

WHEN IS TECHNOLOGY GOOD AND WHY?

Human capacities and predicative goodness in design ethics

Ryan Wittingslow  ryan.wittingslow@anu.edu.au

HASS Digital Research Hub, The Australian National University, Australia

University College Groningen, University of Groningen, The Netherlands

Avigail Ferdman  avigailf@technion.ac.il

Department of Humanities and Arts, Technion-Israel Institute of Technology, Israel

Article type: Research article

Review process: Double-blind peer review

Topical collection: Ethics In/Of/For Design, guest editors: Michael Nagenborg, Deger Özkaramanli

This open-access article is published with a Creative Commons CC-BY 4.0  license

<https://creativecommons.org/licenses/by/4.0/>

DOI: <https://doi.org/10.59490/jhtr.2026.4.8006>

ISSN: 2773-2266

Submitted: 29 January 2025 **Revised:** 04 September 2025, 01 December 2025, 10 February 2026

Accepted: 17 February 2026 **Published:** 13 October 2026

How to cite (APA): Wittingslow, R., & Ferdman, A. (2026). When is technology good and why? Human capacities and predictive goodness in design ethics. *Journal of Human-Technology Relations*, 4, 1-17.

<https://doi.org/10.59490/jhtr.2026.4.8006>

Corresponding author: Ryan Wittingslow

©2026 Authors, published by TU Delft OPEN Publishing on behalf of the authors.

Keywords

Design ethics; Capacities Approach; Predicative Goodness; Attributive Goodness

Abstract

In this paper we examine the limitations of evaluating designed objects solely through 'attributive goodness' and argue for incorporating 'predicative goodness': that is, whether designs contribute to human flourishing. Building on Avigail Ferdman's take on the 'capacities approach' (2024), which links technological affordances to the development of fundamental human capacities, we propose extending her work to assess the predicative goodness of designed things. By analysing how designs support or hinder core human capacities, we can develop more ethical design practices. We outline a framework for applying the capacities approach across design domains and explore its implications for design ethics, arguing that this approach is useful as design is increasingly applied to societal challenges.

Plain Language Summary¹

- This study is relevant for designers, ethicists, and policymakers who judge designs mainly by how well they work. It argues that "working well" is not the same as being good for people, a distinction that matters as design increasingly tackles social problems.
- The study distinguishes designs that are good at their function (like a surveillance system that monitors effectively) from designs that are good for human flourishing, showing that functional success can still cause real harm to well-being and autonomy.
- Building on philosopher Avigail Ferdman's "capacities approach," the article shows how the possibilities for action built into objects, spaces, and technologies shape whether people can develop capacities such as knowledge, creativity, social connection, and independent decision-making.
- The article introduces "privative unflourishing"—three ways design can restrict or gradually weaken human capacities. Using comparisons like fitness trackers versus basketball courts, it shows how designs that succeed at their stated purpose can still differ in whether they help people grow or narrow their development.
- The framework provides a practical basis for comparing design options during early design, helping designers and policymakers weigh whether a design supports multiple human capacities together, spot trade-offs early, and share responsibility for its longer-term effects.

¹ AI-generated; author checked and approved.

1 INTRODUCTION

When we evaluate the ‘goodness’ of designed objects, we typically ask whether they are ‘good’ at what they’re supposed to do. A good hammer drives nails efficiently, for example; so too does a good chair provide comfortable seating. Making these kinds of judgments requires two things.

First, we need to possess some way of grounding these normative standards. The conventional strategy is to adopt some ‘proper’ account of function. According to proper accounts of function, among all of the functions *simpliciter* that an artefact might possess, there exists a subset of functions that are ‘proper’ to those artefacts. The existence of proper functions makes normative judgements of design goodness possible. Second, we need a way of assessing whether or not artefacts satisfy their proper function. Determining whether an artefact fulfils its proper function involves assessing what philosophers call the ‘attributive goodness’ of that artefact: whether something has “the property of being a good K, for some kind K” (per Hanson, 2017, p. 416). So, when we say that some particular hammer is good at driving nails, we are saying a) that it has the proper function of driving nails, and b) that it has the property of being good at driving nails.

This focus on attributive goodness leaves a gap in the evaluation of design because ‘attributive goodness’ refers to assessment relative to kind membership: whether something is good as the kind of thing it is. So, when we say something is attributively good, we evaluate it against standards appropriate to its kind: a good hammer, a good liar, or a good knife, for example. ‘Predicative goodness’, by contrast, assesses whether something is good *simpliciter*: whether it contributes to what is fundamentally valuable without reference to its kind membership. The distinction is important, because attributive success does not entail predicative success. Saying of someone that ‘she is a terrible liar’ does not entail ‘she is terrible’. Similarly, we can evaluate whether an artefact is attributively good at fulfilling its function without determining whether it is good in a broader, predicative sense.

Different theories in normative ethics offer competing accounts of what grounds predicative judgements. While philosophers like Geach (1956) reject predicative goodness as incoherent, arguing that all legitimate uses of ‘good’ are attributive, we claim (following contemporary perfectionist ethicists) that predicative judgements can be grounded more robustly. This paper outlines how such grounding might work for designed objects and systems, building on recent work in perfectionist ethics that connects technological affordances to human capacities.

Building on both Ferdman’s (2024) capacities approach and Wittingslow’s (2024) application of her work to certain theories of proper function, we propose a general foundation for evaluating the predicative goodness of designed objects and systems. By connecting technological affordances to the development and exercise of human capacities (capacities such as knowing, creating, socialising, exercising willpower, and so on), Ferdman’s approach enables us to evaluate whether designed objects and systems are predicatively good for human flourishing based upon the affordances they furnish. Meanwhile, Wittingslow demonstrates that Ferdman’s work can enrich accounts of proper function by providing tools for evaluating whether those proper functions themselves contribute to human capacities and flourishing.

While Ferdman’s work deals primarily with digital technologies, we argue that it can and should be extended to design more broadly. The relationship between affordances and human capacities that she identifies applies equally to all designed things. By understanding how designs either support or hinder the development and exercise of human capacities, we can better evaluate the extent to which they are predicatively good.

This paper thus presupposes a perfectionist account of human flourishing. Readers who reject this basic assumption will find this argument unpersuasive; meanwhile, those willing to

entertain it can assess whether introducing predicative goodness into perfectionism provides a coherent way of evaluating the moral goodness of designed things.

The paper is in four parts. First, we examine the limitations of focusing solely on attributive goodness in design evaluation. Second, we detail Ferdman's capacities approach and show how it enables predicative evaluation. Third, we develop a framework for applying this approach to design practice across domains. Finally, we consider the implications of this paper for design practice.

2 THE LIMITS OF ATTRIBUTIVE GOODNESS

When evaluating designed objects, many philosophers of technology endorse what are called 'proper function' theories of function.² These scholars argue that artefacts have functions that are proper to them: functions that normatively determine what those artefacts are for. This stands in contrast to 'system function' accounts, which hold that an artefact's function simply describes what it does within a given context.

Proper function theories provide robust grounds for evaluating attributive goodness: that is, whether something successfully fulfils its role as the kind of thing it is. When we say a book is 'for reading', we are identifying reading as the function inherent to books, among all their possible uses. This proper function then provides criteria for determining if something is a good instance of its kind. A book that facilitates reading well is good in its capacity as a book. Meanwhile, one that fails at this task is a poor book, regardless of how well it might serve in some other role: for example, as a doorstop, a paperweight, or a weapon.

While proper function approaches have been influential in design on the grounds that they provide clear criteria for evaluating design success, design theorists have long recognised the limitations of purely functional evaluations of design success. Participatory design movements emphasise user agency and democratic involvement (Björgvinsson et al., 2012), while critical design practices interrogate the values embedded in artefacts (Dunne & Raby, 2013). Models like Value Sensitive Design integrate human values throughout the design process (Friedman et al., 2006), and Design Justice explicitly addresses power asymmetries in design processes (Costanza-Chock, 2020). These approaches are united in acknowledging that successfully fulfilling a proper function does not suffice to make a design ethical.³ These approaches, useful as they are, operate at the level of process and method. They excel at identifying stakeholder values and orienting design towards values such as social justice or fairness (Donia & Shaw, 2021) but don't provide substantive accounts of what makes designs good *simpliciter*. This gap becomes particularly obvious when stakeholders express preferences for designs that may undermine their own long-term wellbeing: a phenomenon well-documented in attention economy research (cf. Williams, 2018). A notable exception is the Capability-Sensitive Design approach, which we discuss later.

² Proper function theories are an extremely broad church. While they all endorse the notion of proper functions, they differ regarding the source of normative authority for artifact functions. Selected effects accounts (cf. Millikan, 1984, 1986, 1995; Neander, 1991; Parsons & Carlson, 2008) ground proper functions in historical selection processes, whether biological or cultural, such that an artifact's proper function derives from what past tokens were selected for doing. Intentionalist theories (cf. Vermaas & Houkes 2003, 2006; McLaughlin, 2001) locate proper functions in designer or user intentions, making function attribution dependent on mental states at design or appropriation moments. Social practice approaches (cf. Preston, 1998, 2013; Thomasson, 2014; Wittingslow, 2025) situate proper functions within ongoing community practices, arguing that artifacts acquire functions through processes of social ratification.

³ Here we use 'function' to refer to proper functions rather than mere systemic or descriptive functions. Throughout this paper, unless otherwise specified 'function' refers to proper function.

The limitation points to a deeper issue: whether we can meaningfully evaluate designs beyond their functional success. While Geach (1956) famously argues that predicative goodness does not exist—that is, that all legitimate uses of ‘good’ are attributive—we argue against this view. Following contemporary perfectionist ethicists like Bradford (2015, 2020) and Hurka (2002), we argue that predicative judgements of goodness can be grounded in objective facts about human flourishing and the development of fundamental human capacities. When we say something is ‘good’ in a predicative sense, we are claiming that it contributes to human flourishing by supporting these essential capacities. This allows us to evaluate designs in terms of their broader contribution to human development and wellbeing.

Geach (1956) argues that ‘good’ functions only as an attributive adjective, like ‘large’ or ‘small’. We speak of good knives or large mice, but not of things being simply ‘good’ or ‘large’ without qualification. For Geach, apparent predicative uses either elliptically invoke some sortal concept or merely express approval without truth conditions. Contemporary perfectionist ethicists provide grounds for rejecting this. Bradford (2015, 2020) demonstrates that achievements contribute to wellbeing not through fulfilling some ‘achievement function’ but by actualising human capacities in ways that constitute flourishing. Hurka (2002) grounds predicative goodness in developing our essential human capacities: capacities like rationality, autonomy, sociality, and so on. According to these scholars, predicative goodness tracks contribution to human flourishing through capacity development and exercise.

The attributive/predicative distinction captures a distinction that proper function accounts cannot otherwise address. A superlative surveillance system, for example, is an example of attributive success yet it tells us nothing about its impact on human autonomy and dignity. These predicative dimensions cannot be reduced to or explained in functional terms. The distinction becomes even more clear when we consider whether certain artefacts should exist at all. A Tomahawk missile might excel at fulfilling its proper function (that is, it might excel at blowing things up) while contributing negatively to human wellbeing. Proper function accounts can evaluate the technical performance of the missile but offer no resources for assessing whether such weapons should be designed in the first place.

These limitations point to a gap in design ethics. As Chan (2017) notes in calling for more developed accounts of design ethics, we need practical and practicable approaches that can address the normative dimensions of design. As design methods are applied to addressing societal challenges, it’s important that we develop theories that can evaluate whether those functions contribute to human flourishing and care. This is where Ferdman’s (2024) capacities approach is useful, in that she provides tools for assessing how technological affordances support or hinder the development and exercise of human capacities such as knowing, creating, socialising, and exercising willpower.

3 FERDMAN’S CAPACITIES APPROACH

Ferdman’s approach to evaluating technology builds on recent work in perfectionist ethics, particularly the ‘capacities approach’ to human flourishing, as distinct from the ‘capabilities approach’ of Martha Nussbaum (2011), Amartya Sen (1999), and others (Curzer & Epstein, 2022; Frischmann & Selinger, 2018; Ratti & Graves, 2021). Perfectionism is a normative account of the good life, whereby humans flourish when they develop a set of essential or natural core capacities, exercised in activities of intrinsic value such as knowledge, friendship, morality and beauty.

3.1 AFFORDANCES TO CAPACITIES

Perfectionist approaches generally converge around the following capacities: rational and deliberative capacities, physical embodiment, the social capacities, capacity for morality, and

willpower (cf. Aristotle, 1984; Bradford, 2015; Hurka, 1993; Kraut, 2007, among many others). The fulfilment of the capacities is achieved, according to Hurka (1993), in leading a well-rounded life, or more generally in the self-realisation of a life of activities that embody rational or deliberative control of thought and action (Brink, 2003; Green, 1883/2003).

Working within Gwen Bradford's 'tripartite framework' (2015; 2016; 2020), Ferdman argues that perfectionist human flourishing requires the successful exercise of fundamental human capacities: knowledge, creativity, friendship, practical wisdom, and so on. Unlike virtue-based theories that focus on the development of moral character, this approach emphasises what we actually do with our capacities. The mere having of creative potential, for instance, contributes nothing to flourishing unless that capacity is appropriately developed and applied. Ferdman's account thus provides clear normative standards, in that it can help us evaluate whether something supports or hinders human flourishing by examining how it affects the development and exercise of capacities.⁴

Ferdman operationalises human flourishing through 'affordances'. Affordances, a concept first developed by psychologist James J. Gibson (1979/2014), and introduced into the design and philosophy of technology literature by design theorist Don Norman (1988/2013), refer to the possibilities for action that emerge from the relationship between an agent and their environment. These are opportunities for action that arise when agent and environment meet. For example, a chair affords sitting to a human of appropriate size but might afford hiding to a child or climbing to a cat. In all cases, the affordance emerges from the relationship between object and user.

While traditional virtue ethics accounts in philosophy of technology treat capacity development as primarily a matter of personal responsibility (cf. Vallor, 2016), Ferdman demonstrates how our environments, through affordances, shape our ability to develop and exercise capacities. An environment rich in appropriate affordances supports capacity development. A well-designed library, for instance, affords multiple opportunities for developing our capacity to know thanks to the amenities on offer. Conversely, an environment that constrains or misdirects our capacities through limited or inappropriate affordances can lead to what she calls 'privation of flourishing'. This privation can manifest in several ways. For example, well-developed capacities can go unexercised, capacities cannot be properly developed due to environmental constraints, or impoverished capacities may lead to negative outcomes. These forms of privation often arise from how our environment structures possibilities for action and development through its affordances.

Technologies play a foundational role in this environmental structuring. As Ferdman argues, they "shape what we are able to do, how we are able to do it and what we see ourselves as being able to do" (2024, p. 12). Connecting affordances and capacities in this way offers new tools for evaluating predicative goodness in design, in that we can assess how the affordances furnished by a design impact users' ability to develop and exercise their fundamental human capacities. We can also readily build proper function into this account. Whilst proper function theories tell us what something is 'for', Ferdman's approach helps us evaluate whether that function contributes to human flourishing. A design that supports capacity development contributes positively to human flourishing, regardless of how well it fulfils its narrow functional role. Conversely, a design that impedes capacity development may be predicatively bad even if it excels at its proper function. This creates a bridge between the analysis of proper function and ethical evaluation: we can assess both how well something fulfils its proper function (that is,

⁴ While capacities are universal, i.e. all humans flourish when they exercise their capacities, perfectionism is pluralist regarding how persons express their capacities: capacities can be actionable through affordances in a variety of ways, depending on one's culture, context, personal background, etc.

attributive goodness) and how that function impacts human capacities (its predicative goodness).

A potential objection arises here: if we evaluate designs as good because they support human flourishing, aren't we simply assessing whether they fulfil a 'moral meta-function'? Under this view, supporting flourishing would itself be a proper function, making our supposedly predicative evaluation merely another form of attributive assessment: evaluating how well something performs its flourishing-support function.

This objection misreads the distinction. Attributive goodness evaluates performance against a specific function: a chair's success at providing seating, a Tomahawk missile's excellence at blowing things up, and so on. These evaluations measure the achievement of determinate ends. Predicative goodness, by contrast, assesses the degree to which artefacts contribute to human flourishing. The difference becomes clear through example. A surveillance system that perfectly fulfils the job for which it was designed while degrading privacy and autonomy hasn't failed at some moral meta-function. Instead, its functional success is *itself* what contributes negatively to human flourishing. Of course, surveillance-enabling security might represent a legitimate capacity trade-off where capacity losses from privacy erosion yield greater capacity gains from harm prevention, but this is an empirical question outside the scope of the present argument. In any case, the point stands: predicative evaluation examines how functional performance affects human flourishing, but not whether that has function itself been achieved.

Furthermore, any attempt to treat the support of flourishing as another function would render the concept of proper function itself incoherent. Functions should be understood as specific, *determinate* capacities or roles. Supporting human flourishing, meanwhile, encompasses an open-ended range of activities and outcomes that resist reduction to functional terms. Just as 'being good' cannot coherently be treated as a function that an object might fulfil, so too can't 'supporting flourishing'.

3.2 PRIVATIVE UNFLOURISHING

While Ferdman initially focused on digital technologies, the integration of function and perfectionist ethics applies equally to designed artefacts more generally. A public bench, for instance, has the proper function of providing seating. But through its affordances—viz., its location, design, relationship to surrounding space, and so on—it might support or hinder capacities for social interaction or environmental awareness. A classroom's proper function is to facilitate teaching, but its spatial arrangement and material properties could well create affordances that impact capacities for learning, attention, collaboration, and creative expression.

This broader application becomes particularly important when we consider how designs can lead to 'privative unflourishing'. The term is Ferdman's (2024), and it names not merely the privation or absence of flourishing (which might result from bad luck or circumstance) but the active deprivation of flourishing through environmental constraints on capacity exercise. Privative unflourishing occurs through three mechanisms:

1. Direct impediment: when designs create affordances that immediately restrict capacity exercise;
2. Gradual atrophy: when designs encourage patterns of use that lead to capacity deterioration over time; and
3. Environmental constraint: when designs shape environments in ways that systematically limit capacity development opportunities.

Privative unflourishing helps identify how seemingly benign design choices can have significant negative impacts on human development over time. Instead of waiting to observe negative impacts after implementation, designers can anticipate how their choices about affordances might lead to future capacity degradation. This forward-looking evaluation is important given how technological affordances can become entrenched through use, creating what Ferdman calls ‘practice possibilities’ (2024): patterns of behaviour that can be difficult to alter once established.

This helps resolve apparent tensions in design practice. When faced with competing requirements, designers can evaluate trade-offs in terms of both their functional success and their impact on human capacities. While some constraints on capacity development may prove necessary or even beneficial in specific contexts, the relevant distinction is between capacity degradation and capacity reallocation. For example, the widespread adoption of calculators in schools reallocated students’ cognitive resources from mechanical computation to the higher-order formulation of mathematical problems, but this reallocation did not require the elimination of arithmetic capacity itself: rather, the calculator preserved both options.

The calculator is not an isolated case. The evaluative questions for design in general are a) whether an artefact prevents the exercising of capacities or enables their *reallocation*, and b) whether an artefact contributes to a ‘capacity hostile environment’ (Ferdman, 2025) by creating systemic conditions, via the mechanisms of privative unflourishing, that could inhibit the development and exercise of human capacities. While we cannot predict which capacities future contexts will demand, designs that afford the development of new capacities while maintaining legacy capacities provide much greater adaptive potential should conditions change.

3.3 CAPACITIES AND CAPABILITIES

One design ethics approach comes close to predicative goodness: the Capability-Sensitive Design (CSD) approach (Cenci & Cawthorne, 2020; Jacobs, 2020), grounded in the capability approach (Sen, 1999; Nussbaum, 2000). The capability approach holds that human well-being is understood as opportunity or freedom to achieve good human functionings by possessing central human capabilities (Nussbaum, 2000).⁵ CSD emphasises technology design’s role in enhancing and expanding what people can be and do (what is in the literature called ‘achieved functionings’), operationalised through the central capabilities (Jacobs, 2020). As such, CSD can be viewed as predicative *prima facie*: it evaluates the degree to which technologies succeed in promoting the good, where the good is manifested in all persons being above a minimum threshold of the central capabilities. Nevertheless, a closer look reveals that CSD is more limited as a predicative framework compared to the perfectionist capacities approach.

CSD holds that while technology should be capability-enhancing, it is ultimately up to the user to choose whether or not to actually exercise their capabilities (Jacobs, 2020). This is grounded in the capability approach’s normative commitment to ensure freedoms or valuable options from which one can choose (Sen, 1999; Robeyns, 2005). Consequently, the capability approach, and by extension CSD, are indifferent to whether a person actually exercises their capabilities and realises their potential. As such, CSD can evaluate predicative goodness only to the extent that goodness is defined as freedom of choice in exercising capabilities. The capacities approach, on the other hand, can evaluate predicative goodness more directly, because it assesses how well a technology promotes (or inhibits) the development and exercise of human capacities, thereby promoting (or inhibiting) flourishing.

⁵ Martha Nussbaum lists 10 capabilities as follows: life, bodily health; bodily integrity; senses, imagination and thought; emotions; practical reason; affiliation; connection to other species; play; control over one’s environment.

Key to this argument is that unrealised capabilities can degrade or atrophy (Sypnowich, 2016). As with capacities, maintaining a capability involves exercising it; otherwise it may become too impoverished to be realised. This means that, from a perfectionist point of view at least, CSD's normative commitment to ensure freedom of choice regarding capability exercise might fall short of its aim to ensure a dignified life manifested in achieved functionings, so long as persons freely decide not to realise their capabilities.

However, the problem runs deeper. As Ferdman (2024) shows, the decision whether to exercise a capacity (or capability) is not always a matter of personal choice. On the contrary, the decision whether to exercise a capacity can be the consequence of 'affordance-shrinkage' in the environment. In deleteriously influencing behaviour, affordance-shrinkages reduce the action possibilities for worthwhile activities that constitute worthy functionings, as well as eroding the value of those activities. The capacities approach, which explicitly accounts for affordances, therefore provides a more straightforward approach to evaluating predicative goodness, in that it accounts both for the actual flourishing and for the conditions under which people choose how to flourish.

Finally, a short comment before we move on to the next section. Perfectionism often faces the criticism that steering people towards flourishing constitutes paternalism. What we propose here avoids this objection by evaluating whether designs support or hinder the development of capacities rather than mandating how individuals should exercise whatever capacities they've developed. A musical instrument, for instance, furnishes some of the necessary infrastructure for developing musicality without requiring anyone to become a musician; however, removing access to musical instruments certainly precludes the possibility of developing musical capacities. Design paternalism, at least as we conceive of it, occurs when artefacts either prevent the development of capacities or enforce the specific exercise of developed capacities, not when users could choose to do otherwise.⁶

4 APPLICATIONS FOR DESIGN ETHICS

Assessing predicative goodness requires looking past immediate functional success to developmental impacts that unfold over months and years. For example, a learning platform might successfully deliver content and track student progress, fulfilling its proper function. However, examining this technology through Ferdman's approach invites more demanding questions. How do its affordances affect critical thinking development? Does it encourage genuine curiosity, or mere task completion? The answers to these questions are important because, as we already know, technological affordances can create patterns of behaviour that can be difficult to change once established. Designs generate complex capacity interactions; even relatively straightforward design choices have ripple effects across multiple capacities. The efficiency that makes a communication platform attributively good, for example, might simultaneously impede the development of social and creative capacities that emerge through richer forms of interaction.

Human capacities are interconnected, as Kraut (2007) observes. Walking illustrates this clearly. Whilst primarily a physical capacity, it also involves the cultivation of social capacities in negotiating shared spaces, epistemic capacities in processing environmental information, and creative capacities through the mental states that walking can generate (O'Mara, 2019; Ferdman, 2019b). The interconnectedness between capacities complicates the evaluation of those capacities, but it also reveals important design principles about how predicative goodness emerges.

⁶ For a more extended discussion on how perfectionism addresses the paternalism criticism (see Ferdman (2019a) and Tahzib (2022)).

Consider a Fitbit: a device that is attributively good at increasing your daily step count. It does exactly what fitness tracking devices are properly for, in that it quantifies movement: even where it includes comparative or gamified features, these involve competing against numerical abstractions. It thus encourages focus on a narrow aspect of physical capacity whilst leaving other capacities unexercised. A basketball court, by comparison, also encourages physical activity, but does so by creating a phenomenologically thick infrastructure that requires coordination with other people. As much as playing basketball might be about getting fit, it also involves negotiating shared space, reading other players' movements, understanding your role within a team, and developing collective rhythms of play.

These designs cultivate different modes of capacity development. Playing basketball develops social capacities and practical wisdom through joint action. It demands spontaneous movement and tactical improvisation, and rewards the ability to read the game and respond to opponents. These capacities develop simultaneously and in coordination with each other. The Fitbit, on the other hand, encourages capacities for self-monitoring and numerical self-assessment in isolation from other capacity domains.

We think that predicative goodness of the kind that we're discussing here emerges from supporting the coordinated development of individual capacities. Hurka (1993) argues that flourishing consists in leading a life wherein capacities develop together and reinforce one another. This has practical implications for design decisions. A designer choosing between wearable fitness tracking and community sport infrastructure faces different affordance profiles with distinct capacity development trajectories. The Fitbit succeeds attributively at what fitness trackers are for; the basketball court succeeds predicatively by creating conditions for the cultivation of integrated capacities. Predicative evaluation gives us the vocabulary for articulating this difference and criteria for choosing between these options.

Another example of well-rounded design compared to 'single-target' design can be found in recommender systems. These systems can be understood through two competing metaphors: the 'vending machine' and the 'tool-box'.⁷ Standard recommender systems, including most search engines and content recommendation algorithms, operate like vending machines: users input specific queries or preferences, and the system dispenses closely matching results. This approach optimises for immediate satisfaction at the expense of deeper engagement and discovery. In much the same way that a vending machine presents a fixed menu of clearly labeled options, these systems narrow possibilities to an efficiently curated set of likely matches. Systems that afford serendipity function more like toolboxes: open-ended resources that support exploration and unexpected discoveries. Serendipity, understood as the art of discovering new things through chance encounters with unexpected information (Reviglio, 2019; Danzico, 2010), requires more than mere randomness. A serendipitous recommender system creates an architecture that deliberately introduces elements of surprise and challenge into the user experience, encouraging exploration beyond immediate preferences or familiar patterns. Unlike the vending machine model's focus on efficiency and predictability, serendipitous design embraces 'productive friction' in contexts where capacity development is the goal: moments where users must engage more deeply with the content and their own decision-making processes.⁸

⁷ The analogy is inspired by a distinction in the philosophy of science (see Alexandrova, 2017) but is used here slightly differently.

⁸ Productive friction obviously does not apply universally to all artefacts. Instead, the distinction rests on whether the use-context benefits from capacity development (such as a musical instrument, for example) or requires immediate access (such as the emergency exit on an aeroplane). Toolboxes, like the musical instrument, and vending machines, like the emergency exit, thus serve different capability goals. The evaluative question is whether a given design choice matches its intended capability profile.

5 CHALLENGES AND IMPLICATIONS

Implementing Ferdman's approach raises challenges that range from practical questions about measurement and evaluation to deeper theoretical questions about how we assess a design's contribution to human flourishing.

5.1 EPISTEMIC CHALLENGES

The relationship between proper function and capacity development is difficult to trace across artefact types. While we can assess whether a design fulfils its proper function through relatively straightforward metrics, evaluating its impact on human capacities requires more work. The difficulty lies in capturing how these impacts emerge and evolve over time. These evaluation challenges become more acute when we move beyond digital technologies, where at least some aspects of user behaviour can be tracked. How do we evaluate how a city's street layout longitudinally impacts social capacities? How do we measure whether a classroom's design truly supports the development of creative and intellectual capacities? The long-term, systemic nature of these impacts requires new approaches to evaluation that can capture gradual changes across different contexts and scales while accounting for how capacities develop differently across individuals and contexts.

A second challenge concerns trade-offs and unintended consequences in capacity development. A design choice that supports one capacity might inadvertently hinder others. These interactions create genuine dilemmas: a learning platform that scaffolds novice understanding through structured guidance may at the same time constrain the exploratory capacities that experts need to develop. While Ferdman's approach helps us identify these relationships, weighing different capacity impacts against each other requires a degree of due care. This is particularly true when designs support some capacities while hindering others, raising questions about how to evaluate overall predicative goodness.

Affordances emerge from relationships between designed objects and users, making them context-dependent and difficult to predict. This variability intensifies when making predictive claims about capacity impacts. The epistemic limitation doesn't invalidate the capacities approach, but it does constrain its applicability. Designs can be assessed for the demonstrated impacts they make on capacities and modified accordingly, but obviously perfect prediction will prove impossible. We should note, however, that this is a constraint on any design ethics framework rather than being a weakness particular to what we propose here.

Finally, the cultural specificity of this work constitutes a non-trivial, though non-fatal, limitation. The specific capacities we identify as fundamental to human flourishing emerge from Western philosophical traditions, particularly Aristotelian and perfectionist ethics. Whether these capacities constitute the most appropriate ethical theories across all cultural contexts requires empirical investigation beyond the scope of this paper.

In any case, the basic claim here is that human flourishing involves the development and exercise of capacities. That fact transcends national or cultural boundaries, even if different traditions identify different capacity sets or prioritise them differently. What we outline here thus provides principled grounds for evaluating how designs the development of capacities, even if those specific capacities obviously vary across cultural contexts. Individual variation within cultures, meanwhile, raises an entirely different question. Cultural variation concerns which capacities (say, literacy, craftsmanship, or musicality) belong in a given evaluative context, whereas individual variation concerns whether specific people successfully develop whichever capacities are relevant for that context. This paper addresses the first question (that is, which capacities provide appropriate evaluative criteria for design), not the second question of whether individuals achieve those capacities in practice.

5.2 ASSESSMENT AND RESPONSIBILITY

The capacities approach provides substantive perfectionist grounds for evaluating predicative goodness in design. Designs that systematically degrade human capacities for attention, creativity, autonomous action, and so on, constitute genuine failures of predicative goodness, not merely violations of stakeholder preferences. Perfectionist ethics grounds this evaluation in objective facts about human flourishing through capacity development.

Traditional design metrics lead designers to privilege certain capacities at the expense of other capacities. Ferdman's approach makes these trade-offs explicit: predicative evaluation centres on capacity impacts rather than satisfaction or efficiency alone. However, perfectionist ethics does not provide straightforwardly calculable solutions where capacity trade-offs prove unavoidable. The capacities approach helps instead to identify when such trade-offs exist and what is at stake when we try to resolve them. Applying the framework appropriately also requires deliberation among stakeholders about which capacity impacts matter most in specific contexts. Different stakeholders will legitimately prioritise different capacities depending on their contexts and concerns. The framework enables these actors to make their capacity priorities explicit and comparable, facilitating accountable judgements about objective contributions to human flourishing even where disagreement persists.

Should predicative goodness be measured post-implementation at all? Long-term capacity development proves methodologically difficult to track across contexts and scales. Yet waiting for empirical evidence defeats the practical purpose of the framework we outline here. By the time definitive evidence makes itself known, designs are entrenched and their affordances embedded in practice (Ferdman, 2024). The solution lies in recognising that the framework functions primarily as a tool for comparing design alternatives. Taken this way, questions change from 'does this deployed design support capacity X?' to 'how do competing approaches create different affordance profiles?' The recommender system metaphor is a good example of this: evaluating vending machine versus toolbox approaches early during design processes can reveal capacity trade-offs well before implementation.

This has implications for how we perform design evaluations. Design teams might conduct structured analyses of capacity impacts when comparing alternatives, examining both how each approach creates different affordance profiles and the capacities that are exercised, or stymied, by those affordances. These analyses can contribute to deliberation about trade-offs, grounded in perfectionist criteria about capacity development.

This reframing also transforms how we think about designer responsibility. Traditional accounts of designer responsibility focus on proper functions and foreseeable misuse. ICE-theory ('intentional', 'causal-role', and 'evolutionist'), for instance, locates responsibility in the communication of proper use plans (cf. Vermaas & Houkes, 2006). However, if designs can contribute to privative unflourishing via the gradual degradation of capacities, our assessments of responsibility need to be sensitive to that fact. Making this move also implies that responsibility is distributed across multiple actors. Individual designers shape specific affordances; design teams establish functional parameters; organisations create incentive structures; users exercise agency in engagement choices; policymakers shape regulatory frameworks, and so on. Each contributes to whether designs contribute to capacity development.

All of this suggests that we need a graduated model of designer responsibility. While this obviously requires further development, we can already draw some tentative conclusions. Most immediately, designers bear primary responsibility for capacity impacts that are reasonably foreseeable. When a designer implements autoplay features that hijack attention, for example, or creates interfaces that exploit psychological vulnerabilities, they are directly accountable for the foreseeable degradation of those capacities. Secondary responsibility, meanwhile, concerns

probable interactions with existing systems and practices, and tertiary responsibility encompasses contributions to broader technological trajectories. In short, while designers obviously cannot be held fully accountable for all downstream effects, they do bear some degree of responsibility for considering how their work shapes environments that support or hinder human flourishing. In addition, this graduated model must acknowledge that stakeholders beyond designers share responsibility for capacity assessment.

The framework guides design evaluation by grounding predicative judgements in demonstrable capacity impacts, but applying it requires argument among people who disagree about which capacities matter most. Moral value emerges from specific choices about proper function and affordances: choices that can be assessed against perfectionist criteria even when they cannot be explicitly calculated.

6 CONCLUSION

The framework outlined here offers tools for evaluating predicative goodness in design. Proper function theories provide robust grounds for assessing attributive goodness, but they struggle to tell us whether a design contributes positively to human life. Ferdman's approach bridges this gap by showing how proper function and affordances shape opportunities for capacity development. We have tried to make clear how even seemingly straightforward functional choices can have profound implications for human flourishing through their impact on capacities, how privative unflourishing identifies potential ethical concerns across all domains of design, and how predicative goodness arises from the relationship between affordances and human capacities.

That being said, what remains unsolved is the empirical question how one actually analyses capacity impacts. The structured comparative assessments we discuss in the prior section need testing, and designers and educators who take predicative goodness seriously will need ways for identifying affordance profiles and determining the impacts of those profiles on capacities. Developing those methods is the next step.

Data Access Statement

N/A

Contributor Statement

Ryan Wittingslow: Conceptualisation (lead), Writing – original draft (lead), Writing – review & editing (lead)

Avigail Ferdman: Conceptualisation, Writing – original draft, Writing – review & editing

Use of AI

N/A

Funding Statement

The authors did not receive any financial support for the work on this article.

Acknowledgments

N/A

Conflict of Interest Statement

There is no conflict of interest.

References

- Alexandrova, A. (2017). Is there a single theory of well-being?. In *A Philosophy for the Science of Well-Being* (pp. 27–53). Oxford University Press. <https://doi.org/10.1093/oso/9780199300518.003.0002>
- Aristotle. (1984). *Nicomachean ethics* (W. D. Ross, Trans.). In J. Barnes (Ed.), *The complete works of Aristotle*. Princeton University Press.
- Björgvinsson, E., Ehn, P., & Hillgren P.-A. (2012). Design things and design thinking: Contemporary participatory design challenges. *Design Issues*, 28(3), 101–116. https://doi.org/10.1162/DESI_a_00165
- Bradford, G. (2015). *Achievement*. Oxford University Press.
- Bradford, G. (2016). Perfectionism. In G. Fletcher (Ed.), *The Routledge Handbook of Philosophy of Well-Being* (pp. 24–134). Routledge.
- Bradford, G. (2020). Perfectionist bads. *The Philosophical Quarterly*, 71(3), 586–604. <https://doi.org/10.1093/pq/pqaa055>
- Brink, D. O. (2003). *Perfectionism and the common good: Themes in the philosophy of T.H. Green*. Oxford University Press. <https://doi.org/10.1093/0199266409.001.0001>
- Cenci, A., & Cawthorne, D. (2020). Refining value sensitive design: A (capability-based) procedural ethics approach to technological design for well-being. *Science and Engineering Ethics*, 26(5), 2629 - 2662. <https://doi.org/10.1007/s11948-020-00223-3>
- Chan, J. K. H. (2017). Design ethics: Reflecting on the ethical dimensions of technology, sustainability, and responsibility in the Anthropocene. *Design Studies*, 54, 184–200. <https://doi.org/10.1016/j.destud.2017.09.005>
- Costanza-Chock, S. (2020). *Design justice: Community-led practices to build the worlds we need*. The MIT Press. <https://doi.org/10.7551/mitpress/12255.001.0001>
- Curzer, H. J., & Epstein, A. C. (2022). The virtuous data scientist and the ethics of good science. *Philosophy & Technology*, 35, Article 46. <https://doi.org/10.1007/s13347-022-00541-3>
- Danzico, L. (2010). The design of serendipity is not by chance. *Interactions*, 17(5), 16–18. <https://doi.org/10.1145/1836216.1836220>
- Donia, J., & Shaw, J. A. (2021). Ethics and values in design: A structured review and theoretical critique. *Science and Engineering Ethics*, 27(5), Article 57. <https://doi.org/10.1007/s11948-021-00329-2>
- Dunne, A., & Raby, F. (2013). *Speculative Everything: Design, Fiction, and Social Dreaming*. The MIT Press.
- Geach, P. T. (1956). Good and Evil. *Analysis*, 17(2), 33–42. <https://doi.org/10.2307/3326442>
- Green, T. H. (2003). *Prolegomena to ethics* (D. O. Brink, Ed.). Oxford University Press. (Original work published 1883)
- Ferdman, A. (2019a). A perfectionist basic structure. *Philosophy & Social Criticism*, 45(7), 862–882. <https://doi.org/10.1177/0191453718820891>

- Ferdman, A. (2019b). Walking and its contribution to objective well-being. *Journal of Planning Education and Research*, 43(2), 294–304. <https://doi.org/10.1177/0739456X19875195>
- Ferdman, A. (2024). Human flourishing and technology affordances. *Philosophy & Technology*, 37(1), Article 1. <https://doi.org/10.1007/s13347-023-00686-9>
- Ferdman, A. (2025). AI deskilling is a structural problem. *AI & Society*, 41(4), 3001-3013. <https://doi.org/10.1007/s00146-025-02686-z>
- Friedman, B., Kahn, P., & Borning, A. (2006). Value sensitive design and information systems. In P. Zhang & D. F. Galletta (Eds.), *Human-computer interaction in management information systems: Foundations*. Routledge. <https://doi.org/10.4324/9781315703619>
- Frischmann, B., & Selinger, E. (2018). *Re-engineering humanity*. Cambridge University Press. <https://doi.org/10.1017/9781316544846>
- Gibson, J. J. (2014). *The ecological approach to visual perception*. Psychology Press. <https://doi.org/10.4324/9781315740218> (Original work published 1979)
- Hanson, L. (2017). Artistic value is attributive goodness. *The Journal of Aesthetics and Art Criticism*, 75(4), 415–427.
- Hurka, T. (1993). *Perfectionism* (1st ed.). Oxford University Press. <https://doi.org/10.1093/0195101162.001.0001>
- Hurka, T. (2002). Capability, functioning and perfectionism. *Apeiron*, 35(4), 137–162. <https://doi.org/10.1515/APEIRON.2002.35.4.137>
- Jacobs, N. (2020). Capability sensitive design for health and wellbeing technologies. *Science and Engineering Ethics*, 26(6), 3363–3391. <https://doi.org/10.1007/s11948-020-00275-5>
- Kraut, R. (2007). *What is good and why: The ethics of well-being*. Harvard University Press.
- McLaughlin, P. (2001). *What functions explain: Functional explanation and self-reproducing systems*. *Cambridge Studies in Philosophy and Biology*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511498510>
- Millikan, R. G. (1984). *Language, thought, and other biological categories: New foundations for realism*. The MIT Press. <https://doi.org/10.7551/mitpress/4124.001.0001>
- Millikan, R. G. (1986). Thoughts without laws; Cognitive science with content. *The Philosophical Review*, 95(1), 47–80.
- Millikan, R. G. (1995). In defence of proper functions. In *White queen psychology and other essays for Alice*. The MIT Press. <https://doi.org/10.7551/mitpress/7224.003.0004>
- Neander, K. (1991). The teleological notion of ‘function’. *Australasian Journal of Philosophy*, 69(4), 454–468. <https://doi.org/10.1080/00048409112344881>
- Norman, D. (2013). *The design of everyday things: Revised and expanded edition*. Basic Books. (Original work published 1988)

- Nussbaum, M. C. (2000). *Women and human development: The capabilities approach*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511841286>
- Nussbaum, M. C. (2011). *Creating capabilities: The human development approach*. Harvard University Press.
- O'Mara, S. (2019). *In praise of walking: The new science of how we walk and why it's good for us*. Vintage Publishing.
- Parsons, G., & Carlson, A. (2008). *Functional beauty*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199205240.001.0001>
- Preston, B. (1998). Why is a wing like a spoon? A pluralist theory of function. *Journal of Philosophy*, 95(5), 215-254.
- Preston, B. (2013). *A philosophy of material culture: Action, function, and mind*. Routledge.
- Ratti, E., & Graves, M. (2021). Cultivating moral attention: A virtue-oriented approach to responsible data science in healthcare. *Philosophy & Technology*, 34(4), 1819-1846. <https://doi.org/10.1007/s13347-021-00490-3>
- Reviglio, U. (2019). Serendipity as an emerging design principle of the infosphere: Challenges and opportunities. *Ethics and Information Technology*, 21(2), 151-166. <https://doi.org/10.1007/s10676-018-9496-y>
- Robeyns, I. (2005). The capability approach: A theoretical survey. *Journal of Human Development*, 6(1), 93-117. <https://doi.org/10.1080/146498805200034266>
- Sen, A. (1999). *Development as freedom*. Oxford University Press.
- Sypnowich, C. (2016). *Equality renewed: Justice, flourishing and the egalitarian ideal*. Routledge. <https://doi.org/10.4324/9781315458335>
- Tahzib, C. (2022). *A perfectionist theory of justice*. Oxford University Press. <https://doi.org/10.1093/oso/9780192847119.001.0001>
- Thomasson, A. L. (2014). Public artifacts, intentions, and norms. In M. Franssen, P. Kroes, T. A. C. Reydon, & P. E. Vermaas (Eds.), *Artefact kinds: Ontology and the human-made world* (pp. 45-62). Springer. https://doi.org/10.1007/978-3-319-00801-1_4
- Vallor, S. (2016). *Technology and the virtues: A philosophical guide to a future worth wanting*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780190498511.001.0001>
- Vermaas, P. E., & Houkes, W. (2003). Ascribing functions to technical artefacts: A challenge to etiological accounts of functions. *The British Journal for the Philosophy of Science*, 54(2), 261-289. <https://doi.org/10.1093/bjps/54.2.261>
- Vermaas, P. E., & Houkes, W. (2006). Technical functions: A drawbridge between the intentional and structural natures of technical artefacts. *Studies in History and Philosophy of Science Part A*, 37(1), 5-18. <https://doi.org/10.1016/j.shpsa.2005.12.002>

- Williams, J. (2018). *Stand out of our light: Freedom and resistance in the attention economy*. Cambridge University Press. <https://doi.org/10.1017/9781108453004>
- Wittingslow, R. M. (2024). Building perfectionist ethics into action-theoretic accounts of function: A beginner's guide. *Philosophy & Technology*, 37(1), Article 28. <https://doi.org/10.1007/s13347-024-00711-5>
- Wittingslow, R. M. (2025). Finding function: Exaptation, maintenance, and exemplification. *Techné: Research in Philosophy and Technology*, 29(2/3), 260-291. <https://doi.org/10.5840/techne202584212>