
Teaching and Crafting Human-Robot Relational Ethics

Olya Kudina^{a1}, Joost Mollen^{a2}, Jordi Viader Guerrero^a, Dmitry Muravyov^a and
Juan Pablo Bermúdez^b

^aTU Delft, Delft, Netherlands

^bUniversity of Southampton, Southampton, UK

ABSTRACT

This paper proposes and reflects on a teaching methodology for introducing relational ethics in the engineering curriculum based on a pilot at TU Delft, a technical university in the Netherlands, in a course for robotics engineers. Unlike prevalent models of technology ethics courses, in our pilot, we shifted from having students apply ethical theories to technologies to having them reflect on different aspects of human-robot relations from a relational and more-than-human perspective. Aside from conceptual reasons, this redesign was prompted by practical motivations related to a lack of methodological examples of relational ethics in engineering ethics education, a call for more experiential engineering education, and a push to re-evaluate course assessment due to a rise in generative AI. In combination with lecture content, students explored various dimensions of relational ethics in a thinking-through-doing manner by crafting and interacting with a companion robot in various interactive tutorials. This culminated in an individual essay in which students reflected on the question ‘How to live well with robots?’, drawing inspiration from their developing relations with their robotic companion and supported by visual evidence of their human-robot interactions throughout the course. We offer broader reflections and lessons learned from this experimental course redesign, outline several considerations for those intending to integrate relational ethics into their curricula, and suggest avenues for further work.

KEYWORDS

engineering education, ethics education, robot ethics, engineering ethics, relational ethics, care ethics, robotics

PUBLICATION

Submitted:
20th December 2024

Accepted after revision:
18th July 2025

Acknowledgement: We would like to thank the organizers and attendees of the 2024 Relational Ethics in Technological Societies workshop at ETH Zürich for their feedback and comments on an earlier version of this paper.

¹ Corresponding Author 1, Olya Kudina, University of Delft, Delft, The Netherlands - O.Kudina@tudelft.nl

² Corresponding Author 2, Joost Mollen, University of Delft, Delft, The Netherlands - J.K.Mollen@tudelft.nl

This work is licensed under Attribution-NonCommercial-No Derivatives 4.0 International. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>

Introduction

Ethics is increasingly recognized as a crucial part of engineering education (Bucciarelli 2008; Bowden 2010; Valentine et al. 2020). However, teaching ethics to engineering students is a complex task. It requires balancing the richness of theoretical considerations with the practical utility of applying them to real-world cases. Hence, the teaching curriculum needs to reflect some of the sociotechnical complexity that the students will face upon completing their education. It should also enable them to identify, ideally proactively, a broad range of social and ethical concerns and mitigation strategies regarding the development of specific technologies. To that end, at TU Delft, a technical university in the Netherlands where all of the authors do research and teach at the time of writing, a conventional starting point in most of the ethics-based courses is to position technologies as complex sociotechnical systems that integrate not just a technical, but also a social and an institutional dimension, covering, for instance, the cultural, legal, and organizational context (Bijker et al. 1987; Kroes et al. 2006; Baxter & Sommerville 2011; Kudina and van de Poel 2024). This provides a helpful bridge to the course teachers to suggest that ethics, equally, does not arise only from one specific aspect of technology (e.g., the algorithms and hardware in companion robots) but also its system components (e.g., particular users and non-users, the cultural and normative setting where this robot will be used, etc.). Students are invited to look deeply and broadly within and across each system component to identify ethical opportunities and risks and consider meaningful, practical responses. Against this conceptual framing of technologies, TU Delft offers a broad range of thematic and program-tailored courses that introduce ethics competencies to (engineering) students (TU Delft 2024a).

In this paper, we³ will reflect on one such course in-depth, the MSc course on the societal implications of robotics, and the teaching team's collective effort to embed theoretical and practical considerations of the relational ethics approach in its curriculum. This effort to go beyond traditional ethical theories in the engineering curriculum (e.g., Kantian ethics or consequentialism), their merits notwithstanding, is grounded on various challenges we ran into in previous renditions of the course. These include the need for ethics about technology development that goes beyond a checkpoint assessment, the limits of traditional anthropocentric ethical approaches to account for more-than-human perspectives, a need for more active experiential education to promote neurodiversity and accessibility, and the increased availability and use of students of generative AI technologies. We aimed to address these challenges by turning to relational ethics as a theory that can address these concerns and integrate the background of the sociotechnical systems introduced above (Tronto 2010; de la Bellacasa 2017; Verbeek 2011; Coeckelbergh 2012; Pols 2023). By turning to relational ethics, our goal is to position ethics in our course as an accompaniment rather than an assessment of engineering practice; as something that constantly, however at times unreflectively, accompanies it in the dynamic social and material environment and prioritizes relations, context, and interdependence (Verbeek 2011; Kudina 2023).

This paper aims to capture and reflect on our redesigned version of the Robots and Society course, redesigned in 2023-24 and taught in early 2024. To this end, in Section 2, we will first provide an introductory background to the course and outline the conceptual and practical reasons for its redesign. Then, in Section 3, we will briefly outline how we understand relational ethics and elaborate on how it grounds our educational approach. Specifically, we operationalized relational ethics through

³ Olya Kudina developed the course in 2021 and has since been the course coordinator and main lecturer. Joost Mollen has been teaching in the course since 2021, Jordi Viader Guerrero and Dmitry Muravyov joined in 2023, and Juan Pablo Bermúdez in 2024. This paper focuses on the course's redesign in 2023-24, and its experimental run in 2024, with all the team authoring the paper.

designing our tutorials around various dimensions of caring relations — i.e., identifying, performing, repairing, and sustaining them — and having students craft and interact with a companion ‘robot’ throughout the course. In Section 4, we elaborate on this learning-by-doing approach by outlining and reflecting on the four tutorial sessions in the redesigned course. In Section 5, we zoom out to provide broader reflections and lessons learned from this experimental course redesign, outlining several challenges we ran into and considerations for those intending to integrate relational ethics into their curricula.

Background

The Robots and Society Course

Positioned in the Robotics MSc program and available as an elective across the university since 2021, the Robots and Society course introduces the students to various ethical issues and considerations that typically accompany robot development and their introduction in society. The course focuses, for instance, on the human-like appearance of robots and human-robot interaction, accompanied by theoretical considerations such as technological mediation of human perceptions and actions (Verbeek 2005) or values in the design process (Van de Poel 2013). The course spans seven lectures and four discussion-based tutorials, given yearly in February-April to 125-135 students on average. Because the course does not intend to provide a comprehensive ethical deep-dive and instead strives to introduce the students to the nature of ethical problems related to robots in society, its ethical backbone in 2021-2023 did not steer far from (re)introducing the students to the traditional ethical theories of utilitarianism and consequentialism, deontic and virtue ethics. Only towards the end of the course were other ethical theories, such as Ubuntu and Confucianism, care ethics, and several others, introduced.

Introducing ethical theories to students made them aware of the breadth of moral considerations. It gave them several normative points that could anchor the students in reflecting on robotic development and use in society. To this end, the final assignment of the course throughout 2021-23 did not pursue a ‘correct’ application of an ethical theory to a technological artifact, but rather, through a reflective group essay⁴, asked the students to identify potential or existing ethical issues in a robotic technology of their choice, reflect on them through a prism of any concepts or theories introduced in the course, and finally, provide a range of suggestions for what the students would deem a more responsible design of their chosen technology. The assignment was structured in this way to test the students’ capacity for ethical reflection qualitatively while providing them with some structure to ground it and navigate it. Even though the course has been receiving positive feedback from the students, throughout the years, various developments have called for a principled redesign of the course, which we outline next.

⁴ This course hosts different students: for most of them, the course is mandatory (higher course load), for some, it is a university-wide elective. Because of this, the course needs to have two assessment points: an exam (taken by all) and a reflective group essay (an additional assessment point for the mandatory curriculum students). The exam is open-book and -internet, and is focused on the application of knowledge from the course’s lectures and assigned readings to problem-based questions. This assessment component did not undergo structural change, so we do not mention it further in the reflections on course redesign.

Reasons for Redesign

After the first year's run of the course, we received feedback from the students to include materials on the complex relations between robotic and AI technologies and the environment. Since 2022, this aspect has been addressed in the course by discussing how sociotechnical AI and robotic systems not only often promise to foster sustainable development goals but also have a considerable detrimental environmental impact. Such an impact can be related to, for instance, earth minerals extraction practices, obfuscated supply chains, energy-intensive training of software and data center operations, and water shortages (Couldry & Mejias 2019; Brodie 2020; Brevini 2021; Crawford 2021; Valdivia 2024). Through multiple contemporary examples and case studies in the readings and a designated lecture, the students discussed the current state of robotics development and charted options for a more responsible way forward.

The lecture greatly resonated with the students, but showed the limits of relying primarily on traditional ethical theories. As the students remarked, while having rules and relying on justice principles is essential, they are not enough to help ameliorate the environmental crisis we are facing globally. The complexity of the sustainability problem vividly illustrates the need for ethics in technology development that goes beyond a checkpoint assessment - e.g., a sustainability checkmark if there is an energy calculator in the project or a privacy checkmark if the project is deemed compliant with the data protection regulations. Instead, it pointed to the need for ethics as a practice that constantly accompanies an engineering project's development through formal and informal reflections of engineers and other stakeholders. This type of ethics, as accompaniment (Verbeek 2011; Kudina 2023), is not per se about pre-defined principles or rules but about proactive identification of what may arise when several system components come together, something that robotic engineers have been acutely aware of for a long time in practice, even though such awareness is primarily tacit (e.g., Kaplan 2004; Castañeda and Suchman 2014; Arzberger et al. 2023).

The students' interest in sustainability and robots as complex sociotechnical systems also revealed the limits of traditional ethical approaches, primarily vis-à-vis their anthropocentric focus. Focusing on humans and human-centered value-making leaves technologies as neutral tools in the hands of autonomous and rational human subjects and the environment as a standing reserve for extraction practices. Scholars across disciplines (e.g., Coeckelbergh 2012; Vallor 2016; Nyholm 2020; Cila et al. 2017; Forlano 2017; Wakkary et al. 2018; 2022) increasingly converge on the need to focus on relations and interdependence and engage with an ethics grounded in a more-than-human logic (Haraway 1985; 2016) when researching technological developments. However, this approach is frequently interpreted as speculative, too flexible, and ambiguous to be useful in practice (see discussions of relational ethics proponents on this, e.g., Tronto 2010; Mol 2021; Pols 2023). Thus, prompted by the students' interests in sustainability, the need to incorporate more-than-human ethical concerns, the sociotechnical understanding of technologies, and an idea of ethics-as-accompaniment formed a goal of aligning the course's narrative along the lines of relations, interdependence, and a dynamic understanding of ethics. Even though traditional ethical theories can partly reflect some of these points, the relational ethics theory can accommodate all of them and thus was chosen as the main theoretical backbone for the course redesign.

Parallel to these conceptual considerations, practically, the teachers who had accompanied and graded the students in the group essay in the previous rendition of the course worried that the students interpreted the assignment primarily as a way to demonstrate the 'correct' application of ethical theories. An exaggerated example would be students asking what a morally right action in a technological development should be and appealing to consequentialism to resolve it by comparing

the number of drawbacks and benefits. Integrating additional explanations throughout the lectures before the essay submission, reminding the students to check the essay guidelines regularly, and highlighting these points through the preceding checkpoint assignments (e.g., description of the research questions and methods, draft paper, peer-review) did not significantly help the matters, resulting in high workload for the teachers and the students, and making this final assignment unsustainable in the long run.

Additionally, scholars increasingly criticize the dominant pedagogical focus on textual engagements in engineering ethics education, outlining a need for more active experiential exercises to promote neurodiversity and accessibility (Van Grunsven et al. 2024). Relational ethics is frequently acknowledged as necessary in engineering education and as a theoretical ground that can promote a diversity of relations and experiences (Ley 2023). However, there is a notable lack of methodological examples of how to do so in practice. We were fortunate with a wealth of expertise and the breadth of ethics teaching portfolio at TU Delft, which enabled us to learn from the existing courses, particularly related to environmental ethics by Andrea Gammon (TU Delft 2024b), and the diversity of practice-based teaching approaches, among others from the neurodiversity approaches by Janna van Grunsven and Sabine Roeser (2022) and integrating art and empathy into teaching, proposed by Aafke Fraaije (Fraaije et al., 2022) and Filippo Santoni de Sio (TU Delft, 2024c). In return, we aspired to contribute to the further evolution of the educational content and approaches through the content and structural redesign of the Robots and Society course.

Finally, the teachers had to consider the global technological rise and availability of generative AI technologies. This prompted us to reconsider the essay assignment and whether the essay format would still be the way to achieve it. Various scholars have increasingly drawn attention to the impact of challenges of generative AI on academic education in general (Michel-Villareal et al. 2023) or on engineering education specifically (Johri et al. 2023; Qadir 2023; Yelamarthi et al. 2024). One example of an educational challenge is a concern about academic integrity: students can use the technology to generate content for tests and assignments instead of (completely) writing it themselves, leaving it difficult for teachers to assess whether students have engaged with the education material and have met academic standards (Van Niekerk et al. 2025). In response to students' proliferating use of generative AI tools, we aimed to focus on personal reflection and skills different from those that generative AI tools are allegedly automating, hopefully disincentivizing reliance on generative AI technologies.

All of the above prompted us to redesign the logic underpinning the course, the essay, the lectures, and the tutorials. Jointly, as a course team, we took it as a challenge to do something that would rely less on ethical theories and more on reflections of students' experiences and analysis of current technological affairs. The following section presents our way of doing so.

Crafting relational ethics, from the conceptual to the practical

Relational Ethics

In this section, we will first provide our conceptualization of relational ethics. This is by no means meant as a comprehensive background to the rich history of relational ethics, for which we draw inspiration from scholars across philosophy, science, and technology studies, environmental and indigenous ethics, and more (e.g., Coeckelbergh 2012; Castañeda and Suchman 2014; Whyte and

Cuomo 2016; de la Bellacasa 2017; Neimanis 2017; Nyholm 2020; Mol 2021; Kudina 2023). Instead, it is meant to sketch the theoretical starting points we employed in structuring the lectures and tutorials, reading materials, and the final reflective assignment.

Theoretically, the dominant anchor points we wanted to highlight in the course were the more-than-human emphasis and a dynamic understanding of ethics. Since this course was set in the tradition of ethics of technology, it already featured the moral significance of the material setting (e.g., Verbeek, 2011; Van de Poel, 2013). However, in its attempt to portray the (moral) complexity of robots and AI as sociotechnical systems, the course had not yet emphasized the ethical significance of the environment as another more-than-human element⁵. This was the first conceptual change to the course's narrative, inspired by environmental justice scholarship, indigenous philosophies, and care ethics. The second theoretical inspiration for crafting the relational ethics lens drew on framing ethics as accompanying technological development and use. However, it still relied predominantly on the traditional ethical theories, not entirely satisfying these ambitions. For these reasons, we drew on care ethics (Tronto 2010; Pols 2023; Ley 2023). This ethical theory understands ethics as a dynamic affair arising within interdependent relations of several actors and, hence, is sensitive to the context with more-than-human potential for recognizing ethical concerns and opportunities (de la Bellacasa 2017). All of this positioned the central question of the course as “How to live well with robots?”.

In each lecture, teachers and guest lecturers reminded the students about the overarching course's question, the umbrella of robots understood as complex sociotechnical systems, and ethics accompanying dynamic and embodied relations of people with robots in the cultural, institutional, and environmental world. With this, we aimed to show that ethical opportunities and risks arise within these complex relations and require inter- and transdisciplinary approaches for the responsible development of robots in society. Seven lectures followed this structure, presenting the students with theory and contemporary case studies to test their assumptions. However, it is the tutorials, i.e., more practical and discussion-oriented sessions, where the students could test and challenge the theoretical backbone of the course through a series of workshops.

In the original version of the course, the four tutorials were discussion-oriented and centered around contemporary case studies of robots in society. They required reading journalistic engagements, debating, and making presentations. Even though the students usually found them interesting and relevant, they mentioned that the tutorials appeared disconnected from the lectures and the assigned readings. In the new version, adapting the logic of care (e.g., Tronto 2010; de la Bellacasa 2017) for human-robot-world relations, each tutorial was meant to represent the four dimensions of caring relations: identifying, performing, repairing, and sustaining them.

Crafting Robot Course Companions

Students predominantly explored the above-mentioned caring relations through a personally crafted ‘robot’ companion, which played a central role in the tutorials. Students made these companions by combining an obsolete piece of technology they had to provide themselves (e.g., an old phone, calculator, etc.) with arts and crafts materials provided by the teaching staff. This scrap material

⁵ Note that for practical reasons, in this course, we chose to conceptualize “more-than-human ethics” with the roles of material and ecological environment, not highlighting other potential inclusions in this regard, e.g. non-human animals and other living organisms, etc. For ways of doing so, see Haraway (2003) and de la Bellacasa (2017).

consisted of cardboard, glue, tape, and various recycled materials, such as fabrics, foams, and plastics (see Figure 1 below).

While these personal companion robots do not qualify under more technically oriented definitions as robotic artifacts since they are incapable of sensing or responding to the environment, we adopted a more relational-focused conception of robots for this course. Drawing from postphenomenological approaches to human-robot relations (Coeckelbergh 2011), we proposed understanding robots through how they appear to humans, how we behave towards them, and how they mediate and affect human-human and human-environmental relations. This complemented the general focus of the tutorials, where students explored human-robot relations individually and in groups rather than focusing on their technical or functional complexity. In this manner, the tutorials also invite inquiry into how the term ‘robot’ refers to specific types of technologically-mediated relations, behaviors, and outcomes, rather than mechanical devices.



Figure 1: Example of scrap material used for the crafting exercise.

The intention behind the robotic companion was twofold. First, by creating, interacting, and caring for a robotic companion, the tutorials aimed to attune the students to the experiential dimension of human-robot relations and the different ethical tensions and opportunities this entails. Caring-at-a-distance is a known problem for engineers (Van Grunsven et al. 2023a), as it involves designing socially responsible robotics, while the users of these technologies, with their diverse experiences and values to be affected by these technologies, may not be known to the engineers. Proposals to ameliorate this problem at an early stage include morally attuning the engineering students to the diverse experiences of others who may be affected by these robotic creations through a series of conceptual anticipation exercises (Stone et al. 2020; Van Grunsven et al. 2023b).

In this course, we wanted to explore an alternative strategy by including some experiential dimensions of human-robot relations. While robots are increasingly present in work, care, and home environments, many students might not have had many such interactions, especially over a prolonged period. Even though the experiential dimension with robots is not a prerequisite to becoming a good robotics engineer, some exposure to a direct or testimonial human-robot interaction could be beneficial for students’ recognition of the prevalent role of care in interactions with robots, i.e., the caring work of identifying, performing, repairing, and sustaining human-robot relations. This method allowed students to reflect on abstract issues regarding human-robot relations that others may face (e.g., privacy, anthropomorphism, etc.) and utilize the personal dimension between students and their robotic artifacts (no matter how minimal) as a source for reflection to identify, evaluate, and mitigate potential

ethical issues. These projective and personal insights would form a basis for the student's individual reflective essays while promoting their moral sensitivity and a sense of prospective engineering care at a distance.

Second, by creating a robot companion themselves — instead of, for example, receiving a pre-made one — students had an immediate potential personal connection to the artifact. This further reinforced the idea that robots and human-robot relations are actively created, rather than given by default. It also promoted the idea that engineering students are responsible for co-shaping these technologies and how other people relate to them in their future careers. Here, we drew inspiration from other creative educational approaches with low-tech methods (Van Grunsven et al. 2024) to explore human-robot relations in the field, such as HitchBot, the hitchhiking robot (Smith and Zeller 2017), the couch-surfing BlockBots (Mollen et al. 2023), and the Tweenbot, a 'helpless' cardboard robot on wheels that was only able to drive forward and relied on the help of passersby to reach its destination (Kinczer 2009). Next to actively creating a robotic companion, a connection was further reinforced by having students incorporate an obsolete technology from their lives with which they have had some prior experience and memories. Additionally, the choice to reuse an existing piece of technology and give it a proverbial 'new life' was also a practical expression of our use of more-than-human-ethics, emphasizing that ethical issues of human-robot relations extend beyond the mere object but also include more significant environmental issues regarding the material, ecological, and energy costs of creating and maintaining AI and robotic artifacts.

The final assignment became an individual reflective essay, where the students were asked to answer the question of "How to live well with robots?" underpinned by their activities in the four tutorials. Central here were students' reflections on their relation to their self-made 'robots,' and potential ethical issues emanating from a range of documented experiences with these robots both during and in between classes (e.g., photos, diaries, audio-reflections, etc.). The following section presents four tutorial sessions as the key component to practicing and reflecting on relational ethics in a learning-through-doing manner.

Tutorials

In this section, we describe each tutorial and its intentions, reflect on how it went, and provide visual documentation. These four tutorials reflect (roughly) the four dimensions of care: identifying, performing, repairing, and sustaining relations.

Crafting Human-Robot Relations (Tutorial 1, identifying)

The intention of the first tutorial was twofold. First, it aimed to familiarize students with the ethical dimensions of human-robot relations and facilitate a discussion on the topic. Specifically, the tutorial focused predominantly on personal human-robot relations. The entry of robotics into domains of human life and work, ranging from prolonged collaborative relations in a work setting to leisurely or intimate relations in the privacy of our homes, might bring about particular emotional responses and attachments. Regardless of whether these robots were designed to solicit these emotional responses, these responses raise questions regarding their desirability, how and when they should be avoided or mitigated, and how they differ from other forms of human-human relations and human-object attachments.

Second, the tutorial allowed students to craft their robotic companion from an obsolete piece of technology they would bring to the course, as well as arts and crafts materials provided by the teaching staff. During the crafting, students were instructed to discuss the statement, ‘*Having social and romantic relations with robots is better than having no such relations at all.*’ Additionally, students were encouraged to consider the following questions: (1) *Can you really be friends with a ‘robot’?*, (2) *What is the difference between a human and a robot friend?*, and (3) *Should we ban romantic relations with robots?*. After the crafting session, the tutorial ended with a plenary discussion about the created companions and the earlier prompts.



Figure 2: Examples of robotic companions with anthropomorphic characteristics

Students were positively engaged with the crafting exercise during the tutorial, creating a relaxed and collaborative atmosphere. Students needed or took different amounts of time for the exercise, ranging from 45 to 90 minutes. An apparent benefit of the exercise is that it engages all students in an activity that is easy to start, aligns with the prototyping experiences of engineering and design students, and does not present a potential language barrier, which more traditional discussion or debate formats might face. Tutorial engagement, however, differed between students. Some students had already started building their companions at home, e.g., creating a 3D-printed companion. Other students

planned to finish their robot companion at home because they needed a specific tool (e.g., a 3D printer) or a specific material (e.g., fake fur). Other students did not have a piece of obsolete technology, for example, since they had moved from abroad to study and had not taken any broken or old technology with them. As a solution, other students shared additional obsolete technologies or used pieces from the scrap material as stand-ins.




Figure 3: Examples of robotic companions with zoomorphic characteristics

Despite the limited materials and tools provided by the course and the obsolete technology they brought, students produced a wide variety of robotic companions. However, most constructed companions could be characterized by anthropomorphic or zoomorphic characteristics, including eyes, hands, arms, legs, hair, whimpers, and other human or animal characteristics (See Figures 2 and 3). Some students wrote a text on the robot indicating a name. However, others went in a more abstract or utilitarian direction, using origami, constructing boxes out of cardboard, and creating abstract mechanical shapes (see Figure 4). For example, one student made a large rotating platform for her phone. Another made a transparent hull containing obsolete computer hardware such as transistors and processors. These anthropomorphic tendencies or anti-anthropomorphic tendencies in their companion design were frequently commented on in the reflective essays.

The role that an obsolete piece of technology played in the final companion also differed between students. For example, one student brought an obsolete but working Nintendo DS, which got a prominent place within their companion, displaying a range of photos of her younger brother. For others, the technology they brought was broken and played a more decorative or constructive role. Students brought different pieces of technology, including an old mobile phone, an electric toothbrush, a VHS tape, a CD, an Arduino, and a calculator. Due to the fragile nature of some of the created companions, various forms of repair or alteration of the companion bot were needed throughout the course, which several students commented on in their reflective essays.




Figure 4: Examples of robotic companions with abstract or functional characteristics

Performing Human-Robot Relations (Tutorial 2, performing)

Tutorial 2 presented students with a question: ‘What if the term ‘robot’ refers to technologically mediated relations, behaviors, and outcomes rather than mechanical devices?’ Recalling the original use of the word ‘robota’ in Slavic languages, meaning serf labor, drudgery, or hard work, the tutorial framed robots in pragmatist terms centered on actions: a robot is something we (humans) do. More precisely, robots are expected and performed relations between humans in contexts where mechanical devices organize social relations.

The tutorial consisted of an activity called CAPTCHA (Completely Anthropomorphic and Unautomated Public Turing Test to Make Computers and Humans Alike), followed by a plenary discussion. The exercise aimed to transform the classroom into a data production assembly line, in which students produce visual and textual data about their companion robots on a large scale in a limited amount of time. Inspired by the activity “Text to Art–Art of Text” designed by Lili Lian van Doornick (Heijnen & Bremmer, 2021), this activity could be described as a variation of the telephone game.

The activity unfolded as follows: after taking a good look at their companion robots, leaving them in the back of the room, and making sure that each student was sitting next to each other in an S-shaped distribution (see figure 5), students were handed a deck of blank cards (or input devices) and asked to describe their robot in only three lines on one of them. Students were given a couple of minutes to write this description.

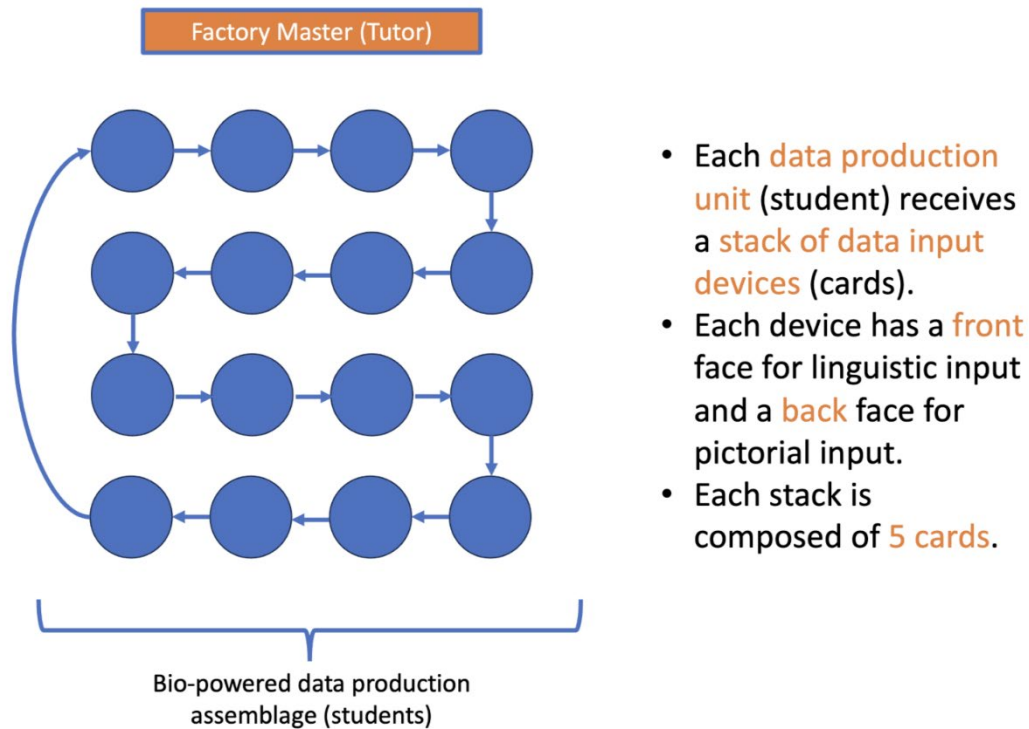


Figure 5: Schematic sketch of the tutorial's instructions.

Once this first round was done, students had to hand their input device with the description to the student next to them, following the S-shaped distribution. Each student then received a card with a short description of a companion robot, and after quickly reading it, they were given a minute and a half to draw this description on the back of the card they received.

After the second round, students passed the used card to the student next to them, who was then asked to quickly observe the picture and write a three-line description of it on a new card from their deck. This time, students were given only one minute. This process was repeated for ten rounds (5 rounds of written descriptions and five rounds of visual depictions). The time to complete the task was reduced for each round, meaning the students' decision-making had to be accelerated. Moreover, the entire activity was accompanied by electronic drum and bass music on the classroom's speakers and video visuals projected on the presentation screen to create a fast-paced environment (see figure 6).



Figure 6: Students produce visual and textual data of their companion robots.

After the activity ended, the students were asked to classify their output as either human or robot-made (see Figure 7). The outcome and the methods for their classification were the starting point for a plenary discussion on the difficulties of this distinction. Discussions in the different tutorials often concerned the context-dependency of the human-robot distinction and the link between robot deployment and industrial expectations of standardized labor and production at large scales.

The main goal of this tutorial was to get students to ‘perform robot relations’ and engage in an activity that involved actions and abilities other than predominantly writing and speaking. While linguistic reflection is crucial and was carried out in the later part of the tutorial, the CAPTCHA activity aimed for experiential learning during the tutorial rather than making them think about, imagine, or reflect on real or hypothetical cases outside of the classroom.




Figure 7: Student outputs classified as human-made or robot-made.

Nevertheless, while students enjoyed this active and performative tutorial approach (the music, visuals, and time constraints successfully framed the activity as a fun game), many seemed to question its relation to their own social life. The activity was conceived to draw parallels between the classroom-as-assembly-line and the everyday activities that have become production sites through technology in which they are almost certainly involved (i.e., social media transforming communication and relationships into productive behaviors). This comparison was also meant to suggest that human-robot relations are also performed in ways that do not necessarily emulate emotional attachments or companionship. While potential objects of moral consideration, robots are often embedded in larger sociotechnical systems of value production that expect specific interactions from humans. However, this link remained implicit in most tutorial sessions as student reflections remained fixed on describing human-robot relations in emotional attachment terms during and after the activity. Yet, it did have the benefit of sparking conversations about their feelings of shock and absurdity when using machines to produce at scale, how this is at odds with their deliberate intentions to produce human-like outputs, and the difficulties of delivering consistent outputs under time constraints.

Navigating Robot Failures (Tutorial 3, repairing)

In the third tutorial, students engaged with the theme of errors and failures in the context of robotics. At its core, this session aimed to expand students' perspective on technological fallibility beyond the vocabulary that exists in traditional robotics training. In robotics, errors are often deeply tied to a discussion of safety; they are to be prevented or anticipated in the design (Bisante et al. 2023). However, beyond notions of malfunction or error rates, technological fallibility is also imbued with social, economic, and political meanings (Appadurai and Alexander 2020; Elish 2019; Mirnig et al.

2017). Through speculative storytelling about possible failures of their companion robots, the students were invited to think about the fallibility of their robot-making endeavors.

Distinguishing ‘successes’ and ‘failures’ in robotics is not always trivial. Does the robot have to be ‘working’ to be a ‘success’? What does it mean that a robot ‘is not working properly’ or is ‘not good enough’? What role do errors play in designing robots? During the session, we asked the students to come up with stories about their companion robots breaking down, erring, and failing, drawing on prompts that could be interpreted by students freely, such as engineering failure, market failure, political failure, failing a social group, or environmental failure.

After creating these stories in small groups, students jointly discussed their differences and similarities. Building on that, we invited them to reflect on the following questions: Even if a robot is technically “perfect,” can it still fail in another sense? What can or should engineers do about that? Is it always desirable to improve a particular robot and make it error-free? Can the errors in robotics design reveal something? Can they bring about social good or have ethical value? Discussing these questions and deriving multiple meanings of “failure” based on them allowed us to have a broad range of conversations about human and machine imperfection, computational modeling, external real-world environments, the limits of calculability of errors, and various demands imposed on technological development.

Discussions on error and failure have often started with students conceptualizing these notions as deviations from designers’ intentions (Parvin and Pollock 2020). However, as the sessions progressed, students have also noticed that it is sometimes precisely when the designers’ intentions seem to be fulfilled that technologies fail certain people. In the discussions, students discussed inviting people from other disciplines and pushing against particular problem framings as possible solutions to the limitations of technical understanding of failure and error.

This tutorial aimed to construe robot fallibility not necessarily and only as an event to account for in the design of robots, but also as a dynamic relation continually orchestrated between multiple parties whose lives are impacted by robotics. Thinking of robot failure in relational ethics terms captures that it is not an accidental property of a technical artifact, but a particular form of relation that emerges between humans and technologies. We hope that connecting such speculative storytelling to their companion robots makes this activity more tangible and engaging. However, moderating such discussions was also not without challenges. While the multitude of meanings attributed to failure and error helped to foster and broaden the conversation, ensuring shared understanding throughout the process required continual attention by the tutor to students’ different uses of terms.

[Keeping Robots in Our Lives \(Tutorial 4, sustaining\)](#)

This tutorial invited the students to consider how to sustain human-robot relations long-term by engaging in a process where groups with different values try to agree on solving breakdowns in said relations. More specifically, in the tutorial, students worked on (1) identifying stakeholders beyond the ‘usual suspects’ (i.e. users and producers); (2) exploring the diverse impacts of human-robot relationships on multiple stakeholders; (3) understanding how maintaining human-robot relations involves navigating value conflicts due to technology affecting stakeholder groups differently; and (4)

grasping the challenges of managing human-robot relations while navigating conflicts between the diversely affected stakeholder groups.

The tutorial was split into two parts. First, students adopted the perspective of different groups affected by a human-robot relationship in a hypothetical future scenario. Then, students participated in stakeholder assemblies, representing their group’s values and trying to reach a consensus with other stakeholder representatives on the future of human-robot relations.

From values to design requirements

This tutorial exercise is inspired by the “technomoral scenarios” method, in which a narrative about a hypothetical future serves to reflect on the moral consequences of technological change (Arnaldi 2019; Swierstra et al. 2009). The exercise started with a story: ‘In the year 2034, Loop Industries launches Loopy, a cute and furry robot endowed with advanced Artificial Intelligence, programmed to serve as an educational tool for children from age five onwards.’ The scenario depicts how Loopy’s arrival affects relationships beyond the direct child-robot relations: it also affects relations between the children and the parents, the school setting, governments, and the workers in the Global South who produce Loopies but cannot afford to provide one for their children.

After students heard the story, they were divided into groups of 4–5. Each group represented one stakeholder: parents, European Union regulators, Loop Industries, the Association for the Ethical Treatment of Robots (of which the students’ robot companions were all members; see Figure 8 below), and a mystery stakeholder. To define who the mystery group should represent, students discussed which other groups impacted by Loopy have not yet been included. Deciding who to represent was up to the students in the mystery group. Student groups decided to represent teachers, manufacturers, Global South communities whose natural environment was affected by Loopy production, and an “anti-robot group” of people fighting for a robot-free society.



Figure 8: Robot companions attending’ the meeting for the Association for the Ethical Treatment of Robots

To sustain the human-robot relationship despite the conflicts it generated, each group was asked to think about how Loopy the robot would have to be redesigned or updated to preserve their group's core values. Following the values hierarchy tool (Van de Poel 2014), each group went through three steps: (1) identifying core *values* of their group that were advanced or threatened by the human-robot relation; (2) picking one threatened value and identifying *norms* the robot should follow to protect it; and (3) selecting one of those norms and proposing *design requirements* that should be built into the next version of Loopy to ensure the chosen value was advanced. After this, the students were informed that there would be a stakeholders' assembly, where all stakeholders would get together and decide whether Loopy should continue to be allowed and what its new version should be like.

Stakeholders' assembly

Groups were reorganized so that each group had at least one stakeholder representative. Each group now constituted one stakeholder's assembly, which was convened by the EU regulators to decide on Loopy's fate. The assembly worked as follows:

- Each stakeholder group presents its design requirement proposal.
- The assembly votes on which proposals should be given priority and votes on each in order of priority. Only unanimous votes are accepted. If there is no unanimity, those voting against it can put forward modified proposals, which are then submitted to a new vote.
- If no proposals are accepted, Loopy and similar technologies are suspended. If agreements are reached, the next version of Loopy must comply with all accepted proposals.

The exercise was designed so that consensus was challenging yet achievable. Indeed, all assembly groups agreed on at least one requirement (less than half reached two or more). This part of the tutorial closed with a general discussion where students reflected together on what made the agreements difficult and what made them possible, as well as on how maintaining human-robot relations required being mindful of the relations (both human and more-than-human) impacted by the technologies, and how tools like value hierarchies and technomoral scenarios could be used to implement these reflections consistently and rigorously.

One limitation of this tutorial design was that it did not involve student interactions with their robotic companions. The robotic perspective was included through the Association for the Ethical Treatment of Robots, which invited students to take the perspective of robots like their companions. However, even that was somewhat removed from the direct interaction with the material objects. This was due to the tutorial leader's choice to focus on one specific aspect of maintenance—the need to update technologies to solve problems that lead to human-robot relationship breakdowns—and the use of design methodologies and consensus-building to navigate value conflicts. While this was aimed as an exercise to put into practice ideas from the course's relational ethics and more-than-human ethical approach, new versions of the course could seek to bring this tutorial more closely in line with the others by including more direct student-robot interactions, as well as exploring other aspects of relationship maintenance.

Discussion

In this section, we reflect on our experimental course redesign, limitations, and lessons learned. We also outline several considerations for those intending to integrate relational ethics into their curricula and suggest avenues for future work. We organize our discussion based on our conceptual and practical

goals for the course redesign as outlined in Section 2. These goals were to (1) provide a concrete example of embedding relational ethics into the engineering ethics curriculum in order to go beyond checkpoint ethics assessments and account for more-than-human perspectives, (2) provide more active experiential education in order to promote neurodiversity and accessibility, and (3) respond to the increased availability and use of students of generative AI technologies. In addressing these concerns, we consider our pilot integrating relational ethics into engineering ethics education a modest success. However, these points also brought about particular challenges and limitations, which we discuss below.

Embedding Relational Ethics in Engineering Ethics Education

Our course redesign provides a methodological example of introducing relational ethics in engineering education. By turning to relational ethics as the theoretical backbone of our course, we were able to incorporate more-than-human ethical concerns, a sociotechnical understanding of technologies, and an idea of ethics-as-accompaniment with our course's narrative and place a larger emphasis on themes such as relations, interdependence, and a dynamic understanding of ethics.

Practically speaking, we incorporated those themes in the tutorials in a learning-by-doing manner. The 'doing' part consisted of students crafting a robot companion for experiential exercises during tutorial sessions, e.g., role-play, the (re-)use of obsolete technologies in the construction of the robot companions, and using a self-crafted robot throughout the course as a focal point for reflection on relational and care dimensions of robotics engineering. The 'thinking' part of our approach invited the students to reflect on robotics from various relational dimensions, such as crafting, performing, repairing, and sustaining human-robot relations throughout the lectures and tutorials and, for example, how the concept of robots and robotic failure are particular relation forms that emerge between human and technologies which are shaped, negotiated and mended by various stakeholders.

Reflecting on our own experiences in the course as teachers, we found that grounding the course on relational ethics offered unique opportunities to invite the students to attune to different (experiential) dimensions of human-robot relations and the different ethical tensions and opportunities this entails, as well as promote the idea that engineering students are responsible for co-shaping these technologies. These opportunities were harder to come by in previous renditions of the course. Additionally, we observed that this theoretical and practical framing in relational ethics brought about an engaged and diverse student experience. In tutorials, students were engaged with practical and creative activities, which activated students who would typically be less inclined to engage in more traditional discussion-oriented formats. This engagement also extended beyond the tutorial activities. Various students commented in their reflective essays that while their journey had started with scepticism, they caught themselves engaging in pro-social, anthropomorphic, or emotional behaviour toward their 'bot' nonetheless. While the purpose of the tutorial(s) was not necessarily to cultivate a human-robot attachment, the fact that these relations did (and in some cases did not) emerge provided a unique opportunity for students to reflect on the ethical questions around human-robot relations based on their own experience.

However, integrating relational ethics theory into our course through our methodological approach also brought about multiple practical challenges. Compared to more traditional tutorial approaches, our approach, in its current form, was labor-intensive. Organizing the tutorial activities came with a heavier workload for the teaching team and required larger material production costs and time

investments. For example, two boxes of scrap material, cardboard, and crafting materials had to be found, ordered, and moved between buildings, and their storage had to be negotiated between different university departments. Additionally, the tutorial involved preparation for the in-class activities (making the cards and writing the technomoral scenario, selecting music and visual supports, acquiring magnets, colored pencils, etc.). The associated workload might not entirely match standardized teaching task descriptions and the hours allocated to them in the way that traditional PowerPoint presentations might. Hence, an effort must be made to streamline and standardize the teaching methodology to ensure that the course is transferable and not dependent on the current teaching staff. So, while the low-tech crafting material and active learning provided a key component in providing an example of how relational ethics can be embedded in the ethics engineering education, it also brought about practical challenges regarding time, effort, logistics, and (material) costs. This places further tension between the effort of developing and giving courses and the resources available within an academic institution.

Experiential education

Our approach also aimed to respond to calls to increase experiential and accessible engineering ethics teaching formats (Van Grunsven et al. 2024). Grounding our approach in relational ethics, we aimed to enrich ethics in the engineering curriculum through enabling active, embodied, and experiential learning opportunities. During the tutorial sessions, we tasked students to reflect on human-robot relations by using their hands, actively reassembling an old electronic device, and, later on, embedding it in different situations to open up alternative forms of engagement. So, rather than only reviewing possible case studies outside the classroom, these sessions aimed at creating embodied experiences that could serve as anchors for reflection on human-robot relations throughout the course. This further reinforced our theoretical backbone: engagement with robots is always relational; never occurring in a decontextualized manner.

During the course, a challenge emerged concerning the role that imagination plays in our experiential and relational methodology. In the tutorial exercises, students used their robotic companions as a reflection point for a particular moral dimension of human-robot relations. This functioned as a reflection point and also extended beyond the tutorials. The students were encouraged to journal or otherwise document their relationship with the robotic companion, leading to their reflective essay. However, the fact that robotic companions were ‘low-key’ mock-ups without interactional properties required students to exert additional (moral) imagination. That is, during the course, students were implicitly asked to use their imagination to ‘complete’ or find the relevance of their relation with their mock-up robot inside and outside the classroom by extrapolating it to ‘real-world’ contexts of engaging with interactive robots, rather than simply focusing on the apparently mundane but potentially illuminating interactions with it. Consequently, this led to the feedback (informally and in reflection essays) that some students perceived they might have failed to engage with the exercise because they had not developed a relationship with their companion. Conversely, others thoroughly let their imagination fly, describing make-believe situations in their essay to contrivedly showcase ethically salient situations with possibly little link to their personal experience.

Thus, some students operated on a seeming misconception that the purpose of the exercises was to form an emotional connection with their companion robot. While it was always explicitly and repeatedly stated that this was not the case, the misconception persisted. This is perhaps at least partly due to calling their artifact a robotic ‘companion,’ the framing of which suggests companionship and

might thus be constitutive of an emotional connection. Students also built their robotic companions in the same tutorial, where they discussed the ethics of human-robot bonding and emotional attachments. Perhaps reframing the robot from ‘companion’ to something else (e.g. ‘robotic device’), or having a different discussion topic for the first tutorial could amend this.

Alternatively, we might take these students' responses as reflecting existing human anthropomorphic tendencies. In other words, students may perceive their absence of a human-robot relation predominantly because experiences with robots are often understood in anthropomorphic terms, leading to difficulties pinpointing other sorts of relations with the robot. This defaulting to anthropomorphized emotional attachment pushed us, nevertheless, to acknowledge a key takeaway of the course: it can be unclear what these relations look like—what sort of actions and emotions they entail and, more broadly, how to conceptualize them. This open-endedness can be interpreted as both a drawback and a strength of the course's design. A strength because, by asking students to craft mock-up robots with e-waste and arts-and-crafts materials that would not be typically used in design and engineering practices, they were given the opportunity to question and expand what robots could be, how to relate to them and, therefore, to reflect on what sort of relations engineers want to foster when designing them. However, it can also be a drawback, leading to the aforementioned confusion among students. This points to the need to better manage students' expectations and better structure and guide the role of moral imagination in future renditions of the course, instead of leaving this reflection outside of lectures and tutorials wholly up to the students. In the future, guiding reflection during off-session times could benefit the development of a more profound and multifaceted consideration of human-robot relationships. Alternatively, future versions could discuss the topic of imagination head-on: explicitly stating that some students might find themselves using imagination to attribute to their robots properties that go beyond the ones they have, while others may not, and discuss with them the role of the imagination in our human tendency towards anthropomorphising and interpreting the behaviour of robots.

Availability of generative AI tools

Finally, our course redesign aimed at responding to the increased availability of generative AI tools for students. We did this in roughly two ways. First, during our tutorials, we offered activities that do not necessarily entail writing: making, drawing, choreographing, role-playing, and storytelling. When writing was involved, we asked them to engage in writing as a self-reflective activity, instead of a traditional academic essay. We expected that pivoting to relational ethics and focusing writing assignments on students' own experiences would limit the potential misuse of generative AI tools. While students may still use generative AI tools for other parts of these tasks, this element of personal reflection would disincentivize excessive reliance on these technologies since students can, and are asked to, draw from personal experience. Additionally, photographic documentation of their activities would add a layer of difficulty to the misuse of AI tools at this moment.

That said, the teaching staff cannot assess whether student reflections or personal experiences are grounded in real experiences or whether documentation is produced to support a pre-conceived narrative. This could be addressed partly by incorporating an intermittent experience sampling method (Csikszentmihalyi & Larson 2014), where students are asked to report on their most recent experiences several times throughout the semester (e.g., once every two weeks), and then to use those reports as materials for the final reflection essay (which they could be asked to write as an in-class exam). Arguably, such a staggered reflection approach would nudge students to stay closer to their everyday

experience. That said, on most occasions, the production of the reflection essays seems to have prompted most students to produce credible reflection points and proactively engage with moral imagination and course literature to write a passing essay, which are ultimately the main learning goals this assessment point seeks to evaluate. Finally, it is also important to mention that we framed generative AI as an inherent challenge that our approach sought to diminish or overcome by reducing the relevance of using the tool. This might have contributed to or reinforced a tendency in students to keep their use of generative AI a secret (producing a kind of ‘secretive’ relation between student and generative AI) rather than bringing this relation to the forefront for students to reflect upon⁶. Making student-generative AI relations a more explicit focus of reflection could be an interesting angle for future renditions of the course.

Conclusion

This paper has proposed and reflected on a teaching methodology for introducing relational ethics in the engineering curriculum based on a pilot redesign of TU Delft’s Robots and Society MSc course. Our redesign specifically responded to provide a concrete example of embedding relational ethics into the engineering ethics curriculum, to go beyond checkpoint ethics assessments and account for more-than-human perspectives, a call to increase experiential and accessible teaching formats, and the need to adapt education to the challenges posed by the availability of generative AI tools. We responded to these challenges through integrating relational ethics theory in a thinking-through-doing manner. The ‘doing’ part consisted of students crafting a robot companion for experiential exercises during tutorial sessions, e.g., role-play, and using a self-crafted robot throughout the course as a focal point for reflection on relational and care dimensions of robotics engineering. The ‘thinking’ part of our approach invited the students to reflect on robotics from various relational dimensions, such as crafting, performing, repairing, and sustaining human-robot relations throughout the lectures and tutorials.

Our work and findings contribute to the larger engineering ethics education community by providing a concrete example of practically integrating relational ethics in engineering ethics education. We have shown that relational ethics can be a solid theoretical basis on which to ground experiential and accessible engineering ethics education, which additionally might lessen the dependence of students on generative AI tools. Our paper offers a reflective case study for practitioners in engineering ethics education interested in incorporating relational ethics within their educational practice. Such a case study can inspire practitioners looking to expand, complement, or adapt their current methods with relational and care-based frameworks. For these practitioners, our paper provides a helpful reflection on the challenges and limitations that we encountered. Going forward, these challenges must be addressed to make the course more robust for future iterations.

One avenue for future research concerns using empirical methods to gain further insights into the pedagogical value of our approach. A natural follow-up could be applying for a research ethics committee approval to be allowed to use the institutional course evaluation forms as an additional source of data, while additionally conducting interviews with the students and teachers. For example, researchers at the University of Toronto involved with integrating ethics education in the computer science curriculum conducted a longitudinal study assessing the impact of their methods through

⁶ Thank you to an anonymous reviewer for bringing this point to our attention.

anonymous surveys that students were invited to participate in before and after the course (Horton et al. 2024). Additionally, we can draw on Van Grunsven and colleagues (2024), who used a triangulation method to assess their tinkering exercise with engineering students by having multiple perspectives (e.g., student, observer, and teacher) assess the same activity concerning the educational goals of the course. Whichever method we choose, we will have to weigh its benefits against inevitable challenges, e.g., how to best assess students' moral attunement and caring attitude during the course and evaluate it over time.

References

- Appadurai, A., & Alexander, N. 2020. *Failure*. Cambridge: Polity Press.
- Arnaldi, S. 2018. "Retooling Techno-Moral Scenarios: A Revisited Technique for Exploring Alternative Regimes of Responsibility for Human Enhancement." *NanoEthics* 12 (3): 283–300. <https://doi.org/10.1007/s11569-018-0329-6>.
- Arzberger, A., S. Menten, S. Balvert, T. Korteland, C. Zaga, E. Verhoef, D. Forster, and D. Abbink. 2023. "Barriers & Enablers to Transdisciplinarity in Practice: Emerging Learnings From a Pilot Project Exploring the Potential for Robotic Capabilities to Improve Nursing Work." *Relating Systems Thinking and Design Symposium. RSD12 Symposium*.
- Baxter, G., and I. Sommerville. 2011. "Socio-Technical Systems: From Design Methods to Systems Engineering." *Interacting with Computers* 23 (1): 4–17. <https://doi.org/10.1016/j.intcom.2010.07.003>.
- Bijker, W., T. P. Hughes, and T. Pinch (Eds.). 1987. *The Social Construction of Technological Systems. New Directions in the Sociology and History of Technology*. Cambridge: MIT Press.
- Bisante, A., A. Dix, E. Panizzi, and S. Zeppieri, S. 2023. "To err is AI." *Proceedings of the 15th Biannual Conference of the Italian SIGCHI Chapter*: 1–11. <https://doi.org/10.1145/3605390.3605414>.
- Bowden, P. 2010. "Teaching Ethics to Engineers – A Research-Based Perspective." *European Journal of Engineering Education* 35 (5): 563–572. <https://doi.org/10.1080/03043797.2010.497549>.
- Brevini, B. 2021. *Is AI good for the planet?* Hoboken: John Wiley & Sons.
- Brodie, Patrick. "Climate Extraction and Supply Chains of Data." *Media, Culture & Society* 42, no. 7–8 (October 2020): 1095–1114. <https://doi.org/10.1177/0163443720904601>.
- Bucciarelli, L. L. 2008. "Ethics and Engineering Education." *European Journal of Engineering Education* 33 (2): 141–149. <https://doi.org/10.1080/03043790801979856>.
- Castañeda, C., and L. Suchman. 2014. "Robot Visions." *Social Studies of Science* 44 (3): 315–341.
- Crawford, K. 2021. *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven: Yale University Press.
- Cila, N., I. Smit, E. Giaccardi, and B. Kröse. 2017. "Products as Agents: Metaphors for Designing the Products of the IoT Age." *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 448–459. <https://doi.org/10.1145/3025453.302579>.

- Coeckelbergh, M. 2011. “Humans, Animals, and Robots: A Phenomenological Approach to Human-Robot Relations.” *International Journal of Social Robotics* 3: 197-204.
- Coeckelbergh, M. 2012. *Growing Moral Relations: Critique of Moral Status Ascription*. Cham: Springer International Publishing.
- Couldry, Nick, and Ulises A. Mejias. “Data Colonialism: Rethinking Big Data’s Relation to the Contemporary Subject.” *Television & New Media* 20, no. 4 (May 1, 2019): 336–49. <https://doi.org/10.1177/1527476418796632>.
- Csikszentmihalyi, M., & Larson, R. (2014). Validity and Reliability of the Experience-Sampling Method. In *Flow and the Foundations of Positive Psychology* (pp. 35–54). Springer Netherlands. <https://doi.org/10.1007/978-94-017-9088-8>
- De la Bellacasa, M.P. 2017. *Matters of Care: Speculative Ethics in More Than Human Worlds*. Minneapolis: University of Minnesota Press.
- Elish, M. C. 2019. “Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction.” *Engaging Science, Technology, and Society* 5: 40–60. <https://doi.org/10.17351/ests2019.260>.
- Fraaije, A., van der Meij, M. G., Kupper, F., & Broerse, J. E. (2022). Art for public engagement on emerging and controversial technologies: A literature review. *Public Understanding of Science*, 31(6), 694-710.
- Forlano, L. “Posthumanism and Design.” *She Ji: The Journal of Design, Economics, and Innovation* 3, no. 1 (2017): 16–29. <https://doi.org/10.1016/j.sheji.2017.08.001>.
- Haraway, D. J. 1985. “A Manifesto for Cyborgs: Science, Technology, and Socialist Feminism for the 1980s.” *Socialist Review* 15 (2): 65–107.
- Haraway, D. J. 2003. *The Companion Species Manifesto: Dogs, People, and Significant Otherness*. Chicago: Prickly Paradigm Press.
- Haraway, D. J. 2016. *Staying With the Trouble: Making Kin in the Chthulucene*. Durham: Duke University Press.
- Heijnen, E., and M. Bremmer, M. (Eds.). 2021. *Wicked Arts Assignments: Practising Creativity in Contemporary Arts Education*, translated by L. Reijnen. Amsterdam: Valiz.
- Horton, D., Liu, D., McIlraith, S. A., Coyne, S., & Wang, N. (2024, March). Do Embedded Ethics Modules Have Impact Beyond the Classroom?. In *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1* (pp. 533-539).
- Johri, A., Katz, A. S., Qadir, J., & Hingle, A. (2023). Generative artificial intelligence and engineering education. *Journal of Engineering Education*, 112(3)
- Kaplan, F. 2004. “Who is Afraid of the Humanoid? Investigating Cultural Differences in the Acceptance of Robots.” *International Journal of Humanoid Robotics* 1 (03): 465–480.
- Kinzer, K. 2009. “Tweenbots by Kacie Kinzer.” <http://www.tweenbots.com/>.
- Kroes, P., M. Franssen, I. van de Poel, and M. Ottens. 2006. “Treating Socio-Technical Systems as Engineering Systems: Some Conceptual Problems.” *Systems Research and Behavioral Science* 23 (6): 803-814. <https://doi.org/10.1002/sres.703>.

- Kudina, O., and I. van de Poel. 2024. “The Sociotechnical Systems Perspective on AI.” *Minds & Machines* 34 (3): 21.
- Kudina, O. 2023. *Moral Hermeneutics: Making Moral Sense Through Human-Technology-World Relations*. Lanham: Rowman & Littlefield Lexington Books.
- Ley, M. 2023. “Care Ethics and the Future of Work: A Different Voice.” *Philosophy & Technology* 36 (1): 7.
- Michel-Villarreal, R., Vilalta-Perdomo, E., Salinas-Navarro, D. E., Thierry-Aguilera, R., & Gerardou, F. S. (2023). Challenges and opportunities of generative AI for higher education as explained by ChatGPT. *Education sciences*, 13(9), 856.
- Mirnig, N., G. Stollnberger, M. Miksch, S. Stadler, M. Giuliani, and M. Tscheligi. 2017. “To Err Is Robot: How Humans Assess and Act toward an Erroneous Social Robot.” *Frontiers in Robotics and AI* 4: 21. <https://doi.org/10.3389/frobt.2017.00021>.
- Mol, A. 2021. *Eating in Theory*. Durham: Duke University Press.
- Mollen, J., P. Van Der Putten, and K. Darling. 2023. “Bonding with a Couchsurfing Robot: The Impact of a Common Locus on Human-Robot Bonding in the Wild.” *ACM Transactions on Human-Robot Interaction* 12 (1): 1-33.
- Neimanis, A. 2017. *Bodies of Water: Posthuman Feminist Phenomenology*. London: Bloomsbury Academic.
- Nyholm, S. 2020. *Humans and Robots: Ethics, Agency, and Anthropomorphism*. Lanham: Rowman & Littlefield.
- Parvin, N., and A. Pollock. 2020. “Unintended by Design: On the Political Uses of ‘Unintended Consequences.’” *Engaging Science, Technology, and Society* 6: 320–327. https://doi.org/10.17351/ests2020.497_
- Pols, J. 2023. *Reinventing the Good Life: An Empirical Contribution to the Philosophy of Care*. London: UCL Press.
- Qadir, J. 2023. “Engineering Education in the Era of ChatGPT: Promise and Pitfalls of Generative AI for Education.” *2023 IEEE Global Engineering Education Conference (EDUCON)*, 1–9. <https://doi.org/10.1109/EDUCON54358.2023.10125121>.
- Smith, D. H., and F. Zeller. 2017. “The Death and Lives of hitchBOT: The Design and Implementation of a Hitchhiking Robot.” *Leonardo* 50 (1): 77–78.
- Stone, T. W., J. van Grunsven, and L. Marin. 2020. “Before Responsible Innovation: Teaching Anticipation as a Competency for Engineers.” In *Engaging Engineering Education: Proceedings of the 48th Annual SEFI Conference*, edited by J. van der Veen, N. van Hattum-Janssen, H.-M. Järvinen, T. de Laet, and I. ten Dam, 1371. Enschede: SEFI.
- Swierstra, T., D. Stemerding, and M. Boenink. 2009. Exploring Techno-Moral Change: The Case of the Obesity Pill. In *Evaluating New Technologies*, edited by P. Sollie and M. Düwell, 119–138. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-90-481-2229-5_9.
- Tronto, J. C. 2010. “Creating Caring Institutions: Politics, Plurality, and Purpose.” *Ethics and Social Welfare* 4 (2): 158-171.

TU Delft. (2024a). Ethics teaching at TU Delft. <https://www.tudelft.nl/ethics/ethics/teaching-activities/ethics-teaching>

TU Delft. (2024b). Environmental ethics course. Coordinator Andrea Gammon. <https://www.tudelft.nl/ethics/ethics/teaching-activities/ethics-teaching/ethics-20/environmental-ethics>

TU Delft. (2024c). Art, empathy, and ethics. Coordinator Filippo Santoni de Sio. https://studiegids.tudelft.nl/a101_displayCourse.do?course_id=48728

Valdivia, Ana. “The Supply Chain Capitalism of AI: A Call to (Re)Think Algorithmic Harms and Resistance (Extended Abstract).” *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* 7 (October 16, 2024): 1466–1466. <https://doi.org/10.1609/aies.v7i1.31738>.

Valentine, A., S. Lowenhoff, M. Marinelli, S. Male, and G. M. Hassan. 2020. “Building Students’ Nascent Understanding of Ethics in Engineering Practice.” *European Journal of Engineering Education* 45 (6): 957-970.

Vallor, S. (2016). *Technology and the virtues: A philosophical guide to a future worth wanting*. Oxford University Press.

Van de Poel, I. 2014. “Translating Values Into Design Requirements.” In *Philosophy and Engineering: Reflections on Practice, Principles and Process*, edited by D. P. Michelfelder, N. McCarthy, and D. E. Goldberg, 253-266. Dordrecht: Springer Netherlands.

Van Grunsven, J., & Roeser, S. (2022). AAC technology, autism, and the empathic turn. *Social Epistemology*, 36(1), 95-110.

Van Grunsven, J., L. Marin, T. Stone, S. Roeser, and N. Doorn. 2023a. “How Engineers Can Care from a Distance.” In *Thinking Through Science and Technology: Philosophy, Religion, and Politics in an Engineered World*, edited by G. Miller, H. M. Jerónimo, and Q. Zhu, 141-163. Lanham: Rowman & Littlefield.

Van Grunsven, J., L. Marin, and T. Stone. 2023b. “Fostering Responsible Anticipation in Engineering Ethics Education: How a Multi-Disciplinary Enrichment of the Responsible Innovation Framework Can Help.” *European Journal of Engineering Education* 49 (2): 283–298.

Van Grunsven, J., T. Franssen, A. Gammon, and L. Marin. 2024. “Tinkering with Technology: How Experiential Engineering Ethics Pedagogy Can Accommodate Neurodivergent Students and Expose Ableist Assumptions.” In *Building Inclusive Ethical Cultures in STEM*, edited by E. Hildt, K. Laas, C. Z. Miller, and E. M. Brey, 289-311. Cham: Springer International Publishing.

van Niekerk, J., Delport, P. M., & Sutherland, I. (2025). Addressing the use of generative AI in academic writing. *Computers and Education: Artificial Intelligence*, 8, 100342.

Verbeek, P.-P. 2005. *What Things Do: Philosophical Reflections on Technology, Agency, and Design*. University Park: Penn State Press.

Verbeek, P.-P. 2011. *Moralizing Technology: Understanding and Designing the Morality of Things*. Chicago: University of Chicago Press.

Wakkary, R., D. Oogjes, and A. Behzad. 2022. “Two Years or More of Co-speculation: Polylogues of Philosophers, Designers, and a Tilting Bowl.” *ACM Transactions on Computer-Human Interaction* 29 (5): 1-44.

Wakkary, R., D. Oogjes, H. W. Lin, and S. Hauser. 2018. “Philosophers Living With the Tilting Bowl.” *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, paper 94: 1-12. <https://doi.org/10.1145/3173574.3173668>.

Whyte, K. P., and C. J. Cuomo. 2016. “Ethics of Caring in Environmental Ethics. Indigenous and Feminist Philosophies.” In *The Oxford Handbook of Environmental Ethics*, edited by S.M. Gardiner and A. Thompson, 234–247. Oxford: Oxford University Press.

Yelamarthi, K., Dandu, R., Rao, M., Yanambaka, V. P., & Mahajan, S. (2024). Exploring the potential of generative AI in shaping engineering education: Opportunities and challenges. *Journal of Engineering Education Transformations*, 37(2), 439-445.