



Delft University of Technology

The distinction between generic and specific concepts and why it matters for conceptual engineering

van de Poel, Ibo

DOI

[10.1080/0020174X.2025.2556750](https://doi.org/10.1080/0020174X.2025.2556750)

Publication date

2025

Document Version

Final published version

Published in

Inquiry (United Kingdom)

Citation (APA)

van de Poel, I. (2025). The distinction between generic and specific concepts and why it matters for conceptual engineering. *Inquiry (United Kingdom)*. <https://doi.org/10.1080/0020174X.2025.2556750>

Important note


To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Inquiry

An Interdisciplinary Journal of Philosophy

ISSN: 0020-174X (Print) 1502-3923 (Online) Journal homepage: www.tandfonline.com/journals/sinq20

The distinction between generic and specific concepts and why it matters for conceptual engineering

Ibo van de Poel

To cite this article: Ibo van de Poel (19 Sep 2025): The distinction between generic and specific concepts and why it matters for conceptual engineering, Inquiry, DOI: [10.1080/0020174X.2025.2556750](https://doi.org/10.1080/0020174X.2025.2556750)

To link to this article: <https://doi.org/10.1080/0020174X.2025.2556750>



© 2025 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 19 Sep 2025.



Submit your article to this journal [↗](#)



Article views: 246



View related articles [↗](#)



View Crossmark data [↗](#)

The distinction between generic and specific concepts and why it matters for conceptual engineering

Ibo van de Poel 

Department Values, Technology & Innovation, TU Delft, Delft, The Netherlands

ABSTRACT

Conceptual engineering is an approach or method for assessing, improving, adapting and disposing concepts. While recent case studies have shown the possibility and success of conceptual engineering, I argue that not all concepts are equally open to conceptual engineering. It is therefore useful to distinguish between generic and specific concepts. While the latter can be the object of conceptual design, I argue that designing generic concepts is problematic for practical and normative reasons. Nevertheless, attempts can be made to change generic concepts through interventions in larger conceptual repertoires. Contrary to specific concepts, generic concepts do not have a well-delineated function that can provide a normative reference for conceptual engineering. Moreover, they are well entrenched in larger conceptual schemes or repertoires, so that they cannot be engineered in isolation. Generic concepts also raise more and larger challenges in terms of implementation and authority than specific concepts. I discuss the consequences of the distinction between generic and specific concepts for conceptual engineering, the ethics of socially disruptive technology and conceptual ethics.

ARTICLE HISTORY Received 24 April 2025; Accepted 30 August 2025

KEYWORDS Conceptual engineering; conceptual disruption; implementation challenge; authority problem; generic concepts; specific concepts

1. Introduction

In the last decade, there has been a growing attention for conceptual engineering, which may be understood as an approach or method for assessing, improving, adapting and disposing concepts (Cappelen 2018; Chalmers 2020; Haslanger 2020; Isaac, Koch, and Nefdt 2022; Koch, Löhr, and Pinder

CONTACT Ibo van de Poel  i.r.vandepoel@tudelft.nl  Department Values, Technology & Innovation, TU Delft, Jaffalaan 5, 2628 BX Delft, The Netherlands

© 2025 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group
This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

2023). Conceptual engineering may be aimed at improving existing concepts (amelioration; cf. Cappelen (2018)), developing new concepts (de novo conceptual engineering; cf. Chalmers (2020)), disposing concepts (for example because they exacerbate inequalities or injustices) or adapting and repairing concepts (for example when they have been disrupted; cf. Hopster and Löhr (2023)).

Recently, several authors have discussed more detailed case studies that show how conceptual engineering can help to revise or improve concepts, for example in response to conceptual disruption, and more generally in response to new challenges and moral problems, e.g. created by new technologies (Hopster and Löhr 2023; Veluwenkamp et al. 2022). These examples suggest that – despite some skepticism about, for example, the implementation of new or revised concepts (cf. Cappelen 2018) – there are contexts in which conceptual engineering can be meaningfully applied and can be successful (Veluwenkamp 2025).

However, this leaves open the question whether all our concepts are equally open to conceptual engineering. In this article, I will argue that they are not and that it is therefore useful for the purpose of conceptual engineering to distinguish between specific and generic concepts. Generic concepts are defined as concepts that are used in different contexts and practices, and that they fulfill different aims, roles or functions in these contexts and practices. Such generic concepts also play central roles in larger conceptual schemes or repertoires to which societies have access. Examples of such generic concepts are NATURE, FREEDOM, CONTROL and RESPONSIBILITY.

Specific concepts fulfill a specific function in a specific context or practice. Often they are proposed and used by specific groups of practitioners or experts, like scientists, philosophers or other professionals to serve specific purposes. Examples are concepts like MEANINGFUL HUMAN CONTROL, GUIDANCE CONTROL, FREEDOM AS NON-DOMINATION and OUTCOME RESPONSIBILITY. Such specific concepts are open to deliberate conceptual design because they have a well-delineated function on basis of which they can be conceptually engineered and conceptual engineers have some degree of control over the use and application of these concepts.

I will argue that whereas specific concepts can be conceptually engineered, the engineering metaphor breaks down when it comes to what I will call generic concepts. That is to say: designing generic concepts is often-times either practically infeasible or normatively questionable, or both. So, while I believe that deliberate attempts can be made to change generic concepts, such attempts are better not understood in terms of conceptual design. Rather they should be understood as attempts to intervene in larger conceptual repertoires, I will argue.

Generic concepts can also get disrupted, e.g. through technological developments (Marchiori and Scharp 2024). In combination with the fact that it is

typically hard or even impossible to (re)design generic concepts, this raises additional (ethical) questions. It suggests that while some conceptual disruptions may be relatively easily repaired with (local) conceptual engineering efforts, others may be more systemic and deep, and may require conceptual interventions at the level of conceptual schemes and repertoires rather than just (re)design of specific concepts. While the need to address conceptual disruption at the level of conceptual schemes and repertoires has been recognized in the literature (e.g. Löhr 2023a; Marchiori and Scharp 2024), the further point that this challenge cannot be met through conceptual design of generic concepts, as I will argue in this paper, has not yet been recognized.

In the next section, I will introduce conceptual engineering and conceptual disruption, and describe how conceptual engineering can help to (re)engineer (specific) concepts. In section 3, I introduce the idea of generic concepts. In section 4, I argue that such generic concepts are harder to conceptually engineer than specific concepts. Section 5 presents a distinction between conceptual design and interventions in larger conceptual repertoires, and it argues that whereas specific concepts can be conceptually designed, generic concepts can only be (successfully) changed through interventions in conceptual repertoires that are ultimately social and political in nature. In section 6, I present conclusions and some implications of the view I have defended for conceptual engineering, the ethics of socially disruptive technologies and conceptual ethics.

2. Conceptual disruption and conceptual (re)engineering

Conceptual engineering is aimed at assessing, improving, adapting and disposing concepts. We can distinguish between a metaphysical and a pragmatic approach to conceptual engineering (Thomasson 2020). The first is aimed at developing concepts that represent how the world really is, to carve nature at its joints, so to speak. The second starts from the assumption that concepts have a function or point that can be used as normative reference point for conceptual engineering. Here I will follow and assume the second functionalist approach to conceptual engineering, which is now followed by many authors on the subject (e.g. Jorem 2022; Nado 2021a; Thomasson 2020; Veluwenkamp 2025) with some exceptions (Capelen 2018).

There have been different proposals for how to understand concept functions. Some authors have proposed specific function theories to understand and describe the functions of concepts (e.g. Köhler and Veluwenkamp 2024; Queloz 2022; Simion and Kelp 2020; Thomasson 2020), others have argued that we should not aim for a substantive account of concept functions, and instead focus on what matters in a specific inquiry (Riggs 2021), or on what a concept is used for (Jorem 2022).

In this contribution, I will assume a broadly functionalist approach to concepts, i.e. I will assume that (the use of) concepts can serve diverse aims and that concepts can be used for a variety of purposes. Concepts can fulfill a variety of functions, ranging from general functions like representation, causal reasoning and moral evaluation to more specific functions like managing the risks of a new technology or distributing responsibility. However, I will argue that proper conceptual engineering requires a clear delineation of the function of a concept because otherwise we lack a clear normative reference point for conceptual engineering.

The function of a concept may sometimes be clear from how the concept is used in a given context or in given practices. In other cases, efforts may be required to find the (current) function of concepts, e.g. through some form of what may be called reverse engineering or through genealogy (Queloz 2021). It is also conceivable that we want a concept to fulfill a new, not yet fulfilled, function, and this may also lead to the development of new concepts (*de novo* conceptual engineering). It could be argued that when we are adapting concepts to fulfill new functions we are effectively ‘changing the subject’ (cf. Strawson 1963), but that as such does not seem to invalidate the enterprise of conceptual engineering if one assumes a functionalist approach to conceptual engineering (Nado 2021a).

Initially, conceptual engineering mainly focused on ameliorating concepts, but recently it has been argued that it can also be used to adapt or ‘repair’ concepts in response to conceptual disruption (Hopster and Löhr 2023). Conceptual disruption has been understood as ‘a challenge to the meaning of a concept that prompts a possible future revision of it’ (Hopster et al. 2023, 143). Although this is a useful definition, I will here adapt a slightly different definition that emphasizes (more explicitly) functional failure, and which understands conceptual disruption as ‘an interruption in the normal functioning of a concept, cluster of concepts, or conceptual scheme’ (Marchiori and Scharp 2024, 4). This definition, moreover, does not restrict conceptual disruption to single concepts but extends it to clusters of concepts and conceptual schemes.

Several authors have recently published case studies of conceptual engineering in a technological context (Himmelreich and Köhler 2022; Hopster and Löhr 2023; Veluwenkamp 2025; Veluwenkamp et al. 2022). In these cases, existing concepts were disrupted by new technologies or they did not function properly, e.g. because of new challenges and moral problems caused by technology. These cases are sometimes historical (e.g. brain death), but most are more recent examples of concepts disrupted by technology or concepts that are at least not fully functional in the light of new technological developments and moral problems, oftentimes related to artificial intelligence. Examples of such concepts are CONTROL, RESPONSIBILITY and AGENCY. The focus is often on the revision of such concepts in a specific, applied

context, although also cases are discussed of more general revisions of concepts, like FREEDOM (Veluwenkamp et al. 2022).

From these cases, we can distill a number of questions (and related phases) that seem central to such conceptual engineering efforts:

1. What is the current (dominant) conception of CONCEPT X?
2. What functions does CONCEPT X currently fulfill? (or do we want it to fulfill)?
3. In what ways is CONCEPT X currently not well functioning? (was CONCEPT X conceptually disrupted? How could it better fulfill its function?)
4. What new conception of CONCEPT X might restore its functionality or might better fulfill its function?

Although currently not typically included in the mentioned descriptions, one might argue that question/phase 4 should be followed by a phase where a new concept is tried out or experimented with to see whether it indeed better fulfills the required function, and on basis of evaluation of this, the new concept might be further adapted or finetuned.

The following gives an example of what the answers to these questions might look like for the case of ‘meaningful human control’, which has been proposed as a new conception of CONTROL in the context of self-driving vehicles:

1. According to Veluwenkamp (2025, 162), the current (dominant) conception of CONTROL in the context of self-driving cars is ‘operational control’ which is defined as ‘Agent A is in *operational control* of outcome O if and only if A is (or has been) able to causally influence O.’
2. Functions that CONTROL fulfills in this context are ‘(1) a decrease of risk we expose others to and (2) to allow for appropriate responsibility attributions’ (Veluwenkamp 2025, 162).¹
3. It is claimed that ‘when an autonomous system is designed with the value of operational control, it does not satisfactorily fulfill the normative function of “control”.’ (Veluwenkamp 2025, 163)
4. The new conception of CONTROL that is proposed is ‘meaningful human control’: ‘Agent A has *meaningful human control* over system S if and only if (1) A’s reasons are being tracked by S and (2) S is designed such that the tracing condition applies to A.’ (Veluwenkamp 2025, 164)

The above descriptions seems to suggest that the target of conceptual engineering is not so much concepts as well as conceptions of such concepts (cf. Löhr 2023b). What the conceptual engineer proposes is not a new

¹Although not explicitly mentioned, a function in this context also seem to be guiding design choices for self-driving cars.

concept but rather a new conception of a concept, at least according to the description above. A conception might here be roughly understood as a specific interpretation and operationalization of a concept (cf. Löhr 2023b). It represents a set of beliefs and typical practices of use or application of the concept.

While I am sympathetic to the idea that conceptions rather than concepts are the target of conceptual engineering and while I am well aware that there is no agreement in the literature on what the proper target for conceptual engineering is (e.g. Löhr 2024), for the purpose of this article, I focus on concepts as the target of conceptual engineering. In doing so, I follow the proposal of Isaac (2023), who not only suggests that concepts are the proper target of conceptual engineering but also that we should employ a psychological rather than a philosophical notion of concepts (cf. Machery 2009). Whereas philosophical accounts take concepts to be the basic elements or building blocks of thought (or propositional attitudes), psychologists understand them in cognitive terms, for example as what has been called ‘categorization devices’ (Löhr 2020). In fact, on such a psychological understanding of concepts, concept are similar to what others have called conceptions (Löhr 2023b). Still, to avoid unnecessary complications and confusion, I will talk here about concepts as the target of conceptual engineering.

For now, the above characterization suffices as a description of the background against with I want to advance my claim that there it is useful to distinguish between two categories of concepts, namely generic and specific concepts and that while specific concepts can be conceptually engineered this is much harder, if not impossible, for generic concepts.

3. Generic versus specific concepts

In the previous section, we have seen how conceptual engineering can be used as a method to adapt concepts such as *CONTROL* in the context of self-driving cars in response to technological developments and (moral) problems raised by such developments. However, it would seem that in this specific example what is conceptually engineered is not so much the concept *CONTROL tout court* but rather its specific use or application in the context of self-driving cars.

To distinguish between such more local, contextual instances of conceptual engineering and more general ones that aim to revise the concept *tout court*, i.e. in all contexts and practices where the concept is used, I will distinguish between what I will call specific concepts and generic concepts.

I propose to understand specific concepts as concepts that are used in specific contexts and practices, and that fulfill a function in these practices that can be delineated and – more or less precisely – be described. As we have seen in section 2, this function can then function as the normative

reference point for conceptual engineering. Oftentimes, specific concepts are developed by specialists like philosophers, scientists, and other professionals like lawyers or medical professionals, and this development typically involves some form of conceptual analysis or definition.

Generic concepts are defined as concepts that are used in a wide variety of practices or contexts and that can and do fulfill diverse functions in these contexts and practices. They are not only used by experts or specialists but also by lay people. Below, I will further elaborate this distinction between generic and specific concepts. First, I will use *CONTROL* as an example of a generic concept that is related to a number of more specific concepts like *DEMOCRATIC CONTROL* (section 3.1). Then, I will explain in more detail what in my view distinguishes generic from specific concepts and how the generic/specific distinction relates to, but is different from two other distinctions that have been made in the literature, namely between thin and thick concepts and between concepts and conceptions (section 3.2). Finally, I will elaborate on the functions that generic concepts play in larger conceptual repertoires (section 3.3).

3.1. *Control as an example*

An example of a generic concept that can and does fulfill different functions in different contexts and practices is *CONTROL*. As we have seen in the context of self-driving cars, it fulfills functions like risk reduction and attributing responsibility, which provides a normative reference point for its re-engineering in that context. However, this combination of functions fulfilled by *CONTROL* seems rather specific to the context of self-driving cars.² In other contexts, *CONTROL* fulfills other (combinations of) functions. For example, in more general discussions about responsibility, it also fulfills the function of attributing responsibility, but not the function of risk reduction. This has also led to proposals for more specific concepts in that context. For example, Fischer and Ravizza (1998) have proposed the more specific concepts of *GUIDANCE CONTROL* and *REGULATIVE CONTROL* to attribute responsibility.

Another context in which the generic concept of *CONTROL* is used is in engineering more generally. One may, for example, look at how the concept *CONTROL* is used in control engineering (e.g. Ogata 2010). Here, more specific concepts have been developed for describing different ways that an engineering system can be controlled like *OPEN LOOP CONTROL* (in which certain input parameters to the system can be changed/controlled) and *CLOSED LOOP CONTROL* (which also contain a feedback mechanism from outcomes of the system to the inputs, like in the case of a thermostat).

²Or more precisely – to a somewhat broader but still delineated context of the engineering of so-called autonomous systems (also including e.g., drones and various forms of AI).

Also more specific concepts for CONTROLLABILITY have been proposed like STATE CONTROLLABILITY and OUTCOME CONTROLLABILITY (Kreindler and Sarachik 1964). These more specific control concepts serve functions like distinguishing between systems that can be controlled and systems which cannot, and ensuring that the outcomes (or states) of a system remain within certain acceptable boundaries.

Yet another context in which the generic concept of CONTROL is applied is democratic decision-making. In this context, control may be understood in terms of the relation between the demos and collective decisions. For example, Pettit (2008) distinguishes between three more specific conceptions for DEMOCRATIC CONTROL: popular causal influence, popular intentional direction and popular institutional control. In this context, the more specific concept of DEMOCRATIC CONTROL fulfills functions like evaluating and shaping democratic systems and institutions.

In still other contexts, CONTROL may be seen as a crucial condition for human well-being, e.g. Nussbaum (2000) has defined 'control over one's environment' as one of the 10 basic human capabilities.³ Here CONTROL is closely related to a number of other concepts like DEMOCRACY, AUTONOMY, AGENCY, CHOICE and WELL-BEING and it fulfills a function like helping to articulate what a worthwhile and dignified human life would consist of.

3.2 What distinguishes generic from specific concepts?

The key defining difference between generic and specific concept is twofold, namely (1) whether the concept is used in multiple practices and contexts (generic concept) or in a specific context or practice (specific concept) and (2) whether the concept is used for multiple aims and functions (generic concept) or for a limited number of well-delineated functions (specific concept). At least theoretically, these two aspect may come apart and I would therefore propose to call a concept generic if it is at least generic in one of these two senses, i.e. either because it lacks a well-delineated function or because it is used in many different contexts and practices.

As this definition already suggests, I take the distinction between generic and specific concepts not to be binary but gradual. Concepts can be more or less generic, and more or less specific. For example, a concept like DEMOCRATIC CONTROL is more specific than CONTROL but at the same time more generic than, for example, DEMOCRATIC ELECTORAL CONTROL. It should also be noted

³Such control according to her has two dimensions: 'A. Political. Being able to participate effectively in political choices that govern one's life; having the right of political participation, protections of free speech and association. B. Material. Being able to hold property (both land and movable goods), not just formally but in terms of real opportunity; and having property rights on an equal basis with others; having the right to seek employment on an equal basis with others; having the freedom from unwarranted search and seizure' (Nussbaum 2000, 80).

that sometimes the same term can be used to refer to both a generic and specific concept. For example, we should distinguish between the generic concept DEATH and the more specific concept DEATH_{MEDICAL} as it is used by medical professionals for certain medical (and legal) purposes.

While the examples in section 3.1 may suggest that specific concepts are derived from generic concepts, this need not necessarily be the case. There may even be specific concepts without a generic concept (and vice versa). An example would seem to be the specific concept of MANSPLAINING that does not seem to have a corresponding generic concept.⁴ Another possibility is that a specific concept is related to two (or more) generic concepts, like the specific concept ARTIFICIAL AGENT that would seem to refer to the generic concept AGENT as well as to the generic concept ARTIFICIAL.

It may also be worth pointing out how the specific/generic distinction relates to two other distinctions that are often made by philosophers, namely that between thin and thick concepts and that between concepts and conceptions.

Thin concepts are evaluative concepts that lack descriptive content, while thick concepts have both evaluative and descriptive content (Väyrynen 2025). A typical example of a thin concept is GOOD, while BRAVE would be a typical example of a thick concept. A first thing to note is that the thin/thick distinction is different from the specific/generic distinction in two ways. First, it is a distinction between different evaluative concepts while the generic/specific distinction also applies to non-evaluative concepts. Second, the defining characterization of the thin/thick distinction is different from the specific/generic distinction. The first distinction is defined in terms of the content of the concepts,⁵ the second in terms of their use(s) and function(s). So the two distinction are clearly different in intension. Nevertheless, they may extensionally map on each other. It seems reasonable to expect that thin concepts are usually more generic than thick ones because their lack of descriptive content means that they can more easily be used in a multiplicity of (very) different contexts and practices, whereas thick concepts may be tied to specific contexts and practices.

The generic/specific distinction also needs to be distinguished from the distinction between concepts and conceptions. A conception might be understood as a specific interpretation and operationalization of a concept (see also section 2).⁶ So 'justice as fairness' (Rawls 1999) is a conception of

⁴I thank an anonymous reviewer for the suggestion and the example. Also note that concepts that start out as specific may become more generic when they start to be used in different contexts or for different purposes than initially.

⁵There is no agreement in the literature how to exactly understand the thin/thick distinction, but the idea that thin concepts only have evaluative content while thick concepts have both evaluate and descriptive content seems a common denominator (Väyrynen 2025).

⁶There is no agreement on how the concept-conception should be exactly understood and some have argued that the distinction is irrelevant to conceptual engineering (Isaac 2023). Here I following Löhr (2023b) in how I understand conceptions.

the concept JUSTICE but it is not a specific concept itself. While a (generic) concept like JUSTICE may both come with different conceptions as well as with different (more) specific concepts, the two should be distinguished. A conception is an interpretation (or conceptualization) of a (generic) concept, while a specific concept is another concept that serves a more specific function for a specific context than the generic concept. (The specific concept may itself have different conceptions.) So, whereas different conceptions of a concept are usually competing interpretations of that concept, specific concepts are not competing but are different specifications of a generic concept, tailored to fulfill specific functions in specific contexts. Specific concepts can usually co-exist without any contradiction, while different conceptions of the same concept often contradict each other. Moreover, as we have seen there may be specific concepts without a corresponding generic concept, while a conception without a corresponding concept is impossible.

3.3 The role played by generic concepts in larger conceptual repertoires

I have tried to make it plausible that in addition to specific concepts that are used in specific contexts for specific purposes or functions, there are generic concepts that can be used in different contexts for a multiplicity of aims or functions. Examples are such concepts as CONTROL, AGENCY, FREEDOM, NATURE and RESPONSIBILITY. When these generic concepts are applied in specific contexts, they may be translated into more specific concepts. In the case of the generic concept CONTROL, examples of such more specific concepts are – as we have seen – GUIDANCE CONTROL, REGULATIVE CONTROL, DEMOCRATIC CONTROL, OPEN LOOP CONTROL and CLOSED LOOP CONTROL.

Given that generic concepts like CONTROL when applied in different contexts seem to get a somewhat different meaning or interpretation, and may be translated in more specific concepts, one might wonder whether (and why) we need generic concepts. It might even be argued that generic concepts can cause a lot of confusion and misinterpretation, and that we are better off with only what I have called specific concepts.

For example, Nado (2011b) has suggested that we may be better off with different specific concepts of TRUTH than with one generic concept.⁷ Similarly, when computer scientists are talking about the AGENCY of an ARTIFICIAL AGENT,

⁷Nado (2011b, 2) argues that ‘many of our ordinary pre-engineering concepts are ‘multitaskers’, playing multiple roles or serving multiple functions. And a concept that is a ‘jack-of-all-trades’ is, arguably, likely to be a master of none. If this is correct, then we may find our purposes better served by a plurality of concepts, each custom-designed to best fill one of the original concept’s roles.’ I do not disagree that having multiple (specific) concepts can be useful, but I think generic concepts (that are ‘jacks-of-all-trades’) are also quite useful and indeed indispensable, as I will argue below.

they mean something quite different than when (moral) philosophers talk about the AGENCY of a HUMAN AGENT. Perhaps in such cases, we should prefer to use the specific concepts of HUMAN AGENCY and ARTIFICIAL AGENCY rather than the generic concept AGENCY.⁸

There seems to be certainly a grain of truth in such an argument. However, it is important to recognize that generic concepts also play important roles or functions in our larger conceptual repertoires and in human practices, and for that reason may be indispensable. One thing is that generic concepts can be applied in new contexts and situations, and therefore allow for cognitive efficiency and flexibility. Without the possibility to apply generic concepts to new situations and contexts, humans may be lost in the world. This is not to say that new situations or contexts can – in the longer run – not lead to new specific concepts, but initially the best reaction may be to apply a generic concept.

Generic concepts thus play a number of roles in our larger conceptual repertoires and in human practices. I want to highlight four of such roles.

First, generic concepts can be used for different purposes or for different functions in different contexts. This flexibility also allows the use of generic concepts in new situations or new contexts that humans have not encountered before. Generic concepts thus allow for cognitive (or mental) flexibility and efficiency.

Second, generic concepts enable the creation of more specific concepts for specific contexts and practices. For example, the generic concept CONTROL has enabled the creation of more specific concepts like GUIDANCE CONTROL, REGULATIVE CONTROL, DEMOCRATIC CONTROL, OPEN LOOP CONTROL and CLOSED LOOP CONTROL.

Third, generic concepts help to connect different human practices. For example, the generic concept of CONTROL allows us to see connections between democratic practices and practices of responsibility attribution. Such connections may help to evaluate and improve human practices. They may also contribute to a more integral picture or understanding of human practices, which may be worthwhile in itself as well as contribute to (individual) human well-being and more just societies.

Fourth, generic concepts are embedded in larger conceptual repertoires and may play key roles in such repertoires. Among others, this means that they have certain (inferential) relations with other (generic) concepts. For example, the generic concept of CONTROL is related to concepts like RESPONSIBILITY, DEMOCRACY, AGENCY, and AUTONOMY. In this way generic concepts may also contribute to coherence in how we perceive and think about the world.

⁸And perhaps we should phrase the concept of artificial agency not even in terms of the word 'agency' to avoid confusion.

4. Why it is much harder (if not impossible) to conceptually engineer generic concepts

I will now argue that generic concepts are much harder to conceptually engineer than specific concepts, even to the extent that we might wonder whether it is meaningful to speak in any straightforward sense about the conceptual engineering of generic concepts. Below, I will present four arguments why it is hard and perhaps even impossible to conceptually engineer generic concepts. First, generic concepts lack any well-delineated function that can serve as normative reference point for their (re)engineering. Second, generic concepts cannot be engineered in isolation because they are part of larger conceptual repertoires. Third, the meaning of generic concepts is beyond our control or at least the control of (individual) conceptual engineers. Fourth, the conceptual engineering of generic concepts is bound to be controversial and will raise questions about who has the authority to engineer such concepts.

I think that each of these arguments shows why conceptual engineering is much more difficult for generic concept than for specific concepts. Perhaps each of the arguments does in itself not show the impossibility of conceptual engineering for generic concepts but it is important to see that they hang together and that they together make a strong case for why it is hard and perhaps even impossible to conceptually engineer generic concepts.

4.1 Generic concepts lack a clearly delineated function that can serve as (normative) basis for their (re)engineering

Generic concepts fulfill – by definition – different roles or functions in different contexts. In addition, they fulfill more generic functions due to their role in larger conceptual repertoires like enabling specific concepts, connecting different human practices and making (inferential) relations to other concepts.

In contrast to generic concepts, specific concepts have a clearly delineated function and – as we have seen – this delineated function can serve as normative reference point for conceptual engineering. Even if generic concepts may be said to fulfill certain functions, it would seem that these functions do not provide a clear normative reference point because they are so diverse and not well-delineated. Conceptual engineering is therefore much harder for generic concepts than for specific concepts.

It might be objected that if generic concepts fulfill functions, these functions can still serve as normative basis for conceptual engineering. One might – for example – suggest that we can derive requirements from the various functions or potential functions that generic concepts can fulfill. However, such an approach suggests that the functioning of generic

concepts can be isolated from the functioning of other (generic) concepts and larger conceptual repertoires. But I think that that is impossible and this brings me to the second argument.

4.2 Generic concepts cannot be engineered in isolation

As we have seen, generic concepts are part of larger conceptual repertoires and they have (inferential) relations with other concepts. This means that if we change a generic concept through conceptual engineering, we will also change its relation to other concepts, and in this way the larger conceptual repertoire or at least a larger cluster of closely related concepts.

Nevertheless, it could be argued that the same problem applies to specific concepts. After all, specific concepts have also inferential relations with other concepts and are also part of larger conceptual repertoires and clusters. However, in the case of specific concepts with clearly delineated functions, it would seem very well possible to isolate the functioning of the specific concept from the larger conceptual repertoire or cluster of which it is part. In terms of conceptual engineering or design, we might simply treat all related concepts as boundary conditions or constraints for the (re)design of a specific concept. In other words we may assume that the surrounding conceptual repertoire remains the same and, on that assumption, it would seem very well possible to engineer the specific concept.

However, this approach does not seem to work in the case of generic concepts because of their tight connection with other concepts and their embedding in larger conceptual repertoires. Even if it might still be possible to do proposals for a revision of generic concepts, these proposals are likely not to be effective (see also my next argument). Moreover even if they would be effective, such proposals are likely to do more harm than good because they neglect the effect of the revision of generic concepts on other concepts and on larger conceptual repertoires. In other words, generic concepts are typically so well entrenched in larger conceptual repertoires that conceptually engineering them in isolation is either impossible or harmful.

As an example consider the case of brain death. Due to the invention of the mechanical ventilator in the 1950s, it became possible to pump air into patients' lungs and to sustain heart activity even when patients had no brain activity. The question arose whether these comatose patients should be considered death or alive (Baker 2019; Nickel, Kudina, and van de Poel 2022). Marchiori and Scharp (2024, 9) interpret this as a conceptual disruption at the level of what they call a conceptual cluster (somewhat similar to what I have called a conceptual repertoire) because it disrupts (seemingly) the (generic) concepts of DEATH and LIFE.

This conceptual disruption was eventually addressed through the (de novo) conceptual engineering of a new specific concept, that of BRAIN

DEATH (see Löhr 2022, 845). BRAIN DEATH is a specific concept because it was designed (by the Harvard Committee that proposed it) for a number of well-delineated functions, like answering questions about what Belkin (2003) calls bed-side considerations, like how comatose patients should be treated, as well as questions about organ transplantation, as it was considered immoral and illegal to transplant organs from patients that are still alive (Giacomini 1997). Moreover, the concept was intended for use in specific medical contexts, namely that of comatose patients.

One might nevertheless wonder how it was possible to design BRAIN DEATH as a new specific concept given its inferential relations with other concepts and the fact that a larger conceptual cluster was disrupted. The answer, I believe, lies in the fact that no direct attempts were made to redesign the generic concepts LIFE and DEATH. The Harvard Committee Report ‘did not provide a worked-out, consistent notion of death, or life, itself. Rather it outlined the permission granted by an irreversible coma’ (Belkin 2003, 349). It did so by introducing the new concept of BRAIN DEATH. This reason why this new specific concept could be successfully introduced was (I claim) that the Harvard Committee successfully decoupled it from broader conceptual questions (and inferential relation) about LIFE and DEATH. If this interpretation of the brain death example is right, it illuminates why it is easier to conceptual engineer a specific concept like BRAIN DEATH than generic concepts like DEATH and LIFE.⁹

4.3 The meaning of generic concepts is beyond the control of conceptual engineers

The third argument is that the meaning of generic concept is beyond the control of conceptual engineers, which makes them inaccessible for conceptual engineering. It might be objected that the same argument applies to specific concepts. Indeed, it has been suggested in the literature on

⁹My interpretation of this case here deviates from the one offered in Marchiori and Scharp (2024, 9–10). They suggest that the response to the conceptual disruption proceeded in two steps, the first implying a change in the meaning of the concept DEATH. However, there is ample empirical evidence that in as far as the concept of DEATH was changed it was a (unintended) consequence of the Harvard Report rather than its goal (Giacomini 1997; Belkin 2003; Nickel, Kudina, and van de Poel 2022). Moreover, what was affected (I would argue) is the specific concept DEATH_{MEDICAL} rather than the generic concept DEATH. In fact, the Harvard Committee explicitly refrained from redefining death. Also, when anesthesiologists turned to Pope Pius XII in 1957 for an answer to the question whether ‘a patient plunged into unconsciousness ... whose ... blood circulation ... is maintained through artificial respiration be considered *de facto* or even *de jure* dead?’ (Baker 2019, 145), Pius XII declined to answer the question, referring back to the medical profession for an answer (Giacomini 1997; Baker 2019) suggesting that what was at stake was not a redefinition of DEATH or LIFE *tout court* (i.e., the generic concepts) but rather their meaning in a very specific medical context (i.e., DEATH_{MEDICAL}). In as far, the Harvard Committee Report was interpreted as a redefinition of the generic concept of DEATH it was not met with general enthusiasm and appraisal, while the specific concept of BRAIN DEATH, and its implications for the treatment of comatose patients were quickly broadly adopted (Nickel, Kudina, and van de Poel 2022).

conceptual engineering that implementing new or revised concepts is a problem due to a lack of control (Cappelen 2018). This has become known as the implementation challenge. In response it has been suggested that the implementation challenge may be addressed by, for example, changing speaker meaning (Pinder 2020) or by implementing new concepts in specific context like engineering design (Veluwenkamp 2025).

I will not to discuss the general challenge of implementation for conceptual engineering, but it seems to me obvious that the implementation of generic concepts is much more difficult than of specific concepts. One reason is that engineered generic concepts affect multiple contexts and are therefore much more difficult to implement than engineered specific concepts that affect only a specific context or practice. Another reason is that the meaning of generic concept is often tacit and implicit. It might often not be easy to explicate in words what we mean with a generic concept and any attempt to do so might at least partly fail. This seems to me quite different for specific concepts which are often deliberately introduced to fulfill specific function in specific contexts. They often result from deliberate attempts to define a concept and to make that concept fulfill a particular function. It would seem therefore that such concepts and their meaning are much more under the control of conceptual engineers than generic concepts.

4.4 Attempts to engineer generic concepts are likely to be (more) controversial and it is (more) unclear who has the authority or legitimacy to engineer them.

Conceptual engineering proposals for generic concepts are likely to give rise to controversy or to assume agreement where we should accept disagreement, plurality and political struggle. Consider, for example, the debate between Williams (2005) and Dworkin (2011) about the generic concept of LIBERTY (see also Queloz 2022). Whereas Dworkin's proposes a generic concept of LIBERTY that alleviates the tensions with JUSTICE, Williams criticizes that attempt because it denies or suppresses underlying political disagreement and struggle.

This suggests that attempts to conceptually engineer generic concepts such as LIBERTY may not only lead to controversy, but that they might also miss the mark in the following sense. We should not aim to solve tensions between values like LIBERTY and JUSTICE through conceptual engineering by philosophers (or other conceptual engineers), but instead we should leave the disagreement about such concepts in the realm where it belongs, the political realm. More generally, we should perhaps accept that at least some generic concepts are essentially contested (Gallie 1955), so that the idea

that we can provide generally accepted generic concepts through conceptual engineering is a fundamental mistake.¹⁰

Related to this is what has become known as the authority problem in conceptual engineering (Queloz 2022, 2025). Why should concept-users recognize the authority of the concepts proposed by conceptual engineers? This question raises further issues about the legitimacy of conceptual engineering. Who has the authority and legitimacy to conceptually engineer concepts? It might be argued that this challenge applies to both specific and generic concepts, but it is clearly more severe for generic concepts. The reason why the authority problem arises less for specific concepts is twofold. First, specific concepts have a clearly delineated function which can serve as a (shared and authoritative) normative reference point for conceptual engineering. Second, specific concepts are often developed by specific professionals or practitioners for a specific context, which gives these professionals or practitioners at least some authority and legitimacy in conceptually engineering the specific concept.¹¹ One might think that for generic concepts such an authoritative role may perhaps be played by democratic institutions like the government or parliament. However, in a liberal democracy, it would seem undesirable for governments to undertake such forms of conceptual engineering (Queloz and Bieber 2022). In other words, conceptual engineering of generic concepts by, for example, governments may not only be technocratic and controversial but ultimately also authoritarian.

5. Conceptual design versus intervention in conceptual repertoires

I have argued that the more generic a concept is, the harder it is to engineer it. But that does not mean that generic concepts cannot change. They certainly can and do. We can only not deliberately redesign them, at least not in a meaningful and straightforward way. But what we can (try to) do is to change them more indirectly; such attempts can best be seen as interventions in larger conceptual schemes or repertoires rather than as the pinpointed redesign of generic concepts.

The distinction I have in mind here between conceptual *design* and *intervening* in larger conceptual repertoires is similar to the distinction that has been made in (the philosophy of) technology between the *design* of technical artifacts and *intervening* in sociotechnical systems of which these artifacts are

¹⁰For an argument why we should not aim or try to define the generic concept of LIFE see Machery (2012).

¹¹Examples are the Harvard Committee in the case of BRAIN DEATH and International Astronomical Union for the case of PLANET.

part (e.g. design of a car versus an intervention in the transportation system) (cf. Vermaas et al. 2011).

Sociotechnical systems cannot be designed (in their entirety) because they already exist and are beyond the control of designers; we can intervene in them, but not completely design them anew (Kroes 2009). Similarly, we can conceptually engineer specific concepts, relatively independent from the larger conceptual repertoires they are part of, but if we want to change generic concepts, we should try to intervene in larger conceptual repertoires and the engineering metaphor breaks down. We are making social interventions rather than designing or engineering a concept.

Below in section 5.1, I will argue in more detail why generic concepts cannot be designed building on the arguments given in section 4. In section 5.2, I will argue that the distinction between design and intervention is also relevant to how we understand the success of conceptual engineering projects and that some earlier conceptual engineering projects (like Haslanger's ameliorative project) can be better understood in terms of intervention than design.

5.1 The impossibility of designing generic concepts

The four arguments given in section 4 point to two more fundamental reasons that make it hard or even impossible to design generic concepts. The second and third argument point to a lack of isolation and a lack of control, which, as I will argue below, are both conditions for the practical feasibility of design. The first and fourth argument both point to the lack of a legitimate normative basis for design. These arguments make design not impossible but normatively questionable. My overall argument may be read as stating that in as far as design of generic concepts (or at least attempts to it) might be practically feasible they are likely to be normatively questionable.

To see why it is practically infeasible to design generic concept consider the following two features of the design paradigm (Kroes 2009, 536):

- *Isolation.* Design assumes that 'it is possible to clearly separate the object of design from its environment' (Kroes 2009, 536). This does not mean that the object of design can function completely independent from its environment, but it should be possible to express the (functional) dependence on the object's environment in terms of boundary conditions or constraints. For example, an electric device requires electricity but its dependence on the electricity net may be captured in terms of boundary conditions.
- *Control.* The designers, or those implementing the design should have enough control to successfully realize the design. Again, this does not

require full control, but at least enough control for the successful implementation of the design.

Both conditions are not met for generic concepts. Generic concepts do not meet the condition of isolation because, as I have argued, generic concepts cannot be engineered in isolation, and the control condition is not met because, as I have argued, the meaning of generic concepts is beyond the control of conceptual engineers.

Now one might argue that a lack of isolation and control does not make it impossible for conceptual engineers to develop design proposals for generic concepts. They make it only less likely that conceptual engineers will succeed. However, there are also normative reasons that make it undesirable or even impermissible to design generic concepts. First, generic concepts lack a well-delineated function that can serve as (shared, accepted) normative basis for design and designing generic concepts therefore raises an authority problem. Due to a lack of a legitimate normative basis, conceptual engineers who design generic concepts without consulting the target audience (which in the case of generic concepts may be the entire population) may de facto misuse their authority, which is normatively questionable and seems democratically illegitimate.

5.2 Understanding the success of conceptual engineering projects: design versus intervention

The distinction between conceptual design and intervention is also relevant for how we judge the success of conceptual engineering projects. Consider, for example, Pinder (2022)'s discussion of whether Haslanger's ameliorative project (Haslanger 2000, 2020) has been a successful conceptual engineering project. He offers three interpretations of the (aims of the) project; a revisionary, a revelatory and a terminological interpretation.

On the revelatory interpretation, 'Haslanger's ameliorative project is to reveal (and analyse) the semantic meaning of terms such as "woman", "man", etc.' (Pinder 2022, 6). However, on that interpretation, the project does not seem to be a conceptual engineering project at all (Pinder 2022, 9) and I will therefore not further consider this interpretation here.

On the revisionary interpretation 'Haslanger is seeking to revise how people use and understand "woman", "man", etc. in the pursuit of social justice. Haslanger's definitions are the end goals: they capture the intended post-amelioration meaning of these terms' (Pinder 2022, 4). As he points out on this interpretation, the project has been unsuccessful, and it is also hard to see how it could succeed.

On this revisionary interpretation, Haslanger's project is aimed – in my terminology – at the conceptual redesign of generic concepts like WOMAN and

MAN. It is not hard to see how all four arguments for why it is hard if not impossible to engineer generic concepts that I have discussed in the previous section apply to it. First, generic concepts like WOMAN and MAN do not have a clearly delineated function. While it seems to make sense to re-engineer them in the light of social justice, it seems naïve (and undesirable) to think we can simply ignore all the other relevant functions they now fulfill (for better or worse) in larger conceptual repertoires. Second, these generic concepts are deeply entrenched in larger conceptual repertoires and it would seem impossible to change them in isolation, without changing these larger repertoires. Third, as also Pinder (2022) points out we lack the required control to revise such generic concepts. Finally, as Haslanger herself is very well aware, appropriating these concepts may be unwarranted and potentially harmful (Haslanger 2020, 236), which also raises questions about the authority to change such generic concepts. For all these reasons, it seems clear that, and why, the project understood as a revisionary project aimed at conceptual design of generic concepts is bound to fail.

However, Pinder (2022) also offers a third interpretation. On this terminological interpretation,

Haslanger is *not* seeking to change how people ... *in general* use and understand 'woman', 'man', etc. The theory she develops ... is better thought of as a *tool* The function of the tool is to expose people to a different way of thinking about gendered and race categories, to cause (or jolt) them to see those categories in a new light – a light that reveals ... injustices. (Pinder 2022, 13)

On this interpretation, Haslanger's project is not an instance of conceptual design but rather an *intervention* in a larger conceptual repertoire. This intervention is moreover not purely conceptual (or linguistic), but it is a social and political intervention. As Haslanger writes:

The thought is that agents who come to understand the historical and political context of their gender and race identification and the role of their identification in perpetuating their own oppression and the oppression of others, will be taking a first step in a process of emancipation. (Haslanger 2020, 236)

This intervention might – ultimately – lead to a change in the generic concepts of WOMAN and MAN, but that effect is the indirect result of an intervention.

6. Conclusion and implications

I have argued that when it comes to the possibilities and limitations of conceptual engineering, it is useful to distinguish between specific and generic concepts. I have provided four arguments why generic concepts are much harder – if not impossible – to conceptually engineer. First, they lack a

well-delineated function that can serve as normative reference point for conceptual engineering. Second, they cannot be engineered in isolation because they are usually deeply entrenched in larger conceptual repertoires. Third, aspiring conceptual engineers lack control over the meaning of generic concepts, and fourth, in as far as conceptual engineering with respect to generic concepts is feasible and can be attempted it is likely to lead to controversy because it is unclear who has the authority to conceptually engineer generic concepts.

Although the distinction between specific and generic concepts may be one of degree rather than an absolute one, it is a distinction that makes a difference, when it comes to the possibilities of conceptual engineering. While generic concepts can be changed, attempts to do so are best understood as interventions in conceptual repertoires rather than in terms of conceptual design.

My argument has three kinds of more general ramifications. First, it has implications for the field of conceptual engineering itself, as it shows certain limitations to what can be subject to conceptual engineering and what not. Perhaps, some would still maintain that also generic concepts can be subject to conceptual engineering and that we should not take the engineering metaphor too literally. So perhaps we can keep talking about both conceptual *design* and the *intervention* in larger conceptual schemes as instances of *conceptual engineering projects* (as Pinder (2022) seems to suggest). However, there also seem to be good reasons to avoid the engineering metaphor altogether when it comes to interventions in larger conceptual schemes that might change generic concepts. First, such interventions are obviously not just conceptual or linguistic, but also social and political in nature. Second, the engineering metaphor suggests that such interventions are instrumental and controlled and are undertaken by a specific class of people, namely conceptual engineers. That assumption may not only be wrong but also lead to potentially harmful appropriations of generic concepts.

A second implication concerns the ethics of socially disruptive technologies, in particular how we can and should react to conceptual disruption (Hopster et al. 2023; Löhr 2023a, 2023b; Marchiori and Scharp 2024; Veluwenkamp et al. 2024). On my view, both specific and generic concepts can get disrupted. But the implications seem different. In the case of disruption of specific concepts, it seems relatively easy to undertake local repair work through conceptual engineering. Of course, implementation of new or adapted specific concepts may not always be easy or uncontroversial, but the task of the conceptual engineer is relatively straightforward.

This seems different for generic concepts. For one, there are reasons to think that generic concepts do not get so easily disrupted. As I have argued, they are typically quite flexible and can often be applied to new situations. Still, it would seem that if technological or other developments disrupt

generic concepts the effects are much bigger. What is disrupted in such cases would typically not be a single generic concept but a conceptual cluster or an entire conceptual repertoire.

Such a possibility has indeed been recognized in the literature on conceptual disruption. For example, Marchiori and Scharp (2024, 6) distinguish disruption at the level of single concepts from that at the level of conceptual clusters, and within the latter, they distinguish between first-order and second-order disruption. In the case of second-order disruption, concepts get indirectly disrupted, i.e. through the disruption of other concepts.

As the earlier given example of brain death suggests conceptual disruption, even it occurs at the level of conceptual clusters and has ripple effects (Löhr 2023a) on other concepts, can sometime be addressed through the (de novo) design of a specific concept (like BRAIN DEATH). In other words, the conceptual engineering efforts may be focused on specific concepts, which in turn may also affect other specific concepts or even generic concepts as a side effect.

However, there are likely also cases in which the disruption cannot be repaired through the conceptual (re)design of specific concepts. This may, for example, happen in the case of deep conceptual disruptions, that ‘challenge basic categories and concepts of thought’ (Hopster 2021, 2). In such cases, conceptual (re)design of specific concepts may not suffice. However, this does not mean that we are empty-handed; we might still try to intervene in conceptual repertoires as I have argued, but this task is more daunting.

If this argument is right, we have reasons to be more concerned about the disruption of generic concepts than of specific concepts, because the latter allow for local repair work. If this is indeed true, it raises further questions about how disruptions of generic concepts come about, how harmful and undesirable (or beneficial and desirable) such disruptions are, and whether we have reasons to try to prevent the disruption of certain generic concepts, for example by putting constraints on technological development.

A final implication concerns the field of conceptual ethics. It is clear that changing either specific or generic concepts can raise ethical questions and hence that there is a need for conceptual ethics, but one might wonder whether interventions in conceptual repertoires with the aim to change or adapt generic concepts do not also raise questions about power, legitimacy and justice, i.e. questions of politics, not just of ethics. Answering such questions may therefore require a political philosophy of conceptual intervention.

Acknowledgements

The author thanks the members of the ESDiT Research Line ‘Conceptual Disruption’, as well as participants to the 4TU.Ethics/ESDiT Conference Rethinking Ethics – Reimagining Technology that was held in October 2024 for comments on earlier versions.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This publication is part of the research program Ethics of Socially Disruptive Technologies, which is funded through the Gravitation program of the Dutch Ministry of Education, Culture, and Science and the NWO [grant number 024.004.031].

ORCID

Ibo van de Poel  <http://orcid.org/0000-0002-9553-5651>

References

- Baker, Robert. 2019. *The Structure of Moral Revolutions: Studies of Changes in the Morality of Abortion, Death, and the Bioethics Revolution*. Cambridge, MA: The MIT Press.
- Belkin, Gary S. 2003. "Brain Death and the Historical Understanding of Bioethics." *Journal of the History of Medicine and Allied Sciences* 58 (3): 325–361. <https://doi.org/10.1093/jhmas/jrg003>.
- Cappelen, Herman. 2018. *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.
- Chalmers, David J. 2020. "What Is Conceptual Engineering and What Should It Be?" *Inquiry* 63: 1–18. <https://doi.org/10.1080/0020174X.2020.1817141>.
- Dworkin, Ronald. 2011. *Justice for Hedgehogs*. Cambridge, MA: Belknap Press of Harvard University Press.
- Fischer, John Martin, and Mark Ravizza. 1998. *Responsibility and Control: A Theory of Moral Responsibility*, *Cambridge Studies in Philosophy and Law*. Cambridge: Cambridge University Press.
- Gallie, W. B. 1955. "Essentially Contested Concepts." *Proceedings of the Aristotelian Society* 56:167–198. <https://doi.org/10.1093/aristotelian/56.1.167>.
- Giacomini, Mita. 1997. "A Change of Heart and a Change of Mind? Technology and the Redefinition of Death in 1968." *Social Science & Medicine* 44 (10): 1465–1482. [https://doi.org/10.1016/S0277-9536\(96\)00266-3](https://doi.org/10.1016/S0277-9536(96)00266-3).
- Haslanger, Sally. 2000. ". "Gender and Race: (what) Are They? (What) Do We Want Them to Be?" *Noûs* 34 (1): 31–55. <https://doi.org/10.1111/0029-4624.00201>.
- Haslanger, Sally. 2020. "Going On, Not in the Same Way." In *Conceptual Engineering and Conceptual Ethics*, edited by Alexis Burgess, Herman Cappelen, and David Plunkett, 230–260. Oxford University Press.
- Himmelreich, Johannes, and Sebastian Köhler. 2022. "Responsible AI through Conceptual Engineering." *Philosophy & Technology* 35 (3): 1–30. <https://doi.org/10.1007/s13347-022-00542-2>.
- Hopster, J., P. Brey, M. B. O. T. Klenk, G. Löhr, S. Marchiori, B. Lundgren, and K. Scharp. 2023. "Conceptual Disruption and the Ethics of Technology." In *Ethics of Socially Disruptive Technologies: An Introduction*, edited by I. van de Poel, L. Frank, J. Hermann, J. Hopster, D. Lenzi, S. Nyholm, B. Taebi, and E. Ziliotti, 141–162. Cambridge: Open Book Publishers.

- Hopster, Jeroen. 2021. "What Are Socially Disruptive Technologies?." *Technology in Society* 67: 1–8. <https://doi.org/10.1016/j.techsoc.2021.101750>.
- Hopster, Jeroen, and Guido Löhr. 2023. "Conceptual Engineering and Philosophy of Technology: Amelioration or Adaptation?" *Philosophy & Technology* 36 (4): 1–17. <https://doi.org/10.1007/s13347-023-00670-3>.
- Isaac, Manuel Gustavo. 2023. "Which Concept of Concept for Conceptual Engineering?" *Erkenntnis* 88 (5): 2145–2169. <https://doi.org/10.1007/s10670-021-00447-0>.
- Isaac, Manuel Gustavo, Steffen Koch, and Ryan Nefdt. 2022. "Conceptual Engineering: A Road map to Practice." *Philosophy Compass* 17 (10): 1–15. <https://doi.org/10.1111/phc3.12879>.
- Jorem, Sigurd. 2022. "The Good, the bad and the Insignificant—Assessing Concept Functions for Conceptual Engineering." *Synthese* 200 (2): 1–20. <https://doi.org/10.1007/s11229-022-03548-7>.
- Koch, Steffen, Guido Löhr, and Mark Pinder. 2023. "Recent Work in the Theory of Conceptual Engineering." *Analysis* 83 (3): 589–603. <https://doi.org/10.1093/analysis/anad032>.
- Köhler, Sebastian, and Herman Veluwenkamp. 2024. "Conceptual Engineering: For What Matters." *Mind; a Quarterly Review of Psychology and Philosophy* 133 (530): 400–427. <https://doi.org/10.1093/mind/fzad064>.
- Kreindler, E., and P. Sarachik. 1964. "On the Concepts of Controllability and Observability of Linear Systems." *IEEE Transactions on Automatic Control* 9 (2): 129–136. <https://doi.org/10.1109/TAC.1964.1105665>.
- Kroes, Peter. 2009. "Foundational Issues in Engineering Design." In *Handbook of the Philosophy of Science. Volume 9: Philosophy of Technology and Engineering Sciences*, edited by Anthonie Meijers, 513–541. Oxford: Elsevier.
- Löhr, Guido. 2020. "Concepts and Categorization: Do Philosophers and Psychologists Theorize about Different Things?" *Synthese* 197 (5): 2171–2191. <https://doi.org/10.1007/s11229-018-1798-4>.
- Löhr, Guido. 2022. "Linguistic Interventions and the Ethics of Conceptual Disruption." *Ethical Theory and Moral Practice* 25 (5): 835–849. <https://doi.org/10.1007/s10677-022-10321-9>.
- Löhr, Guido. 2023a. "Conceptual Disruption and 21st Century Technologies: A Framework." *Technology in Society* 74: 1–9. <https://doi.org/10.1016/j.techsoc.2023.102327>.
- Löhr, Guido. 2023b. "Do Socially Disruptive Technologies Really Change Our Concepts or Just Our Conceptions?." *Technology in Society* 72: 1–6. <https://doi.org/10.1016/j.techsoc.2022.102160>.
- Löhr, Guido. 2024. "If Conceptual Engineering Is a new Method in the Ethics of AI, What Method Is It Exactly?" *AI and Ethics* 4 (2): 575–585. <https://doi.org/10.1007/s43681-023-00295-4>.
- Machery, Edouard. 2009. *Doing without Concepts*. Oxford: Oxford University Press.
- Machery, Edouard. 2012. "Why I Stopped Worrying about the Definition of Life ... and why You Should as Well." *Synthese* 185 (1): 145–164. <https://doi.org/10.1007/s11229-011-9880-1>.
- Marchiori, Samuela, and Kevin Scharp. 2024. "What Is Conceptual Disruption?" *Ethics and Information Technology* 26 (1): 1–14. <https://doi.org/10.1007/s10676-024-09749-7>.
- Nado, Jennifer. 2021a. "Conceptual Engineering, Truth, and Efficacy." *Synthese* 198 (7): 1507–1527. <https://doi.org/10.1007/s11229-019-02096-x>.

- Nado, Jennifer. 2021b "Re-engineering Knowledge: A Case Study in Pluralist Conceptual Engineering." *Inquiry* 64: 1–26. <https://doi.org/10.1080/0020174X.2021.1903987>.
- Nickel, Philip J., Olya Kudina, and Ibo van de Poel. 2022. "Moral Uncertainty in Technomoral Change: Bridging the Explanatory Gap." *Perspectives on Science* 30 (2): 260–283. https://doi.org/10.1162/posc_a_00414.
- Nussbaum, M. C. 2000. *Women and Human Development. The Capabilities Approach*. Cambridge: Cambridge University Press.
- Ogata, Katsuhiko. 2010. *Modern Control Engineering*. 5th ed. Boston: Prentice-Hall.
- Pettit, Philip. 2008 "Three Conceptions of Democratic Control." *Constellations* 15 (1): 46–55. <https://doi.org/10.1111/j.1467-8675.2008.00473.x>.
- Pinder, Mark. 2020 "Conceptual Engineering, Speaker-Meaning and Philosophy." *Inquiry* 63: 1–15. <https://doi.org/10.1080/0020174X.2020.1853342>.
- Pinder, Mark. 2022 "Is Haslanger's Ameliorative Project a Successful Conceptual Engineering Project?." *Synthese* 200 (4): 1–22. <https://doi.org/10.1007/s11229-022-03803-x>.
- Queloz, Matthieu. 2021. *The Practical Origins of Ideas: Genealogy as Conceptual Reverse-Engineering*. Oxford: Oxford University Press.
- Queloz, Matthieu. 2022. "Function-Based Conceptual Engineering and the Authority Problem." *Mind* 131 (524): 1247–1278. <https://doi.org/10.1093/mind/fzac028>.
- Queloz, Matthieu. 2025. *The Ethics of Conceptualization. Tailoring Thought and Language to Need*. Oxford: Oxford University Press.
- Queloz, Matthieu, and Friedemann Bieber. 2022. "Conceptual Engineering and the Politics of Implementation." *Pacific Philosophical Quarterly* 103 (3): 670–691. <https://doi.org/10.1111/papq.12394>.
- Rawls, John. 1999 *A Theory of Justice. Revised Edition*. Cambridge, MA: The Belknap Press of Harvard University Press.
- Riggs, Jared. 2021. "Deflating the Functional Turn in Conceptual Engineering." *Synthese* 199 (3–4): 11555–11586. <https://doi.org/10.1007/s11229-021-03302-5>.
- Simion, Mona, and Christoph Kelp. 2020. "Conceptual Innovation, Function First." *Nous* 54 (4): 985–1002. <https://doi.org/10.1111/nous.12302>.
- Strawson, Peter. 1963. "Carnap's Views on Conceptual Systems versus Natural Languages in Analytic Philosophy." In *The Philosophy of Rudolf Carnap*, edited by P. Schlipp, 503–518. La Salle: Open Court.
- Thomasson, A. L. 2020. "A Pragmatic Method for Normative Conceptual Work." In *Conceptual Engineering and Conceptual Ethics*, edited by Alexis Burgess, Herman Cappelen, and David Plunkett, 435–458. Oxford: Oxford University Press.
- Väyrynen, Pekka. 2025. "Thick Ethical Concepts." In *The Stanford Encyclopedia of Philosophy*. edited by Edward N. Zalta, and Uri Nodelman, Spring 2025 Edition. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2025/entries/thick-ethical-concepts/>.
- Veluwenkamp, Herman. 2025 "Impactful Conceptual Engineering: Designing Technological Artefacts Ethically." *Ethical Theory and Moral Practice* 28 (1): 155–170. <https://doi.org/10.1007/s10677-024-10459-8>.
- Veluwenkamp, Herman, Marianna Capasso, Jonne Maas, and Lavinia Marin. 2022. "Technology as Driver for Morally Motivated Conceptual Engineering." *Philosophy & Technology* 35 (3): 1–25. <https://doi.org/10.1007/s13347-022-00565-9>.
- Veluwenkamp, Herman, Jeroen Hopster, Sebastian Köhler, and Guido Löhr. 2024. "Socially Disruptive Technologies and Conceptual Engineering." *Ethics and Information Technology* 26 (4): 1–6. <https://doi.org/10.1007/s10676-024-09804-3>.

- Vermaas, Pieter, Peter Kroes, Ibo van de Poel, Maarten Franssen, and Wybo Houkes. 2011. *A Philosophy of Technology: From Technical Artefacts to Sociotechnical Systems*. Vol. 6, Synthesis Lectures on Engineers, Technology and Society. Cham: Springer.
- Williams, Bernard. 2005 "Conflicts of Liberty and Equality." In *The Beginning Was the Deed: Realism and Moralism in Political Argument*, edited by Geoffrey Hawthorn, 115–127. Princeton: Princeton University Press.