

FOR BETTER

FAR

FOR HUMANITY

FOR WORSE

Body of Knowledge:

Proportionality Study: Its Importance and Practical Application

Version 1.0 - July 2025

Authors: Esther Y. Chung and Paul Crafer

NEAR

ALL THE TIME

<https://forhumanity.center/>

This document explains the purpose, structure, and implementation of a Proportionality Study as required by ForHumanity's audit frameworks. It explores how proportionality serves as a safeguard for rights and freedoms, and how organisations can apply it to ensure AAA Systems are ethically justified, legally compliant, and socially responsible.

Introduction to ForHumanity

ForHumanity's mission is to mitigate downside risks posed by AI, algorithmic, and autonomous (AAA) Systems to humans. [ForHumanity](#), a 501 (c) 3 tax-exempt public charity, endeavours to be a beacon, examining the impact of AAA Systems on jobs, society, our rights, and our freedoms. We focus on mitigating risk in the areas of ethics, bias, privacy, trust, and cybersecurity at the corporate and public policy levels, always on behalf of and for humanity.

Overview

The decision to design and develop, or procure, an AAA System to achieve a specific purpose is among the key challenges facing organisations today. Much performative hype surrounds AI technologies, and it is expressed in various national strategy documents and frames AI infrastructure investment and rapid adoption as both necessary and inevitable for their respective countries.¹ This hype is now driving many organisations to act without fully considering the implications of those decisions for individuals, communities, and societies. In this process of rationalising the deployment and operation of a developed or procured AAA System, organisations need to understand how their choice of AI model and the purpose assigned to it must be proportional to the risks to and sacrifices of the rights and freedoms of individuals or groups its operation may impact. This is why ForHumanity requires that a Proportionality Study be conducted in each of its system governance (e.g., EU AI Act, CORE AAA System Governance) and privacy protection (e.g., EU

GDPR, DIFC Reg10, DPDP Act) Certification Schemes.

This document provides an understanding of how the principle of proportionality, as a fundamental concept in law, has become a key part of the field of Algorithm Ethics, which itself is a domain of expertise that informs instances of Ethical Choice within the deployment of AAA Systems. It is the ethical dilemmas inherent in the tensions and tradeoffs of deploying an AAA System that a Proportionality Study seeks to address and clarify, and in doing so, inform the decisions taken in other areas of responsible AI governance for an AAA System, such as risk management, bias mitigation, and design choices.

A Proportionality Study shall be performed by organisations before deploying an AAA System and by the Ethics Committee or equivalent specialised and standing body within the organisation.

The principle of Proportionality

*"Proportionality is a general principle of EU law. In the context of fundamental rights, such as the right to the protection of personal data, proportionality is key for any limitation on these rights. More specifically, proportionality requires that advantages due to limiting the right are not outweighed by the disadvantages to exercise the right. In other words, the limitation on the right must be justified."*²

Proportionality, as means-end testing, is a well-established principle in law designed to resolve conflicts between legal norms. Historically, it arose

A Proportionality Study assesses the tensions and tradeoffs between risks to and sacrifices of the rights and freedoms of individuals or groups, balanced against the potential benefits and gains to an individual or group.

out of the Aristotelian concept of justice³ and can now be found in numerous national legal systems around the world. In these frameworks of law, the principle of proportionality serves fundamentally as a safeguard against the excessive exercise of governmental or intergovernmental authority. It is particularly relevant in the context of human rights, specifically those qualified rights that can legitimately be limited to fulfill critical public interests.

In terms of AI development and deployment, reframing the proportionality principle specifically for AAA Systems necessitates ensuring that those systems do not surpass what is strictly required to fulfill legitimate and ethically defensible objectives.

UNESCO was one of the first intergovernmental organisations to create normative guidance on applying ethics in AI when it published *Recommendation on the Ethics of Artificial Intelligence* (2021).

*“The choice to use AI systems and which AI method to use should be justified in the following ways: [...] the AI method chosen should be appropriate and proportional to achieve a given legitimate aim [...]”*⁴

In the Principle of Proportionality and Do No Harm, UNESCO stresses that AI systems should be used strictly in pursuit of their intended objectives, without exceeding what is necessary. In practice, this means carefully weighing harms, such as risks to privacy or environmental impacts from data-intensive methods, against the anticipated

social benefits. To this end, conducting a Proportionality Study offers a structured way to navigate these tensions through a systematic balancing process, ensuring that AAA Systems remain both ethically sound and contextually proportionate. A proportionality-driven approach not only justifies the selection and deployment of specific AAA Systems but also reinforces transparency and accountability in AI governance.

This approach aligns directly with ForHumanity’s audit requirements. By mandating that a Proportionality Study is conducted, these binary certification schemes ensure that the principle of proportionality is integrated into the decision-making process for AAA Systems and is an important resource in the context of other governance decisions, policies, and procedures throughout the AAA System lifecycle. Ultimately, this promotes fairness, strengthens risk controls, and supports the responsible balancing of competing societal interests.

AI systems should be used strictly in pursuit of their intended objectives, without exceeding what is necessary. This requires maintaining a critical balance between the means employed and the ends sought.

Who are the key stakeholders?

Key stakeholders involved in a Proportionality Study across the CORE, EU AIA, and GDPR schemes include Specialty Committees, individual stakeholders, the organisation itself, and multi-stakeholder input. Specialty Committees, such as the Ethics Committee and the Algorithmic Risk Committee, oversee conducting, documenting, and integrating the Proportionality Study into ethical evaluations, risk management decisions, and system designs. Individual stakeholders, specifically the

Data Lead, leverage insights from the Proportionality Study when making essential decisions about handling sensitive data, aiming to reduce bias and enhance fairness, particularly for vulnerable groups.

The organisation is tasked with producing, documenting, and applying the outcomes of Proportionality Studies in critical processes, such as Data Protection Impact Assessments and Necessity Assessments. Additionally, recognising the importance of multi-stakeholder input, ForHumanity Audit Certification schemes encourage engaging external advisory groups, intended users, and other potentially impacted stakeholders. Such engagement facilitates diverse perspectives on possible risks and consequences, ensuring comprehensive identification and management of ethical, societal, and operational concerns

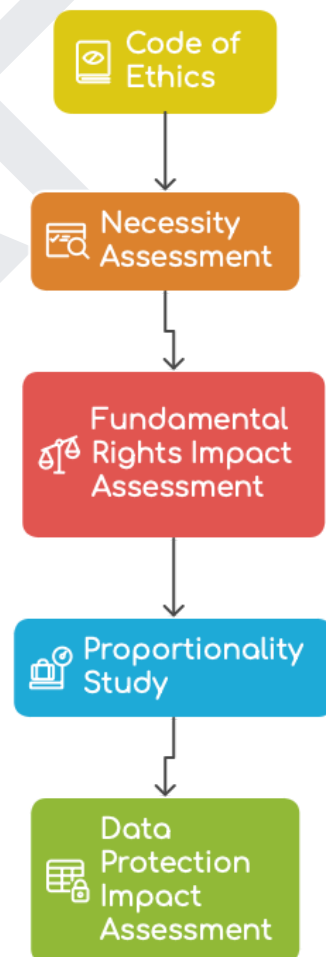
How does the Proportionality Study integrate with other assessments?

A Proportionality Study is conducted in consideration of other assessments that have come before it and informs the Data Protection Impact Assessment that comes after it. The order of operations for conducting the interconnected reports and assessments is as follows:

- The [Code of Ethics](#) documents the shared moral framework of the organisation and is compiled in the Expert Oversight section.
- The **Necessity Assessment** examines if an AAA System is the only or best solution and if each item of Personal Datum is vital. It is conducted by the Algorithmic Risk Committee in the Data Management section.
- The **Fundamental Rights Impact Assessment** examines how an AAA System interacts with the

rights and freedoms guaranteed to AI Subjects. It is compiled by the Ethics Committee in the Ethical Oversight section.

- The **Proportionality Study** is a documented study conducted by the Ethics Committee in the Ethical Oversight section.
- The **Data Protection Impact Assessment** assesses the data protection risks in the processing of personal data. It is conducted by the Algorithmic Risk Committee in the Risk Management section.



How do you conduct a Proportionality Study?

Despite being a staple of jurisprudence in courts worldwide, the proportionality analysis lacks a canonical formulation and carries some variation across jurisdictions. The following section borrows from a four-stage sequential inquiry developed by the German Federal Constitutional Court: (1) a preliminary stage of defining a legitimate objective, followed by a three-pronged proportionality test consisting of (2) Suitability Test, (3) Necessity Test, and (4) Strict Proportionality (Balancing) Test.⁵ These four stages are as follows:

1. Legitimate Objective: The means should promote legitimate ends.

Articulate the legitimate objective, or the end that the chosen means (such as action, decision, or policy) is intended to achieve. The objective must be lawful and legitimate.

Define the means, including its scope, duration, and method of operation.

Identify all expected benefits and potential harms for individuals, society, or the organization from implementing the chosen means. Consider all impacts relating to ethics, bias, privacy, trust, and cybersecurity, as well as infringements on rights, economic burdens, and social disruptions. Take into account both direct and indirect, as well as both short-term and long-term, impacts. Use a structured risk taxonomy and incorporate Diverse Input and Multi-Stakeholder Feedback (DI&MSF) to identify a comprehensive set of risks.

Example use case: A city government proposes deploying an AI-powered facial recognition camera system in public parks.

Legitimate objective: To significantly reduce

violent crime within city park areas and enhance public safety.

Means: The city plans to install specific camera models across 15 high-traffic parks for a two-year pilot, using real-time AI analysis for threat detection.

Benefits: Increased public safety, faster law enforcement response to incidents, potential deterrence of criminal activity.

Harms: Significant privacy infringement for park visitors, potential for algorithmic bias leading to disproportionate surveillance of certain demographic groups, risk of misidentification, potential chilling effect on freedom of assembly, high implementation and maintenance costs.

2. Suitability Test: The means should be suitable for achieving the desired ends.

Determine whether there is a rational connection. Does a clear, rational connection exist between the means's design and the objective? Can this ~~means~~ genuinely contribute to achieving the intended objective?

Example use case: Does the proposed facial recognition system actually help reduce violent crime in parks?

Rational connection: Researchers would need to evaluate the system's technical accuracy in varied conditions (lighting, crowd density) and assess whether similar deployments in other cities genuinely correlated with a reduction in the specific type of crime the city targets, rather than just displacing it. If the system frequently misidentifies people or struggles in common park environments, it might not be suitable.

3. Necessity Test (Least Restrictive Means Test): A less restrictive means should be used if it is equally effective.

Explore alternatives and determine whether there are less restrictive means. Could the objective be achieved utilizing alternative means that are less intrusive, less harmful, or equally effective? For instance, could risks or harms be reduced by using technological alternatives (e.g., other AI models, data processing methods, or system architectures), process alternatives (e.g., non-AI or human-in-the-loop), or mitigation strategies?

Justify the chosen means. If the chosen means is not the least restrictive, provide a rational argument for why less restrictive means are not sufficient or feasible.

Example use case: Are there less intrusive ways to make parks safer than using biometric facial recognition systems?

Alternatives: Increased human police patrols, improved park lighting, community engagement programs, non-AI surveillance cameras with human monitoring for real-time alerts, or public awareness campaigns on crime prevention.

Justification (if facial recognition is chosen): The city might argue that increased human patrols are too expensive or insufficient for the scale of the problem. They might claim non-AI cameras lack the real-time identification capabilities needed for rapid response to violent incidents, making them less effective for their specific objective of immediate crime reduction and apprehension.

4. Strict Proportionality (Balancing) Test: The means should not be disproportionate to the ends.

Weigh benefits vs. harms by balancing, on the one

hand, the *importance* of the objective and the *magnitude* of the expected benefits, and on the other, the *severity*, *likelihood*, and *detectability* of the potential harms. Even if a particular means is suitable and necessary, it must still be proportionate in its overall impact.

Consider the context. Take into account the specific circumstances, including the vulnerability of affected individuals, the duration of the action, decision, or policy, and the potential for unintended consequences.

Finally, determine the acceptability of the residual risk. Determine if the burden imposed by the chosen means, especially the residual risk (the risk that remains after mitigation), is acceptable given the benefits, ensuring that the burden is not excessive in relation to the objective.

Furthermore, your organization can empower end-users to make this determination for themselves by informing them of the residual risk through clear and transparent documentation.



Example use case: Do the public safety benefits of facial recognition in parks outweigh the costs to privacy and civil liberty?

Balancing: The city weighs the high importance of reducing violent crime against the severe privacy infringements and potential for bias-driven harms (e.g., innocent individuals being wrongly flagged, or specific communities feeling targeted). They consider the system's accuracy rates versus the harm of a false positive.

Context: Is this a universally high-crime area, or are particular communities disproportionately affected by both the crime *and* the surveillance? What is the duration of surveillance? Could it lead to unintended effects like reduced park usage or increased social unrest?

Residual risk: Even with strict data retention limits and human oversight for alerts, a privacy risk remains. Is this remaining risk acceptable to the public given the demonstrated reduction in violent crime, and have citizens been clearly informed about these trade-offs?

How do you document the study for transparency and accountability?

Because the documentation for a proportionality study must be verifiable by independent auditors and understood by stakeholders, proper documentation is critical. Key elements include:

1. Executive Summary

Provide a clear, concise overview of the study's purpose, the means (such as action, decision, or policy) assessed, key findings, and the final determination on whether the means is proportionate to the ends or not.

2. Introduction

Describe the purpose of the study, clearly stating why the proportionality study was conducted.

Reference any relevant laws, regulations, or policies that necessitate or guide the study.

Identify all stakeholders (both individuals and groups) who will be affected by the chosen means.

3. Methodology

Describe the steps followed in conducting the study (e.g., data collection methods, analytical frameworks used).

List all sources of data, evidence, and expert opinions considered, including research papers, internal reports, stakeholder consultations, legal opinions, bias audits, and DI&MSF.

4. Analysis

Legitimate Objective: Clearly articulate the objective and the precise nature of the means assessed. Exhaustively list all expected benefits and potential harms, supported by data, evidence, and qualitative assessments. Where possible, quantify impacts (e.g., bias metrics, privacy breach severity).

Suitability Test: Present evidence demonstrating the ability of the means to achieve the intended objective.

Necessity Test: Discuss all considered alternatives and explain why certain alternatives were rejected. Furthermore, explain why the chosen means, after mitigations, is the least restrictive to achieve the objective. If it is not the least restrictive, provide justification for its use.

Strict Proportionality (Balancing) Test: Explain how you weighed the benefits against the harms.

Detail any mitigating measures proposed to reduce potential harms. Discuss the residual risk that remains after mitigation and why it is deemed acceptable.

5. Conclusion

State your final determination on whether the means is proportionate to the ends or not, and why.

6. Recommendations

Offer any recommendations for ongoing monitoring, additional safeguards, or future development to maintain or enhance proportionality.

7. Appendices

Include any supporting documentation, such as:

All data used in the assessment (e.g., bias detection metrics, privacy impact assessments, performance benchmarks).

A comprehensive log of identified risks, their assessment (severity, likelihood, detectability), and proposed or implemented mitigation measures, aligned with ForHumanity's risk management framework.

All stakeholder feedback, such as consultation processes and DI&MSF, and how you addressed it in the study.

Technical documentation, such as relevant system architecture, design specifications, and code documentation that support the claims made in the study.

Audit Trails, such as records of decisions made, changes implemented, and approvals obtained throughout the system's lifecycle, enabling traceability.



Summary of key principles

The rapid spread of AAA Systems calls for a strong ethical framework. As such, the Proportionality Study is a critical inquiry into the trade-offs inherent in deploying AAA Systems, balancing potential benefits against risks to human rights and freedoms.

Ideally conducted by an organization's Ethics Committee, the Proportionality Study integrates with other important governance assessments to provide a holistic view of an AAA System's ethical impact throughout its lifecycle.

The study itself involves four main sequential stages: (1) Define a legitimate objective, (2) Conduct a suitability test, (3) Conduct a necessity test, and (4) Conduct a strict proportionality (balancing) test. Following these four stages leads to a final determination of whether the means is proportionate to the ends or not.

Finally, comprehensive documentation is key for transparency and accountability because it allows

for independent third-party audits, a cornerstone of ForHumanity's mission to build an “infrastructure of trust” for AI. By embedding proportionality into its certification schemes, ForHumanity ensures that it remains a central, verifiable element of responsible AI, ultimately promoting fairness and balancing societal interests on behalf of humanity.

Proportionality in Practice: A ForHumanity FAQ on Auditing AAA Systems

1. How does ForHumanity ensure that the use of AAA Systems does not exceed what is necessary and that individuals are equipped to make informed proportionality assessments?

ForHumanity addresses the practical implementation of proportionality through a framework that embeds [Diverse Inputs & Multi-Stakeholder Feedback \(DI&MSF\)](#) and Expert Oversight mechanisms throughout the lifecycle of an AAA System. DI&MSF pools or risk assessors are composed of individuals with diverse backgrounds, lived experiences, and domain expertise, and are involved at all points of the AI lifecycle including design, development (or procurement), deployment, and decommissioning. These contributors help ensure the proportionality of actions by offering real-world perspectives on necessity, potential misuse, and unintended impacts.

Expert Oversight is delivered via designated standing committees such as the Algorithmic Risk Committee and Ethics Committee. These bodies are trained to evaluate systemic risks such as disproportionate impacts on marginalized groups, automation bias, and ethical concerns. They assess proportionality using tools such as the Proportionality Study, Business Rationale Report,

and specialised tools included with the [ForHumanity Risk Management framework](#) such as [Residual Risk logs](#). These reviews are grounded in both ethical and legal frameworks, ensuring transparency and accountability in decision-making.

2. How does the ForHumanity Independent Audit of AI Systems framework ensure that AAA Systems remain proportionate throughout their lifecycle, and what mechanisms are used to detect when proportionality may be exceeded?

To assist providers and deployers of AAA Systems to maintain proportionality over time, ForHumanity mandates a dynamic system of metrics, measurements, and thresholds embedded in its continuous monitoring and Algorithmic Risk Assessment processes. These include:

- Thresholds for Key Risk Indicators (KRIs) and Key Performance Indicators (KPIs)
- Incident reporting via an Adverse Incident Reporting System
- Feedback loops that feed directly into risk reassessments

When thresholds are exceeded, automated triggers initiate actions such as change impact assessments and comprehensive risk reassessments. The feedback mechanism ensures real-time updates to risk controls, and potential decommissioning procedures are in place if a system no longer aligns with its intended purpose or proportional use.

Importantly, continual monitoring ensures the proportionality assessment is not just a one-time event. The system is designed to evolve, with

regular reviews of metrics and the integration of emergent risks into the governance framework.

3. What is the reason ForHumanity does not recommend a purely metrics-based approach to proportionality assessment?

ForHumanity does not suggest a purely metrics-based or formulaic approach to proportionality because numbers can create a misleading sense of certainty. A single threshold or score cannot capture the fine-grained judgments that arise in the “middle ground,” where intrusions are brief, harms are hard to predict, or benefits and risks are delicately balanced. In such cases, a model might clear the numeric bar while its real-world impact remains disproportionate.

Proportionality also turns on context. Each deployment has a distinct purpose, audience, data environment, and legal or social backdrop; what is justifiable in one setting may be unacceptable in another. Percentages and equations flatten those nuances, masking critical facts that determine whether an intrusion is genuinely necessary and balanced.

Moreover, proportionality must be grounded in evidence, not the elegance of a formula. A traditional deliberative review, first confirming necessity (Necessity Assessment) and then weighing concrete benefits against potential harms (FRIA), guards against the temptation to treat compliance as a box-ticking exercise. Numeric targets can invite “rubber-stamping,” signalling completion the moment a score is achieved even when qualitative red flags remain.

For these reasons, ForHumanity treats quantitative metrics as inputs, never as the sole arbiter. The Proportionality Study rests with its multidisciplinary experts in the Ethics Committee who explicitly consider context and stakeholder perspectives before certifying that an AAA System meets the proportionality requirement.

4. How does ForHumanity address the cumulative effects of multiple AAA Systems or agents, especially in environments where fully autonomous systems are in use?

ForHumanity explicitly excludes fully autonomous systems, those without any form of human-in-the-loop or human-on-the-loop governance, from its certification schemes. This position is grounded in the necessity of human oversight to manage cumulative and systemic risks that arise from the interaction of multiple AAA Systems.


To manage these risks, ForHumanity mandates some form of human oversight responsible for understanding system capacity, monitoring outputs, and intervening when needed. The Algorithmic Risk Assessment process includes mapping systemic riskiness and interdependencies among systems within and across organizations. This includes monitoring growth in data processing and tracking how systems feed into each other.

By compiling and analyzing this information through tools like the [cAIRE Report](#) and [Systemic Societal Impact Analysis](#), ForHumanity enables organizations to detect, assess, and mitigate long-term and compound effects that might otherwise go unnoticed in isolated system reviews.

References

1. Bareis, J., & Katzenbach, C. (2021). Talking AI into Being: The Narratives and Imaginaries of National AI Strategies and Their Performative Politics. *Science, Technology, & Human Values*, 47(5), 855-881.
<https://doi.org/10.1177/01622439211030007>
2. European Data Protection Supervisor.
https://www.edps.europa.eu/data-protection/our-work/subjects/necessity-proportionality_en (Last accessed July 2025)
3. Karliuk, M. Proportionality principle for the ethics of artificial intelligence. *AI Ethics* 3, 985–990 (2023).
<https://doi.org/10.1007/s43681-022-00220-1>
4. UNESCO. 2021. Recommendation on the Ethics of Artificial Intelligence. Paris: United Nations Educational, Scientific and Cultural Organization.
<https://unesdoc.unesco.org/ark:/48223/pf0000381137> (Last accessed July 2025)
5. Lübke-Wolff, G. (2014). The Principle of Proportionality in the Case-Law of the German Federal Constitutional Court, 34 *Human Rights Law Journal*, 12-17.
https://www.researchgate.net/publication/326782433_The_Principle_of_Proportionality_in_the_Case-Law_of_the_German_Federal_Constitutional_Court





Images: <https://betterimagesofai.org>, <https://creativecommons.org/licenses/by/4.0/>
 Cover and back: 'The Omnipresent Tapestry' by Joudy Bourghli & The Bigger Picture
 p.6: 'Surveillance' by Reihaneh Golpayegani & Digit
 p.8: 'AI is...' by Rick Payne and team
 p.11: 'Knowledge Sweatshop' by Leo Lau & Digit

This document is the property of ForHumanity Inc. ©2023, ForHumanity Inc. a 501(c)(3) tax-exempt Public Charity.
 All rights reserved. Creative Commons CC-BY-NC-ND Attribution-NonCommercial-NoDerivs.