

---

# The responsible media representation of artificial intelligence for non-human animal communication (proposal)

---

**Caatje Kluskens**  
Wageningen University and Research  
caatje.kluskens@wur.nl

## Abstract

The body of research on the ethical, legal, and social implications of artificial intelligence for non-human animal communication is growing. However, one area that has not yet been explored is how the outcomes of these technologies should be responsibly represented in the media and how such portrayals influence public awareness, support for conservation efforts, and anthropocentric assumptions about intelligence. This project will examine risks such as speciesism, including human exceptionalism, the polarization of public discourse, Disneyfied stereotypes, and unintentional encouragement for ecotourism. As an alternative, it proposes animal-interest framing and explores multispecies justice cinema as alternative strategies for more ethical portrayals. By proactively developing guidelines for responsible media practices, the research project aims to ensure that the representations of breakthroughs in AI for non-human animal communication serve the interests of animals and broader conservation goals. This proposal will address potential risks of problematic framing of the research outcomes by journalists, discuss alternative representational strategies, and briefly outline the methodology of the PhD research project.

## 1 Introduction

Initiatives like the Cetacean Translation Initiative (CETI) and the Earth Species Project (ESP) aim to decode non-human animal communication through machine learning and various robotic technologies (12)(7). In November 2024, ESP introduced NatureLM-audio, an audio-language model designed specifically for animal vocalizations and sounds (13). It can identify individual animals within recordings and even recognize patterns in the vocalizations of species it has not previously encountered. NatureLM-audio is a significant step towards a future in which humans can decode non-human animal communication with machine learning. The technology has the potential to disrupt existing philosophical paradigms rooted in anthropocentric beliefs about language and intelligence.

Scholars have begun to pay greater attention to the ethical concerns of this technology and its potential risks to animals. Ryan and Bossert raise several issues, including: the risk of anthropomorphism, the right to privacy for animals, the potential emotional and cultural harm to animal communities, "technological solutionism": solving anthropogenic threats to animals with new technological tools rather than addressing the root causes, the potential ineffectiveness for nature conservation, and the risk of projecting gender biases onto other animals (17). Rodríguez-Garavito, Gallant, and Crow warn that while this technology has the potential to be used to understand non-human animals, it can simultaneously be used to exploit and harm them. Artificial intelligence can increase the speed and scale of the exploitation of non-human animals. They propose pillars for an ethical foundation for non-human animal communication technologies, among which is the prioritization of the animal's interest (15). In addition, Gruber and Beguš from the CETI project, together with Rodríguez-Garavito

and Nemeth, recently published a paper exploring the legal impact of AI for non-human animal communication (14).

While the body of research on the ethical, legal, and social implications of AI for non-human animal communication is growing, one area that has not yet been explored is how the outcomes of these technologies should be responsibly represented in the media and how these portrayals impact the public perception of other animals. As Adams argues in *Picturing the Beast* (1), media representations of animals are the only access most people have to animals. The impact of this technology on conservation efforts partially depends on its media portrayal. This proposal will explore three specific concerns, namely, the risk of speciesism, including human exceptionalism, the risk of polarization, and the potential Disneyfied stereotyping by journalists, which could drive ecotourism. In my PhD project, I aim to develop guidelines for the responsible representation of animals in media coverage of AI for non-human communication, grounded in multispecies justice theory. This leads to the following research question: What media practices can maximize the positive impact of AI for non-human animal communication on conservation efforts and public understanding of animal intelligence, while minimizing potential harms to animals and the natural world?

## 2 Conceptual Framework

### 2.1 Human exceptionalism and speciesism

Without careful framing, public discourse is at risk of defaulting to human exceptionalism, positioning human intelligence as the highest standard rather than recognizing that each species possesses forms of intelligence uniquely adapted to its environment. As Nibert (11) argues in *Origins of Oppression, Speciesist Ideology, and the Mass Media*, mass media often reinforce and normalize the existing structures of oppression and speciesist ideologies that marginalize non-human animals. Media representations of research outcomes from ESP or CETI may generate unintended framings, such as reinforcing speciesist hierarchies, e.g. placing whale intelligence "closer" to human intelligence. This framing could fuel the discourse that certain animals can be exploited because their intelligence is perceived as further removed from humans, granting some species more moral consideration than others.

### 2.2 Polarization of public discourse

Although the topic of AI for non-human animal communication does not yet appear polarized, researchers and media professionals should carefully consider how they frame the research to ensure it does not become entangled in existing polarized debates. Online debates around polarized topics often take place within echo chambers, where only a small portion of the content reaches across ideological divides. Even when opposing groups encounter the same visual material, their emotional responses are often opposed, which reflects their pre-existing ideological alignments (16).

Projects such as the Earth Species Project and Project CETI aim to reshape our relationship with nature and encourage reflection on humanity's responsibility to care for the planet and all its inhabitants by showing that we are not the only species with intelligence and communication systems (12)(7). Their research would function as a moral appeal for increased environmental and conservation awareness. Translating scientific findings into a moral message inevitably involves framing. When a message unintentionally triggers a politically charged frame (9), individuals may respond in line with their existing ideological positions, potentially reinforcing polarization. Regarding AI for non-human animal communication, further research is needed to assess the risk that certain framings of the research could activate ideological responses associated with resistance to environmental action or skepticism toward scientific institutions. It is crucial to understand how different narrative strategies influence whether the technology bridges or deepens existing ideological divides.

### 2.3 Disneyfied stereotypes and ecotourism

Another risk that needs to be taken into account when developing narrative strategies for animal communication research is the Disneyfied stereotyping of animals in media (6)(20). The term Disneyfication is derived from the portrayal of animals in Disney movies, but the process can be found within all forms of media (10). Often animals are turned into entertaining characters, and in doing so, are anthropomorphized, stereotyped, and oversimplified. Journalists and other media practitioners

frequently assign strong positive or negative characteristics to particular species in order to provoke emotional responses and create sensational narratives. For example, sharks are often portrayed as evil, and accidents involving them receive disproportionate media coverage. This stereotype overshadows their ecological importance and reduces public willingness to support conservation efforts (6)(20). While Disneyfied framing of animal communication research might generate enthusiasm, it risks flattening scientific nuance. It could produce sensationalist news articles filled with fun facts, which might encourage ecotourism rather than highlighting the need for conservation. Ecotourism can have a negative impact on animals; for instance, beluga whales have suffered heavily due to increased boat traffic associated with ecotourism (3). To avoid these outcomes, narrative strategies must be researched in which the individuality of animals and the issues they face are highlighted, without resorting to stereotypes.

#### 2.4 Animal-interest framing and multispecies justice cinema

To achieve this, I propose the term animal-interest framing, inspired by human-interest framing. A human-interest frame brings a human face to the way a news story is presented (19). With animal-interest framing, I suggest that stories about the development of the technology include narratives about individual animals impacted by anthropogenic threats. This form of framing seeks to address the "collapse of compassion" effect, which refers to the tendency of humans to care more about named individuals than about groups. According to Payne and Cameron (4), the more victims are presented in a story, the less humans care, and the less likely they are to donate money, time, or effort. Alongside narrative choices, animal-interest framing also includes the utilization of visual strategies for placing individual animals at the center of a story. For example, high-angle shots position viewers as if they are literally looking down on an animal, while eye-level shots show the animal's eyes. According to Whitley, Kalof, and Flach (22), visibility of an animal's face and eyes in a portrait-like shot is particularly important for creating human compassion towards animals. An important approach to audiovisual content in this context is multispecies justice cinema, developed by Carbonell (5). Multispecies justice theory critiques society for excluding non-human entities and for failing to acknowledge the interconnectedness of humans, animals, and nature. Multispecies justice theorists argue that knowledge and intelligence are not universal truths, but rather context-dependent, with each species possessing unique forms of intelligence shaped by its relationships and environment (2). In 2024, Carbonell applied multispecies justice theory to filmmaking practices, coining the term multispecies justice cinema. This approach explores new ways of seeing, listening, feeling, and relating to animals, ecosystems, and ecological crises. It seeks to decenter human exceptionalism by foregrounding alternative modes of life, perception, and intelligence through film (5).

### 3 Methodology

Rather than analyzing media representations after an event, as is often done in media studies, my research project will proactively investigate potential risks and opportunities of media representations of the technology and its outcomes. To mitigate risks of misrepresentation and unintended problematic framing, it is crucial to investigate representational strategies in the early stages of the development of the technology, so that future scientific breakthroughs are more likely to have a positive impact on animals and the natural world.

The methodology of the PhD research project is divided into five parts. The first paper, *A Multispecies Approach to Artificial Intelligence for Decoding Non-Human Animal Communication*, will establish the project's theoretical foundation and underlying conceptual framework. The second paper, *Multispecies Justice in Documentaries on Non-Human Animal Behavior Research: Examining My Octopus Teacher, Project Nim, Fathom, and A Life Among Elephants*, will analyze the representational strategies employed in the selected documentaries and reflect on how the strategies support or challenge principles of multispecies justice. The third paper examines the risks and opportunities associated with media representations of this technology and research on animal lives. The fourth paper, *A Phenomenological Approach to Representing Non-Human Umwelten in Audio-Visual Media about Interspecies Communication*, will involve the creation of video essays based on previously researched narrative and visual strategies. The impact of the video essays will be evaluated through audience research. In the final paper, a set of guidelines will be formulated for representing knowledge generated by AI for non-human animal communication based on the insights gathered during the research project.

## 4 Limitations

It is important to acknowledge several limitations of this research proposal. First, various aspects of the study depend on how future technologies and public discourses evolve. Although the study will employ in-depth methods to identify the risks associated with the media representations, the rapid and uncertain development of AI for non-human animal communication means that unexpected risks may emerge over time. Understanding these risks will therefore be an ongoing process, requiring continuous reassessment of the technology and its public discourse. Second, because the project involves the interpretative analysis of media artifacts and the development of normative guidelines, the researcher's positionality, grounded in multispecies justice theory, will shape the framework and recommendations. Third, the study focuses primarily on Anglo-European media contexts, as the research team is based in the Netherlands. This geographical and cultural focus may limit the generalizability of the findings and may overlook media practices from other cultures that frame non-human animals differently. These limitations highlight the importance of an iterative research approach and the need for continued dialogue across disciplines.

## 5 Acknowledgements

This research project is funded by the Dutch Research Council (NWO) through Dr. Bernice Bovenkerk's Vici grant for the project "*The Promise and Perils of Digital Technology for Human–Animal Relationships*." The author is affiliated with Wageningen University and Research.

## References

- [1] Adams, Carol J. 2001. *Picturing the Beast: Animals, Identity and Representation*. University of Illinois Press.
- [2] Barad, Karen. 2007. *Meeting the Universe Halfway*. Duke University Press.
- [3] Blane, Jean M., and Reiner Jaakson. 1994. "The Impact of Ecotourism Boats on the St. Lawrence Beluga Whales." *Environmental Conservation* 21(3): 267-269. <https://doi.org/10.1017/S0376892900033282>.
- [4] Cameron, Daryl C., and Keith B. Payne. 2011. "Escaping Affect: How Motivated Emotion Regulation Creates Insensitivity to Mass Suffering." *Personality and Social Psychology Bulletin* 37(6): 754-768. <https://doi.org/10.1037/a0021643>.
- [5] Carbonell, Isabelle. 2024. "Attuning to the Pluriverse: Multispecies Cinema." *The Minnesota Review* (2024): 103-117. <https://doi.org/10.1215/00265667-11369537>.
- [6] Dettmar, Ute, and Ingrid Tomkowiak. 2023. *On Disney: Deconstructing Images, Tropes and Narratives*. J.B. Metzler Berlin.
- [7] Earth Species Project. n.d. "What we do." Last accessed September 4, 2025. <https://www.earthspecies.org/what-we-do/technology>.
- [8] Falkenberg, Max, Alessandro Galeazzi, Maddalena Torricelli, Niccolò Di Marco, et al. 2022. "Growing Polarization Around Climate Change on Social Media." *Nature Climate Change* 12: 1114–1121. <https://doi.org/10.1038/s41558-022-01527-x>.
- [9] Lakoff, George. 2010. "Why it Matters How We Frame the Environment." *Environmental Communication* 4(1): 70–81. <https://doi.org/10.1080/17524030903529749>.
- [10] Milani, Myrna. 2014. "The Disneyfication challenge to quality veterinary communication." *Canadian Veterinary Journal* 55(10): 997-998. <https://pmc.ncbi.nlm.nih.gov/articles/PMC4187379/>.
- [11] Nibert, David A. 2015. "Origins of Oppression, Speciest Ideology, and the Mass Media." In *Critical Animal and Media Studies: Communication for Nonhuman Animal Advocacy*, edited by Núria Almiron, Matthew Cole, Carrie P. Freeman. Routledge.
- [12] Project CETI. n.d. "About." Last accessed September 4, 2025. <https://www.projectceti.org/about>.
- [13] Robinson, David, Marius Miron, Masato Hagiwara, and Olivier Pietquin. 2024. "Introducing NatureLM-audio: An Audio-Language Foundation Model for Bioacoustics." *Earth Species Project*, November 11. <https://www.earthspecies.org/blog/introducing-naturelm-audio-an-audio-language-foundation-model-for-bioacoustics>.

[14] Rodríguez-Garavito, César, David F. Gruber, Ashley Otilia Nemeth, and Gasper Begus. 2025. "What If We Understood What Animals Are Saying? The Legal Impact of AI-assisted Studies of Animal Communication." *Ecology Law Quarterly* 52(1). <https://doi.org/10.15779/Z383X83N5Q>.

[15] Rodríguez-Garavito, César, Jacqueline Gallent, and Emma Crowe. 2024. "Listening to the more-than-human world: Legal and ethical principles for nonhuman animal communication technologies." *Open Global Rights*, December 6. <https://www.openglobalrights.org/listening-to-the-more-than-human-world-legal-and-ethical-principles-for-nonhuman-animal-communication-technologies/>.

[16] Rossi, Luca, Alexandra Segerberg, Luigi Arminio, and Matteo Magnani. 2025. "Do You See What I See? Emotional Reaction to Visual Content in the Online Debate About Climate Change." *Environmental Communication* 19(3): 449—467. <https://doi.org/10.1080/17524032.2024.2420787>.

[17] Ryan, Mark, and Leonie Bossert. 2024. "Dr. Doolittle uses AI: Ethical challenges of trying to speak whale." *Biological Conservation* 295 (July): 110648. <https://doi.org/10.1016/j.biocon.2024.110648>.

[18] Salomon, Erika, Jesse L. Preston, and Melanie B. Tannenbaum. 2017. "Climate Change Helplessness and the (De)Moralization of Individual Energy Behavior." *Journal of Experimental Psychology: Applied* 23(1): 15–28. <https://doi.org/10.1037/xap0000105>.

[19] Semetko, Hollie, Patti Valkenburg. 2000. "Framing European politics: A content analysis of press and television news." *Journal of Communication* 50(2): 93-109.

[20] Stanton, Rebecca R. 2021. *The Disneyfication of Animals*. Springer Nature Switzerland AG.

[21] Täuber, Susanne, Martijn van Zomeren, and Maja Kutlaca. 2015. "Should the Moral Core of Climate Issues Be Emphasized or Downplayed in Public Discourse? Three Ways to Successfully Manage the Double-Edged Sword of Moral Communication." *Climatic Change* 130: 453–464. <https://doi.org/10.1007/s10584-014-1200-6>.

[22] Whitley, Cameron T., Linda Kalof, and Tim Flach. 2020. "Using Animal Portraiture to Activate Emotional Affect." *Environment and Behavior* 53(8): 837-863. <https://doi.org/10.1177/0013916520928429>.

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The abstract and introduction present claims consistent with the research proposal, as this submission is for the proposal track of the workshop rather than a completed paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: The paper discusses the limitations of the proposal.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The submission is for the proposal track of the workshop “AI for non-human animal communication”. As it is a proposal, it does not yet include research results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [NA]

Justification: The submission is for the proposal track of the workshop “AI for non-human animal communication”. As it is a proposal, it does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: The submission is for the proposal track of the workshop “AI for non-human animal communication”. As it is a proposal for social sciences research, it does not include experiments requiring code.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: The submission is for the proposal track of the workshop “AI for non-human animal communication”. As it is a proposal, it does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: The submission is for the proposal track of the workshop “AI for non-human animal communication”. As it is a proposal, it does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification: The submission is for the proposal track of the workshop “AI for non-human animal communication”. As it is a proposal, it does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn’t make it into the paper).

## 9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The proposal adheres to the NeurIPS Code of Ethics. The research in this submission has not yet been conducted, as it is a proposal.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: The proposal discusses both the potential positive and negative impacts of the technology it will research.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The proposal and future research do not pose such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: The proposal does not and will not use existing assets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](http://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

### 13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The proposal does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The proposal does not involve crowdsourcing nor research outcomes with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The proposal does not involve crowdsourcing nor research outcomes with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

#### 16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard component.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.