

Ethically Aligned Stakeholder Elicitation (EASE): Case Study in Music-AI

Anna-Kaisa Kaila
KTH Royal Institute of
Technology
Stockholm, Sweden
akkaila@kth.se

Petra Jääskeläinen
KTH Royal Institute of
Technology
Stockholm, Sweden
ppja@kth.se

Andre Holzapfel
KTH Royal Institute of
Technology
Stockholm, Sweden
holzap@kth.se

ABSTRACT

Engineering communities that feed the current proliferation of artificial intelligence (AI) have historically been slow to recognise the spectrum of societal impacts of their work. Frequent controversies around AI applications in creative domains demonstrate insufficient consideration of ethical predicaments, but the abstract principles of current AI and data ethics documents provide little practical guidance. Pragmatic methods are urgently needed to support developers in ethical reflection of their work on creative-AI tools.

In the wider context of value sensitive, people-oriented design, we present an analytical method that implements an ethically informed and power-sensitive stakeholder identification and mapping: *Ethically Aligned Stakeholder Elicitation* (EASE). As a case study, we test our method in workshops with six research groups that develop AI in musical contexts. Our results demonstrate that EASE supports critical self-reflection of the research and outreach practices among developers, discloses power relations and value tensions in the development processes, and foregrounds opportunities for stakeholder engagement. This can guide developers and the wider NIME community towards ethically aligned research and development of creative-AI.

Author Keywords

computational creativity, music, ethics, stakeholder, Value Sensitive Design

CCS Concepts

•Human-centered computing → HCI theory, concepts and models; •Social and professional topics → Socio-technical systems; •Applied computing → *Sound and music computing*;

1. INTRODUCTION

Developing responsible AI is a key concern of current machine learning research and data-driven application development. As evidenced by recent controversies, such as litigation

over the data use legitimacy of creative-AI applications [7], or the dilution of artist portfolios through generative pastiches [29], considerations of ethical predicaments and values of development work for data analysis and generation in creative domains often lag behind. In the context of music, several authors have called for increased sensitivity to the social and ethical frameworks in which the applications are being developed (*e.g.* [4, 5, 10, 28, 40, 48]). As Friedman and Nissenbaum were early to warn [24], a lack of diversity and richness of context in the development phase can have serious and widespread consequences in society.

Several high-level principle documents for AI and data ethics have been published in recent years [33], but there is a scarcity of easily-accessible methods directly applicable to the practical needs of AI developers generally [38], and in the area of creative-AI specifically. As argued by Gold *et al.* [25], research communities formed around central conferences of a field, such as NIME, can play a critical role in spearheading the increase of ethical awareness in the industry at large.

We suggest that the exploration of the ethical predicaments in creative-AI development should focus on the real-life implications of individual research projects on actual people rather than on abstract value statements. Furthermore, we argue that such ethical analysis should be sensitive to value tensions and power relations, in order to avoid perpetuating discriminatory structures.

This paper contributes to these purposes in two ways: 1) by proposing a new method, *Ethically Aligned Stakeholder Elicitation* (EASE), that guides developers of creative-AI through a structured, ethically informed and power-sensitive stakeholder analysis, and 2) by advancing a deeper collective understanding of the ethical predicaments of creative-AI development, which will be necessary to inform prospective policy guidelines, ethics standards and analysis methods tailored for the specific needs of the creative-AI sector. With this work, we are taking a step towards helping developers of creative-AI to improve sensitivity towards stakeholder diversity, critically examine the power asymmetries embedded in their work, and align their operational decisions with a framework of ethical principles. Furthermore, we hope the method will inspire engineering and design communities to extend our work on engaging and actionable methods for ethical analysis of creative-AI.

As an initial case study, we evaluated the current version of EASE in a series of workshops with research teams that develop AI in musical contexts (music-AI). The outcomes of the workshops demonstrate that introducing EASE into research projects has the potential to (a) give nuance to critical self-reflections over the social impacts and ethical predicaments of individual research projects, (b) broaden perspectives on the range of impacted stakeholders and their mutual power relations and (c) promote pro-active efforts



Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). Copyright remains with the author(s).

of engagement and empowerment to alleviate value tensions between stakeholder groups.

In the following, we introduce value sensitive design (VSD) as the general context and people-focused care ethics as the specific theoretical framework for our work. We then operationalise this framework through an engaging, low-barrier stakeholder analysis method, and discuss observations gathered from testing EASE in workshops with developers of music-AI. Finally, we connect these insights with reflections on opportunities for further, ethically informed research initiatives on creative-AI applications.

2. BACKGROUND

2.1 Call for ethical creative-AI

Issues of AI and data ethics have been prominently discussed in domains of grave discriminatory or life-threatening implications, such as jurisdiction, policing, or credit allocation [3, 45]. In the domain of creative-AI, and in specific in the context of music, these matters have only recently started to surface in the academic debate (see for instance [2, 9, 18, 26, 31, 32, 39, 55]).

This is, however, no indication that creative-AI would be less susceptible to ethical predicaments. Controversies around training data legitimacy, the emergence of new kinds of style thefts and deep fakes, and collateral damage from biased, fragmentary or aesthetically narrow datasets are just a few examples of issues that may arise in the wake of unreflective deployment of creative-AI technologies [19]. On a wider systemic scale, the proliferation of AI in the cultural and creative sectors is expected to disrupt value chains and job markets, strengthen the existing information monopolies and inequalities, and increase the general ecological impact of the industry [16]. Given the current rush of capital investment in commercial AI applications in domains such as music [13], the conditions of ethically aligned AI must be made more explicit and accessible for developers.

2.2 Scarcity of applied AI Ethics

Along with the increased attention to AI technologies, a number of ethical guidelines for general AI use and development have been established for instance by [17, 57, 59]. In the context of music applications, NIME Principles & Code of Practice on Ethical Research¹ represent pioneering efforts towards formalising ethical frameworks of research and development work, but challenges such as data use in music-AI contexts beyond GDPR-defined personal data are not directly addressed in this code of practice.

In contrast to the multitude of generic, principle-based ethics frameworks, there are much fewer applied analysis methods intuitively accessible even with limited prior expertise in ethics [41]. Some of the existing approaches introduce an extensive series of questions or a range of discussion points as prompts for ethical self-reflection and analysis, *e.g.* [12, 15, 46, 62]. Others translate ethics prompts into graphical or game-like platforms, *e.g.* [8, 20, 60], or focus on a set of specifically chosen principles [47]. Unfortunately, many of the existing tools can be rather overwhelming for less experienced participants to implement. For instance, EC-COLA [60] includes almost a hundred topics to address, and other tools add up to similar figures [12, 15].

Some tools developed in VSD provide routines for stakeholder mapping and analysis, such as envisioning cards [22], stakeholder tokens [65], and interviews with stakeholder

perspectives [50, 56]. These methods are not specifically focused on ethics, but they indicate valuable directions for connecting value awareness with actionable ethical analysis of research projects. Furthermore, none of these tools targets the particular needs of creative-AI applications development, and the impact of its data use on human subjects. In the following, we will introduce an approach that seeks to address these issues and test it in the particular context of creative-AI development.

2.3 Operationalising care ethics

In contrast to the selection of practical ethics methods outlined in the previous section, our EASE method combines two perspectives in a form that – to the best of our knowledge – has not been previously done: the ethics perspectives relevant specifically for data-driven development processes, and the focus on a specific set of ethical issues that directly relate to interactions among *people* and their power relations. The purpose of this focus is to link the analytic exploration with concrete human subjects and real-life situations, and hence provide an intuitive and low-barrier method that does not require an extensive time commitment or previous insights into ethics research from the users.

2.3.1 VSD, Care ethics and the focus on people

Our focus on the *people* of AI development is closely linked with the value sensitive design (VSD) framework and especially the recent explorations in care ethics. VSD is a socio-technical design approach that accounts for values among a diversity of user groups [23]. Care ethics has served several HCI projects as a frame to explore how people in communities relate to one another and to technology [1, 34, 35, 54, 58]. As defined by Held [30, p.10], the epicentre of care ethics is the "moral salience of attending to and meeting the needs of the particular others for whom we take responsibility". Such remarks resonate with Vallor's [61] perspective on the relational understanding of moral obligations as the condition for the development of a virtuous character, and with calls for extended engagement with under-represented populations in the tech policy [66]. Care ethics therefore not only aims at a personal virtue but at the wider questions of the general conditions and values of a community. In particular, care ethics responds to Friedman's call for HCI research to bring in more accounting of power and value tensions in design processes, as well as for developing methods that focus on such goals [23, p.175f].

2.3.2 Stakeholder theory and power analysis

To operationalise the perspectives of care ethics and the focus on power, we integrate an ethically informed stakeholder analysis with a power/interest grid [14], which is one of the constitutive applications of stakeholder theory [21]. The term 'stakeholder' refers here not just to direct users or the financial beneficiaries, but more broadly to an emergent range of individuals and groups that can impact, or are directly and indirectly impacted by the project [42, 49].

While the concept of power has many competing definitions (see [63], or more recently *e.g.* [52]), we have chosen to refer to the practical orientation applied in previous versions of power/interest grids, which define power as "the mechanism through which stakeholders influence the direction and decisions for a project. This power can be used to retain the status quo or to enforce fundamental change" [43, p.842]. Regarding the common distinction between *power-to* and *power-over* [52], this power is a *power-to*, *i.e.* an ability to

¹<https://www.nime.org/ethics/>



Figure 1: General layout of the stakeholder identification. Questions (Q1–Q9) are specified in the main text.

do something, whereas the relational aspect (*power-over*) will be emergent from the present mapping process. Interest is respectively defined here as “[m]easure of how interested each stakeholder group is to impress its expectations on the project decisions” [44, p.322], which can be seen as the gauge of the likelihood for the stakeholder to use the power they have.

3. EASE METHOD

EASE was designed to be run in a workshop format². The format emerged from a series of pilot studies and iterations conducted with researchers at our institution, at a conference workshop, and in the context of doctoral education. While the method in itself may be applicable to data-driven development projects beyond creative-AI, we carried out the first evaluations of it with project teams that develop music-AI applications.

The workshop centres around an actual application development project that the participants agree upon and briefly discuss, shaping a scenario that is familiar to all participants. The workshop process then runs in two steps described in the following subsections: ethically informed stakeholder identification and power-mapping.

3.1 Part 1: Stakeholder identification

In the first exercise of the workshop, participants identify stakeholders related to and impacted by the specified scenario. This process is facilitated by a list of nine question prompts that guide the participants to extend their consideration of stakeholder groups beyond the core development team and immediate collaborators:

1. Who provides data for the project?
2. Who is the data about?
3. Whose permissions do you need to access and use the data?
4. Is someone’s personal or other sensitive data included (whose)?

²The design of the workshop boards can be viewed at <https://doi.org/10.5281/zenodo.7799160>

5. Who do you envision will use your project or its end product?
6. Who sets the rules, regulations and recommendations on your data use?
7. With whom will you communicate about and share information or data on your project?
8. Who could be inadvertently harmed by the project, its data use or its end result? (Additional prompt: “For example: Could it [your project] expose, profile, discriminate or prejudice individuals? Could undetected bias, inclusion/exclusion, data gaps or other limitations cause harm or disadvantage to someone? Could the project and its outcomes unfairly restrict someone’s access or opportunity, or replace another product or service?”)
9. Who does not use your product but could still be positively impacted by it? (Additional prompt: “Are you making things better for the society? For whom (individuals, groups, demographics or organisations?”)

The layout of the first exercise includes a circular board surrounded by these nine question prompts (see Figure 1). Participants conduct the exercise by writing on individual cards a stakeholder or a stakeholder group they find relevant with respect to their project, and then placing the cards on the circular board near the respective question prompt.

The questions are informed by Data Ethics Canvas [46], but have been reformulated to “who”-questions to help participants elicit perspectives related specifically to people, in line with the care ethics framework. Furthermore, some of the questions have been combined and abridged from the original format³ to enable participants to run through the whole process in a reasonable time frame. The last two questions, which are more open in their formulations, have short additional prompts for support.

3.2 Part 2: Stakeholder power-mapping

In the second exercise of the workshop, the previously identified stakeholders are evaluated in terms of power and interest. The definitions of these concepts (as provided in Section 2.3.2) are printed on the board. Stakeholder cards from the first exercise are collected, consolidated in case of overlaps or repetitions, and placed on a 2x2 grid in which interest is on the y-axis and power on the x-axis (see Figure 2). Participants are instructed to locate the cards such that the placement reflects stakeholders’ respective power and interest towards the analysed project, while also being in meaningful locations relative to one another. Through this process, disparities in power distribution between stakeholder groups become apparent.

The resulting grid is then used as a basis for discussing action points for the participants. In the current version of the workshop, we chose to focus on developing action points for very powerful stakeholders with high interest, and for very vulnerable stakeholders (high interest / low power) as we consider the most relevant disparities to exist between such groups. Discussion is stimulated with a series of questions inspired by [46] and [62], which are printed on the workshop board:

³For instance, the prompts in Data Ethics Canvas “Name/describe your project’s key data sources, whether you’re collecting data yourself or accessing via third parties” and “Where did you get the project’s data from?” informed the Question 1 “Who provides data for the project?”

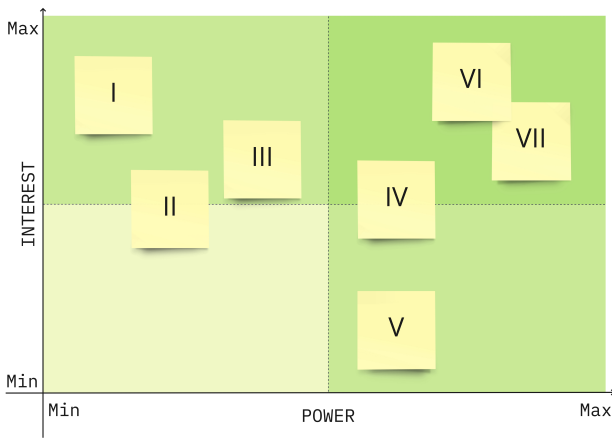


Figure 2: General layout of a stakeholder mapping.

Questions about powerful stakeholders: 1. Are your expectations towards the project aligned with theirs? 2. Do you observe potential for power imbalances? 3. How can you protect your project from negative external influence? 4. Could you harness the power of these stakeholders for a positive overall impact?

Questions about vulnerable stakeholders: 1. What expectations do they have towards your project? 2. How may limitations or risks of the project be communicated to people impacted? 3. How can these stakeholders engage with you with feedback, concerns, requests or appeals? 4. What steps can you take to minimise negative on more vulnerable stakeholders? 5. How could you inform, consult, involve or empower your stakeholders through your project?

3.3 Testing EASE

Current workshop format was tested in April-May 2022 with six volunteering teams. Invitations were sent to in total twelve teams known to be active in music-AI research, in an effort of ensuring some variety in the types of projects included, but there was no specific vetting of participants' prior knowledge of and experience with ethics discussions. The volunteering participant teams came from academic research institutions in Australia, Austria, Finland, France and Great Britain. Each of these teams used a project that they are currently working on or have previously worked on as a scenario for the workshop, as summarised in Table 1.

In total, 17 people took part in the workshops, with the team sizes varying from two to four. Most of the participants self-identified as 'male', only two as 'female' and none as 'other'. This gender ratio, while statistically imbalanced, is consistent with the general gender structure of the NIME community [64], and was thus accepted as sufficient for this initial round of tests.

3.4 Workshop process

Each workshop lasted approximately 90 minutes, with an equal amount of time spent on stakeholder identification and mapping, respectively. The time spent on each of the individual question prompts in the first exercise was not rigorously timed. The total workshop duration also included a brief feedback discussion.

The first author acted as a facilitator and moderator, while the other authors observed the process. For future use of our method, this role of a facilitator could be assumed by someone from the team, who familiarises themselves with

the structure of the exercises and keeps an eye on the time to ensure that each part of the workshop is given sufficient attention. After each workshop, the authors conducted a brief internal feedback session discussing their observations. The workshops were also recorded and analysed in terms of both contents of the discussions and potential problems in the process. All participants provided informed consent.

4. RESULTS

Testing of EASE in six workshops provided various insights into the concrete ethical predicaments music-AI developers are facing in their work. In the following, we present a summary of these insights.

4.1 Stakeholder insights

Workshop participants identified on average 16 stakeholders or stakeholder groups (minimum: 9, maximum: 30). After overlaps and repetitions were removed, on average 13 stakeholder categories were mapped on the power/interest grid (minimum: 9, maximum: 20). The stakeholders most commonly mentioned in the workshops were artists, data providers, users and listeners, research communities, and various companies and employees in the music industry.

Meaningful differentiation, however gradually emerged *within* these broad categories. For example, the category of "artists" included musicians and composers actively involved in developing, testing or using the system, as well as musicians whose works are in the used training data sets used. Some teams additionally included *e.g.* artists whose style is *under-represented* in the training data, and for whom the system might thus not perform well (teams C and F).

In particular, questions 8 and 9 inspired the teams to explore previously undetected use case opportunities such as physically impaired users (teams B, C and F), biases and structural inequalities within the creative industry (teams E and F), and overall positive chain effects that the wider availability of methods for creative work could bring to artists and audiences (team E). Several teams also singled out in their feedback questions 8 and 9 as the most interesting. Overall, these questions were successful in stimulating conversation and debate about a wide range of ethical predicaments, opening up perspectives towards more speculative scenarios of societal and ethical impacts.

4.2 Navigating value tensions

Many teams recognised that the identified sub-groups of stakeholders were dissimilarly affected by the development project. For instance, team C reflected on the possibilities of researchers to empower less known artists to explore the system without validating "artists [who are] trying to have an unfair advantage". Team F discussed value tensions that would arise when companies seek to "make money off the artists". Furthermore, the teams observed that the general public might not necessarily share the values of the research group nor their views towards the desirability of music-AI in general, which could cause unexpected tensions and resistance towards the development projects.

Concerning certain sub-groups of artists that provide data for the project, team E recognised they could consciously influence the power dynamics of the situation themselves: "We can choose to give them more power ... in terms of how their sounds can be used". Some of the teams specifically expressed that in order to better understand the complex power dynamics of the stakeholder networks, they would first need to reflect on and define their internal values and

Table 1: Summary of the workshop participants

Team	Team size	Participants	Project description
A	2	male (2), age 23-49	Real-time deep learning-based sound synthesis and manipulation
B	3	male (3), age 29-39	Symbolic music generation system adapted to a specific instrument
C	2	female (1), male (1), age 28-51	Targeted manipulation of music recommendation systems
D	3	female (1), male (2), age 27-37	Deep generative models as tools for composers
E	2	male (2), age 29-45	Music production tool to select musical elements from a collection
F	4	male (4) age 24-50	Deep-learning based music generation system

expectations towards the project. Such contemplations indicate that the workshop had succeeded in initiating a process of critical self-reflection and heightened sensitivity to the ethical predicaments of the project.

4.3 Developing data management

In most workshops, the teams reflected on the conditions of legitimate and fair data use in AI research. Some teams extended the special care warranted to personal or sensitive data to encompass private artistic expressions, thus extending the legal perspective of GDPR. Others acknowledged that they identified blind spots in their current data management processes, and several expressed frustration in the face of the complexity of setting up responsible data use practices, hoping for more legislative or regulatory support and guidance. Furthermore, many teams contemplated on their capacity to engage with platforms and other commercial entities in promoting responsible data practices. These efforts could include open communication about the observed problems and biases in the systems, or in more extreme cases, pressure towards the platforms through raising media attention.

4.4 Planning further actions

EASE methods aims to help developers identify, protect and empower weaker stakeholders or alleviate harm towards them. The related ethical issues are, however, often complex, open-ended and context-dependent. This was explicitly recognised by teams E and F, who wondered how deep they should continue in the analysis of the down-stream chain of data extraction and preparation networks (see [11]). Furthermore, as observed by teams B and F, the ethical stance of the development project may change over time, for instance, if the project gets commercialised. The exploration initiated in the workshops should therefore fuel a continuous process of proactive ethical reflection including but also extending beyond the stakeholder elicitation. Indeed, our participants spontaneously indicated plans to conduct the process again, possibly with a different case study as a focus, or at different phases of the same project.

There was also an increased appreciation of foregrounding user and other stakeholder engagement in the development work. The teams differed markedly on their previous efforts of stakeholder interaction and consequently, on the range and depth of changes they expressed they would seek to implement in the projects after the workshop. Teams B and C for instance, discussed options for extending their research outreach efforts to motivate a wider range of users to explore their system. Conducting user studies, opening better communication channels and other efforts of direct engagement were often identified as an important step for the project maturation and expansion. Team C, which had identified "user" as a relevant stakeholder group only after being prompted by question 5 of the first exercise, acknowledged their previous detachment from user perspectives and

expressed their intention to change that. This indicates that the workshops were successful in prompting reflection, but it remains an open question whether such plans come to fruition in actual development practices.

5. DISCUSSION

Besides insights on the ethical impacts of the individual music-AI projects themselves, the workshops shed light on some of the benefits and limitations of EASE in facilitating structured ethical reflection in the creative-AI development. In the following, we will discuss these insights and relate them to plans for future work.

5.1 Addressing workshop dynamics

While the workshops were generally successful in their current form, the testing clearly demonstrated the benefits of having at least three participants involved in the process. Active discussions among the participants increased the number and diversity of stakeholders identified, while the insights into the specific stakeholder dynamics grew more detailed over the course of the workshops. Conversely, certain tendencies to a converging agreement were observed in the workshops where the team was small and had strong internal power dynamics, such as between a supervisor and their PhD students.

To alleviate this effect, all participants should be guaranteed to have individual access to the workshop board and feel equally authorised to make changes to it, so that contrasting perspectives can gain a voice in the discussions. Furthermore, for workshops with more than four participants, the stakeholder mapping could be run as two or several parallel break-out groups, allowing for a comparison of the resulting matrices for potentially more diverse insights.

5.2 Feedback and critical reflections

Overall, most participants considered the process "informative" and "revealing", and the set of nine questions on the stakeholder identification board "holistic" and "thorough". They reported "having a better view" of the roles different stakeholders play in the project, and noted that the analysis "helped to elicit values". A certain level of sampling bias should be assumed since the volunteering teams are likely to have a more active interest in ethics discussions than an average research team, but perhaps a limited experience with other ethics analysis methods.

Certain aspects of the workshop format yielded uneven results between the participant groups. For some teams, question 6 ("Who sets the rules, regulations and recommendations for your data use?") mainly related to high-level political entities such as the EU or national governments, with limited brainstorming on other, more local stakeholders. Similarly, some participants found question 2 ("Who is the data about?") confusing. Finally, there was a certain bias in the answers to question 8 towards the issue of autonomous systems replacing human actors, which may have

unnecessarily limited exploration of other potential negative impacts. In the future, we will consider reformulating these question prompts for better clarity and inspiration, and evaluate the option of offering further elective prompts for some or all of the questions.

In one of the workshops (team D), participants felt they were "on a different planet", expressing alienation from the stakeholder perspective and by the strongly guided process of the workshop. This stands in contrast to feedback from other workshops (e.g. teams A, E and F), where participants specifically mentioned a strong appreciation of the guidance and the two-step format of the workshop. As we continue to develop the method, any further cases of alienation will be closely analysed, but currently, it seems advisable to adjust the amount of guidance in correspondence with the preferences of the participant group.

5.3 Future directions for EASE

Since data use is one of the key aspects of AI development, the ODI Data Ethics Canvas provides a solid ground for supporting the process of ethically aligned stakeholder identification. The workshops, however, exposed certain blind spots in the chosen framework. For instance, it was noted that funding institutions exert both active and unintentional power towards research projects in terms of topics, methodology and data management (see also [28]), which was not addressed in the prompt questions. Some of the teams hoped to analyse deeper the intricate and interrelated structures of ownership, access and economic power of commercial platforms or service providers, which the current workshop format did not sufficiently support. The option of extending the human-centered approach of EASE to more-than-human agents, such as ecosystems, was mentioned in two of the workshops (teams D and E) and is also reflected in the recent discussions in the HCI field [23, 34, 36, 37]. It is likely that other aspects specifically relevant for the research or development of creative-AI projects will be uncovered as we continue revising and testing EASE.

To complement the current case study, we aim to continue testing EASE in wider geographic areas (see [51]), in adjacent fields of creative-AI development (e.g. visual, performative), and with developers of commercial creative-AI applications, to strive for greater diversity in gender [64] and other aspects. This will serve to ground the method further to the practical realities of the development work. Alternative ethics frameworks (e.g. [15]), as well as perspectives from critical [27] and feminist [6] stakeholder theories and VSD provide inspirations for subverting any manipulative tendencies that may reside in traditional stakeholder analyses rooted in strategic management (see [49]), and for widening the scope of analysis.

Finally, we aim to accompany some of the participating teams in longitudinal studies to evaluate how action points developed using EASE are implemented (or not) in long term, as well as to observe and analyse the current AI-development culture and practices. These efforts will add to the insights obtained by [53] and others, and provide critical domain-specific context both for the developing EASE and other, possibly creative-AI specific ethical analysis methods and for the drafting of policies and ethics guidelines tailored for creative-AI development.

6. CONCLUSION

The research and development of creative-AI applications is taking place at a rapid pace, with insufficient reflection on the ethics involved. The creative-AI community urgently

needs to proactively contribute to the development of ethically aligned and responsible applications. To support this, we have introduced an analytic method, Ethically Aligned Stakeholder Elicitation (EASE).

Based on the observations during the workshops and the feedback received, EASE provided a structured way to stimulate critical self-reflection, inspired teams discover new aspects of their work and articulate a shared understanding of value tensions and power asymmetries that require their attention, as well as more generally increased their awareness of the range of social and ethical impacts that a music-AI application may have. Hence, EASE paves the way towards sector specific ethical standards that can guide the progress of the creative-AI development in the future. In developing our method further, we hope EASE to be inspiration for more collaborative efforts of co-designing a greater diversity of AI ethics analysis approaches.

7. ETHICAL STANDARDS

This paper complies with the NIME ethical standards. Limitations to the cultural and gender diversity of the participants are addressed and discussed in the article. Research conducted in the paper was evaluated by the institutional ethics board and found to adhere to the national guidelines on research involving human participants. All workshop participants provided informed consent.

8. ACKNOWLEDGEMENTS

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program – Humanities and Society (WASP-HS) funded by the Marianne and Marcus Wallenberg Foundation (Grant 2020.0102), and by the Swedish Research Council (2019-03694).

9. REFERENCES

- [1] J. Bardzell and S. Bardzell. Humanistic HCI. *Interactions*, 23(2):20–29, Feb 2016.
- [2] C. Bauer, A. Ferraro, E. Gómez, and L. Porcaro. Trustworthy MIR: Creating MIR applications with values. Tutorial at the ISMIR Conference, 2022.
- [3] R. Benjamin. *Race After Technology: Abolitionist Tools for the New Jim Code*. Polity, 2019.
- [4] S. M. A. Bin. Discourse is critical: Towards a collaborative NIME history. In *Proceedings of NIME*, 2021.
- [5] G. Born. Diversifying MIR: Knowledge and real-world challenges, and new interdisciplinary futures. *Transactions of the International Society for Music Information Retrieval*, 3(1):193–204, 2020.
- [6] B. K. Burton and C. P. Dunn. Feminist ethics as moral grounding for stakeholder theory. *Business Ethics Quarterly*, 6:133 – 147, 1996.
- [7] M. Butterick. Stable Diffusion litigation. <https://stablediffusionlitigation.com>, 2023. Last accessed Jan 26, 2023.
- [8] A. Calderon, D. Taber, H. Qu, and J. Wen. AI blindspot. <https://aibindspot.media.mit.edu>, 2019. Last accessed Dec 07, 2022.
- [9] M. Clancy, editor. *Artificial Intelligence and music ecosystem*. Routledge, 2022.
- [10] E. H. Clark. A postcolonial MIR? *Resonance*, 3(4):412–432, 2022.
- [11] K. Crawford. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. Yale

- University Press, 2021.
- [12] Digital Catapult. Ethics framework. <https://aibindspot.media.mit.edu>, 2021. Last accessed Dec 07, 2022.
- [13] E. Drott. Copyright, compensation, and commons in the music AI industry. *Creative Industries Journal*, 14(2):190–207, Oct. 2020.
- [14] C. Eden and F. Ackermann. *Making strategy: The journey of strategic management*. Sage Publications, 1998.
- [15] European Commission and Directorate General for Communications Networks, Content and Technology. *The Assessment List for Trustworthy Artificial Intelligence (ALTAI) for self assessment*. Publications Office, 2020.
- [16] European Commission. Directorate General for Communications Networks, Content and Technology. *Opportunities and challenges of artificial intelligence technologies for the cultural and creative sectors*. Publications Office, 2022.
- [17] European Commission. Directorate General for Communications Networks, Content and Technology and High Level Expert Group on Artificial Intelligence. *Ethics guidelines for trustworthy AI*. Publications Office, 2019.
- [18] A. Ferraro, X. Serra, and C. Bauer. What is fair? exploring the artists’ perspective on the fairness of music streaming platforms. In C. Ardito, R. Lanzilotti, A. Malizia, H. Petrie, A. Piccinno, G. Desolda, and K. Inkpen, editors, *Human-Computer Interaction – INTERACT. Lecture Notes in Computer Science*, volume 12933, pages 562–584. Springer, 2021.
- [19] C. Flick and K. Worrall. The Ethics of Creative AI. In C. Vear and F. Poltronieri, editors, *The Language of Creative AI*, pages 73–91. Springer International Publishing, Cham, 2022.
- [20] A. S. Franzke, I. Muis, and M. T. Schäfer. Data Ethics Decision Aid (DEDA): a dialogical framework for ethical inquiry of AI and data projects in the Netherlands. *Ethics and Information Technology*, 23:551–567, 2021.
- [21] R. E. Freeman. *Strategic management: A stakeholder approach*. Pitman series in Business and Public Policy, 1984.
- [22] B. Friedman and D. Hendry. The envisioning cards: a toolkit for catalyzing humanistic and technical imaginations. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 1145–1148, 2012.
- [23] B. Friedman and D. G. Hendry. *Value sensitive design: Shaping technology with moral imagination*. MIT Press, 2019.
- [24] B. Friedman and H. Nissenbaum. Bias in computer systems. *ACM Transactions on Information Systems*, 14(3):330–347, July 1996.
- [25] N. E. Gold, R. Masu, C. Chevalier, and F. Morreale. Share Your Values! Community-Driven Embedding of Ethics in Research. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, pages 1–7, 2022.
- [26] E. Gomez, A. Holzapfel, M. Miron, and B. L. Sturm. Fairness, Accountability and Transparency in Music Information Research (FAT-MIR). Tutorial at the ISMIR Conference, 2010.
- [27] M. Greenwood and R. Mir. Critical management studies and stakeholder theory: Possibilities for a critical stakeholder theory. In J. S. Harrison, J. B. Barney, R. E. Freeman, and R. A. E. Phillips, editors, *The Cambridge Handbook of Stakeholder Theory*, page 35–52. Cambridge University Press, 2019.
- [28] L. Hayes and A. Marquez-Borbon. Addressing NIME’s Prevailing Sociotechnical, Political, and Epistemological Exigencies. *Computer Music Journal*, 44(2-3):24–38, July 2020.
- [29] M. Heikkilä. The Algorithm: AI-generated art raises tricky questions about ethics, copyright, and security. *MIT Technology Review*, Sept 2022.
- [30] V. Held. *The ethics of care: Personal, political, and global*. Oxford University Press, 2006.
- [31] A. Holzapfel, B. L. Sturm, and M. Coeckelbergh. Ethical Dimensions of Music Information Retrieval Technology. *Transactions of the International Society for Music Information Retrieval*, 1(1):44–55, Sept. 2018.
- [32] R. Huang, B. L. T. Sturm, and A. Holzapfel. De-centering the West: East Asian Philosophies and the Ethics of Applying Artificial Intelligence to Music. In *Proceedings of the 22nd ISMIR Conference*, pages 301–309, 2021.
- [33] A. Jobin, M. Ienca, and E. Vayena. The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9):389–399, 2019.
- [34] P. Jääskeläinen, A. Holzapfel, and C. Åsberg. Exploring More-than-Human Caring in Creative-Ai Interactions. In *Proceedings of NordiCHI Conference*, 2022.
- [35] C. Key, F. Browne, N. Taylor, and J. Rogers. Proceed with Care: Reimagining Home IoT Through a Care Perspective. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI ’21, New York, NY, USA, 2021. Association for Computing Machinery.
- [36] S.-Y. Liu. *Posthuman Interaction Design: Designing with, through, and for Human-Nature Interaction*. Phd thesis, Indiana University, 2020. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2022-04-29.
- [37] R. Masu, A. P. Melbye, J. Sullivan, and A. R. Jensenius. NIME and the Environment: Toward a More Sustainable NIME Practice. In *Proceedings of NIME*, April 2021. <https://nime.pubpub.org/pub/4bbl5lod>.
- [38] J. Morley, L. Floridi, L. Kinsey, and A. Elhalal. From what to how: an initial review of publicly available AI ethics tools, methods and research to translate principles into practices. In *Ethics, Governance, and Policies in Artificial Intelligence*, pages 153–183. Springer, 2021.
- [39] F. Morreale. Where Does the Buck Stop? Ethical and Political Issues with AI in Music Creation. *Transactions of the International Society for Music Information Retrieval*, 4(1):105–113, 2021.
- [40] F. Morreale, S. M. A. Bin, A. P. Mcpherson, P. Stapleton, and M. M. Wanderley. A nime of the times: Developing an outward-looking political agenda for this community. In *Proceedings of NIME*, 2020.
- [41] L. Munn. The uselessness of AI ethics. *AI and Ethics*, 2022.
- [42] L. P. Nathan, B. Friedman, P. Klasnja, S. K. Kane, and J. K. Miller. Envisioning systemic effects on

- persons and society throughout interactive system design. In *Proceedings of the 7th ACM Conference on Designing Interactive Systems*, DIS '08, page 1–10, New York, NY, USA, 2008. Association for Computing Machinery.
- [43] R. Newcombe. From client to project stakeholders: a stakeholder mapping approach. *Construction management and economics*, 21(8):841–848, 2003.
- [44] S. Olander and A. Landin. Evaluation of stakeholder influence in the implementation of construction projects. *International journal of project management*, 23(4):321–328, 2005.
- [45] C. O’Neil. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway books, 2016.
- [46] Open Data Institute. Data ethics canvas, 2021. Last accessed 18 May 2022.
- [47] C. O’Neil and H. Gunn. Near-term artificial intelligence and the ethical matrix. In S. M. Liao, editor, *Ethics of Artificial Intelligence*, pages 235–69. Oxford University Press, 2020.
- [48] L. Pardue and S. M. A. Bin. The Other Hegemony: Effects of software development culture on music software, and what we can do about it. In *Proceedings of NIME*, Jun 2022.
- [49] M. Reed and R. Curzon. Stakeholder mapping for the governance of biosecurity: a literature review. *Journal of Integrative Environmental Sciences*, 12(1):15–38, 2015.
- [50] M. Saha, S. Patil, E. Cho, E. Y.-Y. Cheng, C. Horng, D. Chauhan, R. Kangas, R. McGovern, A. Li, J. Heer, and J. E. Froehlich. Visualizing urban accessibility: Investigating multi-stakeholder perspectives through a map-based design probe study. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22, New York, NY, USA, 2022. Association for Computing Machinery.
- [51] N. Sambasivan, E. Arnesen, B. Hutchinson, T. Doshi, and V. Prabhakaran. Re-Imagining Algorithmic Fairness in India and Beyond. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, page 315–328, New York, NY, USA, 2021. Association for Computing Machinery.
- [52] H. Schneider, M. Eiband, D. Ullrich, and A. Butz. Empowerment in HCI – A survey and framework. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–14, 2018.
- [53] N. Seaver. *Computing Taste: Algorithms and the Makers of Music Recommendation*. University of Chicago Press, 2022.
- [54] B. R. Shapiro, A. Meng, C. O’Donnell, C. Lou, E. Zhao, B. Dankwa, and A. Hostetler. Re-shape: A method to teach data ethics for data science education. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, page 1–13, New York, NY, USA, 2020. Association for Computing Machinery.
- [55] B. L. T. Sturm, M. Iglesias, O. Ben-Tal, M. Miron, and E. Gómez. Artificial Intelligence and Music: Open Questions of Copyright Law and Engineering Praxis. *Arts*, 8(3):115, Sept. 2019.
- [56] K. Sun, A. H. Mhaidli, S. Watel, C. A. Brooks, and F. Schaub. It’s My Data! Tensions Among Stakeholders of a Learning Analytics Dashboard. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, page 1–14, New York, NY, USA, 2019. Association for Computing Machinery.
- [57] The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems. *Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems*. Technical Report First edition, IEEE, 2019.
- [58] A. L. Toombs, S. Bardzell, and J. Bardzell. The proper care and feeding of hackerspaces: Care ethics and cultures of making. In *33rd annual ACM conference on human factors in computing systems*, pages 629–638, 2015.
- [59] UNESCO. Recommendation on the ethics of artificial intelligence, 2022. Last accessed 18 May 2022.
- [60] V. Vakkuri, K.-K. Kemell, M. Jantunen, E. Halme, and P. Abrahamsson. ECCOLA — A method for implementing ethically aligned AI systems. *Journal of Systems and Software*, 182:111067, 2021.
- [61] S. Vallor. *Technology and the virtues: A philosophical guide to a future worth wanting*. Oxford University Press, 2016.
- [62] S. Vallor, B. Green, and I. Raicu. An Ethical Toolkit for Engineering/Design Practice. <https://www.scu.edu/ethics-in-technology-practice/ethical-toolkit>, 2018. Last accessed 07 December 2022.
- [63] M. Weber. *The theory of social and economic organization*. Free Press, 1947.
- [64] A. Xambó. Who are the women authors in NIME? — improving gender balance in NIME research. In *Proceedings of NIME*, 2018.
- [65] D. Yoo. Stakeholder tokens: A constructive method for value sensitive design stakeholder analysis. In *Proceedings of the 2017 Conference on Designing Interactive Systems*, pages 280–284, 2017.
- [66] M. Young, L. Magassa, and B. Friedman. Toward inclusive tech policy design: a method for underrepresented voices to strengthen tech policy documents. *Ethics and Information Technology*, 21:89–103, 2019.