

---

# POSITIVE AI: KEY CHALLENGES IN DESIGNING ARTIFICIAL INTELLIGENCE FOR WELLBEING

---

A PREPRINT

**Willem van der Maden \***

Department of Human-centered Design  
Delft University of Technology  
Delft, The Netherlands  
w.l.a.vandermaden@tudelft.nl

**Derek Lomas**

Department of Human-centered Design  
Delft University of Technology  
Delft, The Netherlands  
j.d.lomas@tudelft.nl

**Malak Sadek**

Dyson School of Design Engineering  
Imperial College London  
London, UK  
m.sadek21@imperial.ac.uk

**Paul Hekkert**

Department of Human-centered Design  
Delft University of Technology  
Delft, The Netherlands  
p.p.m.hekkert@tudelft.nl

February 5, 2024

## ABSTRACT

Artificial Intelligence (AI) is a double-edged sword: on one hand, AI promises to provide great advances that could benefit humanity, but on the other hand, AI poses substantial (even existential) risks. With advancements happening daily, many people are increasingly worried about AI's impact on their lives. To ensure AI progresses beneficially, some researchers have proposed "wellbeing" as a key objective to govern AI. This article addresses key challenges in designing AI for wellbeing. We group these challenges into issues of modeling wellbeing in context, assessing wellbeing in context, designing interventions to improve wellbeing, and maintaining AI alignment with wellbeing over time. The identification of these challenges provides a scope for efforts to help ensure that AI developments are aligned with human wellbeing.

**Keywords** artificial intelligence · wellbeing · value-alignment · positive design · positive computing · cybernetics

## 1 Introduction

The rapid advancement and adoption of generative AI (GenAI) technologies like *ChatGPT* signify the dawn of "The Age of AI." (Gates, 2023; Kissinger, Schmidt, & Huttenlocher, 2021) These developments mark a significant leap in the capabilities and adoption of AI systems. However, for many people, the swift and disorienting integration of AI into daily life raises many issues (Cugurullo & Acheampong, 2023; Fietta, Zecchinato, Stasi, Polato, & Monaro, 2022; Qasem, 2023). Concerns include the potential impacts on employment, privacy, and inequality, along with broader societal implications like human rights, mental health, and the preservation of democratic norms (Future of Life Institute, 2023; Prabhakaran, Mitchell, Gebru, & Gabriel, 2022; Shahriari & Shahriari, 2017; Stray, 2020). This article argues for the importance of wellbeing as a key objective in AI and for human-centered design (HCD) as a key methodology. Based on this framing, it shares a set of key challenges that will face designers of AI for wellbeing, or *Positive AI*.

The idea that AI should support wellbeing is not uncommon. In 2018, Zuckerberg (2018) (CEO of Meta, previously Facebook) publicly stated that wellbeing should be the goal of AI. Further, in an interview Jan Leike (Wiblin, n.d.) (head of the 'Superalignment' research lab at OpenAI) said AI optimization should align to "flourishing." Wellbeing,

---

\*Corresponding author: w.l.a.vandermaden@tudelft.nl

however, is complicated. It is not a naturally observable quantity, but rather a multifaceted construct that is based, at least in part, on conscious human experiences (Ruggeri, Garcia-Garzon, Maguire, Matz, & Huppert, 2020). Therefore, designing *Positive AI* requires understanding and shaping human experiences. This situates the challenge squarely in the domain of human-centered design (Auernhammer, 2022). Before reviewing the possibilities for HCD in designing AI for wellbeing, we will briefly address other fields associated with the creation of positive human outcomes in AI. The current article is not the venue for reviewing them in-depth. Yet, we find it important that Positive AI designers are aware broadly aware of their contributions.

## 1.1 Ethical AI

AI ethicists have been formulating ethical principles and frameworks to responsibly guide the development and implementation of AI systems. A key contribution came from Floridi et al. (2018) who synthesized a set of core ethical principles like beneficence, non-maleficence, autonomy, justice, and explicability for cultivating a “Good AI Society.” However, as principles alone they are insufficient (Mittelstadt, 2019): they need to be translated into concrete practices. Much work remains to develop these into practical tools and methodologies (Morley, Floridi, Kinsey, & Elhalal, 2020). Recent work has begun exploring approaches for embedding ethical values directly into AI system design (Klenk & Duijf, 2021; van de Poel, 2020), from which Value-sensitive Design (VSD) has emerged as a candidate to bridge the principle-practice gap (Umbrello & van de Poel, 2021). A recent critical review (Sadek, Calvo, & Mougnot, 2023b) indicates that VSD may be effective but limited. Some limitations include the inadequate elicitation of values, a tendency to depend on pre-established values over context-specific ones, and a lack of precise instructions for embedding values.

## 1.2 AI Alignment

Considering the potential for harm done by AI, some refer to such systems as misaligned with human values. For example, referencing social media, ethicist Tristan Harris says that by optimizing for attention, these platforms are misaligned with human wellbeing and dignity (T. Harris, 2017). “AI alignment” is a field of research that aims to develop systems that are aligned with human values and intent (Christian, 2020). Alignment has been earlier studied as the principal-agent problem in economics and law, where an agent must achieve the objectives and interests of the principal (Hadfield-Menell & Hadfield, 2019). For example, in a car repair scenario, the car owner (principal) expects the mechanic (agent) to fix the car efficiently and affordably. However, the mechanic might suggest unnecessary repairs to increase the bill (misalignment), contrary to the owner’s desire for cost-effective service. Taking this framework to AI systems, alignment means ensuring that AI agents effectively and reliably pursue the goals and preferences set by their designers and users. One successful example of technical alignment work is the use of Reinforcement Learning with Human Feedback (RLHF) (Christiano et al., 2017), which uses human preference data to align the behavior of Large Language Models (LLMs). Related techniques include Constitutional AI (Bai et al., 2022) and inverse reinforcement learning (IRL) (Ng & Russell, 2000). There are many new techniques in the expanding field of technical AI alignment.<sup>2</sup> However, while these technical efforts show tremendous progress, their technology-centered perspective risks missing broader sociotechnical considerations, such as the design of human systems to effectively respond to AI (Dung, 2023).

## 1.3 Why human-centered design?

Given the intrinsic relationship between wellbeing and conscious experience, some scholars have argued for the importance of human-centered design (HCD) in AI (Calvo & Peters, 2014; Desmet & Pohlmeier, 2013). One reason is that, as a field, HCD focuses on understanding and shaping human experiences. However, there are a variety of ways in which Human-centered design (HCD) might complement ethical perspectives and address gaps in the AI alignment field. For instance, HCD might help bring concrete implementation methods and a broader systemic perspective. A core tenet of HCD is to prioritize the needs, values, and capabilities of users, ensuring that the design process is centered around human beings and their interactions with technology.

Designers are trained to attend to—and empathize with—human experiences (Norman, 2013). This means considering the full context surrounding users and technologies, rather than just narrow functionality, as well as prioritizing the understanding of diverse users’ needs and experiences from their point of view (Sanders & Stappers, 2008). They are equipped with the ability to engage in stakeholder participation and reveal the ethical priorities and deeply-held beliefs relevant to design projects (Zhang, Boltz, Lynn, Wang, & Lee, 2023). This encompasses a blend of competencies from engineering design, including problem definition, scoping, and rapid prototyping, combined with methodologies from social sciences like conducting ethnographic research, interviews, deriving understanding from qualitative data, and engaging in empathetic practices (Kramer, Agogino, & Roschuni, 2016).

<sup>2</sup>For a comprehensive overview, readers are referred to the live agenda summarizing ongoing alignment efforts posted to the AI Alignment Forum (Technicalities & Stag, 2023), which has been founded by prominent alignment researcher Eliezer Yudkowsky.

These skills are particularly important for AI because the integration of diverse perspectives ensures that both technical efficiency and societal impacts are considered in AI development (Auernhammer, 2020). For instance, experimentation and prototyping in AI benefit from this blend, allowing for iterative refinement and alignment with human needs and values. Prototyping in AI can be difficult because of the inherent unpredictability and complexity in AI’s capabilities and outputs (Yang, Steinfeld, Rosé, & Zimmerman, 2020). HCD may help by applying user-focused approaches to manage these uncertainties. Moreover, involving end users directly in the design process ensures that AI solutions are tailored to real-world requirements, making the technology more accessible, usable, and effective (Li & Lu, 2021; Zhu, Liapis, Risi, Bidarra, & Youngblood, 2018). In summary, the HCD perspective can complement existing ethical and technical viewpoints in AI development, as it offers methodologies to create systems that balance technical robustness with socially responsible outcomes that benefit people and society at large.

The field of Positive Design focuses on promoting human flourishing. The Positive Design Framework provides a scaffold for solutions that can enhance subjective wellbeing through components like pleasure, meaning, and virtue (Desmet & Pohlmeier, 2013). Grounding positive design in theory and evaluating its effect through controlled studies helps ensure that designed solutions truly contribute to people’s happiness. Similarly, the Positive Computing (Calvo & Peters, 2014, 2019; Gaggioli, Riva, Peters, & Calvo, 2017) movement aims to leverage technology to measurably improve wellbeing and human potential. The emphasis on collaborations between fields like psychology, computer science, and design in positive computing underscores the importance of an interdisciplinary, human-centric approach for developing AI focused on wellbeing objectives (Calvo, Vella-Brodrick, Desmet, & M. Ryan, 2016). In many ways, the tenets of positive design and positive computing have helped lay the foundation for what we now call “Positive AI.”

#### 1.4 Why wellbeing?

A growing movement of scholars advocates for the incorporation of wellbeing metrics into AI systems so that optimization efforts can measurably contribute to social benefit (Schiff, Ayesch, Musikanski, & Havens, 2020; Shahriari & Shahriari, 2017). Specifically, they argue that measures of wellbeing can help manage AI’s effects on society (Musikanski, Rakova, Bradbury, Phillips, & Manson, 2020). Indeed, wellbeing has a strong methodological foundation (Stray, 2020), and there is extensive research on defining and measuring wellbeing; this suggests that algorithmic systems may be able to systematically optimize wellbeing (Havrdá & Rakova, 2020).

Wellbeing’s complexity captures many relevant societal concerns AI systems should address (Stray, 2020). This combination of rich meaning and inherent measurability supports the operationalizing wellbeing as an optimization objective for AI systems. This sentiment is also expressed by a recent IEEE standards review that argues for the adoption of holistic wellbeing frameworks (like IEEE 7010) to guide the design, deployment, and evaluation of AI systems (Schiff et al., 2020). However, significant questions remain regarding whether available wellbeing frameworks are fully sufficient, whether existing metrics are sufficient, what the impacts of wellbeing optimization may be (Musikanski et al., 2020; Schiff et al., 2020; Stray, 2020).

Some argue that wellbeing is a sort of ultimate objective: in *The Moral Landscape*, S. Harris (2010) argues that other values like fairness, transparency, or accountability should be seen as components that contribute to wellbeing, rather than ends in themselves. From this perspective, optimizing for wellbeing involves optimizing for all values that matter, but only insofar as they contribute empirically to wellbeing. In so far as AI systems are able to assess their own impact on human wellbeing, they may be able to potentially maximize all benefits and minimize all harms experienced by users and society (Havrdá & Klocek, 2023). Wellbeing optimization might then allow for the management of complex issues like misinformation and inequality associated with AI systems (Stray, 2020).

#### 1.5 Framing the challenges: human-centered design of AI systems

As a term, ‘Artificial intelligence’ is used to describe both a characteristic of computer systems and the methods employed to develop this feature, such as machine learning (ML) (Gabriel, 2020). Intelligence in both humans and machines has been defined as “an agent’s general ability to achieve goals in a wide range of environments.” (p. 9 Legg & Hutter, 2007) Following this definition, AI researchers Stuart Russel and Peter Norvig define *artificial* intelligence as a *designed* agent that perceives its environment through sensors and acts upon that environment using actuators (Russell & Norvig, 2022). The result of these sensors and actuators is a feedback loop that incorporates system output (e.g., action in its environment) as input for its future actions (e.g., the action had the desired effect). A cybernetic perspective examines these broader feedback loops between AI systems, their environment, and the social context in which they operate. Thus, AI *systems* (in contrast to AI/ML algorithms) can be viewed as sociotechnical systems embedded within a complex network of feedback loops (van de Poel, 2020). This broader and more systemic view of AI has been proposed as an approach to deal with some of the challenges of current and future AI systems (Dobbe, Krendl Gilbert, & Mintz, 2021; Krippendorff, 2021; Van der Maden, Lomas, & Hekkert, 2022).

A human-centered design perspective enables designers to look beyond “the algorithm” to consider how AI interacts within a network of social, ethical, cultural, and political factors (van de Poel, 2020). This means considering how AI influences human behavior, societal norms, and institutional structures—and how AI is, in turn, influenced. This perspective requires engaging with diverse stakeholders to understand their values and needs iterative and reflective design processes that continually assess and respond to these complex dynamics (Sadek et al., 2023b).

This perspective creates new affordances for the design of AI for wellbeing. Shaping the impact of AI can occur through multiple components of the AI system, including the technical artifacts, institutions, practices, etcetera—all in addition to the design of algorithms. For example, consider the role of AI in a video streaming platform like Netflix. A focus on the algorithm is limiting because one can only affect the likelihood of a particular recommendation. In contrast, a broader perspective opens up different kinds of interventions. For instance, in the user interface (e.g., autoplay); the organizational level (e.g., establishing boardroom content acquisition metrics beyond just engagement and growth); data science (e.g., introducing new metrics for optimization that prioritize suggesting content from more diverse voices); or the broader ecosystem (e.g., funding initiatives to broaden representation in the creative industry talent pipeline). In other words, design interventions can occur at multiple levels of the AI *system*, not just in the algorithm. Designers can even consider interventions outside of the control of the AI platform, such as a ‘Netflix Watch Club’ or an alternative YouTube user interface for education (Lukoff et al., 2023).

Why is this broadened view important? Rather than aligning AI to the needs of society, there may be many cases where it may be more appropriate for social institutions to adapt to AI. For instance, while ChatGPT could potentially be aligned with the needs of K12 schools (i.e. so that students are prevented from cheating on their assignments), designers may wish to create new guidelines for positive integration of AI in their courses (e.g., promoting AI literacy in using ChatGPT (Mollick, 2023)). With this systemic perspective, there are expanded opportunities for guiding AI impacts beyond the algorithmic design itself. This shows how there are new opportunities for creating AI for Wellbeing beyond what is typically the scope of AI alignment or ethical AI research.

In this article, we conceptualize AI as a sociotechnical system involving a complex interaction of various feedback loops, each optimized for specific objectives within the system. The essence of ‘Positive AI’ lies in harmonizing these objectives towards a singular, overarching aim: enhancing wellbeing.

## 2 Key Challenges

As experienced design researchers in this domain, we have consistently encountered unique challenges in designing *Positive AI*. The challenges we outline here are intended to inform and guide other designers embarking on similar ventures. This conceptual framework aids in structuring the challenges of designing AI for wellbeing around the following questions:

1. **Modeling** the state of the system: How do we operationally define wellbeing within the context of a particular sociotechnical system?
  - For example: *What wellbeing dimensions are important in the context of Netflix and how do we attribute changes in wellbeing to components of the system?*
2. **Assessing** the state of the system: How do we translate qualitative experiences into assessment metrics?
  - For example: *How can we elicit how people feel about their interactions with TikTok, and how can these experiences be translated into metrics that can be used for assessing future interventions and optimization processes?*
3. **Designing** system actuators: How do we design interventions in AI systems that promote and enhance wellbeing?
  - For example: *How do we know where in the sociotechnical system of ChatGPT we should and can intervene, and how do we know whether our potential interventions will achieve the desired effect?*
4. **Optimizing** the system objective: How do we know whether we are getting close to our desired goal?
  - For example: *How do we manage tradeoffs between autonomy and social connection in designing for wellbeing on Reddit, and how do align immediate outcomes with long-term wellbeing goals?*

We have found it useful to use the concept of cybernetics as a lens for organizing the challenges of designing AI for wellbeing (Dobbe et al., 2021; Krippendorff, 2021; Sato, 1991). Cybernetics allows for a systemic and holistic viewpoint that defines clear mechanisms for impact (Tabari, 2022). A cybernetic viewpoint, for us, naturally accommodates ecological and sociotechnical perspectives on the many feedback loops governing human systems and AI systems in society today. This cybernetic perspective allows us to organize the key challenges of designing AI for wellbeing into

four main categories, as shown in Table 1. This table serves as an overview and guides the structure for the rest of this chapter, with the numbers corresponding to the categories depicted in Figure 1.




Figure 1: Shows a schematic representation of a cybernetic system. The different challenges can be mapped onto this framework: (1) understanding the system context which entails modeling the relation between wellbeing of the systems constituents and its various components; (2) operationalizing said model of wellbeing; (3) designing interventions to actively promote operationalized model of wellbeing; and (4) retaining alignment with the overall goal. The latter refers to both challenges of algorithmic optimization as well as scrutinizing the objective (e.g., is the wellbeing objective still aligned to needs and desires of all relevant stakeholders?)

Table 1: Table 2. Overview of Key Challenges

Category	Challenge	Explanation
How to model wellbeing?	(1a) Identifying relevant wellbeing dimensions	There is a long tradition of wellbeing research spawning many different theoretical paradigms. In any given context, the appropriate theoretical paradigm may vary.
	(1b) Modeling wellbeing and interpreting fluctuations	The multifaceted nature of wellbeing and the richness of the world in which we live, make it difficult to model fluctuations in wellbeing to (specific features of) a system.
How to assess wellbeing?	(2a) Operationalizing wellbeing	In order for a system to respond properly and timely, we have to measure wellbeing in a dynamic, continuous manner that is sensitive to the context.
	(2b) Translating qualitative experiences into system metrics	We currently lack methods for connecting small-scale qualitative research on wellbeing to large-scale, quantitative methods, making it difficult to ensure that metrics of wellbeing are aligned with human experiences.
How to design for wellbeing?	(3) Designing interventions to promote wellbeing	It is not straightforward how an algorithm, the UI of a platform, or its content may lead to wellbeing effects. There’s both a lack of examples and a lack of appropriate design methodologies for Positive AI.
How to optimize for wellbeing?	(4a) Optimization tradeoffs and prioritization	Optimization causes tradeoffs that are difficult to measure, making it difficult to balance different goals and predict the effects of novel metrics on the optimization process and outcomes.
	(4b) Pace layers	Changes in wellbeing are slow but AI optimization cycles are fast, making it difficult to optimize for wellbeing.

## 2.1 Challenges related to modeling wellbeing

### 2.1.1 Identifying relevant wellbeing dimensions

Since the rise of the Positive Psychology movement that gained popularity around the turn of the millennium (M. E. Seligman, 2019) scientific interest in wellbeing has proliferated. Wellbeing is a multifaceted, rich, and dynamic phenomenon, and as such, there are many definitions of it in both scholarly and public contexts (e.g., policymaking). Each definition pertains to different aspects of wellbeing. For instance, the WHO defines health as “a state of complete physical, mental, and social wellbeing, and not merely the absence of disease or infirmity.” (Callahan, 1973) There are also scholarly traditions that call wellbeing “Happiness” which can be distinguished in episodic (“I feel happy”) and attributed happiness (“I am a happy person”) (Veenhoven, 2014). These definitions are akin to *hedonic* and *eudaimonic* understandings of wellbeing (Ryan & Deci, 2001)—which are the two categories typically used to describe wellbeing research.

Contemporary hedonic traditions tend to focus on the degree to which people experience positive and negative emotions and how satisfied they are with their life. Two reviews argue that the most prominent theoretical approach in this tradition is the tripartite model of Diener (Cooke, Melchert, & Connor, 2016; Linton, Dieppe, & Medina-Lara, 2016). This model is typically measured in terms of “life satisfaction” (Diener, Emmons, Larsen, & Griffin, 1985) and the presence of positive and absence of negative emotions (Watson, Clark, & Tellegen, 1988).

In contrast, eudaimonic wellbeing is based on the pursuit of virtue, striving to become the best version of oneself and developing one’s personal strengths (Deci & Ryan, 2008). Psychologists investigating this phenomenon tend to define it in a multifaceted way, such as Ryff’s Psychological Wellbeing (PWB) scale (Ryff, 1989; Ryff & Keyes, 1995) and Seligman’s PERMA model (M. E. P. Seligman, 2010). These theories overlap to a great extent and encompass facets such as positive relations, meaning in life, and personal growth.

Further, there is the tradition of Quality of Life (QoL), which is often used interchangeably with wellbeing. However, the literature on the subject often pertains more to social aspects of the phenomenon and, for example, situations towards end of life, living with a disability, or life after a clinical visit (Bakas et al., 2012; Felce & Perry, 1995; Moons, Budts, & De Geest, 2006).

A final related concept is wellness, which emphasizes a holistic lifestyle (e.g., nurturing emotional and spiritual intelligence) (Corbin & Pangrazi, 2001). In identifying the relevant dimensions of wellbeing, a designer could also investigate non-Western traditions of defining wellbeing such as Ubuntu (Hailey, 2008), Ikigai (Sone et al., 2008), and Gross National Happiness (GNH) (Ura, Alkire, & Zangmo, 2012).

Researchers criticizing the Western viewpoint in wellbeing science also argue that current mainstream approaches tend to be too contextualized (Mead, Fisher, & Kemp, 2021). Historically, psychological research on wellbeing has focused on psychometrics in order to statistically validate generalizable dimensions of wellbeing (Searle, Pykett, & Alfaro-Simmonds, 2021). However, this approach does not necessarily translate to actionable insights. Recognizing this, scholars have developed domain-specific theories of wellbeing for areas like work (Clifton & Harter, 2021) and education (Konu & Rimpelä, 2002).

Contrary to global conceptions of wellbeing, domain-specific theories focus more on aspects that are prevalent in that domain which may not translate to other domains. By contextualizing wellbeing research, the findings become more directly relevant for decision-making such as in policy-making. For example, they can reveal nuances around how wellbeing manifests with a given domain such as stress experience around work relations. However, increased specificity comes at the cost of reduced situational-consistency. This shift underscores the importance of balancing context-sensitive perspectives with more generalized insights that hold over time and place. Both remain indispensable for a comprehensive understanding of wellbeing.

The concept of digital wellbeing illustrates a similar balance. Digital wellbeing refers to wellbeing effects that typically occur in a digital context—i.e., while gaming, surfing the web, or interacting on social media. It goes beyond the mere time spent online, focusing on how digital engagement affects daily life and emotional wellbeing. It is about achieving a harmonious interaction with technology, where the benefits are maximized, and drawbacks such as loss of control are minimized. Furthermore, the IEEE 7010-7020 (IEEE SA, 2020) initiatives represent a pioneering effort to establish a comprehensive understanding of wellbeing within the realm of AI. Other researchers have explored the relationship between AI and community wellbeing (Musikanski et al., 2020) investigating the development of wellbeing metrics, community-centric AI, and applying AI to enhance community wellbeing.

Wellbeing, as highlighted in various perspectives, presents designers of AI systems with a range of dimensions and paradigms to consider. The relevance of wellbeing aspects varies by context; for example, physical health may be central in diet apps, while social media platforms might emphasize social connection. This complexity in people’s

interactions with technology is highly important for designers to consider. It requires careful, context-specific selection of the most relevant wellbeing paradigms and domains. Currently, there is no agreed-upon process for determining the which wellbeing dimensions to prioritize for a given system. This complexity in people’s interactions with technology, being situational and dynamic, is reflective of findings in the literature (Vanden Abeele, 2021). As AI systems become increasingly entangled with daily routines, modeling their impacts on wellbeing becomes difficult. Thus, selecting appropriate theoretical dimensions is crucial yet challenging. Designers, therefore, face the complex task of modeling wellbeing in AI systems within evolving sociotechnical landscapes, a challenge we will explore further.

### 2.1.2 Modeling and attributing wellbeing changes

For AI to effectively promote wellbeing, it requires a deep understanding of what actions and strategies contribute to this goal. This, in turn, requires a contextual model of wellbeing that allows the designer to attribute fluctuations to specific features of the system. Contextual models of wellbeing explain how various aspects of wellbeing manifest situationally and connect to components of a given system. The discussion below focuses on the challenges in developing contextual wellbeing models and attributing wellbeing fluctuations to components of that model, starting with the former.

Considering the complexity of AI systems and contexts, it is crucial to develop an understanding of how different aspects of wellbeing relate to specific features or components of a particular context. Quantitative measures enables researchers to objectively examine explanations and predictions from conceptual models addressing the determinants and outcomes of wellbeing (Diener, 2019). However, developing such models poses significant challenges. The relationships between various causal factors and wellbeing remain unclear, complicated by bidirectional and nonlinear effects. For example, while good sleep benefits mental health, improved mental health also leads to better sleep (Scott, Webb, Martyn-St James, Rowse, & Weich, 2021). Similarly, the connection between mental health and social media use is complicated: it is unclear whether people turn to social media when they are already struggling with mental health issues, or whether social media itself can contribute to these issues (Coyne, Rogers, Zurcher, Stockdale, & Booth, 2020; Hjetland, Schønning, Aasan, Hella, & Skogen, 2021).

The open questions around the relation between specific media and wellbeing have implications for the deployment of AI technologies aimed at promoting wellbeing (Johannes, Dienlin, Bakhshi, & Przybylski, 2022). Specifically, they highlight the need to consider the long-term impacts of these technologies on individuals’ psychological health and wellbeing—e.g., watching Netflix may be conducive to wellbeing in the current moment but how may it shape effects over time? Thus, to design Positive AI, it is necessary to have a thorough understanding of the relationships between wellbeing and its various antecedents. Without this understanding, it is challenging to measure the impact of interventions and determine their effectiveness. The real world is characterized by complex, interconnected patterns that can make it difficult to attribute changes in wellbeing to a single event or intervention (Fokkinga, Desmet, & Hekkert, 2020). This is particularly true when dealing with “narrow” AI systems, such as recommendation algorithms that suggest products a user might like based on their past preferences. These systems may excel at completing specific, limited tasks but struggle to account for the full range of factors that can impact wellbeing in the real world. For example, a product recommendation system would not consider how using that product might affect a user’s sleep, relationships, or long-term wellbeing.

This modeling poses difficulties not just due to AI system opacity (Gabriel, 2020), but because any platform comprises only a fraction of a person’s broader life experience. For example, changes in the dietary practices of a teenager may be a result of the content of their Instagram feed (e.g., only images of people with a specific body type) or because of some other event that occurred in their life (e.g., a breakup or the start of a new fitness program). In order to establish causality, wellbeing has to be measured in a manner that considers the complexity of the context in which the system is deployed. This poses a challenge because existing wellbeing assessments and platforms are not designed to consider the complexity of people’s lives and experiences that extend beyond the platform. For instance, how might Instagram account for life events that occur outside of its platform and influence a user’s wellbeing? And how would it determine which of those external factors are most relevant and impactful? The question of whether Instagram should be held accountable for the wider impact of its user interactions on wellbeing is a matter of ongoing debate. However, for the platform to effectively influence wellbeing, it is imperative that it addresses these interaction effects in some capacity.

In summary, developing contextual models that effectively attribute fluctuations in wellbeing to AI systems poses profound challenges. The interconnected relationships between wellbeing and other life factors resist straightforward causal analysis. This difficulty intensifies for narrow AI platforms, as they comprise limited slices of broader human experience. Although difficult, advancing more contextualized and systemic modeling methodologies promises significant progress toward the goals of Positive AI.

## 2.2 Challenges related to assessing wellbeing

Wellbeing is suited for AI optimization given its history of measurement (Stray, 2020). However, to measure wellbeing in context, we need to model wellbeing in context. For this, we need qualitative inputs to understand subjective personal and community experiences. Translating qualitative insights into quantitative metrics usable for optimization is non-trivial. Here, we discuss challenges related to effective assessment, which requires contextualization and bridging gaps between individual perspectives and system-level scales.

### 2.2.1 Contextually operationalizing wellbeing

As previously discussed, in order to attribute fluctuations in wellbeing to components of a given system, wellbeing must be effectively modeled and measured. While we have explored the theoretical challenges associated with this task, there are also methodological aspects that need to be addressed. This includes the operationalization of chosen wellbeing dimensions.

Wellbeing has traditionally been measured in field of psychology, but since the last decade, there has been an increased interest in also measuring wellbeing for other purposes, such as policymaking (Frijters & Krekel, 2021). The assessment of wellbeing is often done using qualitative surveys and interviews (Alexandrova, 2012). It is important that the assessment instruments employed in these studies are validated, possess temporal stability, and demonstrate cross-situational consistency (Diener & Michalos, 2009). However, literature suggests that these “off-the shelf” wellbeing assessment instruments may not be readily applicable in the context of novel technologies that rapidly change factors influencing wellbeing, such as social media (Kross et al., 2021) and AI (Stray et al., 2023). Traditional instruments aiming for consistency over time and place may be incompatible with contexts that transform quickly and substantially. This also means that present modalities of assessment (e.g., surveys and interviews) may not be sufficient for these emerging contexts (Stray et al., 2023).

Given these limitations, there is a need for measures that are better suited to assess wellbeing in the context of rapidly changing technologies and situations (Vanden Abeele, 2021). Context-sensitive measures, designed to adapt to the specific needs and context of the individuals being assessed (Loveridge, Sallu, Pesha, & R Marshall, 2020; van der Maden, Lomas, & Hekkert, 2023), provide a promising alternative and have been suggested to be more suitable for the operationalization of values (including wellbeing) in concrete applications (Liscio et al., 2021). The value inference process outlined by Liscio et al. (2023) highlights the importance of identifying context-specific values as a crucial step in aligning AI agents with human values. Their methodology demonstrates an approach to elicit relevant, contextualized values that could inform the development of context-sensitive wellbeing measures. These measures allow for a more tailored assessment of wellbeing, as they can be regularly updated to reflect the changing needs of the community. It should be noted that this adaptability poses tradeoffs regarding validation and stability over time and contexts. How to reconcile the discrepancy between the need for sensitive, customized measures and generalizable instruments remains an open question warranting further investigation.

Finally, to develop context-sensitive measures of wellbeing, it is crucial to engage regularly with the community to scrutinize and update the instruments used to assess contextual wellbeing. The literature has identified the importance of community engagement to implement ethical frameworks (Morley, Elhalal, et al., 2021) and articulate shared values in AI systems (Sanderson et al., 2023), though sources acknowledge there are gaps in best practices for stakeholder participation (Sadek, Calvo, & Mougnot, 2023a). While human-centered design methods may provide useful directions for establishing community engagement, they need to be adapted to the particularities of AI systems. Human-centered design excels at developing a deep understanding of contexts, maintaining community participation, and understanding individual and community needs. In conclusion, contextualizing wellbeing measures involves navigating consistency-sensitive tensions and engagement complexities. However, even well-constructed contextual measures struggle to inform AI optimization unless translated to system-level data. This underscores the pivotal challenge of connecting qualitative insights with large-scale, quantitative metrics for AI alignment.

### 2.2.2 Translating qualitative experiences into system metrics

Wellbeing is a highly personal experience that typically requires individuals to report on their own experiences in order to be measured (Diener & Michalos, 2009; Linton et al., 2016). As such, qualitative methods such as 1-to-1 interviews, focus groups, and ecological momentary assessments, are often expected to provide more contextually actionable data. However, currently, there is no agreed-upon process for translating small-scale activities of this nature into large-scale optimization metrics (McGregor, Camfield, & Coulthard, 2015). Instead, easy-to-collect but incomplete or inaccurate metrics are used—such as hours spent on a social media platform to measure satisfaction with content—which do not fully capture users’ experiences and overlook harmful consequences to their wellbeing (Thomas & Uminsky, 2020). Conversely, connecting qualitative research on wellbeing with large-scale optimization methods could not only help

ensure that measures of wellbeing are well-aligned with human experiences, but also help identify areas in which measures of wellbeing can be improved (Camfield, 2016). Aside from issues of scalability, there exists a discrepancy between the types of metrics suitable for on-platform measurement versus those typical for wellbeing assessment.

Whereas self-reporting suggests to be the best way of measuring wellbeing, behavioral data collection is the default method for on-platform optimization. However, behavioral metrics cannot reliably measure wellbeing since research on the relationship between behavioral and self-reported wellbeing measures is limited and inconclusive (Dang, King, & Inzlicht, 2020). Currently, it is unclear whether behavioral metrics can replace self-report measurements. Relying solely on self-reporting may, however, negatively impact the user experience, as users should not be bombarded with wellbeing questions upon engaging with a platform. This presents a challenge for designers and design researchers: how can platforms facilitate user feedback to collect accurate and scalable wellbeing data? This non-trivial challenge has also been acknowledged by other researchers (e.g., Steur & Seiter, 2021).

In conclusion, effectively integrating wellbeing requires translating between qualitative experiences and system metrics. While qualitative insights like self-reports capture personal experiences vital for alignment, convenient behavioral metrics dominate on-platform data collection. Absent mechanisms to translate small-scale activities into optimization inputs, AI risks misrepresenting user needs. Progress necessitates the development of methods that bridge this divide—continuously engaging individuals and communities while surfacing priorities at a systemic level. However, creating participatory channels poses immense practical difficulties around incentivization, standardization, and scalable synthesis. Though an open challenge, pioneering such participatory architectures in ways that meaningfully empower stakeholders promises to actualize AI’s potential for responsibly nurturing human flourishing.

### 2.3 Challenge related to designing for wellbeing

We have discussed the difficulties in conceptualizing and operationalizing wellbeing for AI systems. We have also touched on the need to optimize across individual, community and societal levels of wellbeing. A further question is: given a concept and operationalization of wellbeing, what actions can an AI system take to positively impact wellbeing? And how may we design such actions?

#### 2.3.1 Methodological challenge of designing (AI) actions to promote wellbeing

Designing AI is an incredibly challenging task for designers (Sadek et al., 2023b; Yang et al., 2020). Before tackling the design of AI that promotes wellbeing, it is important to understand the difficulties that designers face when working with any AI-based system in general. Currently, the communication gap between designers, developers, and end-users (Yang et al., 2020; Yu et al., 2020) causes an “AI support vacuum” (Abaza, 2021) where AI neither supports stakeholders nor is supported by them. Despite calls for the more interdisciplinary design of AI-based systems (Harbers & Overdiek, 2022; West, Whittaker, & Crawford, 2019), there is a lack of ‘translational work’ in current interventions that aim to support this collaborative design (Wong, Madaio, & Merrill, 2022). Aside from communication-based challenges, without proper training, it is difficult for designers to understand how to ideate, design and prototype for AI-based systems (van Allen, 2018). It is important to mention and consider these challenges before examining the extra difficulties that a focus on wellbeing might present.

Following the earlier discussion of adopting a systemic perspective, it is currently unclear whether designing Positive AI requires changes to the interface, algorithms, content moderation policies, business models, or otherwise defined components. Interventions could even extend beyond the system itself. This uncertainty stems from a lack of examples and established methods for putting wellbeing at the core of AI design. For example, current strategies for promoting wellbeing through ChatGPT remain largely undefined and untested. Designing Positive AI requires fundamentally rethinking how we approach design to focus on the continuous measurement of wellbeing and alignment with human values. There are well-established frameworks for designing with wellbeing in mind, such as Positive Design (Desmet & Pohlmeier, 2013) and Positive Computing (Calvo & Peters, 2014). These approaches acknowledge the need to consider the impacts on wellbeing, but they were not developed with the unique complexities of AI systems (such as system opacity, unpredictability, and scalability) in mind.

In addition to design methods, alignment approaches like Constitutional AI (Bai et al., 2022), Inverse Reinforcement Learning (IRL) (Arora & Doshi, 2021; Ng & Russell, 2000) and Contestable AI (Alfrink, Keller, Kortuem, & Doorn, 2022) also aim to ensure AI systems remain human-centered and aligned with human values. While these methods primarily focus on modifying the AI system’s internal processes and decision-making mechanisms, they highlight the importance of incorporating human values and wellbeing considerations directly into AI systems. By combining the strengths of these alignment approaches with design methods, we can develop a more holistic and effective AI alignment strategy that addresses user experience, interaction, and overall design, which are crucial for affecting wellbeing. For instance, some platforms, like YouTube and Twitter, have made small changes intended to benefit wellbeing, such as

removing dislike counts and allowing users to ‘unmention’ themselves from conversations. Whether these are cases of ‘ethics washing’ (Floridi & Cows, 2019) or have actual benefit to wellbeing is currently unknown to the public and academia, as research on their effects is not shared publicly.

In this vein, Stray (2020) discuss Facebook’s MSI metrics and YouTube’s satisfaction metrics as cases of AI optimization aligned to wellbeing. However, they criticize that both companies did not involve or get feedback from the people affected by their AI changes. The absence of public assessment and unclear impact on broader aspects of wellbeing, such as social connectedness or life satisfaction, particularly in diverse communities, underscores a crucial shortcoming. This lack of detailed information and engagement hinders a comprehensive understanding of the effectiveness of these interventions in truly aligning AI with community wellbeing. This is a broader concern in the value-sensitive design (VSD) of AI. That is, as discussed in the background section, while VSD methods support the identification and embodiment of values for AI well, they lack support in assessing ‘realized’ values—i.e., whether the designed outcomes in fact achieve the intended effect on said value (Sadek et al., 2023b). This is particularly important for Positive AI, because the core premise lies in empirically confirming that systems positively impact human wellbeing. Without closing the loop between intended and actualized outcomes, the benefits of proposed interventions remain theoretical.

To recapitulate, designing AI actions that support wellbeing is a complex challenge that requires a combination of design methods and alignment approaches. Successful regulation and alignment requires internal diversity matching external complexity. Without design methods and alignment approaches considering the full range of variables shaping wellbeing, AI systems will struggle to effectively promote flourishing. While existing design methods provide guidance, further research and development is needed. By focusing on continuous measurement of wellbeing, alignment with human values, and incorporating aspects of both design methods and alignment approaches, we can work towards creating Positive AI systems that not only avoid harm but actively promote human flourishing.

## 2.4 Challenges related to optimizing wellbeing

### 2.4.1 Optimization tradeoffs

Maximizing user engagement to drive revenue is a central optimization challenge for most platforms, where varied metrics like views, likes, and shares track user interaction with content. Using multiple metrics allows optimization algorithms to take a multi-objective approach to personalize content for optimal user engagement (Trunfio & Rossi, 2021). However, balancing these objectives already involves tricky tradeoffs (Thorburn, 2022). For example, there is tension between showing novel content to pique interest and only showing content matched to the user’s interests to avoid disengagement (Lu, Dumitrache, & Graus, 2020). While some novelty draws users in, too much risks boring them with irrelevant content. Now, considering wellbeing optimization makes this process even more complex.

Firstly, it is difficult to determine which facet of wellbeing should be prioritized—both in the moment and over time. The various facets of wellbeing often compete with one another. For example, social media platforms must balance users’ needs for social connection and personal autonomy. Promoting social interaction may support wellbeing by facilitating relationships, but it could also infringe on users’ freedom to choose their own activities. Additionally, an individual’s priorities may shift over time. What enhances wellbeing in the short-term may differ from long-term needs (e.g. enjoying frequent social activities when moving to a new city versus after settling in). This complexity requires nuanced techniques that can account for tradeoffs between competing needs and evolving individual priorities.

Secondly, because of these optimization tradeoffs there is no optimal solution. This is a common issue in environmental sciences who face complex tradeoffs between, for example, biodiversity and human wellbeing (Daw et al., 2015; McShane et al., 2011). This also goes for optimizing for just human wellbeing in the sense that there is no “right” outcome in balancing, for example, the wellbeing of an individual over that of a family (who are likely to value different things). Here the notions of requisite variety and satisficing become important again. That is, as platforms are faced with these complex tradeoffs they should develop many different strategies to deal with problems as they come up and take small incremental steps to deal with tradeoffs over time.

Thirdly, prioritizing one wellbeing facet over others may yield unintended consequences—beyond just direct tradeoffs. Focusing narrowly on a single aspect of wellbeing can backfire and undermine that very facet over time, given the complexity and interconnectivity of wellbeing. For instance, features meant to enhance social connectedness could hamper autonomy or other unforeseen needs. Wellbeing’s multidimensional nature means myopic solutions risk negative ripple effects from complex causal interrelations that remain poorly understood. Therefore, maintaining broad sensitivity to these complex interactions is critical. By regularly reassessing for subtle harms and tightening feedback loops, platforms can progressively identify and resolve unintentional side effects. An open, responsive systems perspective allows more complete understanding to emerge gradually from ongoing learning. Overall vigilance to complexity and interconnectivity may better serve wellbeing than rigid prioritization of singular facets.

Therefore, when designing AI systems for wellbeing, the models guiding decisions should be continually reassessed through collaboration with relevant stakeholders. This allows for adaptive alignment as understandings of wellbeing, design contexts, and community needs evolve over time. Continual engagement enables updating system priorities to restore balance across wellbeing dimensions. However, even with continual reassessment and stakeholder collaboration, another fundamental challenge persists—the differential pace of change between wellbeing and AI optimization.

#### 2.4.2 Fundamental challenge of pace

Changes in wellbeing typically manifest over longer periods of time. For instance, life satisfaction is not expected to fluctuate dramatically from week to week (Pavot, 2014). While it can be argued that this is due to our assessment instruments (Boschman, Nieuwenhuijsen, & Sluiter, 2018), the argument remains that this is in stark contrast to the pace of AI optimization and, for instance, media consumption, which are much faster. This can make it difficult to reconcile the two, as the wellbeing effects of AI actions should inform the system’s optimization cycles. Note here that system optimization cycles here refer to both analog (e.g., managerial, and designerly) and digital (e.g., algorithmic news feed recommendations) optimization cycles.

In other words, from the perspective of the optimization algorithm, the best it can do is optimize for hedonic wellbeing (momentary pleasure), while the goal is to design AI that supports eudaimonic wellbeing too (long-term wellbeing). An AI system can optimize for what feels good in the moment but not for what is good for you in the future. For example, watching “just one more episode” may be desirable in the moment (hedonic) but regretted the next morning when you feel tired during an important meeting (eudaimonic). Alternatively, from a managerial optimization perspective, linking executive goals to user wellbeing metrics may show progress only after longer periods than quarterly reporting cycles. Yet leaders frequently anticipate prompt, measurable outcomes that can be effectively communicated during shareholder meetings, reflecting a preference for quick, tangible results. This pace mismatch risks reactive changes before initiatives fully play out—such as prematurely disbanding a wellbeing taskforce. Evidently, there is a need to reconcile these pace layers somehow.

Here, we can look to Stewart Brand’s theory of pace layering (Brand, 2018), which suggests that different systems evolve at different speeds (see Figure 2). The pace of platform change far outpaces that of wellbeing. This is a fundamental constraint: the pace of AI optimization processes and the timeline for observable changes in human wellbeing are mismatched, and the rate of these dynamics in either domain is unlikely to significantly change in the future. Acknowledging and making explicit that there are in fact different components in a system that evolve at different paces is essential as we should look for ways to bridge or translate across layers. That is, we need an intermediary layer for measurement—for which qualitative methods may be well-suited. In the example given before, Netflix may choose to incorporate qualitative feedback from users regarding their overall satisfaction with their viewing habits, including how it affects their daily lives, sleep patterns, and long-term goals—bridging the gap between behavior and AI actions.




Figure 2: Different systems operate at different paces, adapted from Stephen (2021)

### 3 Discussion

This paper aimed to outline key challenges designers face when developing AI systems to actively promote human wellbeing, termed Positive AI. It makes two main contributions. First, it proposed adopting a cybernetic, systemic perspective for conceptualizing and addressing these alignment challenges. This viewpoint emphasizes the sociotechnical nature of AI systems, considering potential interventions across multiple system levels. Second, it organized the complex issue of designing Positive AI into four main categories of challenges: 1) modeling wellbeing, 2) assessing wellbeing, 3) designing for wellbeing, and 4) optimizing for wellbeing. These categories provide structure for mapping relevant problems within a systemic framework geared towards continuous improvement through stakeholder participation.

To recapitulate the challenges, designing Positive AI is facing substantial gaps regarding our knowledge on how to do it. The abundance of theoretical wellbeing paradigms makes it difficult for designers to get started on modelling wellbeing within a given context. This leads to issues in measuring wellbeing and attributing changes across different scales, from individuals to communities to sociotechnical systems. Further, the complexity of wellbeing introduces optimization tradeoffs, raising questions around balancing competing wellbeing needs between individuals and groups over time. While these present methodological and technical difficulties, they are fundamentally design challenges requiring creative solutions.

Because of these challenges, in some cases, platforms may opt for simpler derivatives of wellbeing (Pan, Bhatia, & Steinhardt, 2022), or designs that merely mitigate illbeing or improve aesthetics, rather than investing resources in optimizing for wellbeing itself. However, the impending mental health crisis (Organization, n.d.) and the transformative impact GenAI (Chui, Hazan, Roberts, Singla, & Smaje, 2023) make it all the more urgent to move towards Positive AI that systematically prioritizes wellbeing. This requires more than just a focus on mitigating harm; it is about capitalizing on each platform’s potential to foster wellbeing.

The key takeaway is that adopting a cybernetic perspective that places wellbeing assessment at the core can guide this process. It emphasizes the need for design researchers to continuously re-examine contextual models of wellbeing. This requires engaging relevant stakeholders, such as end-users and developers. By framing wellbeing as an ongoing conversation, we can iteratively refine models and measurements as contexts evolve through stakeholder participation. This reflexive, adaptive approach allows designers to navigate complexity and uncertainty when developing AI systems aimed at fostering human flourishing.

#### 3.1 Wellbeing as an objective for AI optimization

The concept of optimizing metrics wellbeing in AI systems emerges as a response to the limitations of traditional metrics like profit or efficiency (Schiff et al., 2020; Stray, 2020). While financial returns have been the primary focus for many companies, this singular pursuit may lead to broader societal harms (Virokannas, Liuski, & Kuronen, 2020). Might wellbeing as a guiding principle for AI design serve as a compelling alternative? Wellbeing, with its quantifiable and multidimensional nature, encompasses various aspects of human life which may make it a more suitable optimization goal. It allows for concrete optimization metrics that are more aligned with human-centered values, potentially reversing the trends of societal harm caused by narrowly focused objectives.

Yet, it seems impossible (and ill-advised) to reduce wellbeing to a single metric. Instead, wellbeing should serve as a multidimensional guide for value-based decisions and a comprehensive principle for moral choices.

Despite arguments for optimizing AI systems for wellbeing, companies may still opt for simpler derivatives due to factors like potential public relations issues, unclear financial rewards, and risks of losing competitive advantages. For example, companies may hesitate to invest in research on the effects of their products on wellbeing, as negative findings could result in bad publicity as exemplified by the 2018 ‘Techlash’ at Facebook for instance (Hemphill, 2019). Further, the benefits of prioritizing user wellbeing over profits are currently ambiguous for technology companies. Specifically, companies grapple with uncertainty about whether prioritizing user wellbeing over profits, which often involves focusing more on ‘doing good’ rather than just ‘preventing harm,’ will yield tangible benefits (Morley, Kinsey, et al., 2021). This ambiguity is rooted in the difficulty of quantifying the return on investment for ethical AI practices that emphasize proactive welfare measures over mere harm avoidance. However, gaining clarity on this trade-off is impeded by corporate reluctance toward transparency and third-party algorithm audits, which are seen as jeopardizing competitive advantage (Stray & Hadfield, n.d.; Stray et al., 2023). This uncertainty, coupled with the perceived risks of optimizing for wellbeing, disincentivizes companies from allocating resources to human-centered design interventions. Overcoming these barriers will require establishing transparent accountability mechanisms, alternative business models not reliant on exploitation, and fostering an ethical culture recognizing that benefiting humanity and profits can be compatible (Di Vaio, Palladino, Hassan, & Escobar, 2020).

This reluctance highlights intricate tensions between public perception, economic incentives, and ethical duties facing corporate decision-makers. Such challenges of power, as also publicly displayed during the OpenAI-Altman debacle late 2023 (Ulanoff, 2023), are highly relevant for Positive AI. These second-order challenges determine system-level goals, shaping whether wellbeing optimization occurs. If companies lack motivation to prioritize wellbeing, alignment is unlikely. Wellbeing as an overarching principle for AI alignment is promising but faces real-world obstacles regarding corporate priorities. While prominent voices endorse human flourishing as the goal, transparency and accountability mechanisms appear necessary to actualize this vision (Morley, Elhalal, et al., 2021).

### 3.2 Future opportunities

AI benchmarking is a popular method for evaluating the capabilities of large language models (LLMs). As AI benchmarking matures and as AI permeates more aspects of life, more sophisticated will be benchmarks required (Burnell et al., 2023). For example, benchmarks for LLM qualities like toxicity are now widely used (Lynch, 2023). Carefully crafted wellbeing metrics could serve as a mechanism for academics and others to indirectly optimize AI systems such as ChatGPT. This is because benchmarks are used to compare different models; as a result, low performance on a benchmark can motivate improvement. Human-centered designers may help attune evaluation methods such as benchmarking better to actual human experiences and ensure that optimization metrics align with these experiences.

### 3.3 Limitations and final remarks

While this paper aimed to provide an overview of key challenges in designing Positive AI, there are inherent limitations in addressing such a complex, transdisciplinary topic. For instance, environmental sustainability is not a core focus of the article, yet it warrants mention given the intricate relationship between sustainability and human wellbeing. As environmental crises increasingly threaten flourishing across communities, sustainability is being recognized as fundamental to comprehensive models of wellbeing (Kjell, 2011; O'Mahony, 2022). That is, wellbeing means more than human wellbeing. Meanwhile, the energy-intensive nature of AI systems presents sustainability challenges (Vinesa et al., 2020). This surfaces an alignment tension between AI benefits and potential unintended harms. Additionally, many relevant issues, from philosophy of technology to data ethics, could only be briefly touched upon given the practical design focus of this paper, but present key areas for further investigation.

Nonetheless, this work aims to spark discussion and research at the intersection of AI, wellbeing, and design. Further interdisciplinary collaboration building on these ideas will develop more pluralistic perspectives on Positive AI. The goal of this work was to identify key challenges that must be addressed by designers in order to develop AI systems aligned with human wellbeing. As the fourth industrial revolution is well on its way, “the time of reckoning for artificial intelligence is now.” (Ozmen Garibay et al., 2023)

## References

- Abaza, A. (2021). *Mlops: Why is it the most important technology in the age of ai?* Retrieved from <https://www.synapse-analytics.io/post/mlops-why-is-it-the-most-important-technology-in-the-age-of-ai>
- Alexandrova, A. (2012, December). Well-Being as an Object of Science. *Philosophy of Science*, 79(5), 678–689. Retrieved 2022-05-17, from [https://www.cambridge.org/core/product/identifier/S0031824800014793/type/journal\\_article](https://www.cambridge.org/core/product/identifier/S0031824800014793/type/journal_article) doi: doi:10.1086/667870
- Alfrink, K., Keller, I., Kortuem, G., & Doorn, N. (2022, August). Contestable AI by Design: Towards a Framework. *Minds and Machines*, 1–27. Retrieved 2022-08-23, from <https://doi.org/10.1007/s11023-022-09611-z> doi: doi:10.1007/s11023-022-09611-z
- Arora, S., & Doshi, P. (2021). A survey of inverse reinforcement learning: Challenges, methods and progress. *Artificial Intelligence*, 297, 103500. doi: doi:<https://doi.org/10.1016/j.artint.2021.103500>
- Auernhammer, J. (2020, 8). Human-centered ai: The role of human-centered design research in the development of ai. In S. Boess, M. Cheung, & R. Cain (Eds.), *Synergy - drs international conference 2020* (p. 11-14). Held online. doi: doi:10.21606/drs.2020.282
- Auernhammer, J. (2022, August). Human-centered AI: The role of Human-centered Design Research in the development of AI. In Boess, S., Cheung, M. and Cain, R. (eds.), *Synergy - DRS International Conference 2020*. Online. Retrieved 2023-03-28, from <https://dl.designresearchsociety.org/drs-conference-papers/drs2020/researchpapers/89> doi: doi:10.21606/drs.2020.282

- Bai, Y., Kadavath, S., Kundu, S., Askill, A., Kernion, J., Jones, A., ... Kaplan, J. (2022, December). *Constitutional AI: Harmlessness from AI Feedback*. arXiv. Retrieved 2023-03-22, from <http://arxiv.org/abs/2212.08073> (arXiv:2212.08073 [cs])
- Bakas, T., McLennon, S. M., Carpenter, J. S., Buelow, J. M., Otte, J. L., Hanna, K. M., ... Welch, J. L. (2012, November). Systematic review of health-related quality of life models. *Health and Quality of Life Outcomes*, 10(1), 134. Retrieved 2022-12-12, from <https://doi.org/10.1186/1477-7525-10-134> doi: doi:10.1186/1477-7525-10-134
- Boschman, J. S., Nieuwenhuijsen, K., & Sluiter, J. K. (2018). Within-person fluctuations in wellbeing and task-specific work ability. *Quality of Life Research*, 27, 437–446. doi: doi:<https://doi.org/10.1007/s11136-017-1713-3>
- Brand, S. (2018, January). Pace Layering: How Complex Systems Learn and Keep Learning. *Journal of Design and Science*. Retrieved 2022-07-12, from <https://jods.mitpress.mit.edu/pub/issue3-brand/release/2> doi: doi:10.21428/7f2e5f08
- Burnell, R., Schellaert, W., Burden, J., Ullman, T. D., Martinez-Plumed, F., Tenenbaum, J. B., ... others (2023). Rethink reporting of evaluation results in ai. *Science*, 380(6641), 136–138.
- Callahan, D. (1973). The WHO Definition of 'Health'. *The Hastings Center Studies*, 1(3), 77–87. Retrieved 2022-12-09, from <https://www.jstor.org/stable/3527467> (Publisher: The Hastings Center) doi: doi:10.2307/3527467
- Calvo, R. A., & Peters, D. (2014). *Positive Computing: Technology for Wellbeing and Human Potential*. MIT Press. (Google-Books-ID: uI6ZBQAAQBAJ)
- Calvo, R. A., & Peters, D. (2019). Design for wellbeing - Tools for research, practice and ethics. *Conference on Human Factors in Computing Systems - Proceedings*, 1–5. doi: doi:10.1145/3290607.3298800
- Calvo, R. A., Vella-Brodrick, D., Desmet, P., & M. Ryan, R. (2016). Editorial for “Positive Computing: A New Partnership Between Psychology, Social Sciences and Technologists”. *Psychology of Well-Being*, 6(1), 4–9. (Publisher: Springer Berlin Heidelberg) doi: doi:10.1186/s13612-016-0047-1
- Camfield, L. (2016). Enquiries into Wellbeing: How Could Qualitative Data Be Used to Improve the Reliability of Survey Data? In S. C. White & C. Blackmore (Eds.), *Cultures of Wellbeing: Method, Place, Policy* (pp. 47–65). London: Palgrave Macmillan UK. Retrieved 2023-04-03, from [https://doi.org/10.1057/9781137536457\\_2](https://doi.org/10.1057/9781137536457_2) doi: doi:10.1057/9781137536457\_2
- Christian, B. (2020). *The Alignment Problem: Machine Learning and Human Values*. WW Norton. (Google-Books-ID: KGCNEAAAQBAJ)
- Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., & Amodei, D. (2017). Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.
- Chui, M., Hazan, E., Roberts, R., Singla, A., & Smaje, K. (2023). The economic potential of generative ai.
- Clifton, J., & Harter, J. (2021). *Wellbeing at Work*. Simon and Schuster. (Google-Books-ID: T5YfEAAAQBAJ)
- Cooke, P. J., Melchert, T. P., & Connor, K. (2016). Measuring Well-Being: A Review of Instruments. *Counseling Psychologist*, 44(5), 730–757. doi: doi:10.1177/0011000016633507
- Corbin, C. B., & Pangrazi, R. P. (2001, December). *Toward a Uniform Definition of Wellness: A Commentary* (Tech. Rep.). President’s Council on Physical Fitness and Sports, 200 Independence Avenue, S. Retrieved 2022-12-12, from <https://eric.ed.gov/?id=ED470691> (Publication Title: President’s Council on Physical Fitness and Sports Research Digest ERIC Number: ED470691)
- Coyne, S. M., Rogers, A. A., Zurcher, J. D., Stockdale, L., & Booth, M. (2020). Does time spent using social media impact mental health?: An eight year longitudinal study. *Computers in Human Behavior*, 104(July 2019), 106160–106160. Retrieved from <https://doi.org/10.1016/j.chb.2019.106160> (Publisher: Elsevier Ltd) doi: doi:10.1016/j.chb.2019.106160
- Cugurullo, F., & Acheampong, R. A. (2023, January). Fear of AI: an inquiry into the adoption of autonomous cars in spite of fear, and a theoretical framework for the study of artificial intelligence technology acceptance. *AI & SOCIETY*. Retrieved 2023-12-20, from <https://doi.org/10.1007/s00146-022-01598-6> doi: doi:10.1007/s00146-022-01598-6
- Dang, J., King, K. M., & Inzlicht, M. (2020). Why are self-report and behavioral measures weakly correlated? *Trends in cognitive sciences*, 24(4), 267–269.
- Daw, T. M., Coulthard, S., Cheung, W. W. L., Brown, K., Abunge, C., Galafassi, D., ... Munyi, L. (2015, June). Evaluating taboo trade-offs in ecosystems services and human well-being. *Proceedings of the National Academy of Sciences*, 112(22), 6949–6954. Retrieved 2024-01-03, from <https://www.pnas.org/doi/abs/10.1073/pnas.1414900112> (Publisher: Proceedings of the National Academy of Sciences) doi: doi:10.1073/pnas.1414900112
- Deci, E. L., & Ryan, R. M. (2008). Hedonia, eudaimonia, and well-being: An introduction. *Journal of Happiness Studies*, 9(1), 1–11. doi: doi:10.1007/s10902-006-9018-1
- Desmet, P. M. A., & Pohlmeier, A. E. (2013). Positive design: An introduction to design for subjective well-being. *International Journal of Design*, 7(3), 5–19.

- Diener, E. (2019). *Assessing Well-being* (Vol. 53). (Issue: 9 Pages: 1791 Publication Title: Fitzpatrick’s Dermatology) doi: doi:10.1017/CBO9781107415324.004
- Diener, E., Emmons, R. A., Larsen, R., & Griffin, S. (1985). The Satisfaction With Life Scale. *Journal of Personality Assessment*, 49(1), 71–75. doi: doi:10.23943/princeton/9780691188959.003.0002
- Diener, E., & Michalos, A. C. (Eds.). (2009). *Assessing Well-Being* (Vol. 39). Dordrecht: Springer Netherlands. Retrieved 2022-05-04, from <http://link.springer.com/10.1007/978-90-481-2354-4> doi: doi:10.1007/978-90-481-2354-4
- Di Vaio, A., Palladino, R., Hassan, R., & Escobar, O. (2020). Artificial intelligence and business models in the sustainable development goals perspective: A systematic literature review. *Journal of Business Research*, 121, 283–314.
- Dobbe, R., Krendl Gilbert, T., & Mintz, Y. (2021, November). Hard choices in artificial intelligence. *Artificial Intelligence*, 300, 103555. Retrieved 2023-04-07, from <https://linkinghub.elsevier.com/retrieve/pii/S0004370221001065> doi: doi:10.1016/j.artint.2021.103555
- Dung, L. (2023). Current cases of ai misalignment and their implications for future risks. *Synthese*, 202(5), 138.
- Felce, D., & Perry, J. (1995, January). Quality of life: Its definition and measurement. *Research in Developmental Disabilities*, 16(1), 51–74. Retrieved 2022-12-12, from <https://www.sciencedirect.com/science/article/pii/S0891422294000288> doi: doi:10.1016/0891-4222(94)00028-8
- Fietta, V., Zecchinato, F., Stasi, B. D., Polato, M., & Monaro, M. (2022, June). Dissociation Between Users’ Explicit and Implicit Attitudes Toward Artificial Intelligence: An Experimental Study. *IEEE Transactions on Human-Machine Systems*, 52(3), 481–489. Retrieved 2023-12-20, from <https://ieeexplore.ieee.org/document/9627595> (Conference Name: IEEE Transactions on Human-Machine Systems) doi: doi:10.1109/THMS.2021.3125280
- Floridi, L., & Cows, J. (2019). A Unified Framework of Five Principles for AI in Society. *Harvard Data Science Review*, 1–15. doi: doi:10.1162/99608f92.8cd550d1
- Floridi, L., Cows, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., . . . Vayena, E. (2018). AI4People—An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations. *Minds and Machines*, 28(4), 689–707. Retrieved from <https://doi.org/10.1007/s11023-018-9482-5> (Publisher: Springer Netherlands) doi: doi:10.1007/s11023-018-9482-5
- Fokkinga, S., Desmet, P., & Hekkert, P. (2020). Impact-centered design: Introducing an integrated framework of the psychological and behavioral effects of design. *International Journal of Design*, 14(3), 97.
- Frijters, P., & Krekel, C. (2021). *A Handbook for Wellbeing Policy-making: History, Theory, Measurement, Implementation, and Examples*. Oxford University Press. (Google-Books-ID: iekrEAAAQBAJ)
- Future of Life Institute. (2023). *Pause Giant AI Experiments: An Open Letter*. Retrieved 2023-12-02, from <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>
- Gabriel, I. (2020). Artificial Intelligence, Values, and Alignment. *Minds and Machines*, 30(3), 411–437. Retrieved from <https://doi.org/10.1007/s11023-020-09539-2> (Publisher: Springer Netherlands) doi: doi:10.1007/s11023-020-09539-2
- Gaggioli, A., Riva, G., Peters, D., & Calvo, R. A. (2017). Positive Technology, Computing, and Design: Shaping a Future in Which Technology Promotes Psychological Well-Being. *Emotions and Affect in Human Factors and Human-Computer Interaction*(April), 477–502. doi: doi:10.1016/B978-0-12-801851-4.00018-5
- Gates, B. (2023, March). The age of AI has begun. *Gates Notes*. Retrieved 2023-07-25, from <https://www.gatesnotes.com/The-Age-of-AI-Has-Begun>
- Hadfield-Menell, D., & Hadfield, G. K. (2019, January). Incomplete Contracting and AI Alignment. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 417–422). Honolulu HI USA: ACM. Retrieved 2023-01-11, from <https://dl.acm.org/doi/10.1145/3306618.3314250> doi: doi:10.1145/3306618.3314250
- Hailey, J. (2008). Ubuntu: A literature review. *Document*. London: Tutu Foundation.
- Harbers, M., & Overdiek, A. (2022). Towards a living lab for responsible applied AI. In D. Lockton, S. Lenzi, P. Hekkert, A. Oak, J. Sádaba, & P. Lloyd (Eds.), *DRS2022: Bilbao, 25 June - 3 July*. Bilbao, Spain. doi: doi:https://doi.org/10.21606/drs.2022.422
- Harris, S. (2010). *The Moral Landscape: How Science Can Determine Human Values*. Simon and Schuster. (Google-Books-ID: VttdxFt4kT4C)
- Harris, T. (2017, Jul). *Technology is downgrading humanity; let’s reverse that trend now*. Retrieved from <https://www.commerce.senate.gov/services/files/96E3A739-DC8D-45F1-87D7-EC70A368371D>
- Havrda, M., & Klocek, A. (2023, April). Well-being impact assessment of artificial intelligence - A search for causality and proposal for an open platform for well-being impact assessment of AI systems. *Evaluation and Program Planning*, 99, 102294. doi: doi:10.1016/j.evalprogplan.2023.102294
- Havrda, M., & Rakova, B. (2020). Enhanced well-being assessment as basis for the practical implementation of ethical and rights-based normative principles for ai. In *2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)* (pp. 2754–2761).

- Hemphill, T. A. (2019, May). ‘Techlash’, responsible innovation, and the self-regulatory organization. *Journal of Responsible Innovation*, 6(2), 240–247. Retrieved 2022-05-04, from <https://www.tandfonline.com/doi/full/10.1080/23299460.2019.1602817> doi: doi:10.1080/23299460.2019.1602817
- Hjetland, G. J., Schønning, V., Aasan, B. E. V., Hella, R. T., & Skogen, J. C. (2021, January). Pupils’ Use of Social Media and Its Relation to Mental Health from a School Personnel Perspective: A Preliminary Qualitative Study. *International Journal of Environmental Research and Public Health*, 18(17), 9163. Retrieved 2023-12-24, from <https://www.mdpi.com/1660-4601/18/17/9163> (Number: 17 Publisher: Multidisciplinary Digital Publishing Institute) doi: doi:10.3390/ijerph18179163
- IEEE SA. (2020). *IEEE 7010-2020 - IEEE Recommended Practice for Assessing the Impact of Autonomous and Intelligent Systems on Human Well-Being*. Retrieved from <https://standards.ieee.org/standard/7010-2020.html>
- Johannes, N., Dienlin, T., Bakhshi, H., & Przybylski, A. K. (2022, January). No effect of different types of media on well-being. *Scientific Reports*, 12(1), 61. Retrieved 2023-01-31, from <https://www.nature.com/articles/s41598-021-03218-7> (Number: 1 Publisher: Nature Publishing Group) doi: doi:10.1038/s41598-021-03218-7
- Kissinger, H. A., Schmidt, E., & Huttenlocher, D. (2021). *The age of ai: and our human future*. Hachette UK.
- Kjell, O. N. (2011). Sustainable Well-Being: A Potential Synergy Between Sustainability and Well-Being Research. *Review of General Psychology*, 15(3), 255–266. doi: doi:10.1037/a0024603
- Klenk, M., & Duijf, H. (2021, November). Ethics of digital contact tracing and COVID-19: who is (not) free to go? *Ethics and Information Technology*, 23(1), 69–77. Retrieved 2022-06-02, from <https://doi.org/10.1007/s10676-020-09544-0> doi: doi:10.1007/s10676-020-09544-0
- Konu, A., & Rimpelä, M. (2002, March). Well-being in schools: a conceptual model. *Health Promotion International*, 17(1), 79–87. Retrieved 2023-12-23, from <https://doi.org/10.1093/heapro/17.1.79> doi: doi:10.1093/heapro/17.1.79
- Kramer, J., Agogino, A. M., & Roschuni, C. (2016, August). Characterizing Competencies for Human-Centered Design. In *Volume 7: 28th International Conference on Design Theory and Methodology* (p. V007T06A026). Charlotte, North Carolina, USA: American Society of Mechanical Engineers. Retrieved 2023-12-22, from <https://asmedigitalcollection.asme.org/IDETC-CIE/proceedings/IDETC-CIE2016/50190/Charlotte,%20North%20Carolina,%20USA/258472> doi: doi:10.1115/DETC2016-60085
- Krippendorff, K. (2021, October). *From Uncritical Design to Critical Examinations of its Systemic Consequences*. Retrieved 2022-07-14, from <https://rsdsymposium.org/professor-dr-klaus-krippendorff/>
- Kross, E., Verduyn, P., Sheppes, G., Costello, C. K., Jonides, J., & Ybarra, O. (2021). Social Media and Well-Being: Pitfalls, Progress, and Next Steps. *Trends in Cognitive Sciences*, 25(1), 55–66. Retrieved from <https://doi.org/10.1016/j.tics.2020.10.005> (Publisher: The Authors) doi: doi:10.1016/j.tics.2020.10.005
- Legg, S., & Hutter, M. (2007, June). *A Collection of Definitions of Intelligence*. arXiv. Retrieved 2022-07-12, from <http://arxiv.org/abs/0706.3639> (arXiv:0706.3639 [cs])
- Li, F., & Lu, Y. (2021). Engaging end users in an ai-enabled smart service design-the application of the smart service blueprint scape (ssbs) framework. *Proceedings of the Design Society, 1*, 1363–1372.
- Linton, M. J., Dieppe, P., & Medina-Lara, A. (2016). Review of 99 self-report measures for assessing well-being in adults: Exploring dimensions of well-being and developments over time. *BMJ Open*, 6(7). doi: doi:10.1136/bmjopen-2015-010641
- Liscio, E., Lera-Leri, R., Bistaffa, F., Dobbe, R. I., Jonker, C. M., Lopez-Sanchez, M., ... Murukannaiah, P. K. (2023). Value inference in sociotechnical systems. In *Proceedings of the 2023 international conference on autonomous agents and multiagent systems* (pp. 1774–1780).
- Liscio, E., van der Meer, M., Siebert, L. C., Jonker, C. M., Mouter, N., & Murukannaiah, P. K. (2021). Axies : Identifying and Evaluating Context-Specific Values. *Aamas 2021*, 799–808.
- Loveridge, R., Sallu, S. M., Pasha, I. J., & R Marshall, A. (2020, December). Measuring human well-being: A protocol for selecting local indicators. *Environmental Science & Policy*, 114, 461–469. Retrieved 2023-04-03, from <https://www.sciencedirect.com/science/article/pii/S146290112030112X> doi: doi:10.1016/j.envsci.2020.09.002
- Lu, F., Dumitrache, A., & Graus, D. (2020). Beyond optimizing for clicks: Incorporating editorial values in news recommendation. In *Proceedings of the 28th acm conference on user modeling, adaptation and personalization* (pp. 145–153).
- Lukoff, K., Lyngs, U., Shirokova, K., Rao, R., Tian, L., Zade, H., ... Hiniker, A. (2023). Switchtube: A proof-of-concept system introducing “adaptable commitment interfaces” as a tool for digital wellbeing. In *Proceedings of the 2023 chi conference on human factors in computing systems* (pp. 1–22).
- Lynch, S. (2023, April). *AI Benchmarks Hit Saturation*. Stanford Institute for Human-Centered AI. Retrieved from <https://hai.stanford.edu/news/ai-benchmarks-hit-saturation> (Accessed: 2024-01-15)
- McGregor, J. A., Camfield, L., & Coulthard, S. (2015). Competing interpretations: human wellbeing and the use of

- quantitative and qualitative methods. *Mixed methods research in poverty and vulnerability: sharing ideas and learning lessons*, 231–260.
- McShane, T. O., Hirsch, P. D., Trung, T. C., Songorwa, A. N., Kinzig, A., Monteferri, B., . . . O'Connor, S. (2011, March). Hard choices: Making trade-offs between biodiversity conservation and human well-being. *Biological Conservation*, *144*(3), 966–972. Retrieved 2024-01-03, from <https://www.sciencedirect.com/science/article/pii/S0006320710001849> doi: doi:10.1016/j.biocon.2010.04.038
- Mead, J., Fisher, Z., & Kemp, A. H. (2021). Moving Beyond Disciplinary Silos Towards a Transdisciplinary Model of Wellbeing: An Invited Review. *Frontiers in Psychology*, *12*. Retrieved 2023-12-18, from <https://www.frontiersin.org/articles/10.3389/fpsyg.2021.642093>
- Mittelstadt, B. (2019). Principles alone cannot guarantee ethical ai. *Nature machine intelligence*, *1*(11), 501–507.
- Mollick, E. (2023, February). *Meet the teacher requiring students to use ChatGPT*. Retrieved 2024-01-17, from <https://www.thirteen.org/metrofocus/2023/02/artificial-intelligence-school-scwlv1/>
- Moons, P., Budts, W., & De Geest, S. (2006, September). Critique on the conceptualisation of quality of life: A review and evaluation of different conceptual approaches. *International Journal of Nursing Studies*, *43*(7), 891–901. Retrieved 2022-12-12, from <https://www.sciencedirect.com/science/article/pii/S0020748906001088> doi: doi:10.1016/j.ijnurstu.2006.03.015
- Morley, J., Elhalal, A., Garcia, F., Kinsey, L., Mökander, J., & Floridi, L. (2021, June). Ethics as a Service: A Pragmatic Operationalisation of AI Ethics. *Minds and Machines*, *31*(2), 239–256. Retrieved 2023-06-05, from <https://doi.org/10.1007/s11023-021-09563-w> doi: doi:10.1007/s11023-021-09563-w
- Morley, J., Floridi, L., Kinsey, L., & Elhalal, A. (2020). From What to How: An Initial Review of Publicly Available AI Ethics Tools, Methods and Research to Translate Principles into Practices. *Science and Engineering Ethics*, *26*(4), 2141–2168. Retrieved from <https://doi.org/10.1007/s11948-019-00165-5> (Publisher: Springer Netherlands) doi: doi:10.1007/s11948-019-00165-5
- Morley, J., Kinsey, L., Elhalal, A., Garcia, F., Ziosi, M., & Floridi, L. (2021). Operationalising AI ethics: barriers, enablers and next steps. *AI & SOCIETY*, *38*(1), 411–423. Retrieved 2023-06-03, from <https://doi.org/10.1007/s00146-021-01308-8> doi: doi:10.1007/s00146-021-01308-8
- Musikanski, L., Rakova, B., Bradbury, J., Phillips, R., & Manson, M. (2020). Artificial Intelligence and Community Well-being: A Proposal for an Emerging Area of Research. *International Journal of Community Well-Being*, *3*(1), 39–55. (Publisher: International Journal of Community Well-Being) doi: doi:10.1007/s42413-019-00054-6
- Ng, A. Y., & Russell, S. J. (2000). Algorithms for inverse reinforcement learning. In *Proceedings of the seventeenth international conference on machine learning* (p. 663–670). San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- Norman, D. (2013). *The design of everyday things*. MIT Press, Cambridge.
- O'Mahony, T. (2022). Toward Sustainable Wellbeing: Advances in Contemporary Concepts. *Frontiers in Sustainability*, *3*. Retrieved 2024-01-04, from <https://www.frontiersin.org/articles/10.3389/frsus.2022.807984>
- Organization, W. H. (n.d.). *World mental health report: Transforming mental health for all* (Tech. Rep.). Retrieved 2023-04-04, from <https://www.who.int/publications-detail-redirect/9789240049338>
- Ozmen Garibay, O., Winslow, B., Andolina, S., Antona, M., Bodenschatz, A., Coursaris, C., . . . Xu, W. (2023, February). Six Human-Centered Artificial Intelligence Grand Challenges. *International Journal of Human-Computer Interaction*, *39*(3), 391–437. Retrieved 2023-01-10, from <https://doi.org/10.1080/10447318.2022.2153320> (Publisher: Taylor & Francis eprint: <https://doi.org/10.1080/10447318.2022.2153320>) doi: doi:10.1080/10447318.2022.2153320
- Pan, A., Bhatia, K., & Steinhardt, J. (2022). The effects of reward misspecification: Mapping and mitigating misaligned models. In *International conference on learning representations*. Retrieved from <https://openreview.net/forum?id=JYtwGwIL7ye>
- Pavot, W. (2014). Temporal Satisfaction with Life Scale (TSWLS). In A. C. Michalos (Ed.), *Encyclopedia of Quality of Life and Well-Being Research* (pp. 6609–6611). Dordrecht: Springer Netherlands. Retrieved 2023-03-03, from [https://doi.org/10.1007/978-94-007-0753-5\\_2993](https://doi.org/10.1007/978-94-007-0753-5_2993) doi: doi:10.1007/978-94-007-0753-5\_2993
- Prabhakaran, V., Mitchell, M., Gebru, T., & Gabriel, I. (2022, October). *A Human Rights-Based Approach to Responsible AI*. arXiv. Retrieved 2023-12-20, from <http://arxiv.org/abs/2210.02667> (arXiv:2210.02667 [cs])
- Qasem, F. (2023, January). ChatGPT in scientific and academic research: future fears and reassurances. *Library Hi Tech News*, *40*(3), 30–32. Retrieved 2023-12-20, from <https://doi.org/10.1108/LHTN-03-2023-0043> (Publisher: Emerald Publishing Limited) doi: doi:10.1108/LHTN-03-2023-0043
- Ruggeri, K., Garcia-Garzon, E., Maguire, Á., Matz, S., & Huppert, F. A. (2020). Well-being is more than happiness and life satisfaction: a multidimensional analysis of 21 countries. *Health and quality of life outcomes*, *18*(1), 1–16.
- Russell, S. J., & Norvig, P. (2022). *Artificial intelligence: a modern approach* (Fourth edition, global edition ed.). Harlow: Pearson.
- Ryan, R. M., & Deci, E. L. (2001). On happiness and human potentials: A review of research on hedonic and

- eudaimonic well-being. *Annual Review of Psychology*, 52, 141–166. doi: doi:10.1146/annurev.psych.52.1.141
- Ryff, C. D. (1989). Happiness is everything, or is it? Explorations on the meaning of psychological well-being. *Journal of Personality and Social Psychology*, 57(6), 1069–1081. doi: doi:10.1037/0022-3514.57.6.1069
- Ryff, C. D., & Keyes, C. L. M. (1995). The Structure of Psychological Well-Being Revisited. *Journal of Personality and Social Psychology*, 69(4), 719–727. doi: doi:10.1037/0022-3514.69.4.719
- Sadek, M., Calvo, R. A., & Mougénot, C. (2023a). Co-designing conversational agents: A comprehensive review and recommendations for best practices. *Design Studies*, 89, 101230.
- Sadek, M., Calvo, R. A., & Mougénot, C. (2023b). Designing value-sensitive ai: a critical review and recommendations for socio-technical design processes. *AI and Ethics*, 1–19.
- Sanders, E. B.-N., & Stappers, P. J. (2008). Co-creation and the new landscapes of design. *Co-design*, 4(1), 5–18.
- Sanderson, C., Douglas, D., Lu, Q., Schleiger, E., Whittle, J., Lacey, J., ... Hansen, D. (2023). Ai ethics principles in practice: Perspectives of designers and developers. *IEEE Transactions on Technology and Society*.
- Sato, K. (1991). From ai to cybernetics. *AI & society*, 5, 155–161. Retrieved from <https://doi.org/10.1007/BF01891721>
- Schiff, D., Ayesh, A., Musikanski, L., & Havens, J. C. (2020). IEEE 7010: A New Standard for Assessing the Well-being Implications of Artificial Intelligence. *Conference Proceedings - IEEE International Conference on Systems, Man and Cybernetics, 2020-October*(March), 2746–2753. doi: doi:10.1109/SMC42975.2020.9283454
- Scott, A. J., Webb, T. L., Martyn-St James, M., Rowse, G., & Weich, S. (2021, December). Improving sleep quality leads to better mental health: A meta-analysis of randomised controlled trials. *Sleep Medicine Reviews*, 60, 101556. doi: doi:10.1016/j.smrv.2021.101556
- Searle, B. A., Pykett, J., & Alfaro-Simmonds, M. J. (2021, June). Introduction to wellbeing research. In *A Modern Guide to Wellbeing Research* (pp. 1–21). Edward Elgar Publishing. Retrieved 2023-12-23, from <https://china.elgaronline.com/edcollchap/edcoll/9781789900156/9781789900156.00008.xml> (Section: A Modern Guide to Wellbeing Research)
- Seligman, M. E. (2019). Positive psychology: A personal history. *Annual review of clinical psychology*, 15, 1–23.
- Seligman, M. E. P. (2010). Flourish: Positive Psychology and Positive Interventions. *The Tanner Lectures on Human Values*, 231–242.
- Shahriari, K., & Shahriari, M. (2017). IEEE standard review - Ethically aligned design: A vision for prioritizing human wellbeing with artificial intelligence and autonomous systems. *IHTC 2017 - IEEE Canada International Humanitarian Technology Conference 2017*, 197–201. doi: doi:10.1109/IHTC.2017.8058187
- Sone, T., Nakaya, N., Ohmori, K., Shimazu, T., Higashiguchi, M., Kakizaki, M., ... Tsuji, I. (2008). Sense of life worth living (ikigai) and mortality in japan: Ohsaki study. *Psychosomatic medicine*, 70(6), 709–715.
- Stephen, D. (2021, July). *Pace layers in experience design — Stabilise innovation by understanding people's needs*. Retrieved 2023-02-07, from <https://duncanstephen.net/pace-layers-in-experience-design-stabilise-innovation-by-understanding-peoples-needs/> (Section: Design)
- Steur, A. J., & Seiter, M. (2021, May). Properties of feedback mechanisms on digital platforms: an exploratory study. *Journal of Business Economics*, 91(4), 479–526. Retrieved 2023-12-13, from <https://doi.org/10.1007/s11573-020-01009-6> doi: doi:10.1007/s11573-020-01009-6
- Stray, J. (2020). Aligning AI Optimization to Community Well-Being. *International Journal of Community Well-Being*, 3(4), 443–463. (Publisher: International Journal of Community Well-Being) doi: doi:10.1007/s42413-020-00086-3
- Stray, J., & Hadfield, G. K. (n.d.). Platforms Can Optimize for Metrics Beyond Engagement. *Wired*. Retrieved 2023-03-03, from <https://www.wired.com/story/platforms-engagement-research-meta/> (Section: tags)
- Stray, J., Halevy, A., Assar, P., Hadfield-Menell, D., Boutilier, C., Ashar, A., ... Vasan, N. (2023, November). Building Human Values into Recommender Systems: An Interdisciplinary Synthesis. *ACM Transactions on Recommender Systems*, 3632297. Retrieved 2023-12-05, from <https://dl.acm.org/doi/10.1145/3632297> doi: doi:10.1145/3632297
- Tabari, P. (2022, September). The Role of Artificial Intelligence in Human-Computer Interaction: Using a Smart Topic Extraction System. In *2022 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)* (pp. 1–3). (ISSN: 1943-6106) doi: doi:10.1109/VL/HCC53370.2022.9833114
- Technicalities, & Stag. (2023, November). *Shallow review of live agendas in alignment & safety — AI Alignment Forum*. Retrieved 2023-12-22, from <https://www.alignmentforum.org/posts/zaaGsFBEDTpCsYHeF/shallow-review-of-live-agendas-in-alignment-and-safety>
- Thomas, R. L., & Uminsky, D. (2020). Reliance on Metrics is a Fundamental Challenge for AI. *Ethics of Data Science Conference*. Retrieved from <https://arxiv.org/abs/2002.08512>
- Thorburn, L. (2022, August). *How to Measure the Causal Effects of Recommenders*. Retrieved 2022-08-24, from <https://medium.com/understanding-recommenders/how-to-measure-the-causal-effects-of-recommenders-5e89b7363d57>

- Trunfio, M., & Rossi, S. (2021, September). Conceptualising and measuring social media engagement: A systematic literature review. *Italian Journal of Marketing*, 2021(3), 267–292. Retrieved 2023-03-03, from <https://doi.org/10.1007/s43039-021-00035-8> doi: doi:10.1007/s43039-021-00035-8
- Ulanoff, L. (2023, November). *OpenAI's reported 'superintelligence' breakthrough is so big it nearly destroyed the company, and ChatGPT*. Retrieved 2023-12-12, from <https://www.techradar.com/computing/artificial-intelligence/openais-reported-superintelligence-breakthrough-is-so-big-it-nearly-destroyed-the-company-and-chatgpt>
- Umbrello, S., & van de Poel, I. (2021, August). Mapping value sensitive design onto AI for social good principles. *AI and Ethics*, 1(3), 283–296. Retrieved 2022-05-11, from <https://link.springer.com/10.1007/s43681-021-00038-3> doi: doi:10.1007/s43681-021-00038-3
- Ura, K., Alkire, S., & Zangmo, T. (2012). GNH and GNH Index. *The Centre for Bhutan Studies*(May), 1–60.
- van Allen, P. (2018, oct). Prototyping ways of prototyping ai. *Interactions*, 25(6), 46–51. Retrieved from <https://doi.org/10.1145/3274566> doi: doi:10.1145/3274566
- Vanden Abeele, M. M. P. (2021, November). Digital Wellbeing as a Dynamic Construct. *Communication Theory*, 31(4), 932–955. Retrieved 2022-09-07, from <https://doi.org/10.1093/ct/qtaa024> doi: doi:10.1093/ct/qtaa024
- van de Poel, I. (2020, September). Embedding Values in Artificial Intelligence (AI) Systems. *Minds and Machines*, 30(3), 385–409. Retrieved 2022-12-02, from <https://doi.org/10.1007/s11023-020-09537-4> doi: doi:10.1007/s11023-020-09537-4
- Van der Maden, W., Lomas, D., & Hekkert, P. (2022). Design for wellbeing during covid-19: A cybernetic perspective on data feedback loops in complex sociotechnical systems. In D. Lockton, S. Lenzi, P. Hekkert, A. Oak, J. Sádaba, & P. Lloyd (Eds.), *Drs2022: Bilbao*. Bilbao, Spain. doi: doi:10.21606/drs.2022.771
- van der Maden, W., Lomas, D., & Hekkert, P. (2023). A framework for designing AI systems that support community wellbeing. *Frontiers in Psychology*, 13. Retrieved 2023-01-15, from <https://www.frontiersin.org/articles/10.3389/fpsyg.2022.1011883>
- Veenhoven, R. (2014). Journal of Happiness Studies. *Encyclopedia of Quality of Life and Well-Being Research*, 3462–3464. doi: doi:10.1007/978-94-007-0753-5\_4049
- Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., . . . Fuso Nerini, F. (2020, January). The role of artificial intelligence in achieving the Sustainable Development Goals. *Nature Communications*, 11(1), 233. Retrieved 2024-01-04, from <https://www.nature.com/articles/s41467-019-14108-y> (Number: 1 Publisher: Nature Publishing Group) doi: doi:10.1038/s41467-019-14108-y
- Virokannas, E., Liuski, S., & Kuronen, M. (2020, March). The contested concept of vulnerability – a literature review: Vulnerability-käsitteen kiistanalaiset merkitykset – systemaattinen kirjallisuuskatsaus. *European Journal of Social Work*, 23(2), 327–339. Retrieved 2022-08-31, from <https://www.tandfonline.com/doi/full/10.1080/13691457.2018.1508001> doi: doi:10.1080/13691457.2018.1508001
- Watson, D., Clark, L. A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: The PANAS scales. *Journal of Personality and Social Psychology*, 54(6), 1063–1070. (Place: US Publisher: American Psychological Association) doi: doi:10.1037/0022-3514.54.6.1063
- West, S., Whittaker, M., & Crawford, K. (2019). *Discriminating Systems: Gender, Race, and Power* (Tech. Rep.). AI Now Institute.
- Wiblin, R. (n.d.). *Jan Leike on OpenAI's massive push to make superintelligence safe in 4 years or less*. Retrieved 2023-12-20, from <https://80000hours.org/podcast/episodes/jan-leike-superalignment/>
- Wong, R., Madaio, M., & Merrill, N. (2022). Seeing like a toolkit: How toolkits envision the work of ai ethics. *Computing Research Repository (CoRR)*. Retrieved from <https://arxiv.org/abs/2202.08792>
- Yang, Q., Steinfeld, A., Rosé, C., & Zimmerman, J. (2020). Re-examining whether, why, and how human-ai interaction is uniquely difficult to design. In *Proceedings of the 2020 chi conference on human factors in computing systems* (pp. 1–13).
- Yu, B., Yuan, Y., Terveen, L., Wu, Z. S., Forlizzi, J., & Zhu, H. (2020). Keeping designers in the loop: Communicating inherent algorithmic trade-offs across multiple objectives. In *Proceedings of the 2020 acm designing interactive systems conference* (pp. 1245–1257).
- Zhang, A., Boltz, A., Lynn, J., Wang, C.-W., & Lee, M. K. (2023). Stakeholder-centered ai design: Co-designing worker tools with gig workers through data probes. In *Proceedings of the 2023 chi conference on human factors in computing systems* (pp. 1–19).
- Zhu, J., Liapis, A., Risi, S., Bidarra, R., & Youngblood, G. M. (2018). Explainable ai for designers: A human-centered perspective on mixed-initiative co-creation. In *2018 ieee conference on computational intelligence and games (cig)* (pp. 1–8).
- Zuckerberg, M. (2018, December). *Facebook Post* [Social Media]. Retrieved 2023-09-27, from <https://www.facebook.com/zuck/posts/one-of-our-big-focus-areas-for-2018-is-making-sure-the-time-we-all-spend-on-face/10104413015393571/>