FG 2025 Ethical Impact Statement Guidelines

History and Rationale

In the past few years, we have seen increasing deployment of artificial intelligence (AI) tools in various societies, which has had real-world impacts on the daily lives of millions of individuals around the world. The scalability of these tools brings the potential for both great benefits and great harms to individuals and groups. Our research community is acutely aware of the possibility that the technologies that we develop and study may be abused, misused, or misunderstood to create harm that we did not foresee or intend. FG 2020/2021 invited papers on the topic of regulation and social impact, and both FG 2023 and FG 2024 invited papers on the topic of privacy and ethical issues; this latter topic was also the conference theme of FG 2023.

This year, at FG 2025, we will be introducing a new requirement that authors submit an Ethical Impact Statement as part of the submission process. (Note that this requirement applies to all short and long papers submitted to the main track. Each special track may have its own rules.) To support authors in meeting this requirement, as well as reviewers in assessing it, we have prepared this document with general guidelines, a checklist for authors and reviewers to complete, and answers to some frequently asked questions. As this is a new policy, we welcome questions and feedback as we refine it and develop a shared understanding.

Length Requirement

Papers must include a dedicated Ethical Impact Statement section at the end of the paper, after the main paper content and before the reference list. It will be published alongside the paper, if accepted. The Ethical Impact Statement **does not count toward the overall page limit** but must not extend beyond **one additional page**; this limit is intended both to encourage conciseness and to minimize reviewer burden. Thus, long papers may have up to 8 pages for main paper content (excluding references), up to the end of a 9th page for the Ethical Impact Statement, and then additional pages for references as needed. Similarly, short papers may have up to 4 pages for main paper content (excluding references), up to the end of a 5th page for the Ethical Impact Statement, and then additional pages for references as needed. These instructions are summarized in the table below, where X means that the section is allowed on that page.

Paper Type	Section	Pages 1-8	Page 9	Page 10+
Long Paper	Main Paper Content	X		
Long Paper	Ethical Impact Statement	X	X	

Long Paper	References	X	X	X

Paper Type	**Section**	**Pages 1-4**	**Page 5**	**Page 6+**
Short Paper	Main Paper Content	X		
Short Paper	Ethical Impact Statement	X	X	
Short Paper	References	X	X	X

Ethical Principles and Guidelines

Ethical Review Boards

The oversight of an ethical review board (e.g., IRB) is primarily required for research involving **human subjects** (e.g., direct interaction with participants, indirect data collection from individuals, and analysis of identifiable private information). However, it may also be required or advisable for research involving animals, environmental impacts, hazardous biological agents, and sensitive data, as well as **"dual-use research"** that has both beneficial and harmful applications (e.g., synthetic biology and cybersecurity). All papers describing studies that were evaluated by an ethical review board should include any rulings of the board (e.g., approval, exemption, or rejection) and study identifiers (e.g., protocol or approval numbers). However, to protect blinded peer-review, the identity of the ethical review board should be **anonymized** during initial submission (e.g., by masking the name of the governing university or organization). Note that standards differ by geographic region and some authors may not have a requirement in their region to seek the oversight of an ethical review board (and may not have access to one).

Potential Harms to Human Subjects

When research involves human subjects and/or their data, there are **potential risks of harm** to the participants themselves. For example, different research projects may expose participants to risks of being physically injured or infected, emotionally distressed, or inconvenienced in a variety of different ways. Participants' **private and sensitive information** may also be revealed to others they would rather not share that information with, which could lead to real-world (e.g., social, legal, or financial) consequences for them. It is thus imperative that participants understand the potential risks and benefits of their participation and provide **informed consent** to the research study and all eventual uses of their data. Researchers must also protect the privacy and **confidentiality** of participants' data unless the sharing of that data was explicitly consented to.

Potential Negative Societal Impacts

When research leads to new products (e.g., knowledge, theory, data, methods, and tools), researchers have a moral obligation to think about potential uses, misuses, and

misunderstandings of those products. Researchers should consider whether/how their research products could be used or misused in applications that (1) limit people's human rights and privacy, (2) negatively impact people's health or livelihoods, (3) deceive people, (4) create/exacerbate discrimination and group inequalities, or (5) destabilize societal functioning. Researchers should consider how the general public may think and feel about their research products and involve relevant stakeholders (e.g., clinicians and patients in medical applications) in the research process to discover their specific needs, feelings, and concerns. Finally, researchers should also consider how their research products could be misunderstood by various audiences and how those mistakes could lead to negative societal impacts (e.g., miscalibrated optimism or pessimism). A notable example of this would be audiences failing to understand the limits of generalizability of your findings (e.g., your study only examined participants from a single region in a single context but your audience may overgeneralize and conclude that the same results would apply in all regions and contexts).

Risk-Mitigation Strategies

Nearly all research involves some potential risk of harm to individuals or negative societal impacts. These risks are typically offset by the potential benefits of the research. However, researchers can and should also use various strategies to mitigate the risks of their research. Strategies to mitigate risk to human subjects include (but are not limited to) ethical review and oversight, public engagement and consultation with relevant stakeholders, data anonymization, and controlled access to data and other research materials (e.g., code). Strategies to mitigate the risk of negative societal impacts include benefit-risk analyses, regulatory compliance, scenario planning, safeguard development, continuous monitoring and adaptation, longitudinal impact studies, and qualitative research to understand the broader implications and societal context. Finally, strategies to mitigate risk of misunderstandings include clear communication plans, the creation of ethical use guidelines, and proactive engagement with policymakers.

Author/Reviewer Checklist

Before submitting, authors and reviewers should be able to check all of the following five (or six, if the paper involves human subjects) boxes by answering "Yes" to the questions they present. Reviewers may refer to checklist items by their number (e.g., concerns about Checklist #6b).

- 1. Did you read the Ethical Impact Statement Guidelines document (provided above)?
- 2. Is it clear that all studies and procedures described in the paper were approved (or exempted) by a valid **ethical review board**? Alternatively, is a valid and sufficient justification provided for why the oversight of an ethical review board was not required?
- 3. Does the ethical impact statement provide a clear, complete, and balanced discussion of the **potential risks** of individual harm and negative societal impacts associated with the research? Note that this includes harm to research participants as well as harm to other individuals that may be affected by use, misuse, or misunderstanding of the research.

- 4. Does the ethical impact statement describe reasonable, valid, and sufficient use of **risk-mitigation strategies** by the authors to lessen these potential risks? Alternatively, if relevant strategies were not used, is a valid and sufficient justification for this provided?
- 5. Does the ethical impact statement provide a valid and sufficient justification for how/why the potential risks of the research are **outweighed** by the risk-mitigation strategies and potential benefits of the research? Note that papers with serious potential risks that are not outweighed by risk-mitigation strategies and potential benefits may be rejected.
- 6. If the paper involves **human subjects**, are *all* of the following sub-boxes checked?
    - 6a. Does the main paper describe whether/how **informed consent** and/or assent were obtained from participants? If consent and/or assent were fully or partially obtained, were the methods used to do so valid? If not fully obtained, does the ethical impact statement provide a valid and sufficient justification for this?
    - 6b. Does the main paper state whether the participants explicitly consented to the **use of their data** in the manner described in the paper? For example, if the data was or will be shared with third parties, does it state that the participants explicitly agreed to this sharing? If some uses were not explicitly consented to, does the ethical impact statement provide a valid and sufficient justification for this?
    - 6c. Does the main paper explain whether/how participants were **compensated**? If participants were compensated, does the ethical impact statement provide a valid and sufficient justification for the form and amount of compensation provided?
    - 6d. If the research involves any special or **vulnerable populations** (e.g., minors, elderly individuals, prisoners, refugees and migrants, individuals with disabilities, individuals with mental illness, or patients in medical settings), does the ethical impact statement provide a valid and sufficient explanation of how the rights, well-being, and autonomy of such individuals were safeguarded in the research?

Frequently Asked Questions (FAQ)

Q. How long should an ethical impact statement be?
A. The appropriate length of an ethical impact statement will vary from paper to paper depending on the complexity and sensitivity of the data and methods used as well as the applications it is relevant to. That said, each ethical impact statement should be at least a paragraph long and no longer than one full page (longer discussions should be standalone papers); a good starting place might be three paragraphs: one for risks, one for strategies, and one for benefit-risk analysis.

Q. Can I spread the information from the ethical impact statement throughout the paper?

A. While this could be a good strategy for longer-form prose (e.g., journal articles), having all the ethical impact information in one place across all papers facilitates the process of verification and peer-review, which is especially helpful at conferences where many papers are reviewed at once.

Q. What if my work has no potential negative applications?
A. It is difficult (probably impossible) to foresee all the potential applications of our work, including those that could lead to individual harm or negative societal impacts. Thus, it is valuable to set aside the time to brainstorm and seek outside input on these topics. Especially in a community like FG, where we tend to work with human behavioral data, potential risks are usually present.

Q. Do I still need to write an ethical impact statement for technical or theoretical work?
A. Many of the examples of ethical impacts that come to mind most readily involve applied work. However, even technical and theoretical work can have ethical implications. A salient example would be the technical and theoretical work used in the development of “deepfakes,” which have the potential to deceive people, spread misinformation, and harass or bully people.

Q. If I used a publicly available dataset to build my model, are there still ethical issues?
A. Using pre-existing data does not automatically satisfy all ethical concerns. Authors still need to reflect on the ethical aspects of their *new work* using that data, as well as the potential limitations of the use of that data (e.g., does that dataset include systematic biases?).

Q. What if I am presenting early, proof-of-concept work?
A. It may feel premature to consider the downstream applications of early work. However, the pace of technological development, especially in AI, is accelerating. It is also the case that many risk-mitigation strategies are most effective when implemented from the beginning. Thus, it is imperative to reflect on the ethical implications of early, proof-of-concept work too.

Q. If I have the approval of an ethical review board, why do I need to write this statement?
A. The evaluation and approval of an ethical review board is a good start for addressing ethical concerns, but such reviews are usually focused on protecting the rights of human subjects rather than preventing negative societal impacts. Thus, an ethical impact statement focused on potential negative applications and societal impacts can be a valuable supplement to ethical board review.

Q. If we remove all personally identifiable information, are there still ethical issues?
A. Removing personal identifiers like names, addresses, and birthdays can be a great start toward protecting participant privacy and confidentiality, but two issues remain. First, it is often possible for research participants to be re-identified from the remaining data and it is difficult to know what will be possible using future technologies. Second, anonymization does not prevent negative applications and negative societal impacts, so these should still be discussed.

Acknowledgments

This document was drafted by Jeffrey Girard and then reviewed/edited by Tempestt Neal. It draws inspiration from a similar one created for the ACII 2023 conference, and I am indebted to the following people who contributed to creating and editing that document: **Desmond Ong**, Javier Hernadez, Rosalind Picard, Fabien Ringeval, Agata Lapedriza, Akane Sano, Rachael Jack, and Shiro Kumano. The following papers are also relevant inspirations, which interested readers might enjoy digging into.

Affective Computing References

- Cowie, R. (2015). Ethical issues in affective computing. *The Oxford Handbook of Affective Computing,* 334-348.
- Hernandez, J., Lovejoy, J., McDuff, D., Suh, J., O'Brien, T., Sethumadhavan, A., Greene, G., Picard, R., & Czerwinski, M. (2021). Guidelines for assessing and minimizing risks of emotion recognition applications. In *Proceedings of the 2021 IEEE International Conference on Affective Computing and Intelligent Interaction.*
- Ong, D. C. (2021). An ethical framework for guiding the development of affectively-aware artificial intelligence. In *Proceedings of the 2021 IEEE International Conference on Affective Computing and Intelligent Interaction*.
- Stark, L., & Hoey, J. (2021). The ethics of emotion in artificial intelligence systems. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*.

Facial Computing References

- Girasa, R. (2020). Ethics and Privacy I: Facial Recognition and Robotics. In R. Girasa (Ed.), *Artificial Intelligence as a Disruptive Technology: Economic Transformation and Government Regulation* (pp. 105–146). Springer International Publishing.
- Raji, I. D., Gebru, T., Mitchell, M., Buolamwini, J., Lee, J., & Denton, E. (2020). Saving Face: Investigating the Ethical Concerns of Facial Recognition Auditing. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 145–151.
- Roundtree, A. K. (2022). Facial Recognition Technology Codes of Ethics: Content Analysis and Review. In *Proceedings of the 2022 IEEE International Professional Communication Conference (ProComm)*, 211–220.