

Over-Researched and Under-Resourced: The ELSA Approach to Transdisciplinary AI Research in Low-Trust Neighbourhoods

Nanou van Iersel¹[0009-0007-0650-4030], Majsja Storbeck²[0009-0003-4584-0046], Marlon Krui-
zinga³[0009-0001-6400-118X] and Michaël Grauwde⁴

¹ Erasmus University Rotterdam, Burgemeester Oudlaan 50, 3062 PA Rotterdam, the Netherlands

vaniersel@law.eur.nl

² Erasmus University Rotterdam, Burgemeester Oudlaan 50, 3062 PA Rotterdam, the Netherlands

³ Erasmus University Rotterdam, Burgemeester Oudlaan 50, 3062 PA Rotterdam, the Netherlands

⁴ Delft University of Technology, Mekelweg 5, 2628 CD Delft, the Netherlands

Abstract. This paper discusses ELSA (Ethical, Legal, and Social Aspects of technology) as an emerging methodology for transdisciplinary AI research, characterized by anticipatory technology assessment through close collaboration with diverse (societal) stakeholders. We offer a methodological reflection based on a 1,5 year-long case study on public safety and AI in Lombardijen, a neighbourhood in Rotterdam, the Netherlands, where we engaged residents as citizen stakeholders. Lombardijen is paradoxically under-resourced, meaning historically neglected and stigmatized as a ‘problem district’, yet over-researched, i.e. scrutinized by countless researchers who engage in what has been called ‘drive-by’ research – driving by, extracting data, and disappearing, often without benefits for the community. The community’s ensuing alienation from governmental and academic institutions means that citizens’ valuable contextual knowledge is often overlooked in public deliberation on AI. This raises our research question: How can citizens in low-trust neighbourhoods be meaningfully and reciprocally engaged in transdisciplinary AI research, and what does an ELSA approach offer in this regard? The paper details our experiences in Lombardijen respectively from ethical, legal, social, and technological perspectives. We candidly discuss our learnings, (modest) successes and limitations, ultimately emphasizing the importance of situated responsibility as a precondition for transdisciplinary AI research.

Keywords: Transdisciplinary Research, ELSA, Citizen Engagement, Artificial Intelligence, Digital Technologies, Public Safety, Neighbourhoods.

1 Introduction

The history of scientific practice and research is also a history of extractivism: the socio-ecologically destructive appropriation and subjugation of human and natural resources through domination without accountability for (grave) consequences [1]. Smith [2] has powerfully described how Indigenous communities have experienced Western

research as a continuation of colonial exploitation, whereby cultural practices, artefacts, even genetic data are extracted with little recognition of Indigenous knowledge practices or the consequences for those being studied.

These critiques are not limited to historically colonized regions or Indigenous contexts. Similar dynamics are at work in cities in the Global North with high urban inequalities, like Detroit, United States [3], or Rotterdam, the Netherlands [4], where under-resourced neighbourhoods are routinely engaged by academic and governmental institutions in ‘drive-by research,’ ‘where researchers drive by, collect the data, and keep going’ [3, p. 2]. In October 2023, the Dutch national newspaper *Volkskrant* published an elaborate account of Rotterdam-south residents feeling worn out by the constant demands of research in the neighbourhood, offering little more than a sense of being scrutinized [5]. This ties in with broader concerns of ‘over-researched’ communities [6] experiencing ‘research fatigue’ [7].

In precisely this context, of research exploitation in Rotterdam South, we situated our research: a 1,5 year-long case study in the sub-district Lombardijen on public safety and artificial intelligence (AI) in the neighbourhood. Significantly, AI itself has been characterized as product of exploitation, built on the extraction of natural resources, war-minerals and low-paid, often invisible and precarious ‘click-work’ [8], a phenomenon termed data coloniality [9]. This dual extractivism – embedded both in research and AI – presents a challenging context. Can and should communities and citizens, who have, with good reason, grown distrustful, be engaged in transdisciplinary AI research?

The alternative of leaving over-researched and under-resourced communities alone risks reinforcing the historical neglect that already affects them. From an egalitarian perspective, these communities should be included in public deliberation on technologies shaping their everyday life. Pragmatically, their participation is also essential: residents bring contextual knowledge of public safety in the neighbourhood [10], with which technological ‘solutions’ should align. There are multiple ways of knowing [11]; knowledge from research and experience contribute equally valuable pieces to the puzzle of public safety and AI in the neighbourhood.

To address how such communities can be engaged, we adopted an ELSA approach: a research methodology for anticipatory technology assessment, defined by close collaboration with diverse (societal) stakeholders with the aim of mutual learning throughout the research cycle [12]. This paper is a methodological reflection on ELSA, based on the case study in Lombardijen. The research question is: *How can citizens in low-trust neighbourhoods be meaningfully and reciprocally engaged in transdisciplinary AI research, and what does an ELSA approach offer in this regard?* In what follows, we begin with an explanation of ELSA, after which we describe the case study in Lombardijen in more detail. The core of the paper reflects on our experiences in Lombardijen from ethical, legal, social and technological perspectives. In the discussion, we consider what worked well and what could be different in our approach, concluding with the importance of situated responsibility in transdisciplinary AI research.

2 ELSA: A Transdisciplinary Research Methodology

What is the ELSA (Ethical, Legal and Social Aspects) methodology and more specifically what is ‘the ELSA attitude?’ [13]. ELSA’s roots are in its American predecessor ELSI (Ethical, Legal and Social Implications), focusing on the societal impact of genetic research [14]. This framework inspired counterparts in Canada and the European Union, marking the beginning of what Zwart and Nelis [15] call the ‘Elsification’ of large techno-science programs. The EU for instance institutionalized ELSA in 1994 under its Framework Programs for Research and Technological Development [16], effectively becoming *the* preferred approach to integrate ethical and societal considerations into technology and science [12, 14].

The term ‘aspects’ (ELSA) is, in the EU context, now preferred over ‘implications’ (ELSI) because ‘implications’ has a negative connotation, whereas ‘aspects’ encapsulates both positive and negative dimensions of the technology-society relationship [14, 17]. Although ELSA experienced a dip in the 2010s, with the preference for Responsible Research and Innovation (RRI), the recent creation of more than 20 Dutch ELSA of AI labs suggests a serious resurgence since 2020 [18]. This revival has also manifested into government frameworks, illustrated by the Dutch Police’s 2025 announcement to integrate ELSA into their governance of AI and data-driven technologies, recognizing the risks they pose for public values [19, 20].

Methodologically, the ‘ELSA approach’ is defined by four elements: (1) Proximity, (2) Anticipation, (3) Interactivity and (4) Interdisciplinarity [15]. *Proximity* originally referred to the close contact between ELSA researchers and exact scientists and engineers, but we now interpret it more broadly as proximity to the practices of various stakeholders; *anticipation* means early-on assessment of the (potential) effects of technology; *interactivity* is the encouragement of diverse stakeholders, including citizens, to co-shape research agendas and articulate results; and *interdisciplinarity* refers to the integration of various research fields.

ELSA’s transdisciplinary nature is reflected in the elements of *proximity* and *interactivity*. Transdisciplinary research engages with societal stakeholders to address complex research questions, bringing together perspectives from research and practice [21]. It necessitates the engagement of relevant stakeholders at the earliest possible stage of AI research and/or development, including from industry, civil society and citizenry, government, research institutions and sometimes (and increasingly) nature [13, 22]. In aiming for shared knowledge production with diverse (societal) stakeholders, ELSA is guided by the notion of epistemic inclusion: the recognition that there is a plurality of equally valid ways of knowing, including research and scientific knowledge, Indigenous knowledge, professional knowledge and contextual or experiential knowledge [11]. The objective is to map values from a broad range of perspectives, which allows for convincing democratic considerations of ethical-legal dilemmas of AI; in turn feeding the desired socially-responsible innovation process [23].

However, ‘involving all stakeholders mentioned in the ELSA methodology – including citizens – is not easy’ [20]. Whereas the participation of industry and government is often immediately assured (in AI), it is often said that ‘ordinary’ citizens do not have enough knowledge about the technology to make meaningful contributions to its design

or desirability. Yet, their experiential knowledge of the direct living environment represents as valuable a puzzle piece as the technical knowledge of engineers or the legal knowledge of policymakers, which is equally partial [10]. Given the context of research fatigue and the undervaluation of citizen expertise, we knew that the engagement of citizens as stakeholders in ELSA would be particularly challenging. The next section outlines our approach to citizen engagement in an under-resourced and over-researched neighbourhood: Lombardijen.

3 Case Study: Lombardijen, the Netherlands

The Lombardijen case study emerged from the ELSA Lab ‘AI-MAPS’, a research project established to assess the impact of AI in public safety through four work packages, covering ethical, legal, social, and technological aspects. In this case study, we engaged with residents from an under-resourced and low-trust neighbourhood. To capture the economic vulnerability implied in ‘under-resourced’, our research was in Rotterdam South, where many socioeconomically vulnerable households are concentrated [24]. The area is marked by debt, unemployment, social exclusion and disproportionate impact of the child’s benefit scandal¹ [25, 26], as well as residents’ sense of neglect by authorities [27].

Over 1.5 years (March 2024 – June 2025), we carried out participatory action research (PAR) in Lombardijen; a relatively small neighbourhood in Rotterdam South, where crime rates are stable but perceived safety has dramatically declined since 2018 to one of the worst scores city-wide in 2024 [28, 29]. This disconnect, between public safety as captured by ‘facts and figures’ and perceived safety of residents, formed part of the reason to go to Lombardijen.

The case study was centred around AI and public safety in the neighbourhood. Specifically, our focus was on ‘public nuisance’ – perceived disruptions in the public order lowering the quality of life in the neighbourhood – and ‘positive safety’: a perspective on public safety that moves beyond policing and crime control, focusing instead on broader factors that contribute to (or detract from) liveability and feeling at home and secure in the neighbourhood [30, 31]. The question, then, quickly became: what concrete factors, of public nuisance or lack thereof, affect the positive safety of residents in Lombardijen? And, what role, if any, is there for (AI-)technologies, actually or potentially, as either a helpful or disruptive factor? We adopted a broad definition of AI, including related digital technologies in our assessment. While AI in the neighbourhood is limited (though present in ‘smart’ doorbells), digital technologies like public camera surveillance are pervasive, and form the infrastructure for AI-capabilities (e.g., automatic violence detection as an AI-enhancement for ‘regular’ camera surveillance).²

Coming to Lombardijen, as we somewhat expected, residents in many cases had pre-established distrust or cynicism towards both local government and researchers. Residents had experiences with not being heard by the municipality, having longstanding problems in public space go unaddressed, and participating in research that they never

¹ This scandal involved the government wrongly targeting thousands of families for childcare benefits fraud.

² Automatic violence detection has in fact been piloted in Rotterdam in 2023 [32].

saw the results from. Lombardijen has at times been designated as a ‘problem district’ in Rotterdam, and all of Rotterdam South has in the past been treated as a sort of ‘living lab’ for research, partially due to its challenging circumstances, and partially because it is perceived as a ‘well of untapped knowledge’ [33]. Such intensive academic attention can backfire, as well-intentioned as the research is [5]. The responsibility of researchers to break through cynicism and distrust, as well as recognizing the many ways in which this distrust is legitimate, occupies the majority of this paper.

We started our case study by attending existing neighbourhood initiatives and groups, to integrate ourselves into the goings-on in the neighbourhood, especially among active members of the community. We made use of ‘gatekeepers’, key persons with pre-existing connections in the neighbourhood, to gain access to the aforementioned neighbourhood initiatives and groups. One was the ‘city marine’, a special civil servant with the municipality, meant to be the bridge between them and the residents. Other gatekeepers included the ‘district manager’ (another municipality figure), the leader of a neighbourhood watch group, and a member of the neighbourhood council. These persons were important not only because they ‘knew people’ in the neighbourhood and could give us contextual knowledge about the community, but also because they had pre-established trust from residents, which could to some degree be conferred onto us when they were the ones introducing us to new potential participants.

With ‘access’ to community through our gatekeepers, we subsequently made it a point to not immediately collect data but to ‘start by giving back’. This was partially in response to warnings from our key persons that immediately asking residents to participate in our research would be met with resistance, as well as due to our own understanding of the over-research problem. Our idea was that, in switching this dynamic around, giving ‘back’ attention, effort, and support before receiving data, we could prove to residents that we had a serious engagement with the neighbourhood, that we took their experiences seriously, and we were ready to centre what they perceived as relevant issues to research.

We started with ‘giving back’ by participating in neighbourhood initiatives: picking up trash, walking along with neighbourhood watch groups, attending neighbourhood meetings with the municipality, baking cakes for neighbourhood initiatives, and letting emergent conversations with residents and other stakeholders inspire the direction of our research. Subsequently, we started to mix ‘giving back’ with data collection: we organized community workshops with residents, students and decisionmakers.³ Participants were asked to send in photos of the neighbourhood beforehand, and we used physical copies to spark discussion during the sessions. Based on what participants defined as the most salient issues and points of pride, we then created visualizations of an ideal Lombardijen using generative AI. The workshop series concluded with a closing meeting in the community centre, where the AI visualizations and residents’ favourite photos were exhibited.

The final workshop in Lombardijen, which served to bookend our case study, took residents, researchers and municipality workers on an ecological tour through the

³ The two community workshops (October 2025) had 12 and 18 participants respectively (residents/students). A follow-up (January 2025) involved 24 decisionmakers from inter alia municipality and social housing.

neighbourhood, sitting them down afterwards to collectively brainstorm about interventions for issues of trash and greenery, which residents had identified as key factors in their perceived (positive) safety. In the next section, we outline the details of our approach for each work package in the ELSA Lab, covering ethical, legal, social, and technological aspects. We discuss the specific research activities undertaken by each work package, key empirical insights, and broader reflections on citizen engagement.

4 ELSA Reflections

4.1 Ethical Takeaway: Fostering Trust Through Hands-On Experience

The following sections detail our experiences in Lombardijen from ethical, legal, social and technological perspectives. We start with the ‘E’ in ELSA: the ethical work package, steered by its philosophy PhD candidate. One of the main sources of empirical insights are the interviews on the ethical reasoning of neighbourhood stakeholders on AI-powered cameras. This, paired with the participatory observation carried out in the neighbourhood by all the work packages, furnished a few unique ethical insights on ELSA research, as well as on inter stakeholder participation and collaboration in research and technological development. The overall theme in these insights is trust, both among societal stakeholders and between stakeholders and researchers.

As mentioned previously, the approach of ‘giving back from the outset’, engaging proactively with existing groups and initiatives in Lombardijen on their terms and helping them, can be framed as a functionalistic approach to overcome research fatigue and distrust. However, it is also worth considering how this distrust, especially on the parts of residents in vulnerable, highly researched neighbourhoods, may be very justified. If a person or a community is in a vulnerable socio-economic position, and if they have experience with giving their time and energy to generate data for a research project of which they themselves never experience significant positive outcomes, it is rational that, gradually, they become more protective of their limited resources.

In this context it is important to realize a potential disconnect in the experiences of researchers and those of other stakeholders, especially residents. Our ‘researcher perspective’ is largely informed by what is useful and appreciated within our own field, and thus it has the tendency to focus on matters like publication, (textual) authorship and recognition of one’s ideas. In lieu of this, it is not uncommon for researchers to be advised, from higher-up, to make stakeholders simply ‘co-authors’ of publications and workshops. However, as we were informed by our gatekeepers, residents are often more interested in concrete, tangible rewards, even if these may seem less significant for us as researchers. Simple things like a free lunch, a social event, physical help with tasks and initiatives, or even potential influence in municipal policy, proved to be more appreciated than ‘credit’ on a research paper.

Falling short of providing these kinds of ‘giving back’ risks making research extractive. Additionally, we must recognize that generating data about certain communities can cause concrete harms. One such case is provocatively introduced by Trottier, Lee, and Boy in their 2022 paper on the methodology and ethics of urban data analytics:

‘At the time of writing, a Google News search of the Dutch neighbourhood Moerwijk revealed terms like ‘rattenoverlast’, ‘rattenplaag’, and ‘ratten zo groot als katten’ (rat nuisance, rat plague and rats as big as cats). For a prospective resident, this may be one of many encounters with digital content that paints an unflattering account of the community. While the rat problem may be a demonstrable fact (as would be the case elsewhere), its prominence online is evidence of this neighbourhood’s embattled reputation. Public records and newspapers have long remained a means to render urban spaces meaningful, generating knowledge that serves to exclude disadvantaged communities’ [34, p. 321].

This kind of reputational damage and ‘data-based disadvantaging’ of already vulnerable communities is obviously made possible not just through public records and newspapers, but also through (open access) academic research. What may be intellectually stimulating (‘sexy’) for researchers about a neighbourhood may be newsworthy for some third parties, politically or strategically useful for others, and outright damaging for the prospects and opportunities of actual residents. All of this makes community members’ distrust of researchers, as well as their concomitant taciturnity, more than reasonable. This, in turn, means that there may not only be a functionalistic, but also a moral reason to engage conscientiously with potential research participants and to earn their trust, rather than presuming it.

We took these kinds of ethical considerations into account in the case study in Lombardijen and were shown first-hand that long-term time investment and engagement with the issues and initiatives of residents, as well as not overpromising what you can deliver, can break through some of the distrust. The workshops we organized would likely not have had the same level of resident-participation if the organizing PhD researchers had not been picking up trash and taking neighbourhood watch walks with many of these same residents. We were transparent about not being able to promise direct policy change, but we were able to bring residents into conversations with representatives from the municipality and the local housing corporation among others. Residents were willing to return for subsequent project events – whether for a free lunch and social gathering, the chance to speak with public officials, or interest in the research itself – indicating that some trust had been established.

The built-up trust, and the functional importance of said trust for successful participatory ELSA research, became even more clear once the interviews with stakeholders on their ethical reasoning about AI-powered cameras started. Almost all participants for these interviews were gained either from previous workshops, or from members of the neighbourhood watch group and the trash-cleaning group, with whom extensive participatory observation was conducted. Participants, especially the residents, evidently had come to believe that we took their and their neighbourhood’s interests to heart, having invested genuine effort into their own initiatives for betterment. Thus, they were not only willing to participate in workshops; they were also willing to engage with interviews which had very little directly to do with their problems and interests, seemingly as a sign of deeper trust and reciprocity.

In the interviews, a pattern emerged in the answers of residents to questions about their ethical reasoning on AI. Many residents indicate that they think they ‘could have a stronger ethical stance’ on technologies like ‘smart’ cameras in their neighbourhood,

but that they do not bother because, to paraphrase, ‘it would not make a difference’. Extrapolating on this, residents say that they think the technologies will enter their communities and living environments no matter what they may feel about it, and will be shaped not by them, but by technologists and decisionmakers. As a result, they see no point in formulating ethical opinions about the technologies in question. From both a broadly ethical and a participatory ELSA research perspective, this should be seen as concerning. It represents a bottleneck for empirical-ethical research, because certain stakeholders, namely residents, are unmotivated to participate fully in shaping the ethical approach towards novel technologies.

This appears related to the lack of trust: residents do not trust that their input would genuinely matter to how new technologies are designed and implemented and pull out of the process altogether, becoming somewhat morally apathetic. This, however, we would argue only confirms the importance of participatory research that puts the emphasis on earning the trust of potential participants, and honouring it by centring their interests and perspectives, including those of residents. Giving back from the outset, being transparent and clear about prospects and eventual results, paired with a strong commitment to generating tangible results for participants within a project’s means, have been our main tools for honouring that trust. The next section reflects on the experiences of the legal work package, linking trust to accountability.

4.2 Legal Takeaway: Unpacking and Practicing Our Role in Accountability

For the legal work package, steered by its legal PhD candidate, empirical insights mostly come from participatory observation at community events and workshops, carried out across all work packages, next to interviews and site visits with local housing corporations responsible for social housing. The key focus, accountability, in many ways, depends on the theme of trust. Accountability is a relationship between two types of actors: the actor whose conduct is being evaluated and the evaluating party or forum. The forum requires access to information and, upon assessment, can impose negative (blame) or positive (praise) consequences [35]. Being a relationship, accountability cannot meaningfully emerge without trust between the involved parties.

In Lombardijen, we observed how fragmented and diffuse accountability for public safety can be. A variety of stakeholders, including housing corporations, municipality, police, and residents themselves, carry some degree of responsibility for pleasant living in the neighbourhood. When numerous stakeholders contribute incrementally to a collective outcome like pleasant living, often without seeing the bigger picture, no single stakeholder may *feel* or be reasonably *held* accountable. This accountability gap, also known as the ‘problem of many hands’ [36], becomes particularly problematic when the collective outcome is ineffective or insufficient, or worse, harmful.

This ‘problem of many hands’ can also be observed in Lombardijen. For instance, residents voice trash, and unclear responsibilities for cleaning it, as a key issue. Housing corporations manage waste within their complexes, while the municipality oversees public space. There can be disputes over grey areas, as one interviewed policymaker from a housing corporation explained: bikes left leaning against building walls but technically standing in public space can result in finger-pointing between the municipality

and the housing corporation, while residents continue to experience nuisance. Another interviewee from a different housing corporation noted that they manage waste not only within their complexes but also in the surrounding greenery, highlighting variations in how duties are interpreted across institutions; this, in turn, may create inconsistent expectations for residents.

Alongside the ‘problem of many hands’, we observed a ‘problem of many heads’. Different stakeholders have varying understandings of what public safety problems exist in the neighbourhood and how these should be addressed (and whether with technology). Policymakers and technology entrepreneurs may operate from entirely different problem conceptions than residents themselves. Consequently, (AI-)technologies may be introduced into the neighbourhood that may not only be unclear or irrelevant but also harmful to residents. To illustrate, residents were often disappointed to learn that public camera surveillance does *not* entail that (i) footage is monitored in real time with immediate responses to disturbances or crime, nor (ii) that footage can be used to sanction any type of nuisance or crime.

Similarly, AI-powered camera surveillance in the form of ‘smart’ video doorbells creates inconsistent expectations for residents. While these are only permitted to film private property, in practice, many record public space too. This constitutes a violation of data protection laws, but the Dutch police routinely request doorbell footage for criminal investigations.⁴ Some residents and housing corporations feel uncomfortable with the video doorbells, but have little recourse when the police effectively accepts the privacy violation. Other stakeholders, however, saw the benefits: a gatekeeper highlighted a case of a neighbourhood crime which was solved using doorbell footage.

The value of an ELSA approach was particularly evident in the discussions about (AI) camera surveillance with residents and representatives from the municipality and housing corporations. In the position of researchers, we were able to highlight and clarify responsibilities related to specific public safety problems and (AI-)technologies. The legal dimension contributed not only to legal literacy (around camera surveillance) but also framed legal questions in relation to issues of desirability, technical feasibility, and social impact. This helped clarify what the law can and cannot do for residents and encouraged them to see the law as part of a broader context shaped by ethical, social and technological aspects. Moreover, by facilitating discussions with different (social) stakeholders, we were able to act as a bridge between residents and institutions. Recognizing that we possessed the capital to connect these groups, we relayed residents’ concerns to institutional stakeholders and helped them better understand local expectations and realities. This bridging role became another way of ‘giving back’.

4.3 Social Takeaway: Developing ‘Out-Of-The-Box’ Methods

For the social work package, led by the sociology PhD candidate, doing transdisciplinary research into AI in an under-resourced yet over-researched neighbourhood meant our methods had to out-of-the-box. As mentioned, we joined trash-cleaning efforts to

⁴ The Dutch police has a voluntary registration system for privately-held cameras called *Camera in Beeld* (Camera in Vision).

ensure our presence had immediate, tangible value, and acted as intermediaries to connect residents with institutions. Beyond this, what truly proved effective was employing more innovative and stimulating methods that would make the community's research experience joyful, 'analytical, proactive, and empowering' [37, p. 179]. To this end, we organized community AI workshops with residents, institutional stakeholders (such as from the municipality and housing corporations) and high schools students, using photos and generative AI to visualize an ideal Lombardijen.

We drew inspiration from participatory action research (PAR) [38, 39] for our methods. PAR, at its core, prioritizes 'knowledge held both within communities among the traditionally oppressed or silenced voices as well as within more commonly recognized institutions' [40, p. 361]. PAR aims to support participants to actively engage in the research process, foster discussion and reflection, co-create knowledge, and amplify the voices of those directly affected [41]. One method we used was photovoice: a PAR-method developed in 1990s China, where female farmworkers documented their health and socioeconomic challenges through photography to influence policy. Participants are invited to capture photos that reflect their experiences or community issues, giving visual narrative to their collective perspectives. Photovoice unfolds in three stages: (1) documenting lived experiences through photography; (2) facilitating reflective discussions on the photos participants produce; and (3) promoting social change, e.g., through engagement with community leaders and policymakers [42].

As it unfolded, photos became the foundation for the community AI workshops. We developed a 'Photovoice-to-AI'-method, inviting participants to submit photos reflecting positive and negative aspects of public safety in the neighbourhood prior to the workshops, which we then used to spark discussion and inspire AI-generated visualizations during the sessions. Visual language proved a welcome alternative to text and speech-oriented methods, which implicitly demand that less-privileged neighbourhoods meet the verbal and written articulation standards of more affluent areas as well as the city's policymakers. Additionally, visual language proved valuable, because having tangibles 'proof' (be it photos or AI-visualizations) gave residents a sense of protection against their complaints constantly being waved away by institutions.

Most importantly, the workshops were designed not as mere 'sites for data collection' but as social community events. We aimed to acquaint high school students with neighbourhood residents and vice versa, provide a fun activity with food and drinks at the community centre, and cultivate digital literacy throughout the process. Interestingly, while the workshops were intended to build digital literacy, an unexpected outcome emerged: a form of digital critical thinking. Participants developed a more nuanced understanding of generative AI through play; they learned to recognize hallucinations in AI-generated content and observed how poorly AI visualizations reflected local conditions in Lombardijen, opening discussions on bias in datasets. What's more, it emerged through laughter and joy. It was a welcome surprise that most residents took genuine pride in their contributions, photographing the AI-visualizations they had helped create to show at home. At the same time, we had a positive experience using generative AI but acknowledge its limitations. We reflect further on its use in research activities aimed at engaging (societal) stakeholders in the next section and discussion.

4.4 Technological Takeaway: Critically Explore AI

Technology is not explicitly captured in the ELSA abbreviation, but we believe it is important to assess it as a stand-alone perspective. The work package on technological aspects, led by the computer science PhD candidate, develops a conversational agent and brings technical insights to discussions on public safety and AI. In this final section, we reflect on our use of generative AI in the workshops described above. Specifically, we made use of ChatGPT 4.0's image-generation function to visualize an ideal Lombardijen. Working with tools like ChatGPT enables researchers to evaluate their usability for specific tasks, but we must also openly and ethically weigh it against its harms.

On the one hand, our experiences show that generative AI can allow stakeholders to imagine different possibilities with respect to their future. Generative AI can serve as thought partner that assists humans in collaborative creation and ideation, and enables participants to co-design how they envision the future of their neighbourhood [43]. Image-generation specifically allows for rapid prototyping, enabling new forms of human creativity [44]; also, its use requires little to no knowledge of the underlying technology, thereby accessible to non-technical users [45]. As such, generative AI tools can provide valuable co-design opportunities [46].

On the other hand, there are specific risks. Amongst many concerns, generative AI relies on underpaid labour in the Global South, is immature in its development, prone to misuse, infringes upon (data) privacy, can facilitate job loss, and is environmentally harmful [47, 48, 49, 50]. Recognizing these risks, we still wanted to explore whether and how generative AI could be used in ethical manner. While eventually ChatGPT became the model of choice, we explored other (open source) models. Yet, there were limitations to choosing other models due to these being less user-friendly in co-design sessions and the lack of Dutch language performance. Both of these issues presented problems for our case study with mostly Dutch-speaking community members. As such, we used ChatGPT, but not without reservation.

One important limitation we observed is (cultural) bias [48]. Many images provided by ChatGPT of a future neighbourhood were mostly reminiscent of areas in the United States, reflective of the data these models are trained on, feeling unfamiliar to the participants in our workshops. Yet, such issues did lead to valuable discussions on the impact of training data on the task at hand; we could explain to participants why it is difficult to create houses that are reflective of their neighbourhoods with generative AI.

Another key limitation is the impact of generative AI on the environment. Residents in the neighbourhood repeatedly stated how important the local environment (like greenery) was for them. As such, it seemed clear that the environmental impact of generative AI would also be important for them to know. Specifics are not yet known on the environmental cost per prompt generation, though there is literature on the impact of training generative AI models as well as that of operating such models, with especially image-generation being a concern [49], as the most energy intensive task when it comes to generative AI [50].

During the workshops, participants prompted ChatGPT to estimate the CO₂ emissions generated by our conversations and visualizations. While it is unlikely that they received a fully correct answer, it drew their attention towards the impact of AI on the

environment regardless. For instance, one resident asked, looking at the laptop and expecting some gas to come out, ‘where [then] is the carbon emission?’, stimulating researchers to explain that the environmental toll comes amongst others from the water needed to cool the data servers that host generative AI models (and that “the cloud” is not so ethereal after all).

Bearing these considerations in mind, the use of generative AI, even with noble goals like citizen engagement, remains ethically and practically fraught. The environmental costs are glaring; and, relinquishing too much of our creativity and imagination to generative AI, with inherent biases and ongoing growing pains, risks disengaging instead of engaging the contextual knowledge of residents. At the very least, each use of generative AI by researchers such as ourselves should be seriously evaluated as to its merits, justifiability and alternatives, which we reflect on more in the discussion.

5 Discussion

5.1 What Worked Well

Having offered ethical, legal, social and technological perspectives on the case study in Lombardijen, we return to the research question: *How can citizens in low-trust neighbourhoods be meaningfully and reciprocally engaged in transdisciplinary AI research, and what does an ELSA approach offer in this regard?* In the discussion, we identify what worked well, and focus even more on our learnings in the section hereafter. We conclude with the need for situated responsibility in transdisciplinary AI research.

The value of ELSA as research methodology lies in its anticipatory and transdisciplinary character. ELSA supports the proactive identification and mitigation of ethical, legal, and social aspects before technologies are fully implemented. This, in combination with engaging diverse stakeholders, can help to avoid the introduction of AI ‘solutions’ that are technically sound, but socially irrelevant or even outright harmful in the community [23]. Engaging citizens fosters a bottom-up dynamic, starting from neighbourhood issues and exploring whether (AI-)technologies can address them, rather than implementing these technologies and then looking for problems to justify their use.

Moreover, being inherently transdisciplinary, ELSA stimulates early-on participation of different types of stakeholders in research, encouraging them to co-shape research questions. It comes naturally in an ELSA context to reflect on how to engage stakeholders meaningfully and reciprocally, since mutual learning is the goal of these interactions; researchers should always ‘give back’ by bringing different stakeholders to the same table – connecting people who might otherwise not meet – and by offering critical allyship with the issues and questions stakeholders raise.

Once in Lombardijen, we learned the importance of ‘giving back from the outset’ by participating in community initiatives, such as the trash-cleaning group and the neighbourhood watch group; not only to establish trust but also to reciprocate the time and efforts residents spent on participating in research activities. We would encourage other researchers focusing on neighbourhood communities to adopt a similar approach. In

addition, the workshops taught us that visual language, through photos and AI visualizations, is valuable; it offered an alternative way for residents to voice their perspectives, distinct from text and speech-based forms of knowledge production.

5.2 What Could Be Different

At the same time, our work is open to critique, a key one being the use of ChatGPT, given the limitations and harms of generative AI, as highlighted in the reflections from technological perspective. Whether these tools should be used in research activities (aimed at citizen engagement), cannot be answered in absolute terms; it should be assessed in context. In our case, we decided in favour of using ChatGPT, because it served as a concrete conversation starter about AI. Considering its extensive media coverage, the first question we often received from stakeholders when introducing our research on AI and public safety was: ‘So, like ChatGPT?’. Building on this curiosity and familiarity allowed us to discuss broader AI issues, such as bias in datasets and the presence of surveillance technologies like (AI-powered) cameras in the neighbourhood.

Yet, our use of ChatGPT reveals a paradox: while it proved effective in engaging the community, it may also illustrate precisely the risk that residents themselves express: that technologies will enter their communities no matter how they may feel about it. Our reasoning was that ChatGPT was already embedded in residents’ worlds: even if not everyone used it directly, most were familiar with it through news coverage or through relatives using it. We hoped that offering residents more hands-on familiarity would positively contribute to their reflections on AI.

Moreover, ChatGPT’s image-generation function proved particularly valuable, but precisely this function is resource-intensive and environmentally harmful [50] – illustrating that AI, in itself, is an extractive industry [8]. Can community-oriented research aiming to move beyond extractivist dynamics legitimately rely on extractivist tools to engage the community? We used the tools, but we do not have a definitive answer to this question. It was valuable to witness the benefits of our approach, yet if we were to do it again, we would re-evaluate alternatives to generative AI, especially image-generation. Our positive experience with visual language could also be achieved through other means; in this case study, we used photos as an inspiration for AI visualizations, but in future research, we would reconsider photos as an inspiration for arts-and-crafts.

Another key critique relates to benefits for the community for participating in our research. We recall one resident’s phone call that left a deep impression; while they acknowledged our good intentions, our co-creative research design, and even our efforts to educate about AI, they remained firmly convinced that the outcome would simply be ‘another report that ends up in a drawer at city hall’, leading to no effective change. They suggested that the funds allocated for our salaries would be far more effectively spent directly on neighbourhood solutions. Although ELSA proved valuable for engaging communities in public deliberation on AI, whether it can achieve (policy) change remains to be seen. Some of our proposed (quite low-hanging fruit) public safety interventions – aligning with the action-ethos of PAR – were met with apprehension over (financial) feasibility by residents and municipal officials alike.

This left us with mixed emotions: we are glad to have made the effort, yet we may have fallen short of our aim to move beyond extractivist dynamics in some respects. The resident's critique is justified in the sense that no effective change came from our research. We offered tangible benefits in other ways, such as volunteering, to reciprocate residents' time and efforts spent on research activities with us. Yet, we could have provided financial compensation. Considering the risks of research exploitation and the slow pace of (policy) change, financial compensation is a justifiable baseline condition for engaging citizens in research. Researchers are paid for their time; other public and private stakeholders typically participate in research during work time. In contrast, citizens contribute their private time. As one resident spoke to us: 'I am leaving now, because I have more things to do; you are paid for this time, but I am not.'

Moreover, it is often in the nature of research to leave once sufficient insights have been gained, which conflicts with the goal of building durable, meaningful relationships with stakeholders. In a way, researchers 'pull out' of the communities once they obtained what they came for. Even our strategies of 'giving back' can be seen as transactional rather than reciprocal, as participation in neighbourhood initiatives also served our research purposes. This can create the impression that communities receive support only when researchers have a need. We sought to mitigate these concerns by continuously reflecting on the power dynamics between us and participants, demonstrating genuine interest and allyship, and being transparent about our intentions and potential outcomes. We also announced our departure through a closing workshop, during which we introduced residents to new researchers continuing our work, focusing on greenery in the neighbourhood.

We also remained intermittently engaged for months after the closing workshop, connecting with current and upcoming research in Lombardijen, giving them our insights to avoid the same things being asked or studied again. An association of researchers in the area seems to have formed, exchanging knowledge and methodology, which appears to be a step in the right direction. We will ourselves be back once more in the spring, to spread 'seed bombs' in the neighbourhood: a subtly rebellious way to give heed to residents' requests for more wildflowers. Nevertheless, in our role as mediators between residents and other stakeholders, such as municipal officials and housing corporations, it is likely that some of the connections we facilitated may not endure after our departure, despite the handover to the new researchers.

5.3 Towards Situated Responsibility

What becomes evident, is that in transdisciplinary (AI) research with citizen stakeholders from low-trust neighbourhoods, we must continually ask ourselves: when does our presence and data collection become another form of extractivism? And, on the other hand, when does our absence contribute to further marginalization – when a neighbourhood, because of its history or challenges, simply drops off the radar? Navigating this tension requires an attitude of flexibility and reflection, what we came to think of as *situated responsibility*. Our goal was not simply to avoid harm, but to actively co-create value – however modest – in a way that respected the neighbourhood's wishes.

In this regard, critical theory, feminist, and Indigenous approaches offer valuable lessons for transdisciplinary AI research, having long highlighted the position of both researcher and researched. Researchers are never neutral outsiders but are always implicated in the contexts they study; they must critically reflect on the power dynamics their presence produces, anticipate harms they could induce in researched communities, and consider how participants' time and efforts are acknowledged, valued and reciprocated within the research process. As Smith eloquently puts it:

'Research in itself is a powerful intervention, even if carried out at a distance, which has traditionally benefited the researcher, and the knowledge base of the dominant group in society. When undertaking research, either across cultures or within a minority culture, it is critical that researchers recognize the power dynamic that is embedded in the relationship with their subjects. Researchers are in receipt of privileged information. They may interpret it within an overt theoretical framework, but also in terms of a covert ideological framework. They have the power to distort, to make invisible, to overlook and to draw conclusions, based not on factual data, but on assumptions, hidden value judgements, and often downright misunderstandings. They have the potential to extend knowledge or to perpetuate ignorance' [2, p. 229].

6 Conclusion

This paper offered a methodological reflection on ELSA: a transdisciplinary research methodology for AI that assesses its ethical, legal and social aspects. We based this reflection on a case study in Lombardijen, about public safety and (AI-)technologies in the neighbourhood. In a community marked by research fatigue and long-standing institutional neglect, our efforts sought to move beyond extractive research practices, by fostering more reciprocal and accountable forms of (citizen) stakeholder engagement. The paper detailed our experiences with citizen participation in Lombardijen from ethical, legal, social and technological perspectives. Ultimately, these perspectives are interconnected and cannot be meaningfully understood without the contextual knowledge of residents. AI and related digital technologies shape the neighbourhood in tangible ways, and for these technologies to align with the realities of the communities affected by them, residents' input is essential.

At the same time, the case study revealed the complex reality and limitations of transdisciplinary research. This underscores the importance of recognizing researchers not as neutral facilitators, but as stakeholders whose presence shapes the dynamics of engagement, interpretation, and knowledge production. Any effort to foster public deliberation on AI must also reckon with the uneven histories and structures that frame public engagement. Rather than presenting ELSA as a finished or rigid framework, we view it as an evolving path and a reflexive practice. It can help shift research practices toward more just, meaningful and reciprocal ways of engaging with AI in places where trust has been fractured. In a way, then, it is not a bug but a feature of the ELSA approach that we reflect on the shortcomings of our engagement with Lombardijen to learn to do better at the next opportunity.

References

1. Chagnon, C.W., Durante, F., Gills, B.K., Hagolani-Albov, S.E., Hokkanen, S., Kangasluoma, S.M.J., Konttinen, H., Kröger, M., LaFleur, W., Ollinaho, O., Vuola, M.P.S.: From extractivism to global extractivism: The evolution of an organizing concept. *The Journal of Peasant Studies* 49(4), 760–792 (2022). Doi: <https://doi.org/10.1080/03066150.2022.2069015>
2. Smith, L.T.: *Decolonizing methodologies: Research and indigenous peoples*, 3rd edn. Bloomsbury Academic, London (2022)
3. Caldwell, W.B., Reyes, A.G., Rowe, Z., Weinert, J., Israel, B.A.: Community Partner Perspectives on Benefits, Challenges, Facilitating Factors, and Lessons Learned from Community-Based Participatory Research Partnerships in Detroit. *Progress in Community Health Partnerships: Research, Education, and Action* 9(2), 299–311 (2015) Doi: 10.1353/cpr.2015.0031
4. Westerhof, S.: Onderzoek onthult grote verschillen in bewonersinitiatieven op Rotterdam-Zuid, <https://www.eur.nl/nieuws/onderzoek-onthult-grote-verschillen-bewonersinitiatieven-op-rotterdam-zuid> (2023)
5. Volkskrant: In deze Rotterdamse wijk zijn bewoners nu wel genoeg onderzocht: ‘Ze denken dat ze hier even data kunnen komen halen’, <https://www.volkskrant.nl/wetenschap/in-deze-rotterdamse-wijk-zijn-bewoners-nu-wel-genough-onderzocht-ze-denken-dat-ze-hier-even-data-kunnen-komen-halen> (2023)
6. Sukarieh, M., Tannock, S.: On the Problem of Over-researched Communities: The Case of the Shatila Palestinian Refugee Camp in Lebanon. *Sociology* 47(3), 494–508 (2013). Doi: <https://doi.org/10.1177/0038038512448567>
7. Ashley, F.: Accounting for research fatigue in research ethics. *Bioethics* 35(3), 270–276 (2021). Doi: 10.1111/bioe.12829
8. Crawford, K.: *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press, London (2021)
9. Posada, G.: *The Coloniality of Data Work: Power and Inequality in Outsourced Data Production for Machine Learning*. University of Toronto, Toronto] (2022)
10. Kruizinga, M.: *Inspraak bij technologische innovaties: De rol van alle stakeholders*, <https://www.eur.nl/nieuws/inspraak-bij-technologische-innovaties-de-rol-van-alle-stakeholders> (2024)
11. Cunningham, A., Williams, P.: De-centring the ‘big picture’: The Origins of Modern Science and the modern origins of science. *The British Journal for the History of Science* 26(4), 407–432 (1993). DOI: 10.1017/S0007087400031447
12. Zwart, H., Landeweerd, L., Van Rooij, A.: Adapt or perish? Assessing the recent shift in the European research funding arena from ‘ELSA’ to ‘RRI.’ *Life Sciences, Society and Policy* 10(1), 11 (2014). DOI: 10.1186/s40504-014-0011-x
13. Zwart, H.: The bioethics of synthetic cells. *Nature Reviews Molecular Cell Biology* 25(3), 157–158 (2024). DOI: 10.1038/s41580-023-00693-w
14. Ryan, M., Blok, V.: Stop re-inventing the wheel: Or how ELSA and RRI can align. *Journal of Responsible Innovation* 10(1), 2196151 (2023). DOI: 10.1080/23299460.2023.2196151
15. Zwart, H., Nelis, A.: What is ELSA genomics? *EUROPEAN MOLECULAR BIOLOGY ORGANIZATION* 10(6), 540–544 (2009).
16. Rodríguez, H., Fisher, E., Schuurbiens, D.: Integrating science and society in European Framework Programmes: Trends in project-level solicitations. *Research Policy* 42(5), 1126–1137 (2013). DOI: 10.1016/j.respol.2013.02.006

17. van Veenstra, A.F., van Zoonen, L., Helberger, N.: Elsa labs: For Human Centric Innovation in AI. NL AI Coalitie, https://pure.eur.nl/ws/portalfiles/portal/44457271/Position_paper_ELSA_Labs_July_2021.pdf (2021)
18. Juhasz, C., Gianni, R., Kamalski, J., Meacham, D.: Navigating many voices: Lessons for ELSA/RRI from the pragmatist perspective of ELSA-by-design. *Journal of Responsible Innovation* **12**(1), 2445308 (2025). DOI: 10.1080/23299460.2024.2445308
19. Knol, J.: Reactie advies WARP – Navigeren in niemandsland, <https://www.politie.nl/informatie/advies-wetenschappelijke-adviesraad-politie---navigeren-in-niemandsland.html> (2025)
20. Wetenschappelijke Adviesraad Politie: Navigeren in Niemandsland. WARP, <https://www.politie.nl/informatie/advies-wetenschappelijke-adviesraad-politie---navigeren-in-niemandsland.html> (2025)
21. Frodeman, R., Klein, J.T., Pacheco, R.C.D.S. (Eds.): *The Oxford handbook of interdisciplinarity*. Oxford University Press, London (2017)
22. Jacobs, G., Van Houdt, F., Coons, G.: Studying Surveillance AI-cologies in Public Safety: How AI Is in the World and the World in AI. *Surveillance & Society* **22**(2), 160–178 (2024). DOI: 10.24908/ss.v22i2.16104
23. KIAST: Ethical, Legal and Societal Aspects (ELSA), <https://www.kia-st.nl/kiast2024-2027/2-kader-omgeving-en-positionering/ethical-legal-and-societal-aspects-elsa-1> (2024)
24. Weck, S., Schmitt, P.: Approaching spatial justice in local development actions. In: Fritsch, M., Kahila, P., Németh, S., Scott, J.W. (eds.) *Spatial Justice and Cohesion*, 1st edn., pp. 49–71. Routledge, [missing city] (2023). DOI: 978-1-00-322968-1
25. Estada, A.: Anderhalf keer meer Rotterdammers door toeslagenschandaal getroffen dan eerder bekend. *Vers Beton*, <https://www.versbeton.nl/2022/12/anderhalf-keer-zoveel-rotterdammers-door-toeslagenschandaal-getroffen-als-eerder-bekend/> (2022)
26. Gemeente Rotterdam: Toeslagen 010, <https://rotterdamraad.bestuurlijkeinformatie.nl/Document/View/b927979b-cc68-4d5c-a5c6-7b17517675fe> (2023)
27. Dol, K., Hoekstra, J., Kleinhans, R.: Case study report ‘Rotterdam South on Course’ Resituating the Local in Cohesion and Territorial Development. Delft University of Technology, Delft (2019)
28. Gemeente Rotterdam: Lombardijen Wijkprofiel, <https://wijkprofiel.rotterdam.nl/nl/2024/rotterdam/ijssemonde/lombardijen> (2024)
29. Meijer, R., de Bruin, H., Schram, K.: *Orde op zaken stellen. Een onderzoek naar veiligheidsbeleving in Lombardijen*. Inholland Rotterdam (2024)
30. Schuilenburg, M., van Steden, R.: Positive security: A theoretical framework. In: *Positive Criminology: Reflections on Care, Belonging and Security*. Eleven, Amsterdam (2014)
31. Suojanen, I., Bayerl, P.S., Jacobs, G.: Citizens' positive safety perceptions in public spaces. In: *Proceedings of the 29th European Safety and Reliability Conference*, pp. 224–231. Research Publishing Services, London (2019)
32. Gemeente Rotterdam: Afdoeningsvoorstel burgemeester toezegging 23bb00330 over ontwikkelingen cameratoezicht 2023, <https://rotterdamraad.bestuurlijkeinformatie.nl/Reports/Item/881c9a11-5203-4175-aa7a-4fc261193895> (2023)
33. EM Magazine: The Rotterdam delta as a living lab for Erasmus and Delft University, <https://www.erasmusmagazine.nl/en/2020/03/13/the-rotterdam-delta-as-a-living-lab-for-erasmus-and-delft-university/> (2020)
34. Trottier, D., Lee, J.-S., Boy, J.: Urban Data Analytics as Research Topic, Method and Ethical Concern. *Digital Culture & Society* **7**(2), 311–328 (2021). DOI: 10.14361/dcs-2021-070215
35. Bovens, M.: Analysing and Assessing Accountability: A Conceptual Framework1. *European Law Journal* **13**(4), 447–468 (2007). DOI: 10.1111/j.1468-0386.2007.00378.x
36. Poel, I. van de: *Moral responsibility and the problem of many hands*, 1st edn. Routledge, New York (2015)

37. Wang, C., Burris, M.A.: Empowerment through Photo Novella: Portraits of Participation. *Health Education Quarterly* 21(2), 171–186 (1994). DOI: 10.1177/109019819402100204
38. Liebenberg, L.: Thinking Critically About Photovoice: Achieving Empowerment and Social Change. *International Journal of Qualitative Methods* 17(1), 1609406918757631 (2018). DOI: 10.1177/1609406918757631
39. McTaggart, R.: *Participatory Action Research: International Contexts and Consequences*. SUNY PRESS, Albany (1997)
40. Derr, V., Simons, J.: A review of photovoice applications in environment, sustainability, and conservation contexts: Is the method maintaining its emancipatory intents? *Environmental Education Research* 26(3), 359–380 (2020). DOI: 10.1080/13504622.2019.1693511
41. Aboulkacem, A.R., Aboulkacem, S., Haas, L.E.: Photovoice 2.0: A Comprehensive Research Framework for the Digital Generation. *TechTrends* 65(5), 874–883 (2021). DOI: 10.1007/s11528-021-00642-4
42. Wang, C., Burris, M.A.: Photovoice: Concept, Methodology, and Use for Participatory Needs Assessment. *Health Education & Behavior* 24(3), 369–387 (1997). DOI: 10.1177/109019819702400309
43. Collins, K.M., Sucholutsky, I., Bhatt, U.: Building machines that learn and think with people. *Nat Hum Behav* 8, 1851–1863 (2024). <https://doi.org/10.1038/s41562-024-01991-9>
44. Gafni, O., Polyak, A., Ashual, O., Sheynin, S., Parikh, D., Taigman, Y.: Make-A-Scene: Scene-Based Text-to-Image Generation with Human Priors. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds) *Computer Vision – ECCV 2022*. ECCV 2022. Lecture Notes in Computer Science, vol 13675. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-19784-0_6
45. Vartiainen, H., & Tedre, M.: Using artificial intelligence in craft education: crafting with text-to-image generative models. *Digital Creativity*, 34(1), 1–21 (2023). <https://doi.org/10.1080/14626268.2023.2174557>
46. Knops, R., Şerban, B., & Houben, S.: Human-human collaboration enhanced with emerging technologies of AI. In: CHI'23: Workshop on Integrating AI in Human-Human Collaborative Ideation (in press)
47. Narayanan, A., & Kapoor, S.: *AI snake oil: What artificial intelligence can do, what it can't, and how to tell the difference*. Princeton University Press (2024)
48. Bird, C., Ungless, E., & Kasirzadeh, A.: Typology of risks of generative text-to-image models. In: *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 396-410 (2023)
49. Weidinger, L., Uesato, J., Rauh, M., Griffin, C., Huang, P. S., Mellor, J., Glaese, A., Cheng, M., Balle, B., Kasirzadeh, A., Biles, C., Brown, S., Kenton, Z., Hawkins, W., Stepleton, T., Birhane, A., Hendricks, L.A., Rimell, L., Isaac, W., Haas, J., Legassick, S., Irving, G., Gabriel, I.: Taxonomy of risks posed by language models. In: *Proceedings of the 2022 ACM conference on fairness, accountability, and transparency*, pp. 214-229 (2022)
50. Luccioni, S., Jernite, Y., & Strubell, E.: Power hungry processing: Watts driving the cost of AI deployment?. In: *Proceedings of the 2024 ACM conference on fairness, accountability, and transparency*, pp. 85-99 (2024)