

Taxonomizing and Measuring Representational Harms: A Look at Image Tagging

Jared Katzman^{*1}, Angelina Wang^{*2}, Morgan Scheuerman³, Su Lin Blodgett⁴, Kristen Laird⁵,
Hanna Wallach⁴, Solon Barocas⁴

¹ University of Michigan

² Princeton University

³ University of Colorado, Boulder

⁴ Microsoft Research

⁵ Microsoft

Abstract

In this paper, we examine computational approaches for measuring the “fairness” of image tagging systems, finding that they cluster into five distinct categories, each with its own analytic foundation. We also identify a range of normative concerns that are often collapsed under the terms “unfairness,” “bias,” or even “discrimination” when discussing problematic cases of image tagging. Specifically, we identify four types of representational harms that can be caused by image tagging systems, providing concrete examples of each. We then consider how different computational measurement approaches map to each of these types, demonstrating that there is not a one-to-one mapping. Our findings emphasize that no single measurement approach will be definitive and that it is not possible to infer from the use of a particular measurement approach which type of harm was intended to be measured. Lastly, equipped with this more granular understanding of the types of representational harms that can be caused by image tagging systems, we show that attempts to mitigate some of these types of harms may be in tension with one another.

Introduction

In recent years, computer vision has gone from a largely academic research topic to a set of technologies that touch the lives of people in a variety of impactful ways. With this success comes the importance of understanding when computer vision systems can cause harms. In this paper, we consider some of the harms that can be caused by a seemingly mundane, yet widespread, computer vision task: image tagging.

Prior work has already documented several problematic cases of image tagging. For example, researchers have identified harms caused by gender classification systems and other systems that sort people into social groups based on their appearance (Keyes 2018; Barlas et al. 2021), demonstrated that the performance of image tagging systems can vary according to the identities of the people depicted in images (Hendricks et al. 2018; Zhao et al. 2017), and revealed a variety of different ways that such systems can perpetuate stereotypes (Schwemmer et al. 2020; Bhargava and Forsyth 2019). These findings have been held up as evidence that image tagging systems can be “unfair,” “biased,” or even “dis-

crimatory.” However, in many of these cases, it can be difficult to identify the full range of normative values at stake.

We aim to bring greater structure and precision to discussions of some of the harms caused by image tagging systems and the many computational approaches that have been proposed to measure them. We first provide relevant background by defining image tagging and revisiting the distinction between allocational and representational harms (Barocas et al. 2017)—emphasizing that problematic cases of image tagging can affect the understandings, beliefs, and attitudes that people hold about particular social groups. We then analyze previously proposed computational approaches for measuring the “fairness” of image tagging systems, clustering them into five distinct categories, each with its own analytic foundation.¹ Next, we identify four types of representational harms that can be caused by image tagging systems, moving beyond generic concerns about “unfairness,” “bias,” or “discrimination.” We consider how different computational measurement approaches map to each of these types of harms, finding that each approach can be used to measure all four types. In light of this, we argue that no single measurement approach is definitive and that, without explicit reasoning, it is not possible to infer from the use of a particular measurement approach which type of harm was intended to be measured. Lastly, we conclude by exploring tensions between attempts to mitigate different types of representational harms—and what these tensions might mean for developing appropriate mitigations.

Background

In this section, we define both image tagging and representational (versus allocational) harms (Barocas et al. 2017).

Image Tagging

Image tagging is the task of applying tags to an image so as to describe salient aspects of its contents. Image tagging differs from the related task of object recognition, where the goal is to identify all objects that are depicted in an image.

¹We restrict our focus to only quantitative measurement approaches, setting aside a survey of qualitative measurement approaches (e.g., Barlas et al. 2019b; Hu and Strout 2018; Kay, Matuszek, and Munson 2015a; Miceli, Schuessler, and Yang 2020; Scheuerman et al. 2020; Offert and Bell 2020) for future work.

Crucially, although both tasks involve applying tags to images, object recognition focuses solely on objects and is concerned with completeness, whereas image tagging focuses on other properties (e.g., actions, facial expressions, social groups) as well as objects and is concerned with salience.

The two tasks also differ in terms of the ways the resulting tags are typically used. The tags applied by image tagging systems are often intended for human consumption. For example, image tagging systems are commonly used to generate alt text descriptions of images for blind and low-vision users (Wu et al. 2017), provide search or grouping functionality in photo-management software, and enable users of search engines to find images based on their contents as well as their accompanying text. In contrast, the tags applied by object recognition systems are often used to automate tasks, obviating the need for human involvement. For example, semi-autonomous vehicles rely on object recognition systems to detect obstacles, but may take autonomous actions to avoid colliding with those obstacles rather than simply communicating their presence to drivers.

Image tagging has a unique normative valence because it is concerned with salience and applies tags that are often intended for human consumption. The tags applied to images by image tagging systems are expressions of what—and who—is important and how they deserve to be described. Because the tags are consumed by people, they affect what people view as salient and how they understand the world.

Representational Harms

Much of the foundational work on the “fairness” of machine learning systems has focused on what Barocas et al. (2017) call allocational harms—that is, when people belonging to particular social groups are unfairly deprived of access to important opportunities or resources. Allocational harms often arise in domains like employment, finance, and housing. However, computer vision systems, like natural language processing systems, are not typically used to allocate opportunities or resources and are therefore less prone to causing allocational harms (Barocas et al. 2017; Crawford 2017; Blodgett et al. 2020). Instead, these systems produce outputs that can affect the understandings, beliefs, and attitudes that people hold about particular social groups, and thus the standings of those groups within society. Barocas et al. (2017) call these representational harms. Concerns about representational harms have motivated research examining, for example, the ways that search engines represent some social groups in less favorable ways than they do others (e.g., returning more images of men than women in response to search queries for “doctor” (Kay, Matuszek, and Munson 2015b) or returning pornographic websites in response to queries for “Black girls” but not “white girls” (Noble 2018)).

Computational Measurement Approaches

Equipped with definitions of image tagging and representational harms, we now analyze previously proposed computational approaches for measuring the “fairness” of image tagging systems. We focus on prior work that self-identifies as measuring “unfairness,” “bias,” or “discrimination.” Al-

though most of the measurement approaches that we describe were proposed for image tagging, some were proposed for other computer vision tasks but can be applied to image tagging. We cluster the approaches into five distinct categories based on their shared analytic foundations.

Incidence-based approaches test whether pre-identified problematic tags, types of images, or pairings of tags and types of images occur in the inputs and outputs of image tagging systems or in the datasets used to train them. These approaches are premised on the belief that there are particular tags, types of images, and pairings that are inherently problematic: problematic tags should not be applied to any images, problematic types of images should not have any tags applied to them, and problematic pairings should not occur.

For example, researchers have used incidence-based approaches to study tags that refer to properties that cannot be determined from visual appearance (Yang et al. 2020), such as gender (Keyes 2018; Scheuerman, Paul, and Brubaker 2019; Scheuerman et al. 2020), character traits (Barlas et al. 2019a), emotion (Bard 2020; Kyriakou et al. 2019; Barlas et al. 2019a; Sedenberg and Chuang 2017), and physical attractiveness (Kyriakou et al. 2019). As another example, researchers have uncovered the presence of lewd tags and images in several widely-used computer vision datasets (Crawford and Paglen 2019; Yang et al. 2020; Prabhu and Birhane 2020). Finally, one of the most widely discussed problematic pairings arose when Google Photos applied the tag *gorillas* to an image of Black people (Simonite 2018). Although neither the tag nor the image would have been considered problematic in isolation, they become so when paired together.

Because problematic tags, types of images, or pairings must be pre-identified, using an incidence-based approach means first answering the question of *which* tags, types of images, or pairings should be considered problematic—a task that requires extensive background knowledge. Yet even with such knowledge, pre-identifying *all* problematic tags, types of images, or pairings is likely impossible.

Distribution-based approaches compare differences in the distributions of tags or types of images in the inputs and outputs of image tagging systems or in the datasets used to train them. One distribution-based approach compares the distribution of images that depict different social groups in a training dataset to some reference distribution, typically the uniform distribution (Buolamwini and Gebru 2018; Kärkkäinen and Joo 2021; Wang et al. 2021; Yang et al. 2020). A more ambitious variant of this approach focuses on the distributions of images that depict attributes (e.g., payot), artifacts (e.g., the Torah), or activities (e.g., lighting the menorah) that are bound up with the identities of different social groups, again comparing them to some reference distribution (Shankar et al. 2017a; Wang et al. 2021; DeVries et al. 2019). Both variants are premised on the idea that image tagging systems may under-perform for social groups, attributes, artifacts, or activities that are not sufficiently well represented in the datasets used to train them.

Another distribution-based approach compares the distributions of tags—either in a training dataset (Barlas et al. 2019a; Otterbacher et al. 2019; Wang et al. 2021; Yang et al. 2020) or output by an image tagging system (Alvi, Zisser-

man, and Nellaker 2018; Kyriakou et al. 2019; Barlas et al. 2019a; Schwemmer et al. 2020)—for images that depict different social groups. Other researchers have instead compared the distributions of social groups depicted in images for different tags (Kay, Matuszek, and Munson 2015a). Both variants of this approach are premised on the belief that it is harmful for images of different social groups to have different distributions of tags, even if these tags are correctly applied, because this casts the groups in different lights.

A final approach compares the distribution of tags applied to images that depict different social groups in a image tagging system’s training dataset to the distribution of tags applied by the system to images that depict these social groups (Wang et al. 2019, 2020; Zhao et al. 2017; Kay, Matuszek, and Munson 2015a; Wang and Russakovsky 2021). This approach is premised on the belief that the outputs of image tagging systems should not exacerbate any differences between social groups that are present in their training datasets.

Performance-based approaches compare the performance of image tagging systems for particular tags or types of images. These approaches are premised on the belief that image tagging systems should perform equally well for different social groups. One widely discussed performance-based approach tests whether image tagging systems exhibit comparable performance for images that depict different social groups when applying tags that indicate membership in those social groups. For example, researchers have investigated whether binary gender classification systems perform better for images that depict men than they do for images that depict women (Bhargava and Forsyth 2019; Buolamwini and Gebru 2018; Hendricks et al. 2018). Other researchers have conducted similar analyses focusing on race, ethnicity, and age (Das, Dantcheva, and Bremond 2018; Kärkkäinen and Joo 2021; Kyriakou et al. 2019). Some researchers have also taken an intersectional lens by focusing on different combinations of social groups (Buolamwini and Gebru 2018; Das, Dantcheva, and Bremond 2018).

Another performance-based approach tests whether image tagging systems exhibit comparable performance for images that depict different social groups when applying *any* tags. For example, researchers have investigated whether image tagging systems perform better when tagging everyday objects in images that depict some social groups than they do for images that depict others, finding performance disparities between genders (Bhargava and Forsyth 2019; Wang et al. 2020), ages (Wang et al. 2020), races (Buolamwini and Gebru 2018; Wilson, Hoffman, and Morgenstern 2019), socioeconomic statuses (DeVries et al. 2019), and geographic locations (Shankar et al. 2017a; DeVries et al. 2019)—as well as different combinations thereof (Buolamwini and Gebru 2018; Khiyari and Wechsler 2016; Klare et al. 2012).

Perturbation-based approaches involve varying aspects of images provided as inputs to image tagging systems to see whether different tags are applied, either correctly or incorrectly. These approaches are premised on two beliefs: first, that the behaviors of image tagging systems should not reflect spurious correlations in their training datasets; and second, that particular aspects of images should not be the basis for particular differences in tagging behaviors (e.g., the

appearance of a person’s face should not affect the occupation with which they are tagged). As a result, they are often used to investigate whether particular aspects of images are driving particular tagging behaviors, providing an opportunity to reflect on whether these behaviors are normatively concerning (Wilson, Hoffman, and Morgenstern 2019; Muthukumar et al. 2018; Klare et al. 2012). For example, if removing occluded faces from images reduces performance disparities between social groups, these disparities are likely due to occlusions rather than something more normatively concerning (Wilson, Hoffman, and Morgenstern 2019).

One perturbation-based approach focuses specifically on invariance. This approach tests whether the tags applied by image tagging systems are unaffected by perturbations of particular aspects of images that are known to be an inappropriate basis for *any* differences in tagging behaviors. For example, researchers have re-taken photos using props, such as wigs (Rodden 2017; Algorithmic Justice League 2021), or used GANs to generate faces with different characteristics (Denton et al. 2021; McDuff et al. 2019), testing whether there are any differences in the tags applied to these images. Other researchers have selected pairs of semantically similar images that depict different social groups (e.g., similarly composed photos of CEOs of different genders), again testing for any differences in tagging behaviors (Otterbacher 2018; Steed and Caliskan 2021; Stock and Cisse 2018).

Internals-based approaches directly investigate the latent representations learned by image tagging systems. Like perturbation-based approaches, these approaches are premised on the belief that the behaviors of image tagging systems should not reflect spurious correlations in their training datasets and the belief that particular aspects of images should not be the basis for particular differences in tagging behaviors. One internals-based approach relies on saliency maps, which highlight the regions of images that are most salient when applying particular tags and are typically generated using gradients. This approach can therefore shed light on the reasoning behind particular tagging behaviors—for example, whether an image tagging system is focusing on the people depicted in an image when tagging objects—regardless of whether the tags in question are correctly applied. Saliency maps have been used to investigate tagging behaviors from quantitative perspectives (e.g., when tagging an object depicted in an image, investigating how much of the object’s bounding box overlaps with the most salient region of the image) and qualitative perspectives (e.g., investigating what is depicted in the most salient region of an image) (Hendricks et al. 2018; Jia, Lansdall-Welfare, and Cristianini 2018; Muthukumar et al. 2018; Singh et al. 2020).

A second internals-based approach tests whether image tagging systems that are not intended to apply tags that indicate membership in social groups still encode notions of social groups in their feature representations. This suggests that they are inappropriately using these notions as the basis for tagging behaviors (Alvi, Zisserman, and Nellaker 2018; Roberts 2018; Song and Shmatikov 2020; Wang et al. 2019).

Representational Harms

Although the computational measurement approaches that we described in the previous section have played an important role in raising awareness of and developing mitigations for some of the harms caused by image tagging systems, their normative motivations are often unclear. These approaches are all premised on beliefs that particular properties or behaviors of image tagging systems or the datasets used to train them are self-evidently problematic, even though these properties or behaviors are rarely connected to concrete harms. This lack of clarity is not unique to image tagging systems or even computer vision systems—Blodgett et al. (2020) found that computational approaches for measuring the “fairness” of natural language processing systems often have similarly unclear normative motivations.

In this section, we draw on the work of Blodgett et al. (2020) to argue that image tagging systems cause representational harms by contributing to the reproduction of harmful social hierarchies—that is, by casting some social groups as inferior to others and confining them to subordinate positions within society. We emphasize that representational harms are not harms to specific individuals; they relate to the standings of particular social groups, affecting *all* members of those groups (Selbst and Barocas 2023). We identify four types of representational harms that can be caused by image tagging systems, delineating how each one reproduces harmful social hierarchies. Having done this, we then introduce an additional normative concern—that is, when image tagging systems deny people the ability to self-identify—and discuss how it can give rise to representational harms.

We do not claim to provide an exhaustive list of all of the types of representational harms that can be caused by image tagging systems. The range of possible variation for tags and images make them remarkably expressive. As a result, anticipating every way in which these variations could reproduce harmful social hierarchies is impossible. Rather, our aim is to bring greater structure and precision to discussions of some of the harms caused by image tagging systems and the many computational approaches that have been proposed to measure them. We leave open the possibility that future work will identify additional types of representational harms as our collective understanding of image tagging evolves.

Four Types of Representational Harms

Reifying social groups. Social groups are socially constructed, so their boundaries are always contested and changing. Despite this inherent fluidity, social groups are often treated as if they are natural, fixed, or objective (Hanna et al. 2020), thereby reifying beliefs about their salience and immutability and beliefs about the boundaries between them. Because image tagging systems rely on visual characteristics and are intended to apply pre-identified tags, they are especially likely to reify social groups, as well as the idea that there are relationships between social groups and visual appearance (Keyes 2018). For example, binary gender classification systems rely on the following beliefs: first, that gender can be determined from visual appearance; and second, that there are only two gender identities. However,

relationships between gender and visual appearance are contested and unstable, as are gender identities themselves.

Stereotyping social groups. Image tagging systems can also perpetuate stereotypes—over-generalized beliefs about social groups that reproduce harmful social hierarchies (Barlas et al. 2021; Bhargava and Forsyth 2019). For example, researchers have demonstrated that some image tagging systems systematically apply the tag *nurse* to images of female doctors. Other researchers have shown that images of female snowboarders are often incorrectly tagged as depicting men (Hendricks et al. 2018; Zhao et al. 2017). In these examples, the latent representations of doctors and snowboarders learned by the image tagging systems are inextricably linked with gender—that is, to be a female medical professional is to be a nurse, while to be a snowboarder is to be a man.

These examples suggest that there are two ways that image tagging systems can perpetuate stereotypes. First, if an image tagging system is intended to apply tags that indicate membership in social groups, then it will necessarily associate particular visual characteristics (e.g., riding a snowboard) with particular social groups (e.g., men). This means that images that depict people with those visual characteristics will be more likely to be tagged as depicting members of those social groups. Second, if an image tagging system is intended to apply other tags, then it may associate visual characteristics that are common among members of particular social groups (e.g., long hair among women) with other properties (e.g., being a nurse). This means that tags that refer to those properties are more likely to be applied to images that depict people with those visual characteristics.

Finally, image tagging systems can also perpetuate stereotypes in less overt ways. For example, some image tagging systems tend to apply tags that refer to professions to images that depict men, while applying tags that refer to appearance to images that depict women (Schwemmer et al. 2020). Even if these tags are correctly applied, they are harmful because they cast men and women in different lights by perpetuating the stereotype that professions are particularly salient for men, while appearance is particularly salient for women.

Demeaning social groups. Image tagging systems can also demean social groups—that is, suggest that members of those groups are less worthy of respect than members of other groups. This can occur in very overt ways, such as when image tagging systems apply tags that contain racial epithets—the very purpose of which is to belittle members of particular racial groups—or tags that are not inherently problematic but are demeaning when applied to images that depict particular social groups. As an example of the latter, applying the tag *gorilla* to an image of Black people is demeaning because there is a long history of diminishing and dehumanizing Black people by likening them to gorillas (Simone 2018). Image tagging systems can also demean particular social groups by applying tags that devalue attributes, artifacts, or activities that are bound up with the identities of those groups. For example, applying the tag *costume* to an image that depicts a religious group’s wedding clothing trivializes the group’s traditions (Shankar et al. 2017b). Evaluative terms can also be demeaning because they can be applied more (or less) often to images that depict some

social groups than to images that depict others. For example, image tagging systems that apply the tag *beautiful* to fewer images that depict a particular social group than they do to images that depict others cast that group as being worthy of less esteem (Levin 2016) and reinforce harmful beliefs about visual characteristics that are considered beautiful.

Erasing social groups. Erasure occurs when image tagging systems fail to correctly apply tags that indicate membership in social groups or tags that refer to attributes, artifacts, or activities that are bound up with the identities of social groups (Keyes 2018; Buolamwini and Gebru 2018; Benjamin 2019). Erasure of a particular social group implies that members of that group are not worthy of recognition and contributes to their further marginalization within society. For example, image tagging systems that do not apply the tag *person* to images that depict people wearing hijabs erase Muslims. Similarly, image tagging systems that fail to apply the tag *menorah* to images that depict menorahs—perhaps because the tag is not available, despite the availability of other tags that refer to Christian artifacts—erase Judaism.

Image tagging systems can also cause harms when they fail to acknowledge the relevance of people’s membership in particular social groups to what is depicted in images—for example, by applying the tags *people*, *walking*, and *street* to an image depicting women suffragists marching. As well as down-playing the gravity of what is depicted, these tags fail to acknowledge the injustices that woman have suffered specifically due to their gender. Erasure of this sort denies the role that membership in particular social groups plays in people’s lived experiences, especially experiences relating to oppression and violence, which are often perpetrated on the basis of people’s membership in particular social groups.

Denying People the Ability to Self-Identify

In this section, we introduce an additional normative concern posed by image tagging systems—denying people the ability to self-identify—and discuss how it can give rise to representational harms. Image tagging systems are often used to sort people into social groups based on their appearance (e.g., gender classification systems), imposing tags on them without their awareness, involvement, or consent. Given the degree to which membership in social groups shapes the way that people are treated within society, including others’ expectations of them (Baker, Hanna, and Denton 2020), this imposition can have profound consequences and may cause people to feel like they have been deprived of a crucial aspect of their autonomy (Hanna et al. 2020).

We consider denying people the ability to self-identify to be a type of individual harm—that is, a harm that affects particular individuals as opposed to the standings of particular social groups. However, patterns in the ways that image tagging systems deny people the ability to self-identify can cause each of the four types of representational harms that we described above, as well as allocational harms. First, it reifies the belief that social groups can be determined from visual appearance. Second, the visual characteristics that image tagging systems rely on when tagging people as members of social groups can perpetuate stereotypes. Third, imposing tags that indicate membership in social

groups onto people can be demeaning—for example by misgendering them or by accurately applying tags (e.g., the tag *women*) to some social groups (e.g., white women) but not others (e.g., Black women). Fourth, if the available tags exclude particular social groups (e.g., people who are non-binary), those groups will be erased. Finally, imposing tags that indicate membership in social groups onto people may increase their visibility and thus their vulnerability to violence, surveillance, loss of employment, and other harms.

Mapping Harms to Measurement Approaches

In this section, we consider how computational measurement approaches belonging to each of the five categories described earlier map to each of the four types of representational harms identified in the previous section, demonstrating that there is not a one-to-one mapping. Due to space constraints, we only provide an in-depth discussion for stereotyping, but include examples for the other types in figure 1.

Stereotyping Social Groups

Stereotyping can be measured using many different computational approaches. One possibility is to use an incidence-based approach to test, for example, whether pre-identified tags that contain stereotyping language, such as *boyish* or *girly*, are applied to any images. Another possibility is to use a distribution-based approach to see whether any differences in the distributions of tags for images that depict different social groups correspond to stereotypes—for example, whether tags that refer to professions are more frequently applied to images that depict men, while tags that refer to appearance are more frequently applied to women (Schwemmer et al. 2020). A third possibility is to use a performance-based approach to compare, for example, the accuracy with which the tag *doctor* is applied to images that depict men to the accuracy with which it is applied to images that depict people of other genders. A fourth possibility is to use a perturbation-based approach to test for invariance in the tags applied to pairs of images that are semantically equivalent, such as similarly composed photos of CEOs of different genders. A final possibility is to use an internals-based approach to investigate whether the regions of an image that are most salient when, for example, incorrectly applying the tag *man* to an image that depicts a woman correspond to stereotypes about gender presentation.

Implications of these Mappings

As demonstrated by the exercise of mapping different computational measurement approaches to different types of representational harms, no single measurement approach is definitive. Different facets of each type of harm will be revealed by different approaches (Jacobs and Wallach 2021). As a result, measuring representational harms does not involve identifying and then using the single best measurement approach. Rather, many different approaches should be used, with the resulting measurements interpreted in the context of what those approaches are capable of revealing. The exercise also demonstrates that without explicit reasoning, it is not possible to infer from the use

	Reifying	Stereotyping	Demeaning	Erasing
Incidence-based	Available tags include different races	Available tags include <i>boyish</i> and <i>girly</i>	Available tags include racial epithets	Available tags include <i>Bible</i> , but not <i>Quran</i> or <i>Torah</i>
Distribution-based	Distribution of gender tags is different for images that depict people with different hair lengths	System more frequently applies tags that refer to professions to images that depict men than to images that depict people of other genders	System more frequently applies the tag <i>beautiful</i> to images that depict white people than to images that depict other racial groups	System tags foods from the Global North by their specific names, but tags foods from the Global South as <i>food</i>
Performance-based	System appears to perfectly apply gender tags, even though this is not possible	System more accurately applies the tag <i>doctor</i> to images that depict men than to images that depict people of other genders	System more accurately applies the tag <i>human</i> to images that depict white people than to images that depict other racial groups	System more accurately tags objects in images from the Global North than objects in images from the Global South
Perturbation-based	System applies different gender tags when varying a person's appearance	System applies different tags to standardized headshots of politicians of different genders	System applies the tag <i>costume</i> when a cheongsam is artificially added to an image	System fails to apply the tag <i>person</i> when an assistive device is artificially added to an image that depicts a person
Internals-based	System's feature representation encodes a notion of gender even though the system is not intended to apply gender tags	Saliency map reveals that the system is focusing on a truck when incorrectly applying the tag <i>man</i> to an image that depict people of other genders	Saliency map reveals that the system is focusing on facial features when applying the tag <i>janitor</i> to an image that depicts a person of color	Saliency map reveals that the system is ignoring people who use wheelchairs when applying the tag <i>people</i> to an image that depicts a group of people

Figure 1: Examples of how computational measurement approaches belonging to each of the five categories described earlier might be used to measure each of the four different types of representational harms that can be caused by image tagging systems. This mapping is by no means exhaustive.

of a particular measurement approach which type of harm was intended to be measured. We therefore recommend explicitly stating the type of harm of interest so as to avoid any ambiguity or confusion. In concurrent work, we provide a concrete demonstration of how to go about this type of analysis, using a variety of computational approaches to measure different types of representational harms that can be caused by image captioning systems (Wang et al. 2022).

Tensions When Mitigating Harms

Lastly, equipped with a more granular understanding of the types of representational harms that can be caused by image tagging systems, we explore tensions between attempts to mitigate some of these types of harms—and what these tensions might mean for developing appropriate mitigations.

The most obvious way to mitigate reifying social groups is simply to remove tags that indicate membership in social groups. By removing such tags, image tagging systems can avoid enforcing boundaries between social groups and allow members of social groups remain unmarked, decreasing their visibility and thus their vulnerability to violence, surveillance, loss of employment, and other harms (Hofmann et al. 2020). In contrast, the most obvious way to mitigate erasing social groups is make sure that tags that indicate membership in social groups exist and are correctly applied. In some cases, marking social groups—that is, calling them out by name—can bring much needed recognition of the injustices they have suffered. In other cases, marking privileged social groups that often go unmarked (e.g., men) can provide a way to draw attention to their privilege.

As another example, removing tags that refer to attributes, artifacts, or activities that are bound up with the identities of social groups can mitigate concerns that such tags might be applied incorrectly, thereby demeaning particular social groups. However, removing tags associated with particular social groups can instead cause those groups to be erased.

In some scenarios, particular types of images are deemed to be inherently problematic—for example, pornographic images that stereotype and demean women—meaning that they should not have any tags applied to them. However, such images cannot be ignored or removed unless they are first identified, usually by some computer vision system explicitly trained for that purpose. Moreover, systems for ignoring or removing particular types of images can erase social groups by incorrectly labeling images that are associated with those groups as problematic. For example, when Tumblr banned “adult content,” they also ended up removing posts from queer communities and sex-positive communities (Pilipets and Paasonen 2020; Bronstein 2020).

As a final example, differences in the distributions of tags for images that depict different social groups often correspond to stereotypes. One straightforward way to mitigate these stereotypes is to ensure that the distributions of tags for different social groups are similar—for example, that the distribution of tags applied to images of Black football players is similar to the distribution of tags applied to white football players (Rhue 2018). However, constraining the distributions of tags for different social groups can be seen as a

denial of the role that membership in particular social groups plays in people’s lived experiences—a form of erasure.

These tensions suggest that developing appropriate mitigations is a complex task, with many considerations. However, in many cases, the scenario in which an image tagging system is intended to be used can (at least partially) determine the appropriateness of different mitigations. For example, the developers of an image tagging system for a social media platform might remove particular tags to avoid stereotyping or demeaning particular social groups. In contrast, the developers of an image tagging system for blind and low-vision users might decide that the consequences of erasing social groups are worse than the consequences of incorrectly applying tags (Bennett et al. 2021; Hanley et al. 2021).

Conclusion

In this paper, we argued that generic concerns about “unfairness,” “bias,” or “discrimination” in the context of image tagging should instead be understood in terms of the reproduction of harmful social hierarchies—that is, representational harms. We analyzed previously proposed computational approaches for measuring the “fairness” of image tagging systems, finding that they cluster into five distinct categories, each with its own analytic foundation. We also identified four types of representational harms that can be caused by image tagging systems, delineating how each one reproduces harmful social hierarchies. By considering how different computational measurement approaches map to each of these types of harms, we demonstrated that each approach can be used to measure all four types, with different facets of each type revealed by different approaches. Lastly, we showed that attempts to mitigate different types of representational harms may be in tension with one another and explored what these tensions might mean for developing appropriate mitigations. Moving forward, we hope that this paper contributes to more precise and actionable discussions of some of the harms caused by image tagging systems.

References

- Algorithmic Justice League. 2021. Drag Vs AI Workshop.
- Alvi, M.; Zisserman, A.; and Nellaker, C. 2018. Turning a Blind Eye: Explicit Removal of Biases and Variation from Deep Neural Network Embeddings. In *Workshop on Bias Estimation in Face Analytics at ECCV 2018*.
- Baker, D.; Hanna, A.; and Denton, E. 2020. Algorithmically Encoded Identities: Reframing Human Classification. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT* ’20*, 681. New York, NY, USA: Association for Computing Machinery. ISBN 9781450369367.
- Bard, J. 2020. Developing a Legal Framework for Regulating Emotion AI. In *University of Florida Levin College of Law Research Paper*.
- Barlas, P.; Kyriakou, K.; Guest, O.; Kleanthous, S.; and Otterbacher, J. 2021. To “See” is to Stereotype: Image Tagging Algorithms, Gender Recognition, and the Accuracy-Fairness Trade-off. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW3): 1–31.
- Barlas, P.; Kyriakou, K.; Kleanthous, S.; and Otterbacher, J. 2019a. Social B(eye)as: Human and Machine Descriptions of People Images. In *Proceedings of the International AAAI Conference on Web and Social Media*.
- Barlas, P.; Kyriakou, K.; Kleanthous, S.; and Otterbacher, J. 2019b. What Makes an Image Tagger Fair? Proprietary Auto-tagging and Interpretations on People Images. In *Proceedings of the 27th ACM Conference On User Modelling, Adaptation And Personalization (UMAP)*.
- Barocas, S.; Crawford, K.; Shapiro, A.; and Wallach, H. 2017. The Problem With Bias: Allocative Versus Representational Harms in Machine Learning. In *Proceedings of SIGCIS*. Philadelphia, PA.
- Benjamin, R. 2019. *Race After Technology: Abolitionist Tools for the New Jim Code*. John Wiley & Sons.
- Bennett, C. L.; Gleason, C.; Scheuerman, M. K.; Bigham, J. P.; Guo, A.; and To, A. 2021. “It’s Complicated”: Negotiating Accessibility and (Mis) Representation in Image Descriptions of Race, Gender, and Disability. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–19.
- Bhargava, S.; and Forsyth, D. 2019. Exposing and correcting the gender bias in image captioning datasets and models. *arXiv preprint arXiv:1912.00578*.
- Blodgett, S. L.; Barocas, S.; Daumé, H., III; and Wallach, H. 2020. Language (Technology) is Power: A Critical Survey of “Bias” in NLP. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, 5454–5476.
- Bronstein, C. 2020. Pornography, Trans Visibility, and the Demise of Tumblr. In *Transgender Studies Quarterly*.
- Buolamwini, J.; and Gebru, T. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. New York City, NY.
- Crawford, K. 2017. The Trouble with Bias. *Conference on Neural Information Processing Systems Keynote*.
- Crawford, K.; and Paglen, T. 2019. Excavating AI: the politics of images in machine learning training sets. *Excavating AI*.
- Das, A.; Dantcheva, A.; and Bremond, F. 2018. Mitigating Bias in Gender, Age and Ethnicity Classification: a Multi-Task Convolution Neural Network Approach. In *Workshop at ECCV 2018*.
- Denton, E.; Hutchinson, B.; Mitchell, M.; Gebru, T.; and Zaldivar, A. 2021. Image Counterfactual Sensitivity Analysis for Detecting Unintended Bias. *arXiv preprint arXiv:1906.06439*.
- DeVries, T.; Misra, I.; Wang, C.; and van der Maaten, L. 2019. Does Object Recognition Work for Everyone? *arXiv:1906.02659*.
- Hanley, M.; Barocas, S.; Levy, K.; Azenkot, S.; and Nissenbaum, H. 2021. Computer vision and conflicting values: Describing people with automated alt text. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, 543–554.

- Hanna, A.; Denton, E.; Smart, A.; and Smith-Loud, J. 2020. Towards a Critical Race Methodology in Algorithmic Fairness. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 501–512.
- Hendricks, L. A.; Burns, K.; Saenko, K.; Darrell, T.; and Rohrbach, A. 2018. Women also snowboard: Overcoming bias in captioning models. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 771–787.
- Hofmann, M.; Kasnitz, D.; Mankoff, J.; and Bennett, C. L. 2020. Living Disability Theory: Reflections on Access, Research, and Design. In *The 22nd International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS)*.
- Hu, Z.; and Strout, J. 2018. Exploring Stereotypes and Biased Data with the Crowd. In *arXiv:1801.03261*.
- Jacobs, A. Z.; and Wallach, H. 2021. Measurement and fairness. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 375–385.
- Jia, S.; Lansdall-Welfare, T.; and Cristianini, N. 2018. Right for the Right Reason: Training Agnostic Networks. In *Lecture Notes in Computer Science*.
- Kay, M.; Matuszek, C.; and Munson, S. A. 2015a. 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, CHI '15*, 3819–3828. New York, NY, USA: Association for Computing Machinery. ISBN 9781450331456.
- Kay, M.; Matuszek, C.; and Munson, S. A. 2015b. 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.
- Keyes, O. 2018. The misgendering machines: Trans/HCI implications of automatic gender recognition. *Proceedings of the ACM on human-computer interaction*, 2(CSCW): 1–22.
- Khiyari, H. E.; and Wechsler, H. 2016. Face Verification Subject to Varying (Age, Ethnicity, and Gender) Demographics Using Deep Learning. *Journal of biometrics & biostatistics*, 7: 1–5.
- Klare, B. F.; Burge, M. J.; Klontz, J. C.; Vorder Bruegge, R. W.; and Jain, A. K. 2012. Face Recognition Performance: Role of Demographic Information. *IEEE Transactions on Information Forensics and Security*, 7(6): 1789–1801.
- Kyriakou, K.; Barlas, P.; Kleanthous, S.; and Otterbacher, J. 2019. Fairness in Proprietary Image Tagging Algorithms: A Cross-Platform Audit on People Images. In *Proceedings of the International AAAI Conference on Web and Social Media*.
- Kärkkäinen, K.; and Joo, J. 2021. FairFace: Face Attribute Dataset for Balanced Race, Gender, and Age. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*.
- Levin, S. 2016. A beauty contest was judged by AI and the robots didn't like dark skin. *The Guardian*.
- McDuff, D.; Ma, S.; Song, Y.; and Kapoor, A. 2019. Characterizing Bias in Classifiers using Generative Models. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Miceli, M.; Schuessler, M.; and Yang, T. 2020. Between Subjectivity and Imposition: Power Dynamics in Data Annotation for Computer Vision. In *ACM Conference on Computer Supported Cooperative Work (CSCW)*.
- Muthukumar, V.; Pedapati, T.; Ratha, N.; Sattigeri, P.; Wu, C.-W.; Kingsbury, B.; Kumar, A.; Thomas, S.; Mojsilovic, A.; and Varshney, K. R. 2018. Understanding Unequal Gender Classification Accuracy from Face Images. In *arXiv:1812.00099*.
- Noble, S. U. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. NYU Press.
- Offert, F.; and Bell, P. 2020. Perceptual bias and technical metaphors: critical machine vision as a humanities challenge. In *AI & Society*.
- Otterbacher, J. 2018. Social Cues, Social Biases: Stereotypes in Annotations on People Images. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*.
- Otterbacher, J.; Barlas, P.; Kleanthous, S.; and Kyriakou, K. 2019. How Do We Talk about Other People? Group (Un)Fairness in Natural Language Image Descriptions. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*.
- Pilipets, E.; and Paasonen, S. 2020. Nipples, memes, and algorithmic failure: NSFW critique of Tumblr censorship. In *New Media & Society*.
- Prabhu, V. U.; and Birhane, A. 2020. Large image datasets: A pyrrhic win for computer vision? *arXiv:2006.16923*.
- Rhue, L. 2018. Racial Influence on Automated Perceptions of Emotions. *SSRN Electronic Journal*.
- Roberts, C. V. 2018. Quantifying the Extent to which Popular Pre-Trained Convolutional Neural Networks Implicitly Learn High-Level Protected Attributes. In *Princeton Masters Thesis*.
- Rodden, K. 2017. Is that a boy or a girl? Exploring a neural network's construction of gender. *Medium*.
- Scheuerman, M. K.; Paul, J. M.; and Brubaker, J. R. 2019. How Computers See Gender: An Evaluation of Gender Classification in Commercial Facial Analysis Services. *Proc. ACM Hum.-Comput. Interact.*, 3(CSCW).
- Scheuerman, M. K.; Wade, K.; Lustig, C.; and Brubaker, J. R. 2020. How We've Taught Algorithms to See Identity: Constructing Race and Gender in Image Databases for Facial Analysis. *Proc. ACM Hum.-Comput. Interact.*, 4(CSCW1).
- Schwemmer, C.; Knight, C.; Bello-Pardo, E. D.; Oklobdzija, S.; Schoonvelde, M.; and Lockhart, J. W. 2020. Diagnosing gender bias in image recognition systems. *Socius*, 6: 2378023120967171.
- Sedenberg, E.; and Chuang, J. 2017. Smile for the Camera: Privacy and Policy Implications of Emotion AI. *arXiv:1709.00396*.

- Selbst, A.; and Barocas, S. 2023. Unfair Artificial Intelligence: How FTC Intervention Can Overcome the Limitations of Discrimination Law. *University of Pennsylvania Law Review*, 171.
- Shankar, S.; Halpern, Y.; Breck, E.; Atwood, J.; Wilson, J.; and Sculley, D. 2017a. No Classification without Representation: Assessing Geodiversity Issues in Open Data Sets for the Developing World. arXiv:1711.08536.
- Shankar, S.; Halpern, Y.; Breck, E.; Atwood, J.; Wilson, J.; and Sculley, D. 2017b. No classification without representation: Assessing geodiversity issues in open data sets for the developing world. In *Proceedings of the NIPS 2017 Workshop on Machine Learning for the Developing World*.
- Simonite, T. 2018. When It Comes to Gorillas, Google Photos Remains Blind. *Wired, January*.
- Singh, K. K.; Mahajan, D.; Grauman, K.; Lee, Y. J.; Feiszli, M.; and Ghadiyaram, D. 2020. Don't Judge an Object by Its Context: Learning to Overcome Contextual Bias. arXiv:2001.03152.
- Song, C.; and Shmatikov, V. 2020. Overlearning Reveals Sensitive Attributes. arXiv:1905.11742.
- Steed, R.; and Caliskan, A. 2021. Image Representations Learned With Unsupervised Pre-Training Contain Human-like Biases. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21*, 701–713. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383097.
- Stock, P.; and Cisse, M. 2018. ConvNets and ImageNet Beyond Accuracy: Understanding Mistakes and Uncovering Biases. arXiv:1711.11443.
- Wang, A.; Barocas, S.; Laird, K.; and Wallach, H. 2022. Measuring Representational Harms in Image Captioning. *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*.
- Wang, A.; Liu, A.; Zhang, R.; Kleiman, A.; Kim, L.; Zhao, D.; Shirai, I.; Narayanan, A.; and Russakovsky, O. 2021. REVISE: A Tool for Measuring and Mitigating Bias in Visual Datasets. arXiv:2004.07999.
- Wang, A.; and Russakovsky, O. 2021. Directional Bias Amplification. arXiv:2102.12594.
- Wang, T.; Zhao, J.; Yatskar, M.; Chang, K.-W.; and Ordonez, V. 2019. Balanced Datasets Are Not Enough: Estimating and Mitigating Gender Bias in Deep Image Representations. arXiv:1811.08489.
- Wang, Z.; Qinami, K.; Karakozis, I. C.; Genova, K.; Nair, P.; Hata, K.; and Russakovsky, O. 2020. Towards Fairness in Visual Recognition: Effective Strategies for Bias Mitigation. arXiv:1911.11834.
- Wilson, B.; Hoffman, J.; and Morgenstern, J. 2019. Predictive Inequity in Object Detection. arXiv:1902.11097.
- Wu, S.; Wieland, J.; Farivar, O.; and Schiller, J. 2017. Automatic Alt-text: Computer-generated Image Descriptions for Blind Users on a Social Network Service. *ACM Conference on Computer-Supported Cooperative Work And Social Computing (CSCW)*.
- Yang, K.; Qinami, K.; Fei-Fei, L.; Deng, J.; and Russakovsky, O. 2020. Towards fairer datasets. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*.
- Zhao, J.; Wang, T.; Yatskar, M.; Ordonez, V.; and Chang, K.-W. 2017. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2979–2989.