


Responsible Use of Artificial Intelligence as Continuous Proportionalization: Fashion Image Generation at OTTO



Simon Sturm, Florian Krause, Benjamin van Giffen

Abstract: The inherent ambivalence of machine learning-based artificial intelligence (AI) technologies makes ensuring their responsible use a pressing concern. While the literature converges on the importance of governance at the organizational level, uncertainty remains about what responsible (use of) AI actually is. We conceptualize responsible AI as both a process and outcome of social evaluation. We propose a model (“continuous proportionalization”) that explains how organizations construct collective interpretations of responsible AI along the dimensions of legitimacy, suitability, necessity, and proportionality. We illustrate the model through a case study of AI-based fashion image generation at Germany’s largest e-commerce company, OTTO.



Keywords: Artificial Intelligence, Responsible AI, AI Ethics, Image Generation, E-Business, Technology Management, Business Ethics, Corporate Governance

Verantwortungsvolle Nutzung Künstlicher Intelligenz als kontinuierliche Proportionalisierung: Fashion-Bildgenerierung bei OTTO



Zusammenfassung: Die inhärente Ambivalenz und die zunehmende Nutzung künstlicher Intelligenz (KI) werfen drängende normative Fragen auf. Die Literatur ist sich einig, dass verantwortungsvolle KI-Nutzung Governance auf der organisationalen Ebene erfordert. Uneinigkeit besteht jedoch darüber, was verantwortungsvolle KI-Nutzung ist und wie sie entsteht. Dieser Artikel betrachtet verantwortungsvolle KI-Nutzung als Prozess und Ergebnis sozialer Evaluation. Er entwickelt ein konzeptionelles Modell („kontinuierliche Proportionalisierung“), das erklärt, wie Organisationen kollektive Interpretationen verantwortungsvoller KI-Nutzung entlang der Dimensionen Legitimität, Geeignetheit, Notwendigkeit und Verhältnis-

mässigkeit konstruieren. Zur Veranschaulichung wird das Modell auf KI-basierte Fashion-Bildgenerierung bei Deutschlands führendem E-Commerce-Unternehmen OTTO angewendet.

Stichwörter: Künstliche Intelligenz, Verantwortungsvolle KI, KI Ethik, Bildgenerierung, E-Business, Technologiemanagement, Unternehmensethik, Corporate Governance

1 Introduction

Advances in machine learning-based artificial intelligence (AI) create new opportunities and risks for humanity (Coeckelbergh, 2020; Taddeo & Floridi, 2018). For example, the same AI technologies that could improve access to mental healthcare (Zhang & Wang, 2024) have also been reported to encourage violent behavior (Gerken, 2024; Kuznia et al., 2025). This ambivalence raises fundamental normative issues (Coeckelbergh, 2020; Mikalef et al., 2022). Consequently, *responsible* (use of) AI has become a central priority for researchers and practitioners across disciplines (Floridi et al., 2018; Garibay et al., 2023; Ryan & Stahl, 2021).

On the one hand, the existing literature agrees that organizations are crucial to ensuring responsible use of AI (Cihon et al., 2021). On the other hand, there is no consensus on what responsible use of AI actually *is*. Rather, the existing body of knowledge seems to presuppose different conceptualizations of responsibility that share structural similarity with the research on the legitimacy of organizational conduct (Suddaby et al., 2017). One stream of research tends to view responsible use of AI as a *property* that can be specified through principles, requirements, or safeguards and assessed against predefined criteria (Bughin, 2025a, 2025b; Heger et al., 2025; Krijger et al., 2023; Minkkinen et al., 2023). Another stream of research treats the phenomenon as a *process*, suggesting that the responsible use of AI emerges from context-sensitive interpretation and negotiation (Elia et al., 2025; Hagendorff, 2022; Kallina & Singh, 2024; Mittelstadt, 2019; Yilma, 2025).

The issue is that these perspectives rest on conflicting assumptions about the nature of responsibility. Property-oriented views assume that responsibility can be specified independently of situated interpretation, whereas process-oriented views emphasize that responsibility comes into being through such interpretation. Choosing one perspective over the other would simplify the debate, but at the cost of discarding insights that capture important aspects of how responsibility is understood and enacted in organizational practice.

This study addresses this problem by conceptualizing responsible use of AI through the lens of social evaluation as an integrative perspective. Specifically, we combine Bitektine and Haack's (2015) process theory of legitimacy with the principle of proportionality (Karliuk, 2023) to theorize responsible use of AI as both process and outcome of social evaluation across individual and collective levels. We propose a conceptual model ("continuous proportionalization") that captures how organizations form, enact, and revise collective interpretations of responsible use of AI through the convergence of individual propriety judgments along four discursive dimensions: legitimacy, suitability, necessity, and proportionality. For illustration, we apply our model to a case study of fashion image generation at Germany's largest e-commerce company, OTTO.

The remainder of this paper is structured as follows. First, we situate our study in the literature on responsible use of AI and motivate social evaluation as an integrative perspective. We then conceptualize responsible use of AI as continuous proportionalization. Next, we apply the model to the OTTO case for illustrative purposes. Finally, we discuss the contributions, limitations, and implications of this research and offer concluding remarks.

2 Conceptual Background

2.1 Property and Process Perspectives on Responsible Use of AI

Research on responsible use of AI has expanded rapidly in response to the diffusion of AI technologies across organizational contexts (Bach et al., 2025). Private companies, public organizations, and research institutions have issued a growing number of principles, guidelines, and governance frameworks aimed at ensuring ethical or responsible AI use (Jobin et al., 2019). Despite this proliferation of guidance, considerable uncertainty remains about what responsible use of AI means in organizational practice.

As outlined in the introduction, the literature tends to mirror broader debates on the legitimacy of organizational conduct (Suddaby et al., 2017), advancing either property or process perspectives on responsible use of AI.

From a property perspective, responsible use of AI is treated as an attribute of technologies, decisions, or organizational arrangements (Bughin, 2025a, 2025b; Heger et al., 2025; Krijger et al., 2023; Minkkinen et al., 2023). AI use is considered responsible if it satisfies predefined criteria, such as compliance with legal requirements, adherence to ethical principles, or the implementation of technical and organizational safeguards. Responsibility thus appears as something that can be specified *ex ante* and assessed *ex post*. This understanding underpins many regulatory and governance approaches that seek to classify AI applications according to their risks and acceptable uses. In this view, responsibility functions as an evaluative label that can be attached to organizational conduct.

In contrast, a process perspective conceptualizes responsible use of AI as something that emerges through ongoing interpretation, deliberation, and justification (Elia et al., 2025; Hagendorff, 2022; Kallina & Singh, 2024; Mittelstadt, 2019; Yilma, 2025). Rather than being a stable property, responsibility is enacted in practice as organizations continuously negotiate what responsible AI use means in light of evolving technologies, shifting regulations, and changing societal expectations. From this perspective, responsibility cannot be fully specified in advance but remains context-sensitive and provisional.

Importantly, these perspectives on the nature of responsible use of AI coexist in tension. Organizations are expected to demonstrate responsibility by pointing to concrete properties (e.g. compliance documentation, safeguards, or compliance structures). Yet, these properties tend to remain incomplete, contested, or insufficient (Mittelstadt, 2019). Especially where established standards lag behind technological development, responsibility cannot be conclusively established through predefined criteria alone (Gogoll et al., 2021). Instead, it must be continuously constructed, justified, and stabilized through processes of deliberation and evaluation (Coeckelbergh, 2024; Gogoll et al., 2021; Watson et al., 2025).

This tension is particularly pronounced in the context of AI, where technological change outpaces the stabilization of normative expectations (Coeckelbergh, 2024; Floridi et al., 2018; Garibay et al., 2023). As a result, responsible use of AI cannot be reduced either to fixed properties or to open-ended processes alone. What remains underexplored is how property and process characteristics are connected in practice. Therefore, we turn to social evaluation as an integrative perspective on the responsible use of AI.

2.2 Social Evaluation as Integrative Perspective

In this paper, we draw on legitimacy theory (Bitektine & Haack, 2015; Suchman, 1995; Suddaby et al., 2017) to integrate the property and process characteristics of responsible use of AI. In legitimacy research, organizational conduct is evaluated not solely based on intrinsic qualities, but from the perspective of relevant audiences who assess whether actions are acceptable, appropriate, and justifiable within a given social context (Suchman, 1995; Suddaby et al., 2017). Legitimacy thus emerges through social evaluation.

Adopting an evaluator perspective (Bitektine & Haack, 2015) allows us to conceptualize responsible use of AI in a way that accommodates both process and property characteristics. Social evaluation foregrounds the processual dimension of responsibility by focusing on how judgments are formed, contested, and revised. At the same time, it explains how the outcomes of these evaluations stabilize responsibility in observable and communicable forms.

It is important to note that legitimacy and responsibility are not equivalent. Organizational conduct may be legitimate without necessarily being responsible in a stronger normative sense. For example, if it complies with regulations or aligns with prevailing norms. Responsibility entails an additional demand: actors must be able and willing to justify their actions by providing reasons that withstand critical scrutiny. The key question is therefore not merely whether AI use is accepted, but which reasons are considered valid in justifying it.

Responsibility thus becomes visible through deliberation, understood as the practice of articulating, weighing, and contesting reasons for action. Crucially, such deliberation is not arbitrary. It is structured by recognizable procedures, standards, and formats that define who evaluates AI use, which considerations are relevant, and how competing reasons are balanced. These structures shape how responsibility is constructed and stabilized in organizational contexts.

Social evaluation provides an integrative lens precisely because it captures this duality. Properties of responsible use of AI can be understood as provisional outcomes of evaluative processes, while these processes themselves are oriented toward producing evaluative outcomes that can guide action. In this sense, responsibility is neither fully given nor endlessly fluid; it is continuously produced, stabilized, and revised through social evaluation.

By adopting an evaluators perspective (Bitektine & Haack, 2015), we can therefore analyze responsible use of AI without reducing it to either fixed normative criteria or open-ended sensemaking processes. Instead, responsible use of AI appears as the outcome of structured social evaluation under conditions of uncertainty, through which organizations seek to render their AI use justifiable to relevant audiences.

In the next section, we build on this perspective to conceptualize responsible use of AI as both process and outcome of social evaluation.

3 Conceptual Model Development

3.1 Core Constructs and Initial Conceptualization

As mentioned above, we derive the core constructs of our conceptualization from the theoretical body on the legitimacy of organizational conduct (Suddaby et al., 2017). We do so because responsible use of AI, like legitimacy, involves a social evaluation of action.

The notion of responsible use of AI presupposes that the decision to (not) use AI is an action that requires justification, and justification requires an audience. Justification entails both the expectation and the obligation to show that an action is taken for good reason. Whether an action is taken for good reason depends on whether it can withstand the scrutiny of those to whom it must be justified. The social nature of action evaluation creates a structural similarity to the legitimacy (Suchman, 1995), which allows us to draw on legitimacy theory (Suddaby et al., 2017) to conceptualize responsible use of AI. While responsibility remains the normative point of reference for our conceptual work, legitimacy theory provides the framework.

We base our work on the multi-level theory of the legitimacy process by Bitektine and Haack (2015). The authors conceptualize the evaluation of organizational conduct as recursive interplay between *individual propriety* and *collective validity* judgments. We adopt several other core constructs from Bitektine and Haack (2015) that shape the recursive interplay between individual propriety and collective validity judgement. Which we summarize as follows.

Individual evaluators execute propriety judgments based on two perceptual inputs (Bitektine & Haack, 2015, p. 51). First, they draw on their *perceptions of organizational properties and behavior* (Bitektine & Haack, 2015, p. 51). Second, they draw on their perceptions of what the collective judgement is (*validity belief*) (Bitektine & Haack, 2015, p. 51). Individual evaluators express their propriety judgement through discourse and *action* (Bitektine & Haack, 2015, p. 53). As evaluators attend to the behavior of other evaluators (Bitektine & Haack, 2015, p. 51), their actions aggregate into collective effects, and individual propriety judgments converge into validity as a general consensus about the appropriateness of action (Bitektine & Haack, 2015, p. 51; Suddaby et al., 2017, p. 468). This general consensus becomes “institutionalized” (Bitektine & Haack, 2015, p. 53) and reenters the perceptions of individual evaluators as a validity belief, where it will shape and be shaped by future propriety judgment. *Judgment validation institutions* influence this cyclical process by arbitrating among conflicting evaluations and informing the validity beliefs held by individual evaluators (Bitektine & Haack, 2015, pp. 51–52).

Bitektine and Haack (2015) offer a general theory of normative evaluation that is designed to apply across a wide range of evaluative contexts. For the purpose of this research, we introduce several modifications to account for the specificities of responsible use of AI.

First, on the macro level of analysis (Suddaby et al., 2017), we distinguish *organization* and its *environment* as two evaluative spheres. These spheres are connected through observable organizational properties and behavior, which in our case concern the *responsible use of AI*. Individual *perception(s) of organizational AI use* therefore become one of two inputs that *inform individual propriety judgement* and subsequently affect *instance(s) of AI use*.

Second, we introduce *advancement(s) in AI technology* as an additional macro-level construct. This addition reflects that responsible use of AI unfolds within a sociotechnical environment in which normative concerns are continuously reshaped by rapid technological development (Mikalef et al., 2022; Stahl, 2012).

These specifications yield an initial conceptualization of responsible use of AI as process and result of social evaluation (Figure 1).

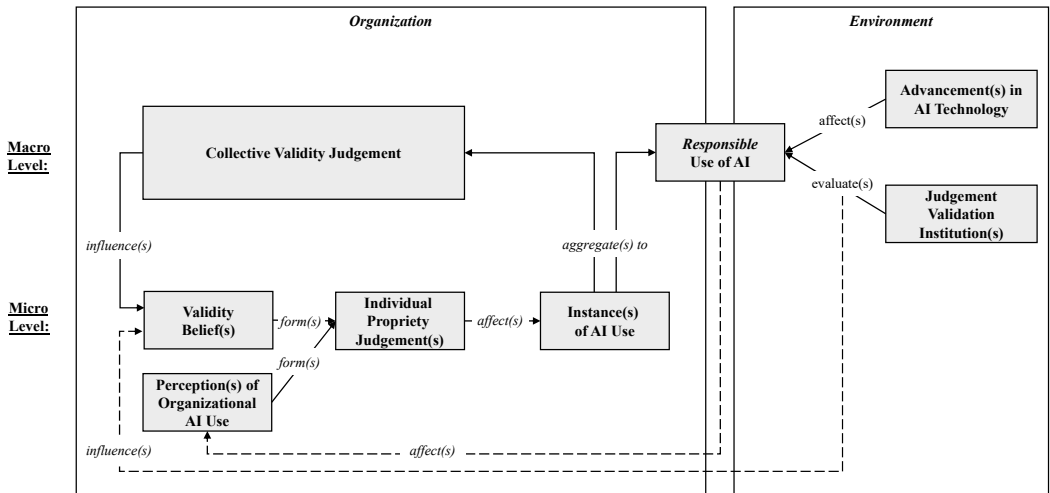


Figure 1: Initial Conceptualization of Responsible Use of AI as Social Evaluation

While this initial conceptualization clarifies how individual and collective evaluations interact and change over time, it leaves open an essential question: how do individual propriety judgments converge into a collective validity judgement? Bitektine and Haack (2015) provide the important insight that this aggregation happens discursively. Building on this, we adopt the principle of proportionality (Karliuk, 2023) to identify discursive dimensions through which individual propriety judgements translate into collective validity judgment of responsible use of AI.

3.2 Leveraging the Principle of Proportionality

The principle of proportionality is a long-standing heuristic for evaluating the validity of action. It originated in Prussian administrative law to assess the legitimacy of government intervention in economic and social affairs (Cohen-Eliya & Porat, 2010). It was based on the emerging understanding that citizens are legal entities with individual rights, and that state actions restricting these rights can only be carried out if (and to the extent to which) they are justified (Cohen-Eliya & Porat, 2010). To this end, Prussian court judges tested state action for *Legitimacy*, *Suitability*, *Necessity*, and *Proportionality* in the strict sense. Progressing through this sequence of tests (Figure 2) enabled them to determine whether state action is justified in the sense that its end(s) are legitimate and its mean(s) adequate (Cohen-Eliya & Porat, 2010; Sobek & Montag, 2018).




Figure 2: Proportionality Test Components

Today, proportionality is a cornerstone of constitutional and human rights law across the globe (Alexy, 2014). This expansion was possible, because the principle of proportionality and the proportionality test are sufficiently abstract. Conceptually, there is no reason not to apply it beyond the legal domain to evaluate actions that affect others.

In this spirit, Karliuk (2023) suggested reconceptualizing proportionality for AI ethics. Specifically, for decisions regarding “both to using AI as such and, if a decision is made to use it, to choose the right AI method” (Karliuk, 2023, p. 989). Applying the proportionality principle for the responsible use of AI means “addressing AI systems in a way that [...] their use do[es] not exceed what is necessary to achieve legitimate aims” (Karliuk, 2023, p. 988). It follows to consider the use of AI responsible if (and only to the extent to which) it is *suitable*, *necessary*, and *proportionate* to achieve a *legitimate* aim.

We agree that the proportionality test offers valuable “structural guidance [...] to reach [...] justifiable decision[s]” and could therefore “play an important role in AI ethics” (Karliuk, 2023, p. 987). Its usefulness for our work lies in the fact that it makes the basic structure of practical reasoning explicit. When people scrutinize the legitimacy of action, they intuitively ask questions like: Was it done for the right reasons? Was this the appropriate way to pursue the goal? Were the means and ends in balance? The proportionality principle captures these intuitions in a systematic way.

If we assume that the organizational use of AI is, like any other action, subject to justification, then these same questions naturally arise in deliberation about AI. The proportionality principle therefore provides a set of shared evaluative dimensions that are likely to surface in discursive exchanges, and through which evaluators compare and scrutinize their propriety judgments against those of others. If Bitektine and Haack (2015) are correct that collective validity judgement emerges discursively, then this convergence will plausibly be based on shared understandings of legitimacy, suitability, necessity, and proportionality.

In this sense, the proportionality principle offers a plausible account of how individual propriety judgments may converge into a collective validity judgement of responsible use of AI. We refer to this discursive convergence as *proportionalization*.

Proportionalization fills the conceptual gap in our initial conceptualization of responsible use of AI as social evaluation. We present our extended conceptual model in the next section.

3.3 Responsible Use of AI as Continuous Proportionalization

We conceptualize responsible use of AI in organizations as a process and result of recursive interplay between individual propriety and collective validity judgment in a dynamic, socio-technical environment (Figure 3).

propriety and collective validity judgment) can be traced in organizational practice. While capturing the full contextual range of continuous proportionalization would require longitudinal research, the present case allows us to observe key dynamics and evaluative dimensions in situ.

4.1 Case Study Research Method

In line with the illustrative function, we conducted a propositional single case study (Yin, 1994). We chose to study the use of AI at Germany's largest e-commerce company OTTO. We selected OTTO because its retail platform business relies heavily on AI (Christophersen & Pärn, 2021), and the company positions itself as a responsible AI user. As the CIO explained: *"We are very much determined by the values of our owner family. They have given us a Code of Ethics on how we want to act as a company [...]. One sub-category [...] is the responsible use of technology. [...] We are already making extensive use of artificial intelligence and [...] are giving this a lot of thought"* (Otto, 2021).

We decided to focus our case study on OTTO's AI use for fashion image generation. This is because fashion image generation promises significant business value for e-commerce companies like OTTO. However, fashion articles are typically displayed on human bodies, which raises normative concerns and subjects fashion image generation to a certain level of public scrutiny. As evidenced by the recent controversy regarding Guess's decision to use AI models in an international fashion campaign (Rufo, 2025). In addition, fashion image generation raises general concerns of intellectual property, bias, and environmental impact (Katirai et al., 2024). Hence, OTTO faces a difficult question: how to use AI in fashion image production responsibly?

Data Type	Σ	Duration
Key Informant Interviews	4	196 min
(ID-1) Digital Content Management Expert	1	49 min
(ID-2) Teamlead Content Operations	1	47 min
(ID-3) Digital Content Production Specialist	1	46 min
(ID-4) Product Owner CGI Production Service	1	54 min
Participant Observation	3	270 min
Internal AI Use Case Assessment Workshops	3	270 min
Documents & Archival Records	13	39 min
Internal Company Documents & Intranet Websites	12	-
Podcasts	1	39 min

Table 1. Case Data by Data Type

We collected data on OTTO's use of AI in fashion image generation from interviews, participant observation, documents & archival records (Table 1). We purposefully selected interviewees who spearheaded or managed AI initiatives in OTTO's image content production because they represented different functions along the fashion image production

and AI-use-case evaluation process (content operations, production, product ownership, and digital content management). All interviews were conducted via Microsoft Teams, recorded, and transcribed. They lasted between 46 and 54 minutes and were conducted in German. We used the speech recognition system “Whisper” for transcription (Radford et al., 2022) and corrected the output as needed. We also collected data by observing three workshops in which OTTO assessed potential AI use cases for image generation. In these workshops, we acted as non-intervening participant observers and did not influence discussions or decision-making processes. Observing the workshops allowed us to capture firsthand how discursive interactions between individual propriety judgments and collective validity judgments unfold. In particular, how participants articulated, contested, and aligned reasoning. We triangulated interview accounts and workshop observations with information from internal documents (e.g., PowerPoint presentations, online whiteboards) and archival records (e.g., intranet pages, podcasts).

For data analysis, we stored the case evidence from all sources in an online database (ATLAS.ti). We used qualitative methods and analyzed the data in four steps. First, we marked all case data that related to fashion image generation and responsible use of AI. Which resulted in over 120 quotes. Then, we coded these quotes deductively using the main constructs of our conceptual model following Yin (1981, 1994). This deductive coding allowed us to group case evidence from different sources under common categories (e.g. individual propriety judgement[s]). Next, we applied inductive coding to abstract and capture additional insights from the quotes. Finally, we combined the coding results and interpreted them using our case knowledge. Table 2 exemplifies how we marked, grouped, coded, and interpreted our data.

Data Source	Data Excerpt	Deductive Code(s)	Inductive Code(s)	Interpretation
Interview (ID-1)	<i>“And that’s exactly what it is: subjective perception. When you start discussing something like this with a lot of people, everyone brings their own subjective view to the table. And stepping outside your own perception and trying to discuss things objectively is the biggest challenge for everyone.”</i>	Individual Propriety Judgement(s)	Subjective Perception Plurality of Perspective Difficulty: Discursive Objectivation	The codes indicate that the democratization of AI image generation increases the number of <i>individual propriety judgement(s)</i> . This plurality of subjective perspectives seems to generate normative uncertainty and complicate the discursive convergence.
Interview (ID-3)	<i>“[...] it was a bit like the Wild West because content production is becoming totally democratized. I’m a trained expert, I studied it and worked my way through the whole subject area. And now there’s someone who comes from the finance sector, but has a good imagination and knows how to write a prompt – they can produce just as good content.”</i>	Individual Propriety Judgement(s)	Democratization of Image Production Normative Uncertainty (“Wild West”)	

Data Source	Data Excerpt	Deductive Code(s)	Inductive Code(s)	Interpretation
Interview (ID-1)	<i>“The inspiration comes from both sides. You could say: Okay, there’s the competition. OTTO wants to be part of the competition. The competition is pretty fast when we measure ourselves against it.”</i>	Perception(s) of Organizational AI Use	Monitoring of Competition Benchmarking	The codes indicate that <i>perception(s) of organizational AI use</i> are shaped by constant monitoring of the organizational environment. In addition to technological advancements, close observation of competitors and benchmarking against their practices play a central role.
Interview (ID-4)	<i>“A lot of it is driven internally. We have highly motivated colleagues who are constantly monitoring what is happening out there, what new technologies are available, what we can do, and what the competition is doing.”</i>	Perception(s) of Organizational AI Use	Continuous Monitoring of Competition Monitoring of Technology	
Interview (ID-1)	<i>“But our utility analysis [...] really helps us prioritize: What do we tackle next, and what don’t we tackle? [...] Otherwise, you don’t know how to decide. At least, I wouldn’t know.”</i>	Proportionality of AI Use	Institutionalization of Proportionality Utility Analysis as Decision Instrument	The codes indicate that a shared understanding of the <i>proportionality of AI use</i> tends to be institutionalized through formal decision instruments, such as utility analyses, which help decision makers compare AI use cases and impose order under conditions of limited organizational resources. At the same time, institutionalized criteria remain open to revision and may change over time.
Interview (ID-4)	<i>“There are lots of great things you could do. And, of course, everyone has their own baby or pet project. To help us bring a little order to the process, we also have our utility analysis.”</i>	Proportionality of AI Use	Utility Analysis as Decision Instrument Ordering of AI Use Cases	
Participant Observation [Workshop]	<i>During a workshop, we observed practitioners applying OTTO’s utility analysis to evaluate and select AI use cases. However, the predefined scale(s) for assessing business value did not allow for meaningful differentiation between the potential AI use cases in fashion image generation. The participants discussed and adapted the scale to better reflect the specific business context.</i>	Proportionality of AI Use	Institutionalization of Proportionality Utility Analysis as Decision Instrument Adaptation of Evaluation Scheme	

Table 2. Example of Case Data Analysis

We made sense of the data until the research team agreed that explaining the use of AI in fashion image generation as a process and result of continuous proportionalization is free from obvious absurdities and thus sufficiently plausible (Polanyi, 1962). Finally, we discussed our findings with OTTO practitioners for external validation.

Having described the research method and introduced the case company, we now provide additional background to understand OTTO’s use of AI in fashion image production.

4.2 How OTTO Uses and Produces Image Content

Digital images are essential to e-commerce retailers like OTTO, as they primarily engage with customers through electronic channels. Unlike in brick-and-mortar retail, where customers can usually look at or touch products beforehand. As an interviewee explained: *“As humans, we are visual beings. We need information about a product to be presented*

not only in text form, but also visually, for us to be able to imagine it. [...] Hence, image content is particularly important in facilitating purchasing decisions.” (ID-3)

Type	Marketing Content	Product Content
Function	Awareness & Inspiration	Activation & Decision
Use (e.g.)	Landing Pages, Campaign Hubs, Inspiration Widgets, Editorial Content	Product Detail Pages (incl. Variations, Bundles, Configurators)

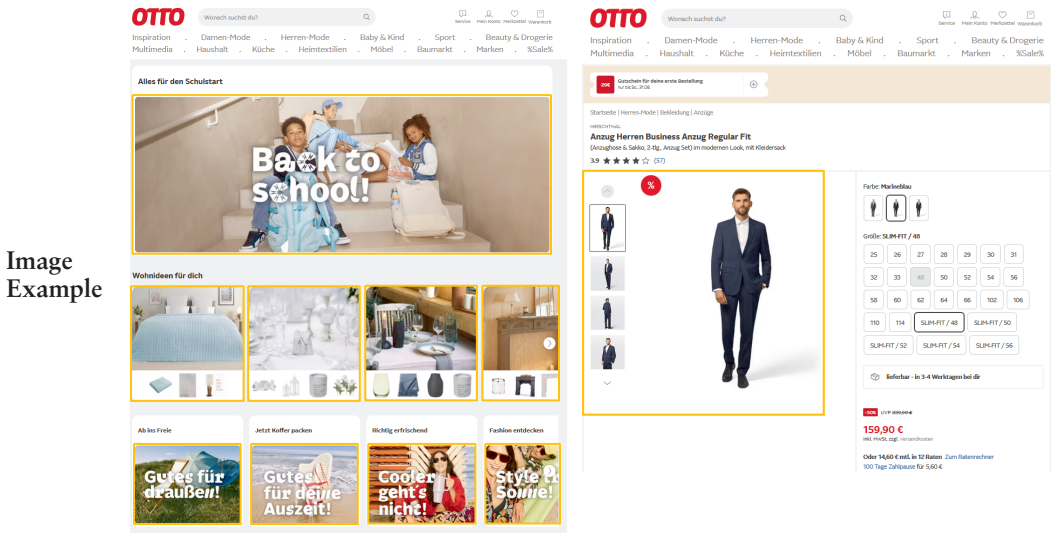


Figure 4: Image Content Types

OTTO differentiates between *marketing* and *product* images (Figure 4). Marketing images are used in the early phases of the customer journey to position the OTTO brand, create desire, and inspire customers. Product images are used in later stages to help customers find specific items, gain orientation, and compare alternatives.

The main difference between marketing and product images is the degree of exactness with which they must represent the products sold on OTTO’s website (“article fidelity”). Unlike marketing, product images require a high level of article fidelity. After all, product images are supposed to help customers imagine the product without creating unrealistic expectations.

Until recently, OTTO used either photography or computer-generated imagery (CGI) to produce digital images. Generative AI models such as Midjourney and Stable Diffusion added a third possibility. They can create digital images from textual inputs by predicting likely pixel sequences (Saharia et al., 2022). Unlike photography and CGI, image generation models do not require real or digital objects. Which promises new opportunities to produce image content faster, cheaper, and more flexible. In the next section, we apply our conceptual model to illustrate the responsible use of AI in OTTO’s fashion image production.

5 Responsible Use of AI for Fashion Image Production at OTTO

For OTTO, the question of responsible use of AI in fashion image production emerged with the rise of image generation models in 2023. Since then, OTTO progressed through what can be described as two *proportionalization cycles*. Below, we outline these cycles and provide examples of AI use cases that OTTO put into production based on its proportionalized understanding of responsible use of AI.

5.1 First Proportionalization Cycle

5.1.1 Formation of Individual Propriety Judgement(s)

At OTTO, individual propriety judgements initially formed through observation and perception of the socio-technical environment in which the organization operates. These perceptions were shaped not only by rapid technological developments in generative AI, but also by close monitoring of competitors' activities. As an interviewee noted: *"We have highly motivated colleagues who are constantly monitoring what is happening out there, what new technologies are available, what we can do, and what the competition is doing."* (ID-4)

Based on these emerging propriety judgements, AI image generation developed organically within OTTO, giving rise to multiple initiatives across the organization. The low entry barriers to the technology enabled many employees to experiment with AI image generation, including individuals outside the units traditionally responsible for content production: *"[...] it was a bit like the Wild West because content production is becoming totally democratized. I'm a trained expert, I studied it and worked my way through the whole subject area. And now there's someone who comes from the finance sector, but has a good imagination and knows how to write a prompt – they can produce just as good content."* (ID-3)

5.1.2 Formation of Collective Validity Judgement(s)

As increasing numbers of employees at OTTO engaged with AI image generation and observed how others expressed propriety judgements, normative questions emerged regarding how the technology should be used. These questions triggered extensive discursive exchanges across hierarchical levels: *"Discussions that were held naturally included questions like, 'How perfect do our people actually look?' I still remember that when we were with the executive board, this question was raised as well. [...] From my perspective, this is a secondary issue, because even today we do not photograph the average person next door, but professional models who go through casting processes beforehand."* (ID-3)

Through these discursive exchanges, individual propriety judgements began to converge into collective validity judgements regarding the responsible (use of AI) in fashion image generation. Interviewees emphasized, however, that this convergence was neither immediate nor straightforward. The large number of involved actors and the diversity of perspectives made it difficult to establish common understanding. Nevertheless, over time, discussions increasingly crystallized around questions of the legitimacy, suitability, necessity, and proportionality.

(1) *Legitimacy of AI Use*: The purpose of AI use in fashion image production is to obtain *"inspiring content for customers cheaply and quickly"* (ID-1). OTTO assessed the

legitimacy of this purpose with respect to internal expectations and regulatory validity clues. As for regulatory validity clues, the EU AI Act classifies fashion image generation as a “limited risk” application, which entails transparency and labelling requirements but does not constitute a prohibited use of AI (European Commission, 2021). Internally, fashion image generation aligns with OTTO’s general business interests. Therefore, fashion image generation was considered legitimate, provided it complies with OTTO’s guidelines for image production, which would prohibit, for example, the generation of visuals that are offensive or discriminatory.

(2) Suitability of AI Use: OTTO determined the suitability of AI use for fashion image production primarily through experimentation: *“We explored different AI services and asked ourselves, what happens when we input images? [...]. What kind of results do we get? [...] And so, step by step, we keep sorting things out: This works, that doesn’t. This works, that doesn’t...”* (ID-1).

This experimentation took place in the context of rapid technological development: *“AI [...] gives us new possibilities every week [...]. The pace is enormous, and I have rarely seen an exponential growth curve as clearly.”* (ID-2). Early experiments revealed limitations such as anatomical inaccuracies in depictions of human figures, which were gradually reduced as the technology improved. At the same time, OTTO encountered situations in which technological possibilities shifted in unexpected ways: *“We were surprised to find that the quality of lingerie images was better than that of ordinary textiles. [...] However, it is almost impossible to produce lingerie images with current models from Google and others. Because they implemented filter functions [...]. So, things looked very promising regarding lingerie images for a while, but possibilities are currently very limited again.”* (ID-2)

Through such experimentation, OTTO developed a general understanding of the types of problems for which AI is particularly suitable. As one interviewee summarized: *“Generative AI excels at predicting [...] which pixel is most likely to appear in a particular spot. However, when it comes to representing articles accurately, [...] you must ‘bend’ AI models for this purpose. Because they are not built to reproduce something exactly [...]. Which is why we still face clear limitations when it comes to article fidelity.”* (ID-2). OTTO therefore concluded that current image generation models are more suitable for producing content with low article fidelity (i.e., marketing images) and less suitable for content that requires high article fidelity (i.e., product images).

(3) Necessity of AI Use: OTTO assessed the necessity of AI use by comparing it with alternative means of image production such as photography and CGI. In this comparison, AI was perceived as *“almost unbeatable in terms of efficiency”* when creating marketing content (ID-2). At the same time, AI was not seen as merely replacing existing processes but also as enabling new applications that had previously been infeasible. As one interviewee explained: *“There are uses that would not have been possible [...] before. There simply would have been no image. That would have been the solution.”* (ID-2). Accordingly, OTTO developed the understanding that AI may be necessary in *“problem cases”* (ID-4) where photography or CGI fall short.

(4) Proportionality of AI Use: At OTTO, the use of AI needed to be justified in terms of the company’s top priority to *“do the things that take us furthest forward”* (ID-4). To operationalize this maxim, OTTO developed a *“utility analysis”* (ID-1) that enables deci-

sion-makers to weigh all relevant considerations for a specific AI use case and compare it to others.

The utility analysis consists of more than 20 criteria (e.g., business benefits, strategic alignment, and risks) that OTTO identified as relevant for evaluating AI use cases. These criteria were assigned individual weights by systematically comparing their relative importance. Based on the resulting scores, use cases were classified within a portfolio and compared against one another to support prioritization decisions.

The potential AI use cases for fashion image generation were also supposed to be assessed using this instrument. Because “*There are lots of great things you could do. And, of course, everyone has their own baby or favorite project. We have the utility analysis to bring a little order to things*” (ID-4). During one workshop, however, we observed that the predefined scales for assessing business value did not allow for meaningful differentiation among fashion image generation use cases. In response, participants discussed and adapted the evaluation scale to better reflect the specific business context.

Based on the utility analysis, OTTO put several AI use cases for fashion image generation into production. Three of which we describe below.

5.1.3 Examples of AI Use in OTTO’s Fashion Image Production

The first AI use case that OTTO deemed justifiable concerned the generation of marketing images. Because article fidelity plays only a secondary role in this context, AI was assessed as suitable for producing visuals whose primary purpose is to inspire customers, encourage engagement, and increase the likelihood of sales. Traditional production processes, such as model and location photography, were highly demanding in terms of effort, cost, and scalability. The use of AI was therefore considered both necessary and proportionate, as it enabled the efficient creation of inspiring content at approximately 60 percent lower cost than conventional image production. Based on this assessment, OTTO established a fully digital workflow in which marketing images are generated entirely by AI.



Figure 5: AI-Generated Marketing Image on OTTO Landing Page

In contrast to marketing visuals, product images require high article fidelity. OTTO therefore uses AI as a complementary tool alongside established production methods, rather than as a full replacement (Figure 6).

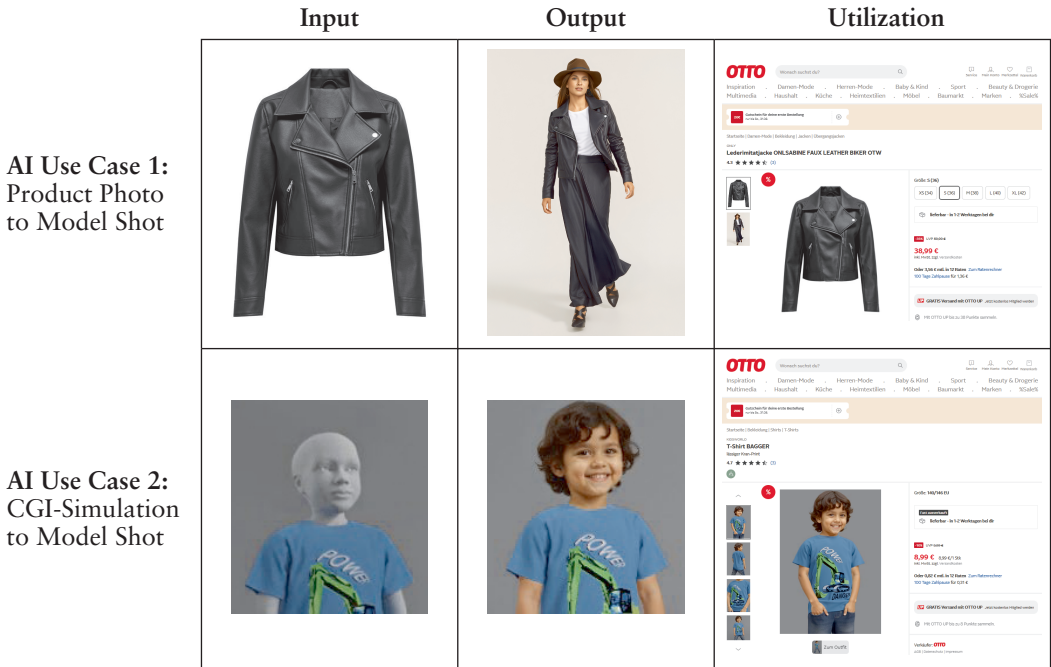


Figure 6: AI-Enhanced Product Images

One such use case involves generating model images based on existing product photos. AI is applied where model shots are unavailable due to logistical or economic constraints. As one interviewee noted, traditional model photography requires an “*insanely high level of effort*” (ID-1) in terms of sample logistics, studio costs, and personnel coordination. In this context, OTTO found AI suitable for generating front-view model images from product photos, while manual post-production ensures that no counterfactual or misleading images are created.

A related use case builds on CGI simulations from OTTO’s fashion product development process. While CGI enables “*maximum product fidelity*” (ID-3), it has “*limitations when it comes to depicting people*” (ID-3). Here, AI allows OTTO to generate realistic model images for fashion articles that have not yet been physically produced. This combination enables earlier visualization of products while maintaining control over article fidelity.

5.1.4 Institutionalization of Collective Validity Judgement(s)

Once established, the collective understanding of responsible use of AI in fashion image generation became institutionalized at OTTO. One prominent form of institutionalization is the expansion of the utility analysis as an agreed decision instrument across multiple initiatives. As one interviewee explained: “*We now use the utility analyses across different initiatives. But we have also noticed that [...] not all teams in the company are using it yet.*” (ID-1)

Beyond formal decision instruments, we observed additional forms through which the collective understanding of responsible AI use becomes institutionalized in everyday practice. These include the development of standardized prompts and curated pools of AI models for image generation. As one interviewee noted: *“For example, what do future prompts need to look like so that we really only get these kinds of models? [...] Is there also an AI model pool that people can draw on? [...] These are the steps we are thinking about now after this MVP phase... How can we bring more stability into this and more control over what is generated?”* (ID-1)

At the same time, ongoing changes in the socio-technical environment have prompted OTTO to enter a subsequent proportionalization cycle that is still unfolding.

5.2 Second Proportionalization Cycle

5.2.1 Expansion of the Legitimacy Object(s)

The second proportionalization cycle emerged from organizational learning about the technical possibilities of AI image generation in the first cycle. Compared to the initial phase, the legitimacy object shifted in a fundamental way. While AI had initially been discussed primarily as an alternative means for established purposes, such as replacing traditional model photography, OTTO began to consider whether AI could also be used for new purposes that had not previously been part of image production practices.

One example concerns the exploration of AI image generation for baby fashion. After successfully generating model images for adult fashion items based on product photos, OTTO began examining whether a similar approach could be applied to baby clothing. Because babies had not been photographed before, this exploration marked an expansion of the legitimacy object and opened a qualitatively new field of application.

In contrast to adult fashion, however, generating synthetic images of babies raised not only technical challenges related to article fidelity but also heightened ethical and normative concerns. A central issue concerned compliance with internal guidelines on nudity, which stipulate that *“too much skin should not be shown”* (ID-1). At the same time, the absence of established production practices for baby images limited the availability of reference points. Additional concerns emerged regarding the potential misuse of synthetic baby images. As one interviewee explained, no collective consensus has yet emerged regarding responsible use of AI in this context: *“It is definitely a case. There is a need and we are moving forward. But only within the guidelines. And this is where the ethical aspect comes in: not everything at all costs, not just focusing on the business case [...] but even if they are synthetic people, portraying them with dignity. [...] We cannot yet say what this will look like in the end, because we are only just starting.”* (ID-1)

5.2.2 Socio-Technical Dynamics in OTTO’s Environment

The ongoing proportionalization cycle is further shaped by judgement validation institution(s) in OTTO’s environment. In addition to a dynamic regulatory landscape, anticipated customer reactions are particularly salient for how OTTO reflects on what it may hold to be a legitimate, suitable, necessary, and proportionate use of AI over time.

Under current legal conditions, OTTO does not yet visibly mark AI-generated images as such. This will change as the EU AI Act enters into force. At present, however, *“for the*

customer, it is not distinguishable whether an article was photographed on a real model or generated by AI” (ID-3).

In line with our main proposition, the collective understanding of responsibility at OTTO that has emerged through proportionalization and currently informs the use of AI in fashion image generation remains provisional. It guides organizational action in the present, yet remains open to revision as new evaluative inputs emerge from the socio-technical environment. As one interviewee concluded: “*And this is now the point where we say: We generate the [image] content, try to design it as well as possible in the sense of OTTO. And once the content is on the website, we go into surveying [...] to include our customers’ perceptions in further development. [...] It could be that everyone says: ‘Oh God, we don’t want this at all.’ Then we would have to stop the project and say: we no longer generate images, we go back to photography. [...] But to even get into this try-and-error, the images have to go live*” (ID-1).

6 Discussion and Conclusion

We conducted this research to advance our understanding of responsible use of AI in organizations. To this end, we integrated Bitektine and Haack’s (2015) process theory of social evaluation with the principle of proportionality (Karliuk, 2023). This integration resulted in a conceptual model (“continuous proportionalization”) that captures how organizations form, enact, and revise collective interpretations of responsible use of AI in a dynamic sociotechnical environment. For illustration, we applied the model to a case study of AI use in fashion image generation at Germany’s largest e-commerce company, OTTO.

With this research, we make three interrelated contributions to the growing body of literature on governance and responsible use of AI in organizations.

First, we offer an integrative perspective on the nature of responsible use of AI. Prior research has tended to conceptualize responsible use of AI either as a relatively stable property of systems, principles, or organizational arrangements (Bughin, 2025a, 2025b; Heger et al., 2025; Krijger et al., 2023; Minkkinen et al., 2023), or as an ongoing process of deliberation, reflection, and adaptation (Elia et al., 2025; Hagendorff, 2022; Kallina & Singh, 2024; Mittelstadt, 2019; Yilma, 2025). This has left a theoretical puzzle as to how responsibility can be both dynamic and, at least temporarily, stabilized in organizational practice. By conceptualizing responsible use of AI as a process of social evaluation (Bitektine & Haack, 2015), we provide an integrative stance that accounts for both its processual and property-like characteristics without collapsing into either.

Second, we theorize how shared organizational understandings of responsible use of AI emerge. Prior work has highlighted the importance of discourse and deliberation in responsible AI (Coeckelbergh, 2024; Gogoll et al., 2021; Watson et al., 2025). We extend this literature by introducing a process model that explains how individual propriety judgments converge discursively into collective validity judgments within a dynamic sociotechnical environment. Our model further shows that this convergence is not arbitrary, but crystallizes around interpretations of legitimacy, suitability, necessity, and proportionality of AI use. In doing so, we elevate the principle of proportionality from a situational decision test (Karliuk, 2023) to a set of shared discursive dimensions through which collective interpretations of responsible use of AI are formed, stabilized, and revised over time.

Third, we illustrate the value of drawing on legitimacy and institutional theory for research on responsible use of AI. As other scholars have noted (de-Lima-Santos et al., 2025; Horneber, 2025), legitimacy and institutional theory provide useful lenses for conceptualizing responsible use of AI. While beyond the scope of this paper, studying dynamic sociotechnical phenomena such as responsible use of AI may offer opportunities not only to draw on, but also to contribute back to legitimacy and institutional theory.

We argue that our findings are transferable beyond the empirical context examined in this study (i.e., e-commerce). This is because our conceptual model is grounded in general theory of the legitimacy of organizational conduct (Bitektine & Haack, 2015) and a broadly applicable heuristic of action validation (Karliuk, 2023). It follows that challenging the explanatory power of continuous proportionalization in other settings offers promising avenues for future research. On the one hand, the model could be applied across different industries and regional contexts to examine how proportionalization is enabled or constrained under varying regulatory, competitive, or cultural conditions. On the other hand, the model could be extended to high-stakes domains of AI use, such as healthcare or military applications. While we expect the recursive interplay between individual propriety and collective validity judgments to persist, the dynamics of proportionalization may differ as the severity of potential consequences increases.

Additional research opportunities arise from the limitations of our study. While our case study illustrates the recursive interplay between individual propriety judgments and collective validity judgments, it does not capture proportionalization over an extended period or in all its contextual nuances. Future research could therefore adopt a longitudinal perspective, to trace how proportionalization unfolds over time (e.g. using ethnographic methods). Moreover, our study did not assess the effects of proportionalization on measurable outcomes (e.g., regulatory compliance and organizational performance). Future research could develop and test metrics that capture both the effectiveness and the ethical implications of responsible AI use. Finally, although we identify practical challenges in the discursive convergence of propriety judgments, we did not examine how such challenges might be addressed. This opens avenues for design-oriented research to develop tools, processes, or governance mechanisms that support organizations in forming shared understandings of responsible use of AI.

Our research offers actionable implications for managers. We suggest that responsible use of AI is neither a one-time compliance exercise nor something that can be “solved” by issuing principles or guidelines alone. Instead, it requires continuous attention and the active organization of discourse through which collective interpretations of responsible use of AI will be formed, challenged, and revised. Creating and sustaining such discursive spaces then becomes a central managerial task. At the same time, our findings imply that responsible use of AI cannot simply be imported from external standards or best practices. While external cues such as regulation or industry norms matter, a shared understanding of responsibility must be worked out internally, in relation to an organization’s specific goals, constraints, and technologies. To support this process, our conceptual model provides practitioners with a set of orienting questions around which responsible use of AI crystallizes: Is the aim for using AI legitimate? Is AI suitable for achieving it? Is the use of AI necessary? And are means and ends proportionate? In discussions with OTTO practitioners, these dimensions were perceived as intuitive and useful for structuring delib-

eration. They may help managers navigate the difficult task of creating conditions under which continuous proportionalization can unfold.

In conclusion, this research contributes to a more nuanced understanding of what responsible use of AI is and how it can be organized. Which we trust other researchers and practitioners may find useful in their shared efforts to ensure that AI technologies are used for good.

References

- Bach, T. A., Kaarstad, M., Solberg, E., & Babic, A. (2025). Insights into suggested Responsible AI (RAI) practices in real-world settings: a systematic literature review. *AI and Ethics*, 5(3), 3185–3232. <https://doi.org/10.1007/s43681-024-00648-7>
- Bitektine, A., & Haack, P. (2015). The “macro” and the “micro” of legitimacy: Toward a multilevel theory of the legitimacy process. *Academy of Management Review*, 40(1), 49–75. <https://doi.org/10.5465/amr.2013.0318>
- Bughin, J. (2025a). Doing versus saying: responsible AI among large firms. *AI & Society*, 40(4), 2751–2763. <https://doi.org/10.1007/s00146-024-02014-x>
- Bughin, J. (2025b). The role of AI assets and capabilities in shaping responsible AI deepening: a random forest machine learning view. *AI and Ethics*, 5(6), 6313–6327. <https://doi.org/10.1007/s43681-025-00802-9>
- Christophersen, T., & Pärn, J. (2021). Data Science bei OTTO. In P. Buxmann & H. Schmidt (Eds.), *Künstliche Intelligenz* (pp. 101–115). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-662-61794-6_6
- Cihon, P., Schuett, J., & Baum, S. D. (2021). Corporate governance of artificial intelligence in the public interest. *Information (Basel)*, 12(7), 275. <https://doi.org/10.3390/info12070275>
- Coeckelbergh, M. (2020). *AI Ethics*. MIT Press.
- Coeckelbergh, M. (2024). Artificial intelligence, the common good, and the democratic deficit in AI governance. *AI and Ethics*, 5(2), 1491–1497. <https://doi.org/10.1007/s43681-024-00492-9>
- Cohen-Eliya, M., & Porat, I. (2010). American balancing and German proportionality: The historical origins. *International Journal of Constitutional Law*, 8(2), 263–286. <https://doi.org/10.1093/icon/moq004>
- de-Lima-Santos, M.-F., Yeung, W. N., & Dodds, T. (2025). Guiding the way: a comprehensive examination of AI guidelines in global media. *AI & Society*, 40(4), 2585–2603. <https://doi.org/10.1007/s00146-024-01973-5>
- Elia, M., Ziehmman, P., Krumme, J., Schlögl-Flierl, K., & Bauer, B. (2025). Responsible AI, ethics, and the AI lifecycle: how to consider the human influence? *AI and Ethics*, 5(4), 4011–4028. <https://doi.org/10.1007/s43681-025-00666-z>
- Floridi, L., Cowsls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People-An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations. *Minds and Machines*, 28(4), 689–707. <https://doi.org/10.1007/s11023-018-9482-5>
- Garibay, O. O., Winslow, B., Andolina, S., Antona, M., Bodenschatz, A., Coursaris, C., Falco, G., Fiore, S. M., Garibay, I., Grieman, K., Havens, J. C., Jirotko, M., Kacorri, H., Karwowski, W., Kider, J., Konstan, J., Koon, S., Lopez-Gonzalez, M., Maifeld-Carucci, I., ... Xu, W. (2023). Six Human-Centered Artificial Intelligence Grand Challenges. *International Journal of*

- Human-Computer Interaction*, 39(3), 391–437. <https://doi.org/10.1080/10447318.2022.2153320>
- Gerken, T. (2024, December 11). Chatbot “encouraged teen to kill parents over screen time limit.” *BBC News*. <https://www.bbc.com/news/articles/cd605e48q1vo>
- Gogoll, J., Zuber, N., Kacianka, S., Greger, T., Pretschner, A., & Nida-Rümelin, J. (2021). Ethics in the Software Development Process: from Codes of Conduct to Ethical Deliberation. *Philosophy & Technology*, 34(4), 1085–1108. <https://doi.org/10.1007/s13347-021-00451-w>
- Hagendorff, T. (2022). A virtue-based framework to support putting AI ethics into practice. *Philosophy & Technology*, 35(3), 1–24. <https://doi.org/10.1007/s13347-022-00553-z>
- Heger, A. K., Passi, S., Dhanorkar, S., Kahn, Z., Wang, R., & Vorvoreanu, M. (2025). Towards a Responsible AI Organizational Maturity model. *Proceedings of the ACM on Human-Computer Interaction*, 9(7), 1–33. <https://doi.org/10.1145/3757514>
- Horneber, D. (2025). Understanding the implementation of responsible artificial intelligence in organizations: A Neo-institutional theory perspective. *Communications of the Association for Information Systems*, 57, 8. <https://doi.org/10.17705/1cais.05708>
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- Kallina, E., & Singh, J. (2024). Stakeholder involvement for responsible AI development: A process framework. *Proceedings of the 4th ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–14. <https://doi.org/10.1145/3689904.3694698>
- Karliuk, M. (2023). Proportionality principle for the ethics of artificial intelligence. *AI and Ethics*, 3(3), 985–990. <https://doi.org/10.1007/s43681-022-00220-1>
- Katirai, A., Garcia, N., Ide, K., Nakashima, Y., & Kishimoto, A. (2024). Situating the social issues of image generation models in the model life cycle: a sociotechnical approach. *AI and Ethics*, 5(2), 1769–1786. <https://doi.org/10.1007/s43681-024-00517-3>
- Krijger, J., Thuis, T., de Ruiter, M., Ligthart, E., & Broekman, I. (2023). The AI ethics maturity model: a holistic approach to advancing ethical data science in organizations. *AI and Ethics*, 3(2), 355–367. <https://doi.org/10.1007/s43681-022-00228-7>
- Kuznia, R., Gordon, A., & Lavandera, E. (2025, November 6). “You’re not rushing. You’re just ready:’ Parents say ChatGPT encouraged son to kill himself. *CNN*. <https://www.cnn.com/2025/11/06/us/openai-chatgpt-suicide-lawsuit-invs-vis>
- Mikalef, P., Conboy, K., Lundström, J. E., & Popovič, A. (2022). Thinking responsibly about responsible AI and “the dark side” of AI. *European Journal of Information Systems: An Official Journal of the Operational Research Society*, 31(3), 257–268. <https://doi.org/10.1080/0960085x.2022.2026621>
- Minkkinen, M., Zimmer, M. P., & Mäntymäki, M. (2023). Co-shaping an ecosystem for responsible AI: Five types of expectation work in response to a technological frame. *Information Systems Frontiers: A Journal of Research and Innovation*, 25(1), 103–121. <https://doi.org/10.1007/s10796-022-10269-2>
- Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence*, 1(11), 501–507. <https://doi.org/10.1038/s42256-019-0114-4>
- Otto. (2021, September 16). *Ethik vs. Wirtschaftlichkeit, können wir KI vertrauen? | MAIN Session – OTTO [Ethics vs. economics, can we trust AI?]*. Youtube. <https://www.youtube.com/watch?v=aiay2hfDiOg>

- Polanyi, M. (1962). The Republic of science: Its political and economic theory. *Minerva*, 1(1), 54–73. <https://doi.org/10.1007/bf01101453>
- Radford, A., Kim, J. W., Xu, T., Brockman, G., McLeavey, C., & Sutskever, I. (2022). Robust Speech Recognition via Large-Scale Weak Supervision. In *arXiv [eess.AS]*. arXiv. <http://arxiv.org/abs/2212.04356>
- Rufo, Y. (2025, July 27). What Guess's AI model in Vogue means for beauty standards. *BBC News*. <https://www.bbc.com/news/articles/cgeqe084nn4o>
- Ryan, M., & Stahl, B. C. (2021). Artificial intelligence ethics guidelines for developers and users: clarifying their content and normative implications. *Journal of Information Communication and Ethics in Society*, 19(1), 61–86. <https://doi.org/10.1108/jices-12-2019-0138>
- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E. L., Ghasemipour, S. K. S., Ayan, B. K., Mahdavi, S. S., Lopes, R. G., Salimans, T., Ho, J., Fleet, D. J., & Norouzi, M. (2022). Photo-realistic text-to-image diffusion models with deep language understanding. *Neural Information Processing Systems*, *abs/2205.11487*, 36479–36494. <https://doi.org/10.48550/arXiv.2205.11487>
- Sobek, T., & Montag, J. (2018). Proportionality Test. In *Encyclopedia of Law and Economics* (pp. 1–5). Springer New York. https://doi.org/10.1007/978-1-4614-7883-6_721-1
- Stahl, B. C. (2012). Responsible research and innovation in information systems. *European Journal of Information Systems: An Official Journal of the Operational Research Society*, 21(3), 207–211. <https://doi.org/10.1057/ejis.2012.19>
- Suchman, M. C. (1995). Managing legitimacy: Strategic and institutional approaches. *Academy of Management Review*, 20(3), 571. <https://doi.org/10.2307/258788>
- Suddaby, R., Bitektine, A., & Haack, P. (2017). Legitimacy. *Academy of Management Annals*, 11(1), 451–478. <https://doi.org/10.5465/annals.2015.0101>
- Taddeo, M., & Floridi, L. (2018). How AI can be a force for good. *Science*, 361(6404), 751–752. <https://doi.org/10.1126/science.aat5991>
- Watson, D. S., Mökander, J., & Floridi, L. (2025). Competing narratives in AI ethics: a defense of sociotechnical pragmatism. *AI & Society*, 40(5), 3163–3185. <https://doi.org/10.1007/s00146-024-02128-2>
- Yilma, K. (2025). From principles to process: the principlist approach to AI ethics and lessons from Internet bills of rights. *AI and Ethics*, 5(4), 4279–4291. <https://doi.org/10.1007/s43681-025-00719-3>
- Yin, R. K. (1981). The Case Study Crisis: Some Answers. *Administrative Science Quarterly*, 26(1), 58–65. <https://doi.org/10.2307/2392599>
- Yin, R. K. (1994). *Case Study Research: Design and Methods*. SAGE Publications.
- Zhang, Z., & Wang, J. (2024). Can AI replace psychotherapists? Exploring the future of mental health care. *Frontiers in Psychiatry*, 15, 1444382. <https://doi.org/10.3389/fpsy.2024.1444382>

Simon Sturm, M.A. HSG, is a Research Associate and a PhD candidate at the Institute of Information Systems and Digital Business at the University of St.Gallen, Switzerland. His research focuses on the management and governance of machine learning-based AI innovation in electronic commerce.

Address: University of St.Gallen, Institute of Information Systems and Digital Business (IWI-HSG), Müller-Friedberg-Strasse 8, 9000 St.Gallen, Switzerland,
E-Mail: simongabriel.sturm@unisg.ch
ORCID: <https://orcid.org/0009-0008-9670-2547>

Florian Krause, Dr., is a PostDoc and Lecturer at the Institute for Business Ethics at the University of St.Gallen, Switzerland and at the Institute for Interdisciplinary Work and Employment Studies at the Leibniz University of Hanover, Germany. His research addresses ethics of digital transformation as well as foundational questions in economics and business ethics, with a particular interest in their normative underpinnings.

Address: University of St.Gallen, Institute for Business Ethics (IWE-HSG), Blumenbergplatz 9, 9000 St.Gallen, Switzerland, E-Mail: florian.krause@unisg.ch
ORCID: <https://orcid.org/0000-0002-9981-165X>

Benjamin van Giffen, Prof. Dr., is an Associate Professor at the University of Liechtenstein, Liechtenstein and an associate researcher with the European Research Center for Information Systems (ERCIS). His research focuses on the organizational adoption of artificial intelligence, AI business value, digital platforms and ecosystems, and human-centered design innovation (e.g., Design Thinking). Prior to joining academia, he held various roles in IT project and compliance management at a global pharmaceutical manufacturing company, where he also served as a manager of the sales and operations planning processes for six sales regions and five manufacturing plants in Europe.

Address: University Liechtenstein, Information Systems and Digital Innovation, Fürst-Franz-Josef-Strasse, 9490 Vaduz, Liechtenstein, E-Mail: benjamin.vangiffen@uni.li
ORCID: <https://orcid.org/0000-0001-6753-8399>