

5-2016

# Principles of Anti-Discriminatory Design

D. E. Wittkower

*Old Dominion University*, [dwittkow@odu.edu](mailto:dwittkow@odu.edu)

Follow this and additional works at: [https://digitalcommons.odu.edu/philosophy\\_fac\\_pubs](https://digitalcommons.odu.edu/philosophy_fac_pubs)

 Part of the [Graphics and Human Computer Interfaces Commons](#), and the [Philosophy Commons](#)

---

## Repository Citation

Wittkower, D. E., "Principles of Anti-Discriminatory Design" (2016). *Philosophy Faculty Publications*. 28.  
[https://digitalcommons.odu.edu/philosophy\\_fac\\_pubs/28](https://digitalcommons.odu.edu/philosophy_fac_pubs/28)

## Original Publication Citation

Wittkower, D. E. (2016, 13-14 May 2016). Principles of anti-discriminatory design. Paper presented at the 2016 IEEE International Symposium on Ethics in Engineering, Science and Technology (ETHICS).

This Conference Paper is brought to you for free and open access by the Philosophy & Religious Studies at ODU Digital Commons. It has been accepted for inclusion in Philosophy Faculty Publications by an authorized administrator of ODU Digital Commons. For more information, please contact [digitalcommons@odu.edu](mailto:digitalcommons@odu.edu).

# Principles of Anti-Discriminatory Design

D.E. Wittkower  
Department of Philosophy  
Old Dominion University  
Norfolk, USA  
dwittkow@odu.edu

***Abstract***—Technical design can produce exclusionary and even discriminatory effects for users. A lack of discriminatory intent is insufficient to avoid discriminatory design, since implicit assumptions about users rarely include all relevant user demographics, and in some cases, designing for all relevant users is actually impossible. To minimize discriminatory effects of technical design, an actively anti-discriminatory design perspective must be adopted.

This article provides examples of discriminatory user exclusion, then defining exclusionary design in terms of disaffordances and dysaffordances. Once these definitions are in place, principles of anti-discriminatory design are advanced, drawing upon a method of phenomenological variation employed in the context of standpoint epistemology.

***Keywords***—*design; affordances; discrimination; race; gender; religion; engineering*

## I. INTRODUCTION

Technological design cannot possibly accommodate every user, but there are clear cases where the line is crossed from mere loss of functionality to discriminatory user exclusion, as recognized legally through e.g. ADA compliance and as recognized socially through e.g. popular rejection of so-called “flesh colored” crayons. This paper constructs a theory of disaffordances and dysaffordances that identifies forms of exclusion which are materially or socially exclusive in ways which can rise to the level of discrimination, then recommending principles which may aid in proactively anti-discriminatory design processes.

## II. USER EXCLUSION AND THE DISCRIMINATORY THRESHOLD

The exclusion of some users seems inevitable in many design contexts, and in many of these cases user exclusion is not discriminatory. A clothing designer, for the most part, designs an article of clothing as conforming with male or female gendered norms for self-presentation, and it would clearly be a misconstrual of an appropriate understanding of discriminatory design to fault a woman’s dress or blouse for failing to afford men’s gender-typical self-presentation. Here, patterns of user behavior require design targeted at a distinct user demographic, and require such targeting to take place in a way which tends to be exclusionary of other demographics.

It is less clearly non-discriminatory if a designer chooses to produce, for example, clothing not only for women in

particular, but for very thin women in particular, since here norms of user behavior do not delimit the design space: making a particular dress design in a large variety of sizes may require creativity and artfulness in order to give the right effect for all wearers, but it does not present either the same degree or the same kind of problem as designing a dress to be gender-typical clothing for a variety of gender presentations. Clothing, however, can be obtained from a great variety of designers without significantly compromising its functionality for the user, and there are designers enough to provide reasonable access to the ability to clothe oneself to the variety of kinds of user embodiment, with the possible exception of those persons whose bodies are at the upper limits of humanly possible size. It may be that designers cluster in the design space that caters to privileged body types, producing more variety and availability to some rather than others, but it seems perhaps an overstrong claim to describe this effect as discriminatory.

By contrast, were a company to require job applications to be submitted via an online system which was incompatible with screen readers—assuming that being sighted was not a bona fide occupational qualification (BFOQ)—this would be clearly discriminatory, as it would remove eligibility for employment for a non-employment-related reason. The difference (1) is not merely in ease of access, but in access itself; (2) does not reflect a difference in user demographics that requires distinct provision of services to distinct demographics; and (3) cannot be accommodated through multiple equivalent providers of service, since employers provide substantially different opportunities to employees, unlike e.g. clothes which perform at least their primary functions with little difference between available brand names.

These examples illustrate a threshold which separates exclusion which is (a.i) pragmatically necessary or (a.ii) optional but unproblematic from exclusion which is (b.) discriminatory. Here, we see the discriminatory impact arise even while considering the technical interface with the user in isolation from compounding social factors. When we consider compounding social factors as well, we can see that discriminatory effects can be created even in the absence of all three of the factors initially identified above.

A classic example of discriminatory design is the Band-Aid. The original and standard color of the Band-Aid seems designed to appear inconspicuous when placed on the skin of some but not all users. It could be argued that this is akin to designing clothes for either male or female gender

### III. NON-AFFORDANCES

presentation: the user demographics provide a forced choice, since the product function will be broken for either light- or for dark-skinned users, depending on whether a dark or light skin tone is used for the product's color. Given the forced choice, it seems then unproblematic to choose to accommodate the majority rather than a minority of users. However, the fact that all brands of commercially available adhesive bandages made this same individual determination to cater to the majority creates an emergent problem: access to this function is denied upon racial lines.

This itself is not a major concern. The primary function of the adhesive bandage is not related to its color, and remains accessible for all users, and the disadvantage of diminished skin color similarity presents little substantial impact on anyone's life, considered in isolation. Moments in our lives are not, however, experienced in isolation, and neither can they be properly understood in isolation. The widespread and historical prejudices against those with darker skin provide a context which gives this denial of functionality greater weight, as does the specific practice—now thankfully in the past—of calling a specific white-normative pinkish-grey color “flesh color,” making dark-skinned persons unable to name the color of their own flesh as flesh. The exclusionary effect of the adhesive bandage not only appears within these contexts, but creates a further context: it is yet another small moment in a system of separate and unequal accommodation and access. We can see the impact of social and historical context by imagining that Band-Aid had chosen an umber color instead—this design choice may have seemed curious, but would not have been grounds for any serious complaint of discrimination. I suspect that it may not have occurred to many white users of the umber bandage that it had been intended to be flesh-colored at all.

In the initial case of the job applicant, we saw (b.i) a direct substantial discriminatory effect. In this second case, we might say that the substantial discriminatory effect is (b.ii) indirect and emergent; it only becomes substantial when we consider the design choices throughout the population of service providers, and within the social and historical context of user communities.

These considerations do not give us any clear list of necessary and sufficient conditions for determining which forms of user exclusion are properly regarded as discriminatory, but instead are intended only to give us a sense of scope and scale. Any abstract theory might be able to establish limit cases, but is very likely, by its nature, to be unable to address all the factors which would go into determining whether a given questionable case is unproblematically exclusionary or problematically discriminatory. It is sufficient to our purposes if this section has established that thresholds exist past which exclusionary design choices are no longer necessary or acceptable but problematic and discriminatory, and that these thresholds vary according to the social context of user groups. This allows us to proceed to a consideration of non-affordances which produce exclusionary effects, and to do so with special attention to those such effects which are exclusionary to groups subject to normative exclusion on the basis of e.g. race, gender, sexual orientation, religion, or disability.

The language of “affordances” has become widespread in engineering as well as interdisciplinary fields including Design Studies, Human-Computer Interaction (HCI), Science and Technology Studies (STS), and Internet Studies, but there has been little discussion of non-affordances. Gibson [1] is the originator of the conceptual structure of “affordances,” and Norman is its most influential popularizer [2][3], but neither discusses non-affordances in a systematic or theorized way. Among the great many authors working with this conceptual structure, only a very few make use of the terms “non-affordance” or “disaffordance,” and none provides either a full definition or a discussion of the scope and variety of cases falling under either term. Gee mentions “disaffordances” in conjunction with “affordance” without providing a definition for the term or any discussion of the concept in isolation, but usefully pointing out that “affordances and disaffordances do not reside in the world alone, but in the combination of the specific mind/body [the player-character] brings to the that world and the way in which that world encourages or discourages that specific mind/body in terms of possible actions” [4][5]. Marcus introduces the term on his own, but credits Gibson [1] for tacitly putting forth the idea, and gives us a bit more of a definition: “As Gibson notes ([1], p. 37), this can also concern the creation of obstacles in the environment, a form of ‘disaffordance’, to protect from or exclude other species or members of the own species, why we here also can sense a potential political dimension to the concept” [6].

No doubt it will be objected that, while the discussion of non-affordances and disaffordances in the scholarly record may be vanishingly infrequent, there is an obvious and implicit understanding of the idea. For example, consider the entire first chapter of Norman's most famous book [2], or even just Carelman's “Coffeepot for Masochists”—the coffeepot with its spout pointing backwards toward its handle that Norman discusses in the first pages of the book, and which even appears on the cover of my edition. Norman's discussion, however—both in this first chapter and elsewhere—looks only at failed affordances in order to inform his analysis of successful affordances.

A merely negative account of non-affordances cannot distinguish between kinds of failures—those which for example arise between the product and its function versus those which arise between the user and the product—and a positive account of various disaffordances is needed to fully address the challenges of design, and, in particular, to address our topics in this paper: when design problems become discriminatory, and how to avoid discriminatory design.

### IV. A TYPOLOGY OF NON-AFFORDANCES

Thus far, following the two scholars I have found to have previously used the term, I have used the term “disaffordance” for the generic lack of affordances in design. We will now give this term a specific definition within a typology of non-

affordances which is targeted toward gaining an understanding of how some forms of user exclusion can be discriminatory.

Still & Dark [7] provide a very good opening to this discussion. They discuss Gaver's [8] example of the cat door, and note that, although we do not perceive it as an affordance insofar as we are under no illusion that it would be a viable pathway whereby we could go outside, "the affordance of passage to another user (the cat) [is] recognized" [7]. They go on to ask, "Is this an example of a perceived affordance? Although an interesting philosophical discussion could be generated in response to this question, our answer in this context is no. In the design context, the action must apply to the perceiver" [7]. It is this philosophical conversation which I seek to generate, not to disagree with Still & Dark, but to expand upon and draw distinctions within their "no," and to make clear that this and other non-affordances are of interest in a design context.

To make sense of how the cat door can be perceived as *affordancing* without being perceived as an affordance, we need an idea of the normative construction of the user through design. We see the cat door as a door either because we recognize conventions that map design onto function [3] or through imagining ourselves as differently embodied, just as we might look at a scale model of a building and imagine it inhabited by ants rather than humans [9]. The functions are recognizable, and we can readily imagine from their form the characteristics of the user necessary to encounter those functions as *affordancing*, even though the tiny door affords us no passage. This can occur when design normatively implies afforded users from which we are physically excluded, like the cat door, or when we are excluded through other constraints, for example, social, cultural, or cognitive. Still & Dark [7] provide another nice example of this latter variety: the child's understanding of chopsticks' intended affordance of eating, even when she has not learned to use them. Here, the child sees the chopsticks as providing an in-principle affordance; the design does not present an affordance to her, but it does present an affordance to her imagined potential self.

Design presents functions relative to normatively implied users, and our relation to the normatively implied user varies in consequential ways. In the cat door and chopsticks cases, the normatively implied user is unlike us in unproblematic ways, for there are either alternate objects oriented to us that provide us equivalent affordances (the human door), or we have clear pathways to changing design intention into affordance which are open more-or-less equally to all users like us (the chopsticks to a skilled user). In the problematic cases with which we are here concerned, the normatively implied user constructed through design is exclusionary of some potential users who are not otherwise appropriately accommodated and who are not different from accommodated users in relevant ways.

#### A. Non-Affordances and Poor Affordances

Let "non-affordance" refer to a general lack of an affordance in question, where we understand "affordances" as the perceptual presence of potential actions as such to an actor by virtue of a given material context of practical action. Non-

affordances, then, can include lacks of affordance due to user interface connection failures, e.g. using aural cues for deaf users. This is importantly different from design failures—in non-affordance, the object either does not appear in the user's phenomenological horizon at all, or it does not appear as equipment [10] which carries with it a set of possible actions and functions related to our projected action(s). But non-affordances are not necessarily discriminatory, and the term applies equally to the non-affordance of aural cues to the deaf person's receiving instructions and windows' non-affordance of infrared vision, or to a rock's non-affordance of freshly-brewed coffee.

Let us then use the term "poor affordance" to refer to what we most often call "bad design." In this case, the affordance *does* appear within the horizon of the user within the actor-network [11]—in fact, it is only by virtue of this appearance and with reference to it that the design of the object can be poor! The poorness of the design consists in the presence of the affordance, without, however, a clear and unobstructed interface between its user or its object (or, in Floridi's useful terminology, its "prompter" [12]). Here, too, poor affordances are not necessarily discriminatory, and this is where Carelman's ponderously designed coffeepot belongs in our typology. As related to the previous example of aural cues for deaf users, we might consider here written prompts for pre-literate children as a relevant variation: written words are present to toddlers, and are understood to be in-principle communicative and informational in function, even though these users may not personally be able to make them work.

#### B. Discriminatory Non-Affordances and Poor Affordances

Both non-affordances and poor affordances may be discriminatory according to the understanding of 'discriminatory design' previously outlined. If, for example, a website contains vital information in an image with no alt-text, this is discriminatory to those with visual impairments, just like having a sign containing safety warnings with no braille translation or non-linguistic signalling of the danger—in either case, the discriminatory effect arises from a non-affordance; the entire non-appearance of the technical object in the user horizon. The adhesive bandage provides an example of discriminatory design arising from poor affordances: the "flesh coloredness" of the bandage appears as an intended affordance<sup>1</sup>

---

<sup>1</sup> I use the circumlocution 'appears as an intended affordance' in order to avoid a contentious nest of issues which we will not be able to untangle in the space of this article, but which should be at least acknowledged in passing. The intentions of the designer to create this or that affordance are not necessarily relevant to the creation of discriminatory effects—it is perceived and effective affordances that are consequential, and these are subject to multistability [13] and interpretive flexibility [14] within a cultural horizon [15]. Material implications cannot be ignored, though, even on a constructivist account of technical function, and I think a notion of a technical proper function—adapted from Millikan's notion of a biological proper function [16]—is needed to give a good account of the social-material interplay which gives rise to correctly perceived intended affordance. The account would approximately be that a technical proper function is an affordance the existence of which materially brings about the adoption, use, and replication of the technical object. This account would have the additional advantage of locating intended uses within the material conditions and genetic history of the technical object rather than referring to a designer's mental states, allowing for circumstances where functional design may be replicated by

even though companies now know that they shouldn't use the term "flesh color" to refer to white people's average skin tone, and it is the phenomenal presence of this intended function, in conjunction with its demonstrable failure to interface with darker-skinned users, which constructs the user's color as non-normative.

Floridi's discussion of the 'Janus-headedness' of the tool [12] also helps to identify a distinction between discriminatory design and unproblematic non-affordances and poor affordances. The god Janus has two faces looking in different directions, and Floridi compares the tool to the Janus head: one (inter)face looks to the user; the other (inter)face looks to the prompter—i.e. the element of the world which the tool is responding to or modifying. It is the interface with the user whose gaze can become more easily and obviously discriminatory. If the bandage does not fit all or most wounds, we say the design is poor, or that the design problem is difficult (e.g. bandages interfacing with prompting wounds appearing on knuckles and knees). The "flesh-coloredness" of the bandage, however, gazes upon the user and recognizes whites only; its affordances are designed to interface with users, but in such a way that privileges one racial grouping at the exclusion of others.

### C. *Disaffordances and Dysaffordances*

The way that discriminatory design seems to arise from user-(inter)facing poor affordances and non-affordances indicates that we can gain additional clarity by looking at how the object's design fails to recognize its user. This distinction is not of a kind with the distinction between non-affordance and poor affordance: in that distinction, we looked at how the object appears within the user's horizon—here, by contrast, we look at how the user appears within the object's horizon.

Let "disaffordance" refer to an object which fails to recognize relevant aspects of relevant users, resulting in either non-affordances or poor affordances. We may think of the stony gaze which the staircase directs toward the wheelchair user, but without seeing her, or seeing her as unworthy of responding to and interfacing with. We could also think of how kitchen counters are designed for women's average height, and baby strollers are made so that a tall man (or an exceptionally tall woman) must lean down, and must take care not to take too long of a stride, since the brakes are often placed on the back of the back wheels. These and other forms of female-oriented design make many male homekeepers subtly but consistently uncomfortable and frustrated. We can also think of VR systems

---

producers who do not fully understand the utilities and adaptations of traditional designs which are reflexively and uncritically replicated.

The need for such an account is not "merely academic" in the colloquial sense. Without a robust account of what affordances can be said to be rightly and objectively perceived as intended, it is unclear how we can justifiably distinguish between a poorly designed affordance and someone just using something wrong. Hand-waving toward 'perceived affordances' may be fine for our purposes here, but it does not properly establish the objective basis for these judgments. If we are to give a complete account of how some but not all logically possible affordances can be objectively identified as relevant when not provided, a desk in a public office which cannot accommodate wheelchair users must be describable as broken according to reasoning which is simultaneously able to explain why the desk's failure to accommodate napping hyenas or naval military offensives does *not* make it broken.

which produce perceived three-dimensionality through parallax-motion systems rather than shape-from-shading: the former system works well for most men but produces nausea in many women, whose brains are more likely than men's to use the latter pattern to construct perceived three-dimensionality [17].

Disaffordances, in this novel and specific definition, can follow from *poor* affordances with discriminatory effects, as in the examples in the previous paragraph, but we have also already seen examples of disaffordances which follow from *non*-affordances: for example, the sign without braille or the image without alt-text.

Using a related but distinct prefix to name a related but distinct phenomenon, let "dysaffordance" refer to an object which not only fails to recognize relevant aspects of relevant users, but which also requires users to misidentify themselves in order to gain access to an object's functions or products. The Latin prefix "dis" usually names a lack or a separation; by contrast, the Greek prefix "dys" usually names a malfunction or problem—here, "disaffordance" refers to a design-based separation from an object's functionality based on user identity, while "dysaffordance" refers to a design-based requirement of certain users to misidentify themselves in order to gain access to an object's functionality.

Examples of dysaffordances abound, but seem necessarily to fall entirely within the realm of poor affordances rather than non-affordances. We may think of binarisms in user data entry fields, requiring gender non-binary persons to choose between male and female gender identities, and requiring bisexual dating site users to (mis)identify themselves as either gay or straight. We may think of SNS real-name policies, which often require users to identify themselves in accordance with governmental documentation rather than in accordance with their lived identity—a particularly consequential misidentification for transgender persons.

## V. PRINCIPLES OF ANTI-DISCRIMINATORY DESIGN

The various kinds of discriminatory design problems identified above can give us a sense of the scope and variety of issues that the anti-discriminatory designer must guard against. Consideration of this scope and variety will demonstrate the necessity of taking a proactively anti-discriminatory approach and the insufficiency of merely not being actively discriminatory to the prevention of discriminatory effects. A methodology for anti-discriminatory design can then be recommended and described: producing discrimination impact assessments through phenomenological variation oriented by standpoint epistemology.

### A. *The Scope and Variety of Discriminatory Design Problems*

In section II, we identified (b.i) direct and (b.ii) indirect discriminatory effects, where in the former the discriminatory effect arises from user exclusion from object functionality, and in the latter the discriminatory effect arises from differential user access to functionalities rather than outright exclusion. In section IV.A. & B., we distinguished non-affordances from poor affordances, where in the former some affordances do not appear within some relevant users' experience *at all*, and in the

latter *intended* affordances do not provide *actual* affordances. In section IV.C., we distinguished discriminatory disaffordances from discriminatory dysaffordances, where in the former the object fails to recognize relevant user identities in a way that gives rise to direct (b.i) or indirect (b.ii) discriminatory effects through either poor affordances or through non-affordances, and in the latter the object both fails to recognize some relevant user identities and also requires some relevant users to actively misidentify themselves in order to gain access to intended affordances.

In these overlapping distinctions we see three primary loci of design problems: 1. The user's phenomenal experience, in which the object may be either present in a more-or-less obstructive form, or simply absent; 2. The object's implicit construction of a user in its interface design, in which some users may not be recognized or may be recognized poorly; and 3. The social context in which the user-object connection, considered in either directionality (1. or 2.), takes place.

### B. The Need for Anti-Discriminatory Design

The identification of these loci demonstrates the likely insufficiency of a mere lack of discriminatory intent to prevent discriminatory design. In (1.) and (2.), we can easily anticipate how difficult it is, at each stage of the design process, to imagine how the user interface will appear to all kinds of users—persons of all races, language abilities, gender identities, sexual orientations, religions, disabilities, ages, body types, etc.—and how users among those groups and within the intersectionalities that they present will experience the interface. In (3.), we can see, further, how, even were this kaleidoscopic imagining of all relevant users' experience of and through the interface successful, it would still be insufficient to anticipate discriminatory effects which follow from social histories and prior lived experience actuated by but not contained within the interface. To use a blatant example: thinking about what it's like to be black can tell you that your photobooth application ought to include brightness correction to that a dark skin tone can be treated as the subject of the portrait rather than being treated as "shadow," balancing exposure and contrast so that black people do not appear as silhouettes. But, valuable as that consideration of how the interface interfaces with users is, it won't tell you why it would upset some users if your app prompted the user to smile by saying "Say fried chicken and watermelon!"—for this, we need a history of racialized experience in the larger context in which this interface appears.

The overwhelming diversity of user experiences, along with the overwhelming difficulty of anticipating and accounting for these user experiences, impresses upon us the insufficiency of simply following established practices for non-discriminatory design, such as ADA/508 compliance, to prevent discriminatory design from occurring. Hence, we must not be merely non-discriminatory designers, but actively anti-discriminatory designers. Even so, avoiding discriminatory effects entirely during the design process seems impossible, and being responsive to user complaints and requests in post-release iterative redesign certainly seems to be a necessary remediation. But still, we can become better at anticipating discriminatory effects, and we have an obligation to do so even

if perfect anticipation of and design for all relevant user experiences is impossible.

### C. Methodology for Anti-Discriminatory Design

The importance of diversity in developer teams cannot be ignored. Those who are marginalized or excluded by systems of power and interpretation are better able to see those systems at work than those who are privileged by them—privilege, indeed, consists largely in the freedom from being required to notice these systems at all. This point from feminist standpoint epistemology—that marginalization produces locations of epistemic advantage—is obviously relevant in a design context. It seems common-sensical, for example, to assume that very few, if any, black developers were involved in programming a facial recognition system that was bad at recognizing dark-skinned users [18], or an automated image tagging system that identified black photographic subjects as "gorillas" [19].

Diverse developer teams, however helpful, cannot present a complete solution or a foolproof fix, for several reasons. First, even those in locations of epistemic advantage may fail to notice their own marginalization, since marginalizing interpretations tend to be internalized by those marginalized by them, this being, in fact, one of the primary pathways by which marginalization occurs. Just because someone is a woman, for example, it cannot be assumed that she will notice all instances of sexism—for example, use of the generic "he" in documentation discussing users of technical or professional software—or that she will feel her voice legitimated enough within the design context to express and insist upon addressing these issues when they *are* noticed. Second, even if we could count on all marginalized persons to be fully aware of their own marginalization, and fully politicized such that they would always insist on addressing such issues, no developer team could be expected to be fully representative of the diversity of identity groups and intersectionalities within the relevant user base. Third, even if a developer team somehow were absolutely diverse, absolutely self-aware, and absolutely politicized, such a strategy for anti-discriminatory design wrongly implies that the socially privileged designer has neither ability nor responsibility to engage in anti-discriminatory design. In other words, a developer working on her own should not excuse herself from the responsibility to attempt to think through diverse user experiences, and "Yeah, well I'm not Asian" would not be a satisfactory apology for e.g. designing image recognition software that causes cameras to insist on repeatedly asking many East Asian users "Did someone blink?" [20].

Aside from the obviously valuable but no-silver-bullet method of increasing diversity in developer teams, anti-discriminatory design methods can be identified at micro- and macro-levels, in order to address the both the user-object interface (V.1. & 2. above) and the larger historical and social context and impact (V.3. above).

At the micro-level, phenomenological variation, as described by Ihde [13] will help in proactively anti-discriminatory design. In the process of phenomenological variation, we begin with a given human-technics relation and successively alter different aspects of the user and of the device

in order to identify and articulate where and when alterations of the human or of the device alter the human-technics relation and its functionality. This cannot be reduced to ‘imagining what it would be like to be a user, except [gay or latino or deaf or etc. etc.]’. When we simply imagine being this or that identity group we import our own sense of our own normality, producing a bias against alternate user experiences, and we are not likely to think through the concrete details of different social and physical embodiments. The systematic control/experiment group process of phenomenological variation has a much better chance of noticing e.g. that using contrast patterns in facial mapping may require different programming for different skin tones, or that e.g. hosting professional activities on Second Life may create troubling dynamics given the relative unavailability of office-appropriate clothing for female avatars in that virtual environment.

At the macro-level, Bush [21] has suggested conducting gender impact assessment reports, akin to the environmental impact assessment reports to which businesses have (sometimes reluctantly) become accustomed. Here, I will suggest a more general diversity impact assessment report, to include the range of relevant diverse user groups rather than gender alone. A population-wide and longitudinal impact analysis of this kind is needed to identify emergent discriminatory effects which may not be phenomenally present in the interface itself in an isolated user interaction and, thus, which phenomenological variation is unlikely to highlight. We can find a fine example in the classic case of the FBI “Taco Circuit” [22]. The FBI recognized and valued the Spanish-language skills that Hispanic agents often brought to the agency, and assigned them tasks where those skills were recognized and put to use. So far, so good! But as a result, Hispanic agents were more often placed in secondary and support roles, where they translated, interpreted, and transcribed for their white colleagues who, because they had lesser language skills on average, were placed in primary and leadership roles where they were more able to gain recognition and promotion. We can also consider another related but very everyday example: calendar design. Scheduling all office holidays and team deadlines so that they match with Christian faith observance reduces opportunities from observant members of non-Christian faiths, since their observance, unlike that of their Christian coworkers, requires them to miss work, to miss opportunities to play more central roles, and possibly to be perceived as unreliable.

These techniques, though, are no substitute for robust engagement with marginalized users and user communities, and, although worth pursuing, are strongly limited by the difficulty of anticipating and understanding the lived experiences of others. When we try to imagine the experience of others, we tend to imagine someone else as ‘ourself + [variation]’. The constancy of what we perceive as the ‘core’ or ‘true self’ fails to consider and account for the determination of values and goals related to varying life experiences in ways we fail to appreciate. These limitations are displayed clearly in the discredited [23] practice of “disability simulation,” where, for example, a nondisabled person is asked to navigate an environment in a wheelchair in order, supposedly, to gain a better understanding of the experiences of disabled persons.

These “simulations” produce an unrealistic understanding of the life experience of disability for a number of reasons: the nondisabled person does not have the alternate skill sets developed by disabled persons, and thus overestimates the loss of function which disability presents, and is furthermore likely to think of able-normative solutions rather than solutions more attuned to a disabled person’s life experience. Similarly, even when they proactively try to imagine and accommodate architectural engagements of disabled persons, nondisabled designers tend to enforce a bipedal-normative ablist agenda in accommodation design, while by contrast, as Shew [24] points out, conversations with wheelchair users may reveal less interest in finding technologies which allow bipedal navigation of able-normative architectures and more interest in altering architectures to accommodate non-normative embodiments which better represent disabled persons’ preferred modes of mobility and interaction.

Phenomenological variation and longitudinal impact assessments can help to identify discriminatory design, but the value of these practices is not only limited by, but is even potentially significantly undermined by overestimation our ability to understand and anticipate the life experiences of others.

To conclude: Best practices for implementing anti-discriminatory design methodology—discussed here as phenomenological variation and diversity impact assessment as oriented by standpoint epistemology—will necessarily vary widely, depending on what is being developed, and which groups are most relevant and most at risk. I hope this discussion is sufficient, though, to establish these guiding principles of anti-discriminatory design:

- Seek out diversity in developer teams, and give weight and respect to marginalized perspectives *especially* when they do not accord with normative (e.g. white/male/straight/Christian/nondisabled/etc.) perceptions.
- Engage in systematic phenomenological variation to identify problem areas in the user interface. Think through differences relevant to social and physical embodiments in areas of race, disability, gender, religion, sexual orientation, transgender status, and other areas as prompted by intended affordances.
- Conduct population-wide and longitudinal diversity impact assessments of potentially marginalized groups, especially emphasizing groups identified as relevant through phenomenological variation.

And, given the particular histories of many of the cases brought up in the course of this discussion, we should also add:

- Test with diverse users and be clear that you value and are seeking out feedback related to their particular experiences from their particular social and physical embodiments.
- Keep in mind that even all this will not catch every problem. When something goes wrong, admit that it’s a real problem, that the user is not wrong or unimportant, and try to fix it.

What's most distressing about some of these cases is not the discriminatory impact, or the designers' failure to think through user diversity, but instead the dismissive responses to user complaints which explicitly restate what discriminatory design says implicitly: You don't count.

#### REFERENCES

- [1] Gibson, J. (2014). *The ecological approach to visual perception: Classic edition*. New York: Psychology Press.
- [2] Norman, D. (2013). *The design of everyday things: Revised and expanded edition*. New York: Basic Books.
- [3] Norman, D. (1999). Affordance, conventions, and design. *Interactions*, 6(3), 38-43.
- [4] Gee, J. (2007). Pleasure, learning, video games, and life: The projective stance. In M. Knobel & C. Lankshear (Eds.), *A new literacies sampler*. New York: Peter Lang.
- [5] Gee, J. (2008). Video games and embodiment. *Games and Culture*, 3(3-4), 253-263.
- [6] Marcus, L. (2015, 13-17 July). Ecological space and cognitive geometry. In K. Karimi, L. Vaughan, K. Sailer, G. Palaiologou, & T. Bolton, *Proceedings of the 10<sup>th</sup> International Space Syntax Symposium*. London: Space Syntax Laboratory.
- [7] Still, J.D. & Dark, V.J. (2012). Cognitively describing and designing affordances. *Design Studies*, 34(3), 285-301.
- [8] Gaver, W.W. (1991). Technology affordances. In *Computer-human interaction conference proceedings* (pp. 79-83).
- [9] Stiller, B., Rudin, S., & Cornfeld, S. (Producers), & Stiller, B. (Director). (2001). *Zoolander* [Motion picture]. USA: Paramount.
- [10] Heidegger, M. (1996). *Being and time*. (J. Stambaugh, trans.) Albany: SUNY Press. (Originally published 1927.)
- [11] Latour, B. (1999). A collective of humans and nonhumans: Following Daedalus's labyrinth. In B. Latour (Ed.), *Pandora's hope: Essays on the reality of science studies*. Cambridge: Harvard University Press.
- [12] Floridi, L. (2014). *The fourth revolution*. Oxford: Oxford University Press.
- [13] Ihde, D. (1990). *Technology and the lifeworld: From garden to Earth*. Bloomington: Indiana University Press.
- [14] Pinch, T. & Bijker, W. (1984). The social construction of facts and artefacts: or How the sociology of science and the sociology of technology might benefit each other. *Social studies of science*, 14(3).
- [15] Feenberg, A. (1991). *Critical theory of technology*. New York: Oxford University Press.
- [16] Millikan, R. (1987). *Language, thought, and other biological categories*. Cambridge, MA: MIT Press.
- [17] boyd, d. (2014, Mar. 28). Is the Oculus Rift sexist? *Quartz*. Retrieved from <http://qz.com/192874/is-the-oculus-rift-designed-to-be-sexist/>
- [18] *HP computers are racist* (2009), video, wzamen01, 10 Dec., viewed 10 August 2015, <https://www.youtube.com/watch?v=t4DT3tQqgRM>.
- [19] Mullen, J. (2015, July 2). Google rushes to fix software that tagged photo with racial slur. *CNN*. Retrieved from <http://www.cnn.com/2015/07/02/tech/google-image-recognition-gorillas-tag/>
- [20] Rose, A. (2010, Jan. 22). Are face-detection cameras racist? *TIME*. Retrieved from <http://content.time.com/time/business/article/0,8599,1954643,00.html>
- [21] Bush, C. (1983). Women and the assessment of technology. In: J. Rothschild (Ed.), *Machina ex dea*. New York: Teacher's College Press.
- [22] Kennedy, J. & Overend, W. (1988, Oct. 1). FBI discriminated against Latino agents, judge rules. *Los Angeles Times*. Retrieved from [http://articles.latimes.com/1988-10-01/news/mn-4259\\_1\\_fbi-agent](http://articles.latimes.com/1988-10-01/news/mn-4259_1_fbi-agent)
- [23] Flower, A., Burns, M.K., & Bottsford-Miller, N.A. (2007). Meta-analysis of disability simulation research. *Remedial and Special Education*, 28(2), 72-79.
- [24] Shew, A. (2016, May). *UpStanding norms*. Paper presented at IEEE International Symposium on Ethics in Engineering, Science, and Technology, Vancouver, Canada.