. The verbs 'Classifying' and 'Classified' was tagged from an open vocabulary of -ing and -ed ending verbs to describe 'Situations.' Agents in 'Situations' could be 'Characters' like the Annotator in *Suspicious Behavior*, Entities like 'Humans in General' or 'Images,' or machine vision 'Technologies.' for further exploration of the database the archived pages can be accessed using QR codes in the diagram.

Figure 2 is a network visualisation depicting all agents and actions ('Verbs') in the dataset as nodes in a network, but with the verb 'Classifying' and the agents Image linking to this action highlighted. Agents and actions that are not related to classifying are shown as grey. In the network visualisation nodes cluster together when they share many links, so if two technologies are close on the network, that means they act in

Who is Classified in 'Situations'

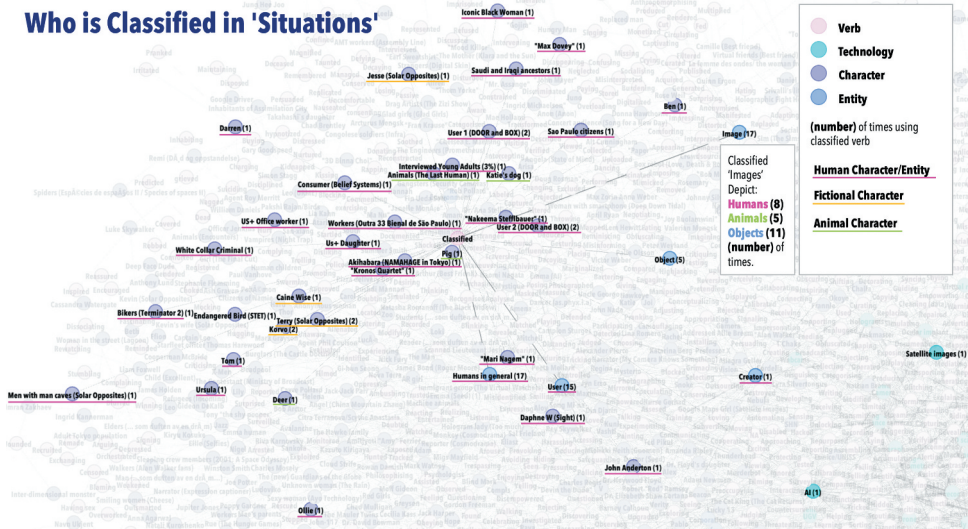


Fig. 3. Who is Classified in 'Situations'? The same network visualisation as in Figure 2, but with the action 'Classified' and agents that are labelled as being 'Classified' highlighted. As images are frequently shown as being classified in the works, a table is included showing counts of what the classified images in the different machine vision situations depict.

Other verbs describing interactions between classified humans and machine vision

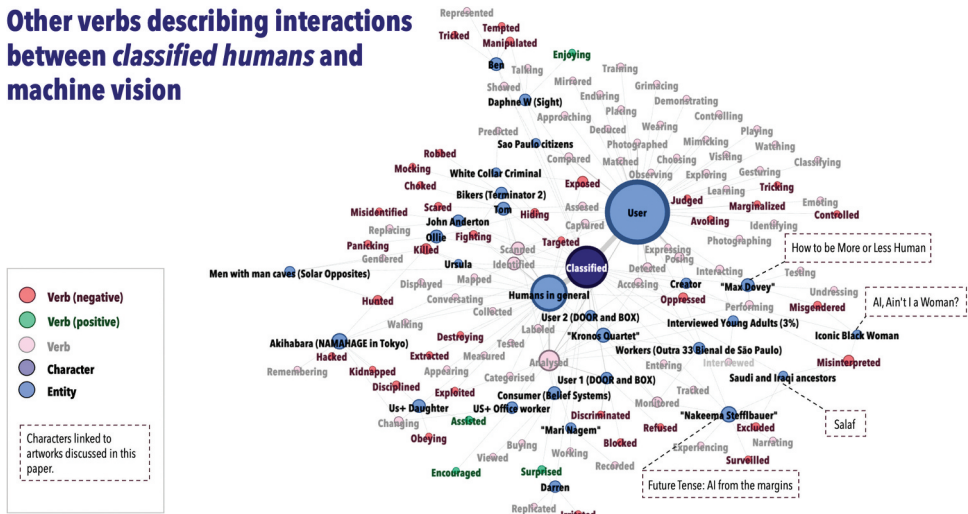


Fig. 4. A network visualisation made in Gephi highlighting other verbs used in machine vision situations where human agents are classified by machine vision.

artworks, games and narratives about machine vision are also Killed (4), Hunted (3), Oppressed (3), Judged (2), Excluded (2), Discriminated (2), or Marginalized (1). Many of these verbs also imply that the classified human finds themselves in an unjust situation. Verbs like Misinterpreted (4), Misgendered (2) and Misidentified (1) indicate that the machine vision is flawed and fails in its task to classify humans. Only Surprised (2), Enjoying (1), Assisted (1) and Encouraged (1) express a positive experience when being classified (green dots in [Figure 6](#)). It is notable that in comparison only a few actions express positive experiences suggesting that cultural conceptions of machine vision present a critical stance towards machine vision deployed to classify humans.

To conclude, data from the Machine Vision Database shows that machine vision deployed to classify humans is a reoccurring interaction depicted in art, video games, movies and novels. Notably, being classified by machine vision is overwhelmingly portrayed as a negative experience. This broader analysis confirms the impression from the four artistic audits I discussed earlier. The discussion of the four artistic audits and the data analysis demonstrates how machines operate humans through images in clearly harmful ways. In the next section I change perspective to human bias and delve deeper into how humans operate machine classifiers through images.

Artistic research – *suspicious behavior* classifying humans in cybernetic loops of indirect reverse operativity

As a fictional annotation tutorial *Suspicious Behavior* takes a unique point of view by examining visual datasets from the perspective of data annotation. Even though automated perception is shaped by image datasets, they have been given little value in discourses about model building.³⁶ Likewise, image annotation is given little attention in computer vision, and when annotation practices of image datasets are researched, they mainly focus on how bias gets embedded into datasets through the individual subjective worker.³⁷ In contrast, *Suspicious Behavior* ([Figure 5](#)) demonstrates ways in which interpretation is imposed on annotators and thus on datasets. As an example of indirect reverse operativity, *Suspicious Behavior* demonstrates that the ways outsourced annotators operate on images is actually ‘informed by the interests, values, and priorities of other actors.’³⁸

Artistic research for the project involved investigating the layers of labour that goes into assembling visual datasets. Artists have investigated aspects of outsourced labour often by utilizing platforms like Amazon’s Mechanical Turk (AMT) to produce artworks with the aspiration of rendering visible this type of hidden human labour.³⁹ However, *Suspicious Behavior* differs from these projects because on-demand labour was not used to create the artwork. Instead, the work of an on-demand annotator is simulated, thus offering the reader an opportunity to experience the conditions under which images are classified and prepared for machine learning. In what follows, I will use examples from *Suspicious Behavior* to discuss how data curators enforce certain interpretations of images through instructions and interface design.




Fig. 5. *Suspicious Behavior* installation setup. © esc medien kunst labor CYBORG SUBJECTS, Photo: Martin Gross.

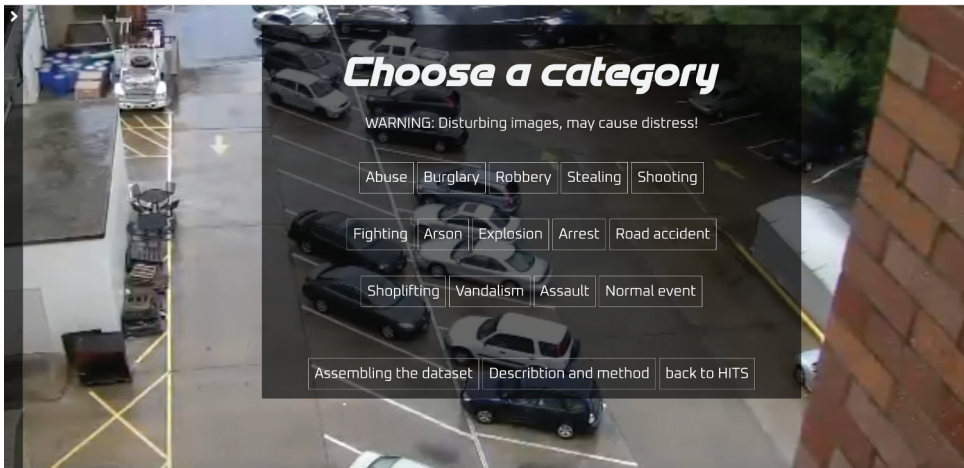


Fig. 6. In *Suspicious Behavior* the reader can traverse 12 anomaly categories and one normal category of the UCF-Crime Dataset.

Data acquisition and annotation can be done in many ways. ImageNet, an influential dataset that has become an integral part of AI infrastructure, established a model practice of scraping images from the internet and outsourcing annotation. This model has been criticized as ‘data extraction without consent and labelling by underpaid crowdworkers.’⁴⁰ Even though automated annotation as a cost-efficient



Fig. 7. Screenshot of annotation interface in Suspicious Behavior.

option is increasingly popular, such automation still involves an annotation apparatus which at some point of its development depended on what I call classifying humans. Typically, humans involved in the annotation of visual datasets are curators and annotators. Dataset curators are the clients. They usually work at companies or universities and are in a need of a visual dataset to develop machine learning models and deploy machine vision. To sustain the illusion of machine automation, human labour is often intentionally hidden.⁴¹ This is also the case with annotators who typically belong to a low-valued and invisible crowd of outsourced on-demand workers.⁴² Annotators are typically hired through platforms like AMT or through companies specialising in annotation services for artificial intelligence. In addition to curators and annotators, project managers, consultants, and quality assurance personnel might be involved in classification.

In *Suspicious Behavior*, the reader is invited to step into the role of an on-demand annotator-trainee. The annotator-trainee is guided through a fictional tutorial that teaches how to navigate ‘Human Intelligence Tasks’ (HIT’s). HIT’s are micro-tasks, for example, tagging a bundle of images or videos and facilitated by often custom made annotation interfaces. In *Suspicious Behavior*, the annotator’s task is to decide if a video contains suspicious behaviour or not. On-demand annotation is casually called ‘clickwork,’ a term that implies work that is easy and requires little thought. However, assembling, and annotating millions of images or videos is tedious labour usually involving two tasks: labelling and segmenting. Segmenting involves separating the objects that are visible in an image and classifying them as different from each other. Labelling is the act of naming an image or a segment. In *Suspicious Behavior*, gradually the annotator-trainee learns that the video dataset that is being labelled is assembled in order to train machine vision to spot suspicious human behaviour. This setup exemplifies indirect reverse operativity in which the annotator classifies images, and the images are then used to train machine vision to classify humans.

Milagros Miceli, Martin Schuessler, and Tianling Yang have studied hierarchies and power relations in image annotation, which they define as a ‘sense-making practice’ in which meaning is assigned to data by adding labels to images.⁴³ This ‘sense-making practice’ is a step in translating images into

machine-readable formats. When images become machine-readable, they shift from being visual representations to becoming non-representational processes of calculation. Thus, operative images are perceived as disappearing into black boxes and becoming invisible.⁴⁴ In this process of translating the visual into data, labelled images become mediators of meaning. Aud Sissel Hoel suggests considering ‘operative images’ as interfaces, ‘in the epistemological and ontological sense as intermediaries.’ If this idea is situated in the machine learning context, then representational images in the datasets can be understood to operate as interfaces.

The images that are annotated function as containers of what Bowker and Leigh Star have called ‘boundary objects,’ a means of translating human-image-machine interpretation. According to Bowker and Leigh Star, boundary objects have multiple meanings that differ between various social worlds, yet they are recognizable in more than one of those worlds.⁴⁵ When an interpretation of an object is not questioned by a community, it becomes a naturalized part of that social world. In a similar way, when images and the objects (or actions) in them are labelled for machine perception, the interpretations of those objects are naturalized, and consequently, alternative interpretations are denied, like in the case of binary gender classification. As those interpretations are built into machine vision, they become naturalized standards in more than one community. An example of this can be observed in *Suspicious Behavior*’s advanced module HIT01:Explorer where the reader can traverse videos and learn more about the UCF-Crime Dataset (Figure 6). In an article accompanying the UCF-Crime Dataset, data curators acknowledge the difficulty of defining anomalous behaviour, ‘since it is quite subjective and can vary largely from person to person.’⁴⁶ Nevertheless, they selected 13 anomaly categories they decided have ‘significant impact on public safety.’⁴⁷ The dataset also contains one category for normal behaviour. Without much contemplation these categories become naturalized, ‘disappearing into the realms of what is considered common sense,’⁴⁸ even though they impose specific interpretations of anomaly behaviour. Thus, those who classify image datasets hold enormous power, because they make decisions which embed certain ideologies affecting how images indirectly reverse operate upon humans. How is that power distributed among classifying humans?

The annotation interface is yet another layer in the sense-making of images as it facilitates the labelling of images. According to Christian Ulrik Andersen and Søren Bro Pold, layers of ‘[i]nterfaces organize how data are stored, translated, exchanged, distributed, used, and experienced.’⁴⁹ In the process of translating human perception into machine perception annotation can be understood as operating on and through images as interfaces. When investigating the annotation apparatus behind ImageNet, Nicolas Malevé calls to our attention that ‘[t]he interfaces of annotation are designed to control workers’ productivity, to find the optimal trade-off between speed and precision.’⁵⁰ In *Suspicious Behavior*, when the reader chooses to ‘become a clickworker’ they receive their first HIT. The annotation interface is simple (see Figure 7): videos are presented in a sequence overlaid with the question: ‘Spot anything suspicious?’ The question is followed by a yes and a no button. The

annotator is given 10 seconds to decide, and a countdown of remaining seconds is placed next to the buttons.

The videos, taken from VIRAT Video Dataset, are designed to assess action detection algorithms in video surveillance. In their documentation, the VIRAT dataset curators describe how they prepared hours of video material for annotation by breaking it up into 10 seconds segments.⁵¹ To replicate the experience of annotating short segments of videos, the interface in *Suspicious Behavior* only allows the annotator-trainee a 10-second glance at each video. The binary yes or no question is an adaption of another annotation interface presented in a paper describing the assembly of the MIT Moments in Time action dataset.⁵² Binary labelling is considered the most efficient annotation method to determine if an object or action appears in an image.⁵³ The reader browsing through the tutorial might ask, what is suspicious behaviour? Instructions for what the annotator-trainee should look for is presented through the

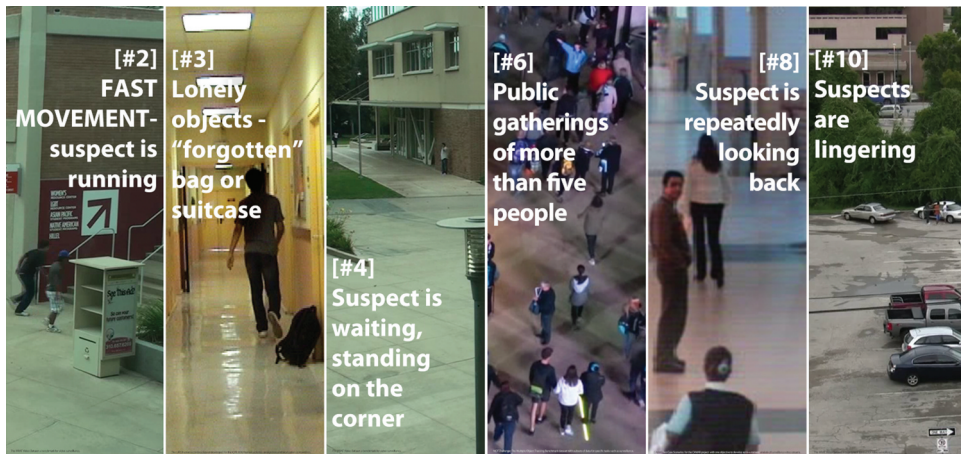


Fig. 8. Examples of posters depicting suspicious behaviour.

Time's up! HIT complete!

You managed to annotate 13 videos in 60 seconds. 7 / 13 are annotated correct, efficiency rate = 53.84615384615385%. An accuracy rate <80% means that you will get rejected and less HITs offered in the future. If you can hold up the pace for 1 hour your result would be:

CLICKWORKER paycheck	Euro/hour	Projected annotated videos
Projected earnings	7.0200000000000005€	780 (each pays 0.009 cents)

Your hourly salary of 7.0200000000000005€ is LESS than the minimum wage paid in Canada.

Fig. 9. Screenshot of the report in end of HIT03: Speed Master.

12 posters (examples depicted in Figure 8). In addition, an instructive YouTube video montage collects examples of what various security authorities define as suspicious: ‘anyone quickly leaving an area after dropping a package,’ ‘abandoning a vehicle;,’ ‘lurking in corners,’ or ‘seeming nervous or acting in a disturbing manner’ are just a few things the reader is instructed to look for. Written definitions of how to name objects and actions, as well as images of ‘good’ and ‘bad’ examples, are embedded in annotation interfaces to assist the annotator to interpretate images correctly.⁵⁴ By mimicking the design choices of real-world annotation interfaces, *Suspicious Behavior* recreates an annotation environment in which a specific reading of the image is imposed on the annotator.

On demand annotators are seldomly provided with additional information about given tasks. They often lack labour rights and struggle to make a living wage. Thus, they need to optimize their work pace and meet quality thresholds set by the client.⁵⁵ Dataset curators have developed methods to evaluate worker quality to check that the interpretations of individual annotators are in line with given definitions.⁵⁶ Only ‘consensus’⁵⁷ among several annotators validates that the labelled video contains what the dataset curator is looking for, in this case, suspicious behaviour. This means that several annotators are expected to have the same interpretation of suspicious behaviour for a video to enter the dataset. These are mechanisms dataset curators implement in order to test and filter out those whose performance do not reach a set threshold.⁵⁸ In *Suspicious Behavior*, ‘Advanced HIT03: Speed Master’ represents such a test. It exemplifies Nicolas Malevé’s observation that ‘[f]or the annotators, structurally, the “glance” is the norm’⁵⁹ as the annotator-trainee is given 60 seconds to annotate as many videos as possible. At the end of this test a ‘report’ (Figure 9) shows how many videos were ‘correctly’ labelled. If their accuracy rate is higher than 80%, the annotator-trainee qualifies for more work. The annotator-trainee learns a rhythm to maintain a balanced ratio between speed and accuracy while interpreting actions, and this does not allow questioning of how categories are made, or classes defined.⁶⁰

Previous examples from *Suspicious Behavior* expose the hierarchical power dynamics at play among classifying humans. The interface is a technical layer that obscures the fact that dataset curators impose certain interpretations of images on annotators and ultimately manipulate the ways annotators operate on images. The role of annotators is reduced to matching labels with images, thus, the power to define classes and the control over which kind of images those classes contain remains with dataset curators. Therefore, it is not solely the subjectivities of annotators that introduce bias into image datasets. Often the decisions made by curators have a greater effect on how images indirectly reverse operate.

Although dataset curators keep control of image interpretation, there are operations out of their control before annotation even begins. When images are collected into a candidate pool for annotation, search queries on chosen platforms determine what images end up annotated in the first place. For example, describing the construction of the earlier mentioned UCF-Crime Dataset, the dataset curators used variations of search words, like ‘car crash’ for the category of ‘road accident,’ in order to find videos for the dataset.⁶¹ The choices of platform and of the

interaction mode between users, content creators, advertisers and the platform algorithms are absorbed into the annotation apparatus. Indirect reverse operativity comes into play when platforms or content producers use automated image recognition to manage the labelling of videos. For example, Google Clouds Video AI claims to ‘recognize over 20,000 objects, places, and actions in video.’⁶² Such automated labelling in turn only operates because a dataset was once classified by another annotation apparatus which now indirectly reverse operates in cybernetic loops. This means that indirect reverse operativity constantly feeds an evolving internet-sized ‘mass image’ with labelled images. In such loops of indirect reverse operativity classifications become naturalized standards and earlier human decisions are taken for granted, embedded in machine vision infrastructure.

New media theorist Jussi Parikka notes how histories of photography and technical images are constructed of ‘[o]perations built on operations’ meaning that instruments, infrastructures, practices, and techniques build upon upon each other.⁶³ In a similar sense, the operations of an annotator build upon practices and tools such as search engines that are built upon other instruments, practices, and infrastructures. Classification systems embedded in machine vision in turn serve as infrastructure that, for example, city security management systems build upon. Among layers of operations involved in assembling an image dataset, automated interpretation also influences image annotation. The classifying human and the machine classifier merge, and in cybernetic loops of indirect reverse operativity, norms and worldviews embedded in machine vision classification propagate.

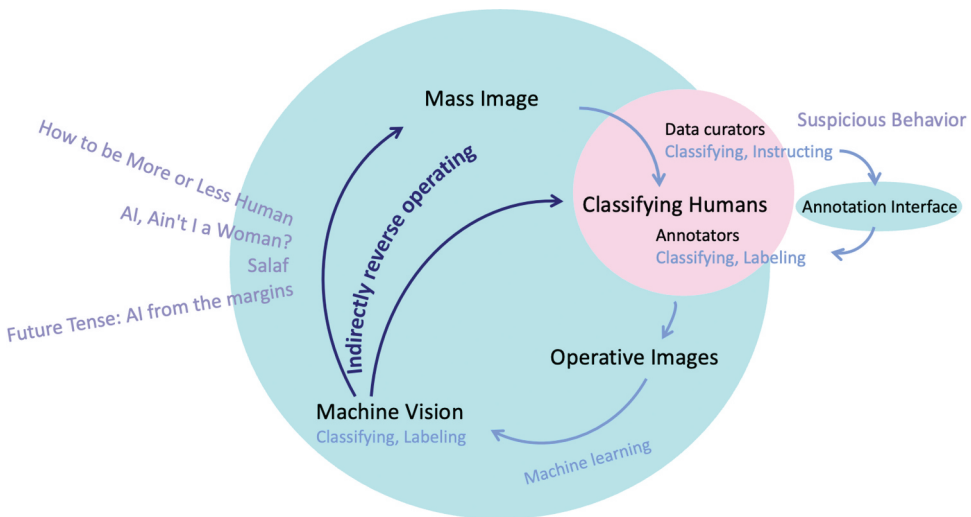


Fig. 10. Indirect reverse operativity.

Conclusion

In this article I examined cultural conceptions of machine vision bias. I particularly draw attention to machine vision bias in instances when humans operate machines and machines operate humans through images. Figure 10 shows how I understand the artworks I engage with to depict different instances of what I call indirect reverse operativity. This entails that data curators and annotators classify images which operate back on humans when a machine vision classifier trained on these images is deployed to classify humans. What is made explicit in this diagram is that classifying for machine learning is ultimately a human-centred process. Although cost-efficiency and the demand for internet-sized datasets drives the development of automating dataset curation to a constantly evolving ‘mass image,’ we must also recognise that the ‘mass image,’ as Sean Cubitt writes, is a ‘post-human hybrid of work and labour performed by natural human and technical agencies.’⁶⁴

I have used what I call artistic audits as examples to describe how images indirectly reverse operate in ways which are unfair and discriminatory particularly towards marginalized populations. In these artworks machine vision deployed to classify humans is represented as troublesome: it is discriminating, misjudging, and oppressing systematically marginalized populations. This portrayal aligns with scholarly audits of popular algorithmic systems.⁶⁵ What the artistic audits by Dovey, Buolamwini and Aljowaysir show is that binary gender classes as well as harmful stereotypes are naturalized as they become part of machine vision infrastructure. Nakeema Stefflbauer & Nushin Yazdan’s work explains that machine vision bias is experienced differently at the intersections of race, gender, and class, and it is harmful because it perpetuates systems of discrimination against those already marginalised by society. My data analysis of machine vision situations in a dataset of 500 works of art, games and narratives demonstrated that the cultural conceptions of harmful machine vision bias we saw in the artistic audits are echoed in this larger sample of creative works.

Artworks have been particularly influential in communicating the link between bias in datasets and bias in machine vision to a broader audience. This has pushed the computer vision community to revisit influential datasets and respond with efforts to mitigate harmful bias. Because attention has mainly been directed to representation and historical bias in datasets involving problematic image classes, responses to mitigate bias has involved ‘technical fixes’ such as balancing the number of images in a dataset to equally represent intersections of gender and race, adding data, or to erase problematic classes.⁶⁶ Some methods of diversifying datasets might mitigate bias, while others end up doing the opposite.⁶⁷ However, sanitation of datasets in hindsight does not magically fix harmful machine vision bias. Machine vision classifiers already trained on biased datasets still propagate those biases. Increasingly automated data curation and labelling practices create cybernetic loops of indirect reverse operativity in which the ways humans are classified is easily taken for granted even though disproportionately harmful for those in the margins.

Suspicious Behavior presents a different path to confront dataset bias. By acknowledging that datasets are always biased the artwork simulates the experience of

performing bias. By reflecting upon artistic research for *Suspicious Behaviour* I demonstrate how designing annotation interfaces and the instructions they embed is one instance of performing bias. Interfaces are designed to keep dataset curators in control of image interpretation. Even if an annotator's individual subjectivity is often said to introduce bias into datasets, on-demand annotators are given only limited agency and are reduced to following instructions and matching labels to images. Thus, it is important to acknowledge that the technical layer of the interface obscures the fact that dataset curators hold the power to define classes and certain interpretations are imposed on annotators and ultimately affecting on how images indirectly reverse operate. This demonstrates that the power to perform bias, in other words to interpret and classify images, is distributed among annotators and data curators and facilitated by annotation interfaces. This means that bias is introduced into machine vision through a multiplicity of instances.

The point at which bias is introduced into machine vision is further complicated because images do not reverse operate instantaneously nor directly back on those who classify them. When machine classifiers automatically label images which are then scraped, annotated, and fed into a new machine learning model, cybernetic loops of indirect reverse operativity occur. Such cybernetic loops further propagate and naturalize problematic classification. Therefore, for those classifying humans — developers, designers, data curators and annotators — who operate on images in the assembly pipeline of visual datasets, it becomes crucial to acknowledge that images indirectly reverse operate, and instances of performing bias should be accounted for.

Acknowledgments

I am grateful to the Machine Vision team for cooperation on the MV database, feedback and support, as well as, to Andreas Zingerle for collaboration on *Suspicious Behaviour*, both projects contributed in many ways and made this article possible. In addition, I am grateful for to the anonymous peer-reviewers who provided helpful feedback to develop this article.

Notes

1. Bowker and Star, *Sorting Things Out*, 1.
2. Mackenzie, *Machine Learners*, 10.
3. Klein and D'Ignazio, *Data Feminism*, 103.
4. Kate Crawford, *Atlas of AI*.
5. Crawford, Kate, and Paglen. "Excavating AI." <https://www.excavating.ai>.
6. For how facial recognition is particularly oppressive towards already marginalized populations see: Shoshana Amielle Magnet, *When Biometrics Fail*; or Kelly Gates, *Our Biometric Future*. More on bias in image datasets are discussed for example in: Joy Buolamwini and Gebru, "Gender Shades: Intersectional Accuracy Disparities

- in Commercial Gender Classification,” In *Proceedings of the 1st “Gender Shades,”* 77–91; and Kate Crawford, *Atlas of AI*.
7. Crawford and Paglen.
 8. Malevé, *Algorithms of Vision. Human and Machine Learning in Computational Visual Culture*, 193.
 9. Kronman, “Intuition Machines,” 54–68.
 10. Several artists have critically scrutinized and audited datasets, mostly for facial recognition, but also video datasets for action detection. Joy Buolamwini and Timnit Gebru audit gender classification products in their project Gender Shades (<http://gendershades.org/>); Adam Harvey and Jules LaPlace raise concerns about privacy issues in Exposing.ai (<https://exposing.ai/>); Kate Crawford and Trevor Paglen describe their media archaeological approach to image datasets in ‘Excavating AI: The Politics of Images in Machine Learning Training Sets’ (<https://www.excavating.ai>); Everest Pipkin worked with MIT’s Moments in Time video dataset to describe experiences of exhaustive watching in ‘On Lacework: Watching an Entire Machine-Learning Dataset’ (<https://unthinking.photography/articles/on-lacework>). These are only a few examples of an increasing number of artists working with image datasets. Projects like Data/Set/Match by The Photographers’ Gallery have also brought together several artists to interrogate image datasets and produce artworks and essays: <https://thephotographersgallery.org.uk/data-set-match>. Scholar and artist Nicolas Malevé’s writing for the Photographers’ Gallery and his article “On the Data Set’s Ruins,” *AI & SOCIETY* 36 (2021): 1117–31 has been especially influential for developing this article.
 11. Walker Rettberg et al., “Database of Machine Vision in Art, Games and Narratives,” May 3, 2022, <https://doi.org/10.5281/zenodo.6514729>.
 12. KairUs, *Suspicious Behavior*, 2020, <http://kairus.org/suspicious/>.
 13. Farocki, “Phantom Images,” 12–22.
 14. For different definitions of operative images and degrees of operativity see for example: Hoel, ‘Operative Images. Inroads to a New Paradigm of Media Theory’ or Eder and Klonk eds., *Image Operations*. This article uses Volker Pantenburg’s narrow definition of operativity from “Working Images.” 58.
 15. Hoelzl and Marie, *Softimage*, 84.
 16. With ‘aiming at’ Hoelzl and Marie reference Farocki, “Phantom Images”, 17. And in this passage, Farocki actually responds to Paul Virilio’s writing in *War and Cinema*.
 17. Hoelzl, “The Operative Image — an Approximation.”
 18. Thylstrup, “The Ethics and Politics of Data Sets in the Age of Machine Learning.”
 19. I use quotation marks to indicate that I speak of the artist as a persona which is portrayed in a particular way in an artwork.
 20. Davis, “A Dataset Is a Worldview.”
 21. Safiya U, *Algorithms of Oppression*; for gender stereotypes in images searches on occupations see Kay, Matuszek, and Munson, “Unequal Representation and Gender Stereotypes in Image Search Results for Occupations,” 3819–28.
 22. Prabhu and Birhane, “Large Image Datasets,” 14.
 23. Os Keyes, “The Misgendering Machines: Trans/HCI Implications of Automatic Gender Recognition.” 1–22.

24. Buolamwini and Gebru, "Gender Shades."
25. For more on Americentric and Eurocentric bias in image retrieval, see for example Shankar et al., "No Classification without Representation."
26. Costanza-Chock, "Introduction: #TravelingWhileTrans, Design Justice, and Escape from the Matrix of Domination," 20.
27. The examples in *Future Tense: AI from the margins* refer to Joy Buolamwini's artwork *The Coded Gaze: Unmasking Algorithmic Bias*, 2016, <https://youtu.be/162VzSzzoPs>. As well as to cases published in scholarly blogs and publications namely: Noble, *Algorithms Oppression*; Sharp, "Nikon Camera Says Asians: People Are Always Blinking — Sociological Images" and Rhue, "Racial Influence on Automated 685 Perceptions of Emotions."
28. Benjamin, *Race After Technology*.
29. Patricia Hill Collins, *Black Feminist Thought: Knowledge, Consciousness, and the Politics of Empowerment*.
30. Stefflbauer and Yazdani, *Nakeema Stefflbauer & Nushin Yazdani — Future Tense* (time code 17:03).
31. Sarah Myers West, Meredith Whittaker, and Kate Crawford, *Discriminating Systems Gender, Race, and Power in AI*.
32. Eubanks, *Automating Inequality*.
33. The dataset situations.csv is the main unit for my analysis. Further characters.csv was used to identify if 'Characters' were ascribed as Humans, Machines, Animals or Fictional Species. Published datasets are available here: Walker Rettberg et al., "A Dataset Documenting Representations of Machine Vision Technologies in Artworks, Games and Narratives," <https://doi.org/10.18710/2G0XKN>.
34. Walker Rettberg et al., "Representations of Machine Vision Technologies in Artworks, Games and Narratives," 108,319.
35. Rettberg et al., "Representations of Machine Vision Technologies in Artworks, Games and Narratives."
36. Hutchinson et al., "Towards Accountability for Machine Learning Datasets: Practices from Software Engineering and Infrastructure."
37. Examples of this are: Bhavya Ghai et al., "Measuring Social Biases of Crowd Workers Using Counterfactual Queries."; and, Hube, Fetahu, and Gadiraju, "Understanding and Mitigating Worker Biases in the Crowdsourced Collection of Subjective Judgments," 1–12.
38. Miceli, Schuessler, and Yang, "Between Subjectivity and Imposition: Power Dynamics in Data Annotation for Computer Vision."
39. Joanna Zylińska names several projects in addition to her own project *View from the Window* (2018) utilising on-demand work from platforms like AMT in "Undigital Photography," 231–52; and Gabriel Pereira and Bruno Moreschi discuss their artwork *Recoding Art* (2019) and their aspiration to render visible the on-demand work behind AI in: "Artificial Intelligence and Institutional Critique 2.0," 1201–23.
40. Crawford, *Atlas of AI*, 109.
41. Atanasoski and Vora, *Surrogate Humanity Race, Robots, and the Politics of Technological Futures*.
42. Gray and Siddharth, *Ghost Work*.
43. Miceli, Schuessler, and Yang, "Between Subjectivity and Imposition," 1–2.

44. For example Rebecca Uliasz in ‘Seeing like an Algorithm: Operative Images and Emergent Subjects’ describes how images of our faces disappear into black boxes only to reappear with labelled bounding boxes. Artist Trevor Paglen, who investigated machine vision in several of his artworks, observes in ‘Operational Images’ *E-Flux*, no. #59 (2014) that operational images are increasingly invisible to the human eye. Adrian MacKenzie and Anna Munster in ‘Platform Seeing: Image Ensembles and Their Invisibilities’ call these non-representational forms of images ‘invisual images.’
45. Bowker and Leigh Star, *Sorting Things Out*, 297.
46. Sultani, Chen, and Shah, “Real-World Anomaly Detection in Surveillance Videos,” 6479–88.
47. Sultani, Chen, and Shah, “Real-World Anomaly Detection in Surveillance Videos,” 6483.
48. Miceli, Schuessler, and Yang, “Between Subjectivity and Imposition,” 4.
49. Christian Ulrik Andersen and Søren Bro Pold, “Interface.”
50. Malevé, “On the Data Set’s Ruins,” 1121.
51. Sangmin et al., “A Large-Scale Benchmark Dataset for Event Recognition in Surveillance Video,” 3153–60.
52. Monfort et al., “Moments in Time Dataset: One Million Videos for Event Understanding,” 502–8.
53. Deng et al., “Scalable Multi-Label Annotation,” 3099–102.
54. Deng et al. in ‘Scalable Multi-Label Annotation’ gives an example of ‘good’ and ‘bad’ examples of images were embedded in the annotation interface for ImageNet.
55. Gray and Siddharth write in *Ghost Work*, xxix about the insufficient information on-demand workers receive from clients and Kotaro Hara et al., analyse the earnings of on-demand workers in “A Data-Driven Analysis of Workers,” 1–14.
56. Hata et al., “A Glimpse Far into the Future: Understanding Long-Term Crowd Worker Quality.”
57. Hara et al., ‘A Data-Driven Analysis of Workers’ Earnings on Amazon Mechanical Turk’.
58. Deng et al., “Scalable Multi-Label Annotation.”
59. Malevé, “On the Data Set’s Ruins.”
60. Linda Kronman and Andreas Zingerle, “Suspicious Behavior,” 318–26.
61. Sultani, Chen, and Shah, “Real-World Anomaly Detection in Surveillance Videos.”
62. “Video AI — Video Content Analysis | Cloud Video Intelligence API,” <https://cloud.google.com/video-intelligence>.
63. Parikka, “Operational Images — Preface in the Forthcoming Book.”
64. Cubitt, “Mass Image, Anthropocene Image, Image Commons,” 26.
65. Bandy, “Problematic Machine Behavior: A Systematic Literature Review of Algorithm Audits,” 74:1–74:34.
66. For examples of technical fixes see Crawford, *Atlas of AI*, 129., or Scheuerman et al., “How We’ve Taught Algorithms to See Identity.” 1–35.
67. Crawford and Paglen, “Excavating AI.”

Funding

The work was supported by the HORIZON EUROPE European Research Council [771800].

Disclosure statement

No potential conflict of interest was reported by the author(s).

ORCID

Linda Kronman  <http://orcid.org/0000-0001-7699-4505>

Bibliography

- Andersen, Christian Ulrik, and Søren Bro Pold. "Interface." In *Uncertain Archives: Critical Keywords for Big Data*, edited by Nanna Bonde Thylstrup, Daniela Agostinho, Annie Ring, Catherine D'Ignazio, and Kristin Veel, Cambridge, Massachusetts: The MIT Press, 2021.
- Atanasoski, Neda, and Kalindi Vora. *Surrogate Humanity Race, Robots, and the Politics of Technological Futures*. Durham and London: Duke University Press, 2019.
- Bandy, Jack. "Problematic Machine Behavior: A Systematic Literature Review of Algorithm Audits." *Proceedings of the ACM on Human-Computer Interaction* 5, no. CSCW1 (April 22, 2021): 74:1–74:34. doi:10.1145/3449148.
- Benjamin, Ruha. *Race After Technology: Abolitionist Tools for the New Jim Code*. Polity, 2019.
- Bowker, Geoffrey C., and Susan Leigh Star. *Sorting Things Out: Classification and Its Consequences*. First Paper Back Edition. Cambridge, Massachusetts; London, England: MIT Press, 2000.
- Buolamwini, Joy. "The Coded Gaze: Unmasking Algorithmic Bias." 2016. Video Poem. <https://youtu.be/162VzSzzoPs>
- Buolamwini, Joy, and Gebru. Timnit. "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification." In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, edited by Sorelle A. Friedler and Christo Wilson, 81:77–91. Proceedings of Machine Learning Research. New York, NY, USA: PMLR, 2018. <http://proceedings.mlr.press/v81/buolamwini18a.html>
- Collins, Patricia Hill. *Black Feminist Thought: Knowledge, Consciousness, and the Politics of Empowerment*. 2nd ed. Routledge, 1990. doi:10.4324/9780203900055.
- Costanza-Chock, Sasha. "Introduction: #travelingwhiletrans, Design Justice, and Escape from the Matrix of Domination." In *Design Justice*. The MIT Press, 2020. <https://design-justice.pubpub.org/pub/ap8rgw5e/release/1>
- Crawford, Kate. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. Yale UP, 2021.

- Crawford, Kate, and Trevor Paglen. "Excavating AI: The Politics of Images in Machine Learning Training Sets." *The AI Now Institute* (2019). <https://www.excavating.ai>.
- Cubitt, Sean. "Mass Image, Anthropocene Image, Image Commons." In *Photography off the Scale : Technologies and Theories of the Mass Image*, edited by Jussi Parikka and Tomáš Dvořák, 25–40. Edinburgh: Edinburgh University Press, 2021.
- Davis, Hannah. "A Dataset is a Worldview." *Medium* (blog), March 7, 2020. <https://towardsdatascience.com/a-dataset-is-a-worldview-5328216dd44d>
- Deng, Jia, Olga Russakovsky, Jonathan Krause, Michael S. Bernstein, Alex Berg, and Li Fei-Fei. "Scalable Multi-Label Annotation." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 3099–3102. CHI '14. New York, NY: Association for Computing Machinery, 2014. [10.1145/2556288.2557011](https://doi.org/10.1145/2556288.2557011).
- Eder, Jens, and Charlotte Klonk, edited by. *Image Operations: Visual Media and Political Conflict*. Manchester: Manchester University Press, 2017.
- Eubanks, Virginia. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*: Picador, 2017.
- Farocki, Harun. "Phantom Images." *Public*, January 1, 2004. <https://public.journals.yorku.ca/index.php/public/article/view/30354>
- Gates, Kelly. *Our Biometric Future: Facial Recognition Technology and the Culture of Surveillance*. NYU Press, 2011.
- Ghai, Bhavya, Q. Vera Liao, Yunfeng Zhang, and Klaus Mueller. "Measuring Social Biases of Crowd Workers Using Counterfactual Queries." *ArXiv:2004.02028* [Cs], April 4, 2020. <http://arxiv.org/abs/2004.02028>.
- Google Cloud. "Video AI - Video Content Analysis | Cloud Video Intelligence API." Accessed June 30, 2022. <https://cloud.google.com/video-intelligence>
- Gray, Mary L., and Suri Siddharth. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*: Houghton Mifflin Harcourt, 2019.
- Harvey, Adam, and Jules LaPlace. "Exposing.Ai." *Exposing.ai*, 2021. <https://exposing.ai/>.
- Hoel, Aud Sissel. "Operative Images. Inroads to a New Paradigm of Media Theory." In *Image – Action – Space: Situating the Screen in Visual Practice*, edited by Luisa Feiersinger, Kathrin Friedrich, and Moritz Queisner, 11–28. Berlin, Boston: De Gruyter, 2018. <https://www.degruyter.com/view/books/9783110464979/9783110464979-002/9783110464979-002.xml>
- Hoelzl, Ingrid. "'The Operative Image - an Approximation'. The new everyday a media commons project." Accessed February 3, 2014. <http://mediacommons.org/tne/pieces/operative-image-and-approximation>
- Hözl, Ingrid, and Remi Marie. *Softimage: Towards a New Theory of the Digital Image*. Intellect, 2015.
- Hube, Christoph, Besnik Fetahu, and Ujwal Gadiraju. "Understanding and Mitigating Worker Biases in the Crowdsourced Collection of Subjective Judgments." In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–12. Glasgow Scotland Uk: ACM, 2019. doi:[10.1145/3290605.3300637](https://doi.org/10.1145/3290605.3300637).
- Hutchinson, Ben, Andrew Smart, Alex Hanna, Emily Denton, Christina Greer, Oddur Kjartansson, Parker Barnes, and Margaret Mitchell. "Towards Accountability for Machine Learning Datasets: Practices from Software Engineering and Infrastructure." *ArXiv:2010.13561* [Cs], Accessed January 29, 2021. <http://arxiv.org/abs/2010.13561>

- KairUs. *Suspicious Behavior*. 2020. Net-art. <http://kairus.org/suspicious/>
- Kenji, Hata, Ranjay Krishna, Li Fei-Fei, and Michael S. Bernstein. "A Glimpse Far into the Future: Understanding Long-Term Crowd Worker Quality." *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, February 25, (2017): 889–901. doi:10.1145/2998181.2998248.
- Keyes, Os. "The Misgendering Machines: Trans/HCI Implications of Automatic Gender Recognition." *Proceedings of the ACM on Human-Computer Interaction* 2, no. CSCW (November 1 2018): 88:1-88:22. doi:10.1145/3274357.
- Klein, Lauren F., and Catherine D'ignazio. *Data Feminism. Strong Ideas*: MIT Press, The MIT Press, 2020.
- Kotaro, Hara, Abigail Adams, Kristy Milland, Saiph Savage, Chris Callison-Burch, and Jeffrey P. Bigham. "A Data-Driven Analysis of Workers' Earnings on Amazon Mechanical Turk." In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 1–14. Montreal, QC: ACM, 2018. doi:10.1145/3173574.3174023.
- Kronman, Linda. "Intuition Machines: Cognizers in Complex Human-Technical Assemblages." *A Peer-Reviewed Journal About* 9, no. 1 (4 August 2020): 54–68. doi:10.7146/aprja.v9i1.121489.
- Kronman, Linda, and Andreas Zingerle. "Suspicious Behavior." In *XCoAx 2022: Proceedings of the 10th Conference on Computation, Communication, Aesthetics and X*. Combría, Portugal, upcoming.
- Mackenzie, Adrian. *Machine Learners: Archaeology of a Data Practice*. MIT Press, 2017.
- MacKenzie, Adrian, and Anna Munster. "Platform Seeing: Image Ensembles and Their Invisibilities." *Theory, Culture & Society* 36, no. 5 (September 2019): 3–22. doi:10.1177/0263276419847508.
- Magnet, Shoshana Amielle. *When Biometrics Fail: Gender, Race, and the Technology of Identity*. Duke University Press, 2011. <https://www.dukeupress.edu/When-Biometrics-Fail/>
- Malevé, Nicolas. "On the Data Set's Ruins." *AI & SOCIETY* (November 11, 2020). doi:10.1007/s00146-020-01093-w.
- Malevé, Nicolas. *Algorithms of Vision. Human and Machine Learning in Computational Visual Culture*. London: London South Bank University, 2021.
- Matthew, Kay, Cynthia Matuszek, and Sean A. Munson. "Unequal Representation and Gender Stereotypes in Image Search Results for Occupations." In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 3819–3828. Seoul Republic of Korea: ACM, 2015. doi:10.1145/2702123.2702520.
- Miceli, Milagros, Martin Schuessler, and Tianling Yang. "Between Subjectivity and Imposition: Power Dynamics in Data Annotation for Computer Vision." *ArXiv:2007.14886 [Cs]* (July 30, 2020). <http://arxiv.org/abs/2007.14886>
- Monfort, Mathew, Carl Vondrick, Aude Oliva, Alex Andonian, Bolei Zhou, Kandan Ramakrishnan, Sarah Adel Bargal, et al. "Moments in Time Dataset: One Million Videos for Event Understanding." *IEEE transactions on pattern analysis and machine intelligence* 42, no. 2 (February 1, 2020): 502–508. doi:10.1109/TPAMI.2019.2901464.
- Nakeema Stefflbauer & Nushin Yazdani – *Future Tense: AI from the Margins // Spy on Me #2 Online*, 2020. <https://www.youtube.com/watch?v=HbyPcRGN2uM>

- Noble, Safiya U. *Algorithms of Oppression: How Search Engines Reinforce Racism*. E-Book: New York University Press, 2018.
- Oh, Sangmin, Anthony Hoogs, Amitha Perera, Naresh Cuntoor, Chia-Chih Chen, Jong Taek Lee, Saurajit Mukherjee, et al. "A Large-Scale Benchmark Dataset for Event Recognition in Surveillance Video." *CVPR* 2011, (2011): 3153–3160. doi:10.1109/CVPR.2011.5995586.
- Paglen, Trevor. "Operational Images." *E-Flux*, no. #59 (November 2014). <https://www.e-flux.com/journal/59/61130/operational-images/>
- Parikka, Jussi. "Operational Images – Preface in the Forthcoming Book." *Machinology* (blog), February 16, 2022. <https://jussiparikka.net/2022/02/16/operational-images-preface-in-the-forthcoming-book/>
- Pereira, Gabriel, and Bruno Moreschi. "Artificial Intelligence and Institutional Critique 2.0: Unexpected Ways of Seeing with Computer Vision." *AI & SOCIETY* (September 14, 2020). doi:10.1007/s00146-020-01059-y.
- Pipkin, Everest. "On Lacework: Watching an Entire Machine-Learning Dataset." July 2020. <https://unthinking.photography/articles/on-lacework>
- Prabhu, Vinay Uday, and Abeba Birhane. "Large Image Datasets: A Pyrrhic Win for Computer Vision?" arXiv, Accessed July 23, 2020. doi:10.48550/arXiv.2006.16923.
- Rettberg, Jill Walker, Linda Kronman, Ragnhild Solberg, Marianne Gunderson, Stein Magne Bjørklund, Linn Heidi Stokkedal, Gabriele de Seta, Kurdin Jacob, and Annette Markham. "Database of Machine Vision in Art, Games and Narratives: Archival Version in HTML and CSS." 3 (May 2022). doi:10.5281/zenodo.6514729
- Rettberg, Jill Walker, Linda Kronman, Ragnhild Solberg, Marianne Gunderson, Stein Magne Bjørklund, Linn Heidi Stokkedal, Gabriele de Seta, Kurdin Jacob, and Annette Markham. "A Dataset Documenting Representations of Machine Vision Technologies in Artworks, Games and Narratives." *DataverseNO* (March 22, 2022). doi:10.18710/2G0XKN.
- Rettberg, Jill Walker, Linda Kronman, Ragnhild Solberg, Marianne Gunderson, Stein Magne Bjørklund, Linn Heidi Stokkedal, Kurdin Jacob, Gabriele de Seta, and Annette Markham. "Representations of Machine Vision Technologies in Artworks, Games and Narratives: A Dataset." *Data in Brief* 42 (June 1, 2022): 108319. doi:10.1016/j.dib.2022.108319.
- Rhue, Lauren. "Racial Influence on Automated Perceptions of Emotions." SSRN Scholarly Paper. Rochester, NY: Social Science Research Network, Accessed November 9, 2018. doi:10.2139/ssrn.3281765.
- Scheuerman, Morgan Klaus, Kandra Wade, Caitlin Lustig, and Jed R. Brubaker. "How We've Taught Algorithms to See Identity: Constructing Race and Gender in Image Databases for Facial Analysis." *Proceedings of the ACM on Human-Computer Interaction* 4, no. CSCW1 (May 28, 2020): 1–35. doi:10.1145/3392866.
- Sharp, Gwen. "Nikon Camera Says Asians: People Are Always Blinking - Sociological Images." *The Society Pages* (blog), Accessed May 29, 2009. <https://thesocietypages.org/socimages/2009/05/29/nikon-camera-says-asians-are-always-blinking/>
- Shreya, Shankar, Yoni Halpern, Eric Breck, James Atwood, Jimbo Wilson, and D. Sculley. "No Classification without Representation: Assessing Geodiversity Issues in Open Data Sets for the Developing World." *ArXiv:1711.08536 [Stat]*, Accessed November 22, 2017. <http://arxiv.org/abs/1711.08536>

- Sultani, Waqas, Chen Chen, and Mubarak Shah. "Real-World Anomaly Detection in Surveillance Videos." In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6479–6488. Salt Lake City, UT: IEEE, 2018. doi:10.1109/CVPR.2018.00678.
- Ullasz, Rebecca. "Seeing like an Algorithm: Operative Images and Emergent Subjects." *AI & SOCIETY* 36, no.4 (Accessed September 16, 2020): 1233–1241. doi:10.1007/s00146-020-01067-y.
- Virilio, Paul. *War and Cinema: The Logistics of Perception*. London: Verso, 1989.
- West, Myers, Meredith Whittaker Sarah, and Kate Crawford. "Discriminating Systems Gender, Race, and Power in AI.": AI Now Institut, Accessed April 2019. <https://ainowinstitute.org/discriminatingystems.pdf>
- Zylinska, Joanna. "Undigital Photography: Image-Making beyond Computation and AI." In *Photography off the Scale : Technologies and Theories of the Mass Image*, edited by Jussi Parikka and Tomáš Dvořák, 231–252. Edinburgh: Edinburgh University Press, 2021.

Linda Kronman is a media artist, designer, and is currently completing a PhD at the University of Bergen, researching how machine vision is represented in digital art as a part of the Machine Vision in Everyday Life project. Kronman is a member of the artist duo KairUs, where she creates art together with Andreas Zingerle exploring topics such as surveillance, smart cities, IoT, cybercrime, online fraud, electronic waste, and machine vision. KairUs has recently been recognized by the Austrian Federal Ministry of Arts and Culture, Civil Service and Sport with the BMKÖS/Mayer Outstanding Artist Award 2022 in the category of Media Art.