2023 IEEE International Conference on Robotics and Automation (ICRA 2023)
May 29 - June 2, 2023. London, UK

Ethical Assessment of a Hospital Disinfection Robot

Conor McGinn[1,2] and Robert Scott[2] and Niamh Donnelly[2] and Michael F. Cullinan[1,2]
and Alan Winfield[3] and Pat Treusch[1]

Abstract— **Robots have the potential to deliver very positive impacts for society, however, it's critical that in preparing for real-world deployments, we recognize and take steps to mitigate against the potential harms, both direct and indirect, that they may cause. In this paper, we explore how the ethics canvas (EC) and the ethical risk assessment (ERA) methodology defined in British Standard 8611 can be combined to better align robot technologies with ethics and their socio-cultural context of operation. We illustrate this through a practical case-study involving the real-world introduction of a disinfection robot to a radiology department in a European hospital. Using the EC, we identified 49 distinct ways that the technology was likely to impact key stakeholders and 11 ways that failure or misuse of the technology was likely to impact service provision. From this data, 8 mitigating measures were identified. Then, using the ERA tool, 9 risks were identified that were considered to represent a high likelihood of occurrence. From these insights, a further 8 mitigation measures were proposed. The combined use of both tools was found to be complementary, since the EC fostered a bottom-up, subjective critical thinking process whereas the ERA provided a broader, more top-down objective view. This example provides a practical template for robotics practitioners to better understand and manage the ethical and socio-cultural dimensions of their work, and contributes towards the standardization of ethical assessments in robotics with an emphasis on the move from principles to practice.**

I. INTRODUCTION

Hospital disinfection robots represent a fast growing new category of service automation [1][2]. Their adoption within clinical settings involves high operational and socio-cultural complexity, especially given their potential for labor displacement and the requirement to adapt workflows. These factors require that the robot's deployment is carefully planned to mitigate potential harms and risks beyond technological operability, necessarily involving the socio-cultural and ethical impact of such deployment. However, standards and best practices in guiding the ethical use of hospital disinfection robots, as well as other comparable robotic and artificial intelligence (AI) technologies, is lacking given that the rate of technological development in these areas tends to be faster than the speed at which professional standards are formed and harmonized. While companies using robots are required to comply with certain legal requirements (GDPR, HIPAA, etc.), these extend only to relatively narrow domains, such as physical safety and personal data collection, and provide less coverage against the potential for broader socio-cultural impacts and even harms.

Although it's widely accepted that the research community has an important role in the establishment of guidelines concerning the use of robots and AI technologies, detailed discussion of applied ethical issues often remains absent from technical publications [3]. Calls for greater ethical oversight of the design and use of AI have been increasing, driven by growing discourse around ethical dilemmas of AI [4] and awareness of harms arising from misuses of AI due to factors such as unbalanced data [5] and bad data curation [6]. Examples of harms include structural inequality [7], labor devaluation [8] and unemployment [9], and even weaponizing of AI systems, including robots [10]. It is particularly concerning that vulnerable groups and populations tend to be disproportionately affected by misuse of AI [11]. This problem extends into robotics where it has been observed that the socio-cultural and ethical implications of use-cases have often not been fully considered [12].

The lack of best practices to align the design of robots with socio-cultural values and ethical principles presents a significant barrier to the responsible development of robotics technology. In this paper, we set out to address this problem by exploring how two established tools for guiding ethical technology development, the ethics canvas and ethical risk assessment, can be harnessed in a systematic way to help robotics development teams identify and mitigate against a wide range of ethical risks, including socio-cultural factors. Through a case-study based on a real-world hospital disinfection robot application, we explore how the framework can be used to identify and mitigate against a wide range of ethical risks, including socio-cultural factors, while remaining accessible to practitioners that may not identify as experts in applied ethics. The paper is therefore of direct interest to researchers and practitioners whose work involves the deployment of robots, and of broad interest to research ethics committees, scientific publishers, and standards organizations with responsibilities of overseeing ethical research and development in robotics.

II. PRIOR WORK

In recent years, there has been considerable research involving to the development of ethics guidance for robotics practitioners. For example, a code of ethics has been suggested by Ingram et al. which outlines seven principles that robotics engineers should follow in their work [13]. A more comprehensive code of ethics has been proposed by Riek et al. based around the overall principle of "*respect for human persons, including respect for human autonomy, respect for human bodily and mental integrity, and the affordance of all*

---

[1] School of Engineering, Trinity College Dublin , Ireland
[2] Akara Robotics, Dublin, Ireland
[3] Bristol Robotics Lab, Bristol, UK

979-8-3503-2365-8/23/$31.00 ©2023 IEEE

rights and protections ordinarily assumed in human-human interactions" [14]. These codes of ethics, while useful in establishing an overarching ethical framework, provide limited practical support and guidance to researchers involved with explicit design, planning and execution of robots and robot-based studies.

Some, albeit limited, efforts have been made to develop more applied ways to incorporate ethical considerations into the development of robotic systems. The first of these was proposed by Van Wynsberghe and involved a method of incorporating ethics into the development of care robots from the early design stages [15] using a framework that involves consideration of the context, practice, stakeholders involved, the type of robot, and manifestation of moral elements. While the method is said to offer a 'step-by-step', guide to incorporate ethics during the design process, using the framework mandates the involvement of a dedicated ethicist which may not be feasible for many real-world robotics teams. Furthermore, the only demonstration of its use in the published literature involves hypothetical and retrospective design cases, making its utility difficult to evaluate.

More recently, O'Brien et al. explored how the Ethics Canvas, a collaborative brainstorming tool with the overall aim to foster ethically informed technology design, could be adapted to inform the development of a therapeutic robot for children [16]. The Ethics Canvas is structurally similar to Van Wynsberghe's framework while being more accessible to non-ethicists. However, since the tool focuses strongly on the use context, ethical considerations pertaining to the robot's design that are less application specific (i.e. environmental sustainability) may not be easily identified using this approach.

There has been recent progress in the development of professional standards concerning the ethical design of robot systems [17]. Arguably the first is British standard BS 8611-2016 (Robots and robotic devices: Guide to the ethical design and application of robots and robotic systems) [18]. A key component within the standard is guidance on how to perform an ethical risk assessment. This is elaborated in [19], where a case-study of an ethical risk assessment is provided based around a hypothetical robotic toy based loosely on a robotic character in a sci-fi movie which, according to the authors "*is far more advanced in intelligence, mobility and longevity than we can presently contemplate*".

Another standard that is of relevance to this research is IEEE 7007-2021 (Ontological Standard for Ethically Driven Robotics and Automation Systems). Despite being a new standard, since it deals with ontology, the task of incorporating it within practical robot design workflow is not straightforward.

III. METHODS

A. Methodology

We postulated that using a combination of the ethics canvas and the ethical risk assessment defined in British Standard 8611 (BS8611) offered a systematic way to gain a deep understanding of the ethical and social impacts of a robotic use case. The two methods were deemed to be complementary since completing the ethics canvas is a bottom-up process that focuses heavily on affected stakeholders and use context, whereas the ethical risk assessment is a more top-down approach which explores explicit aspects of the robot's design and aspects of operational deployment which may not normally appear in an ethics canvas (i.e. environmental sustainability).

Our ethical assessment was carried out in three stages. In stage one, we completed an ethics canvas. Then in stage 2, we identified ethical hazards, harms, and risks using the framework outlined in BS8611. In the third stage, we examined the applicability of each hazard identified in the risk assessment and considered whether (a) it had previously been identified using the ethics canvas, and (b) a mitigation measure proposed through the ethics canvas would mitigate the risk associated with the hazard. New mitigation measures were then proposed if they had not already been suitably addressed in completing the ethics canvas.

The team conducting the study included two engineers and a social scientist. One of the team members had prior experience using the ethics canvas. To ensure that the ethical risk assessment was implemented as intended in BS8611, we familiarized ourselves with a worked use-case [19] and corresponded with one of the co-authors of the standard, who provided additional guidance on how to carry out the assessment.

B. Ethics Canvas

The ethics canvas (EC) is a visual-linguistic tool that captures key social and ethical impacts of a technology or technology application. The EC (ver. 1.9 [20]) is structured into nine 'building blocks' and is orientated to help answer three basic questions: (1) Who might be affected by the technology? [blocks 1,2], (2) What are the potential ethical impacts for these people and groups? [blocks 3-8], and (3) how can we address these ethical impacts? [block 9] [21].

C. Ethical Risk Assessment (BS8611)

The ethical risk assessment (ERA), as defined in BS8611, is a formal method for identifying, analyzing, and taking steps to mitigate against potential ethical hazards that may cause harm to individuals, groups, or society. BS8611 comprises a set of 20 distinct ethical hazards[1] and risks, grouped under four categories: societal, application, commercial and financial, and environmental.

D. Use-case

The application we considered concerned the real-world use of a disinfection robot in a hospital. This is a use-case that the authors have prior experience with through current research activities [22], [23]. The robot used in the application was developed by Akara Robotics (Dublin, Ireland). It decontaminates surfaces using germ-killing UV

---

[1] BS8611 defines an ethical harm as "*anything likely to compromise psychological and/or societal and/or environmental well-being*" and an ethical hazard as "*a potential source of ethical harm*"

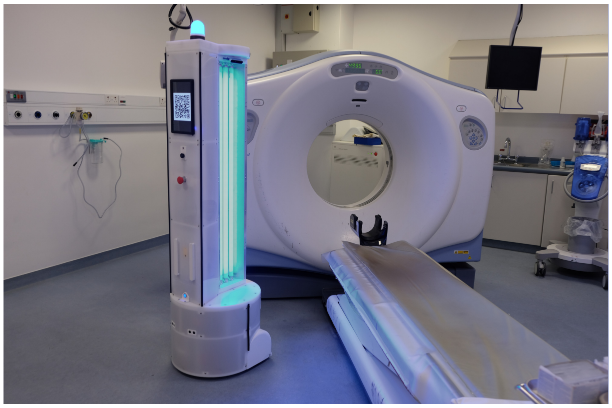

Fig. 1. Photo of the disinfection robot being used to decontaminate a radiology treatment room.

lamps housed in a rotating vertical column. During operation, the robot can move around the room autonomously and by regulating the orientation of the vertical column, it can control which areas of the room are irradiated. A photo of the robot is given in Figure 1. The following assumptions were made concerning the underlying use context:

- The robot is located in the radiology department and used primarily to decontaminate CT or Xray rooms after they have been occupied by an infectious, or potentially infectious patient.
- The hospital allocates specific treatment rooms for infectious patients. This is a practice which was introduced during the COVID-19 pandemic and remains in place at many hospitals at the time of writing [24]. This means that, during the course of a normal day, the robot is typically only required in 1-4 rooms.
- Prior to the introduction of the robot, room decontamination was normally performed by an environmental services worker employed directly by the hospital and based within the radiology department.
- It is standard practice for environmental services workers not to disinfect medical equipment in the room; this is done by clinical staff, normally the radiographer. sev
- Radiographers and environmental services staff receive training on how to use the robot and the standard operating procedures to be followed.
- Only parts of the room that are within line of sight to the exposed UV lamps receive UV irradiation. This means that at any time during a disinfection procedure, background UV levels are negligible in certain parts of the room. This has been verified through empirical testing.
- It is safe for staff to be present in the room when the robot is operating provided they wear suitable (i.e. UV rated) skin and eye coverings and follow standard operating procedures.
- When being moved between rooms, the robot is escorted by a member of the environmental services staff
- The hospital under consideration is a public hospital based in the republic of Ireland.

IV. ETHICS CANVAS

A. Stage 1: Identifying Stakeholders

The first stage of the Ethical Canvas is concerned with identifying key individual (block 1) and group stakeholders (block 2). The purpose of this stage is to better understand the people and groups who are likely to be most affected by the introduction of the technology. We identified 13 individual and 15 group stakeholders. These included a range of direct users (i.e. environmental services staff, radiographers, robot support staff), indirect users (i.e. radiology services manager, infection prevention team, facilities team), and tertiary users (i.e. hospital CEO/CFO, insurance company, healthcare standards authority). Direct users were defined as those that were required to directly operate the technology on a regular basis. Indirect users were people that did not work directly with the technology, but whose workflow was affected by its use. And tertiary users were stakeholders with limited day-to-day exposure to the robot, but would have an interest and influence over its use.

B. Stage 2: Identifying Ethical Impacts

During the second stage, which considers blocks 3-8 of the EC, potential ethical impacts for the different stakeholders are captured. This involves firstly identifying the ‘micro’ impacts that influence the everyday lives of people using and living with the robot. This part captured how peoples behaviour may change because of the robot (block 3) and also how relations between individuals and groups may change (block 4). The analysis of the use-case revealed two encouraging behavioural changes. The first was that since the robot could autonomously decontaminate the room, it may be possible to bypass the need for the environmental services worker and reallocate their time to something else. This would be beneficial to the hospital because it would reduce downtime associated with waiting for cleaning staff to arrive, and beneficial for the environmental services staff as it would alleviate some of the pressure they experience. The second benefit arises from real-time access to cleaning tasks performed. Since the robot can digitally log when it was used, facilities and infection prevention staff can conveniently monitor cleaning activities at the hospital, replacing the current approach which involves collecting and analysing paper check-sheets. Remarkably, the implementation of the technology in this way is likely to lead to increase responsibilities for radiographers, since they will need to learn how to use the robot, adopt new operating procedures and wear more PPE than may presently be the case. These anticipated changes clearly mark a shift of workflow that involves a restructuring of work practices on a micro level and that are propelled by the deployment of the robot. Thus, they are an ethical, socio-cultural concern that requires to be not only addressed but also negotiated between the different stakeholders involved.

Blocks 5 and 6 deal with potential ‘macro’ impacts. These impacts extend beyond the level of an individual’s everyday life and pertain to broader social structures. They consider

TABLE I

(SUMMARY OF PROPOSED MITIGATING MEASURES USING EC AND ERA (ASTERISK '*' INDICATES MEASURE ADDRESSED HAZARD SEPERATELY IDENTIFIED USING ERA)

ID	Proposed Mitigating Measure	Source
1	Perform environmental sampling on a periodic basis to demonstrate efficacy	EC*
2	Share feedback on improved standards and occupational safety with radiographers	EC
3	Involve staff, especially direct users, in developments of new work practices	EC*
4	Develop robot operator qualification to incentive up-skilling of radiographers and environmental services staff	EC*
5	Develop and maintain clear communications pathway to unions, regulators, insurance agencies and other external stakeholders	EC*
6	Monitor patient experience, especially during initial robot roll-out	EC
7	Build trust by validating technology with independent labs	EC*
8	Provide a service-based offering which addresses potential group conflicts and workflow disruptions	EC*
9	Implement sensors to measure UVC output of lamps	ERA
10	Don't record videos or images containing people	ERA
11	Restrict port access and perform regular pen tests to ensure security	ERA
12	Maintain transparency over what data is recorded by robot and make easily searchable	ERA
13	Require operation authentication to ensure only trained operators use robot	ERA
14	Record operator ID in digital cleaning logs	ERA
15	Ensure Akara retains ownership of robot at end-of-life to ensure recyclability	ERA
16	Implement EMS (BS ISO 14001)	ERA

how the use of the robot may influence people's worldviews (block 5) as well as inter-group conflict that may arise as a direct or indirect result of the technology (block 6). Our analysis indicated that the technology had the potential to affect the perspectives of users in a positive way. Firstly, it could augment how environmental services staff view automation, helping see it more as an enabler to change work in a welcomed manner and less as a threat. Second, the high value that could be unlocked by the technology could help hospital management see the value in engaging in innovation, motivating them to be more proactive in the adoption of new technologies. Interest groups such as unions and insurers may be confronted with the realization that their traditional ways of doing things may not be consistent with the changing times, and new procedures will be needed to co-exist with technology and automation. Due to the extent of workflow changes that the introduction of robotics requires, we anticipated numerous potential conflicts between the robot provider and different hospital stakeholders. Some internal conflicts were also identified, such as tensions between hospital management and labour unions, and tensions between infection prevention and radiology departments.

Finally, the remaining sections in stage 2 concern the negative side-effects of robot failure (block 7) and misuse (block 8). We identified that in the event of a critical failure, this could lead to two adverse operational outcomes: (1) procedures being cancelled or postponed arising from increased room downtime (i.e. to accommodate longer manual cleaning methods), and (2) high unscheduled demand for environmental services staff arising from the technology breakdown. We envisioned an alternative response scenario where the manual cleaning process was shortened, which may prevent cancellations, but would lead to a significant occupational safety risk.

We also envisioned cases where the presence of the robot might give staff a false sense of security, leading to a reduction in the standard and frequency that necessary manual cleaning was performed[2]. In addition, we identified quality and safety concerns associated with (1) the prospect of untrained staff operating the robot, and (2) the creation of robotic protocols for using the robot in new parts of the hospital without prior consultation with an expert in robotic UV disinfection. In this view, the system breakdown is only one potential failure while the circumstances of operating, i.e. the need of appropriate training and context-specific guidance, amount to an at least equally high risk of potential failures.

C. Stage 3: Addressing Ethical Impacts

The final stage, presented in block 9, involved outlining possible strategies to address the ethical and socio-cultural issues that emerged in previous stages. These responses typically involve adapting the design of the system, or making changes to how it is deployed and used in the application. Based on the ethical impacts identified, several proposed mitigating steps were identified (Table I).

V. ETHICAL RISK ASSESSMENT

When completing the ERA, we first evaluated the applicability each of the ethical hazards listed in BS8611 relative to our use-case, and the perceived likelihood of occurrence (LO) if no action was taken. We rated the likelihood of occurrence as being either low (L), medium (M), or high (H). These findings, grouped under the categories of societal, application, commercial and financial, and environmental, are given in Tables II-V.

In the societal category (Table II), we determined that job replacement was unlikely, since there will continue to be a strong demand for manual cleaning to remove physical particulates, however, the introduction of the robot would likely lead to some level of job disruption with the introduction of new practices and workflows. We also recognized that the robot's sensors created a data security risk, especially considering the fact that it may be co-located in a treatment

---

[2] Friction-based manual cleaning is still required to remove larger particulate matter i.e. dust, blood, dirt.

room with a patient. Finally, since the UVC radiation is not in the visual spectrum, an operator could be easily deceived if the wrong lamp was fitted to the robot.

TABLE II
SOCIETAL HAZARDS AND THEIR PERCEIVED LIKELIHOOD OF OCCURRENCE (LO)

ID	Ethical Hazard	LO
1	Loss of trust	L
2	Deception	H
3	Anthropomorphization	M
4	Privacy and confidentiality	H
5	Lack of respect for cultural diversity & pluralism	M
6	Robot addiction	M
7	Employment	H

In the application grouping (Table III), we identified a high possibility of misuse in the event that untrained operators were tasked with using the robot. We also anticipated divergent use and inappropriate "trust" if the hospital decided to expand its use beyond the radiology treatment rooms it had been optimized for without consulting an expert in automated UV disinfection.

TABLE III
APPLICATION HAZARDS AND THEIR PERCEIVED LIKELIHOOD OF OCCURANCE (LO)

ID	Ethical Hazard	LO
8	Misuse	H
9	Unsuitable divergent use	H
10	Dehumanization of humans	L
11	Inappropriate "trust" of a human by robot	H
12	Self-learning system exceeding its remit	L

When considering commercial and financial hazards (Table IV), in addition to recognizing further impacts to employment, we identified that while the potential for the robot to optimize its disinfection procedure over time using machine learning was desirable, current environmental auditing procedures (i.e. visual inspection) were not adequate to quantify the direction of performance improvement, creating a risk that performance could degrade over time.

TABLE IV
COMMERCIAL AND FINANCIAL HAZARDS AND THEIR PERCEIVED LIKELIHOOD OF OCCURANCE (LO)

ID	Ethical Hazard	LO
13	Approbation of legal responsibility & authority	M
14	Employment issues	H
15	Equality of access	L
16	Learning by robots with behavioral autonomy	H
17	Informed consent	M
18	Informed command	L

In the environmental category (Table V), we identified a relatively small risk due to lamps breaking since they contain small amounts of mercury. A second, and more likely risk, was that at their end-of-life, the robots would be sent to landfill rather than appropriately recycled.

For each hazard that was determined to have a high risk of occurrence, mitigation measures were identified. These are listed alongside recommendations produced by the ethics canvas and illustrated in Table I.

TABLE V
ENVIRONMENTAL HAZARDS AND THEIR PERCEIVED LIKELIHOOD OF OCCURANCE (LO)

ID	Ethical Hazard	LO
19	Environmental awareness (robot and appliances)	H
20	Environmental awareness (operations)	M

VI. DISCUSSION

The application of the EC and the ERA led to sixteen distinct measures to mitigate ethical risks associated with the use-case under investigation. The first eight of these recommendations originated from the EC and the remainder came from ERA. Of the eight recommendations proposed through the EC, six of them also addressed hazards that were deemed to have a high likelihood of occurrence based on the ERA. We do not believe this to mean that the ERA is the better instrument, since it is not clear whether the ERA would have led to the generation of the same mitigation measures. A more appropriate interpretation is that there are important areas where the ethics canvas does not provide coverage. This is a natural consequence of a bottom-up approach which has the benefits of obtaining a richer understanding of relevant contextual factors but may miss broader or more context independent issues such as environmental and data security hazards. Similarly, the BS 8611 standard identifies both likelihood of occurrence and severity as an important criterion for ethical risk identification, however, it only provides practical guidance on how to apply the former to the ERA. Future work should address this limitation.

There are also more specific hazards which are somewhat unique to robotics and AI technologies which were captured using the ERA, such as the potential for user deception and the negative impacts associated with autonomous systems capable of self-learning. As such, the data clearly demonstrates the two methods in their diverging approaches to be highly complementary. While the ERA highlights the structural framework and addresses general fields of risk of a mismatch, failure or even harm, the EC allows to open up ways of reflecting that forces the researcher to take on the perspective of the potential user and to ask the highly important questions of 'who', 'with', 'what', and 'how' that are central to any kind of deployment of a robot in a real-world setting. Further, it allows us to grasp with socio-cultural complexities within one specific context and given one use case example by working at the hinge between socio-cultural structure (top-down method) and socio-cultural practice (bottom-up method). Thus, this identified complementary effect, indicates the need for such cross-methodological approaches as presented in this paper.

On reflection, the biggest barrier to adoption for this use-case may be resistance from radiologists, who would be required to take on the most additional responsibilities and training. Furthermore, radiologists have been reported to experience lower job satisfaction than other clinical roles, and often report burnout [25]. There are three recommended measures to mitigate this risk, namely (1) clearly communicating how the technology will improve their occupational safety. This will help address concerns shared by many radiographers that their personal well-being as front-line workers has been overlooked [26], (2) their involvement with the development of new work practices helps ensure workflow integration of the technology, and (3) the development of a formal training program and qualification to help to motivate and incentive the necessary up-skilling required.

The proposed ethical mitigation measures are multi-faceted, involving updates to hardware (i.e. adding sensors to empirically measure intensity and wavelength of UV lamps), software (i.e. implement user authentication, update data logging) and operations (i.e. designing of staff training programs, validate efficacy on a periodic basis, involve staff in development of new work practices). It was observed that given the large proportion of operational changes needed, this would pose a significant challenge for a company involved with commercializing UV disinfection robots using traditional capital equipment purchasing models. This could potentially be overcome, however, if the company was to offer a service-based model, as proposed in mitigating measure 8, which provided additional supports in these areas.

One way in which this paper builds on previous studies is through grounding the application of the methods to a real-world use-case in a sector that the authors have prior experience and knowledge. While this has clear advantages over considering hypothetical examples, there may be some concerns that the findings may not generalize to other use-cases. However, this is unlikely to be the case provided there are no categorical differences in the type of robot or application. For example, it's likely that it would take relatively little effort to update the respective ECs and ERA to cover different application of a disinfection robot (i.e. use in other departments or in another health systems). Where categorical differences do exist, it's possible that large components of the EC (i.e. sections dealing with group relations, worldviews, problematic use of resources) and ERA (i.e. sections concerning topics of privacy, inappropriate trust, environmental awareness) are broadly generalizable to different use-cases, robot embodiments, and contexts.

We acknowledge that the proposed ethical risk mitigation measures do not address all possible hazards. Furthermore, in completing the ethics canvas, the data produced could not fit within the single page that is recommended. Nonetheless, should a new hazard emerge in the future, or a current hazard cease to be one, both the EC and the ERA can be updated, thus the proposed framework can be integrated as part of a fluid and continuous development process, readily adapting to changing circumstances over time. As such, it's use overcomes a big limitation of traditional ethics reviews carried out by institutional review boards in a 'snapshot' fashion before deployment of the technology, and which fail to incorporate valuable insights that might be generated from on-going critical evaluation or from reflection.

Unlike traditional safety risk assessments, it can be challenging to determine the likelihood of occurrence of an ethical hazard using purely objective measures. For example, the ERA identifies the tendency for a user to 'anthropomorphize' the robot as an ethical hazard. Ensuring ethical hazards are appropriately accounted for therefore depends heavily on the expertise, experience and make-up of the team performing the ethical review, and regular periodic revision of the EC and ERA based on insights from both qualitative and quantitative evaluation. The process of applying and refining the EC and ERA over a longitudinal period remains important future work.

VII. CONCLUSIONS

This study set out to investigate if the combined use of the ethics canvas and ethical risk assessment could support robotics practitioners to ensure the ethical deployment of a disinfection robot in a hospital setting. The use of both methods facilitated deep consideration of the stakeholders affected, potential ethical risks of socio-cultural harms (at both micro and macro levels), potential consequences of robot failure and resource demand, as well as possible remediation measures.

The method proved to be accessible to non-ethicists, which was evidenced by the fact that analysis was predominately undertaken by engineers and social scientist with no formal ethics training. Assembling insights through the application of both methods permitted a holistic evaluation of the problem, and ultimately cumulated in a comprehensive representation of the ethical landscape along with a number of recommendations to address key ethical issues. Findings from this research could ideally contribute directly to the establishment of best-practices concerning the use of disinfection robots in hospitals and provides a useful demonstrator that motivates adoption of the underlying approaches by other robotics researchers, institutional review boards, funding agencies and other stakeholders concerned with the ethical development of AI technologies.

ACKNOWLEDGEMENTS

This research was partially funded by European Union's Horizon 2020 Research and Innovation programme project no. 825003: Digital Innovation Hubs in Healthcare Robotics (DIH-HERO) BOSCUR project. It was further partially funded by and European Union's Horizon 2020 European Innovation Council Accelerator, project no. 190106381, AVRICS and a grant from the SFI ADAPT and from the European Union's Horizon 2020 Research and Innovation Programme under the HUMAN+ COFUND Marie Skłodowska-Curie grant agreement No. 945447.

ADDITIONAL MATERIALS

Copies of the completed ethics canvas and ethical risk assessment can be downloaded from `https://github.com/Conor-McGinn/ethical_assessment`

REFERENCES

[1] R. Scott, L. T. Joshi, and C. McGinn, “Hospital surface disinfection using ultraviolet germicidal irradiation technology: A review,” *Healthcare Technology Letters*, vol. 9, no. 3, pp. 25–33, 2022.

[2] I. Mehta, H.-Y. Hsueh, S. Taghipour, W. Li, and S. Saeedi, “Uv disinfection robots: A review,” *Robotics and Autonomous Systems*, p. 104332, 2022.

[3] A. Birhane, P. Kalluri, D. Card, W. Agnew, R. Dotan, and M. Bao, “The values encoded in machine learning research,” in *2022 ACM Conference on Fairness, Accountability, and Transparency*, 2022, pp. 173–184.

[4] K. Crawford and R. Calo, “There is a blind spot in ai research,” *Nature*, vol. 538, no. 7625, pp. 311–313, 2016.

[5] J. Buolamwini and T. Gebru, “Gender shades: Intersectional accuracy disparities in commercial gender classification,” in *Conference on fairness, accountability and transparency*. PMLR, 2018, pp. 77–91.

[6] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, “On the dangers of stochastic parrots: Can language models be too big?” in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 2021, pp. 610–623.

[7] I. F. Ogbonnaya-Ogburu, A. D. Smith, A. To, and K. Toyama, “Critical race theory for hci,” in *Proceedings of the 2020 CHI conference on human factors in computing systems*, 2020, pp. 1–16.

[8] A. Kshirsagar, B. Dreyfuss, G. Ishai, O. Heffetz, and G. Hoffman, “Monetary-incentive competition between humans and robots: experimental results,” in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2019, pp. 95–103.

[9] D. Acemoğlu and P. Restrepo, “Robots and jobs: Evidence from the us,” *NBER Working Paper No*, vol. 23285, 2017.

[10] A. Krishnan, *Killer robots: legality and ethicality of autonomous weapons*. Routledge, 2016.

[11] A. Birhane, “Algorithmic injustice: a relational ethics approach,” *Patterns*, vol. 2, no. 2, p. 100205, 2021.

[12] P. Lin, K. Abney, and G. A. Bekey, *Robot ethics: the ethical and social implications of robotics*. The MIT Press, 2014.

[13] B. Ingram, D. Jones, A. Lewis, M. Richards, C. Rich, and L. Schachterle, “A code of ethics for robotics engineers,” in *Proceedings of the 5th ACM/IEEE International Conference on Human-robot Interaction*. IEEE Press, 2010, pp. 103–104.

[14] “A code of ethics for the human-robot interaction profession,” *Proceedings of We Robot*.

[15] A. Van Wynsberghe, “A method for integrating ethics into the design of robots,” *Industrial Robot: An International Journal*, 2013.

[16] C. O’Brien, M. O’Mara, J. Issartel, and C. McGinn, “Exploring the design space of therapeutic robot companions for children,” in *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, 2021, pp. 243–251.

[17] A. Winfield, “Ethical standards in robotics and ai,” *Nature Electronics*, vol. 2, no. 2, pp. 46–48, 2019.

[18] B. BS8611, “Robots and robotic devices, guide to the ethical design and application of robots and robotic systems,” *British Standards Institute*, 2016.

[19] A. F. Winfield and K. Winkle, “Roboted: a case study in ethical risk assessment,” *arXiv preprint arXiv:2007.15864*, 2020.

[20] W. Reijers, K. Koidl, D. Lewis, H. J. Pandit, and B. Gordijn, “Discussing ethical impacts in research and innovation: The ethics canvas,” in *IFIP International Conference on Human Choice and Computers*. Springer, 2018, pp. 299–313.

[21] D. Lewis, W. Reijers, H. Pandit, and W. Reijers, “Ethics canvas manual,” 2017.

[22] C. McGinn, R. Scott, N. Donnelly, K. L. Roberts, M. Bogue, C. Kiernan, and M. Beckett, “Exploring the applicability of robot-assisted uv disinfection in radiology,” *Frontiers in Robotics and AI*, vol. 7, p. 590306, 2021.

[23] C. McGinn, R. Scott, C. Ryan, N. Donnelly, M. F. Cullinan, and M. Beckett, “Rapid disinfection of radiology treatment rooms using an autonomous ultraviolet germicidal irradiation robot,” *American Journal of Infection Control*, 2022.

[24] M. J. Tuite, “What preparedness advice for covid-19 did radiology departments follow?” *Radiology*, vol. 304, no. 1, pp. 126–127, 2022, pMID: 35258377. [Online]. Available: https://doi.org/10.1148/radiol.220362

[25] T. N. Akudjedu, N. A. Mishio, W. Elshami, M. P. Culp, O. Lawal, B. O. Botwe, A.-R. Wuni, N. Julka-Anderson, M. Shanahan, J. J. Totman *et al.*, “The global impact of the covid-19 pandemic on clinical radiography practice: A systematic literature review and recommendations for future services planning,” *Radiography*, vol. 27, no. 4, pp. 1219–1226, 2021.

[26] S. J. Foley, A. O’Loughlin, and J. Creedon, “Early experiences of radiographers in ireland during the covid-19 crisis,” *Insights into imaging*, vol. 11, no. 1, pp. 1–8, 2020.