ZE Zentrum
VE verantwortungsbewusste
DI Digitalisierung

Centre Responsible Digitality

Research Ethics for AI Research Projects

Guidelines to Support the Work of Ethics Committees at Universities

Version 1.0
February 2023

ZE Zentrum
VE verantwortungsbewusste
DI Digitalisierung

Centre Responsible Digitality

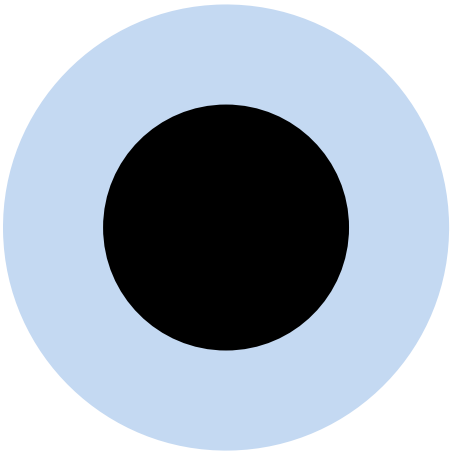

Research Ethics

The ethical evaluation of research projects that involve test subjects, generate/purchase/use personal or otherwise critical data, or that are expected to involve a questionable use of research results ("dual use"), is common in science. In many cases, research funding institutions or publication outlets even explicitly prescribe ethics reviews. Research institutions set up ethics committees for this purpose and define suitable review procedures. In principle, the review is voluntary. Nevertheless, obtaining an ethics vote is typically part of the methodological standard. It should be noted that ethics votes are no substitute for the individual responsibility of the person or persons for the quality and also the ethically relevant aspects of their concrete scientific work.

Research ethics can be seen as a form of quality management. However, it also includes aspects of good scientific practice, professional ethical standards and prevailing law.[1] It can also include the institution's or university's own regulations (e.g. an explicit mission statement or similar or a civil clause) as well as agreements with (and thus legitimate expectations on the part of) cooperation partners as well as normative requirements of a research funding body (e.g. to counteract the so-called gender data gap or to refrain from animal experiments).

"AI" as a subject of assessment

Research in the field of so-called artificial intelligence (AI) is not entirely new. In recent years, however, it has gained a new depth of intervention. This is due to the combination of sophisticated algorithms on the one hand and very large data sets, usually acquired from everyday life, on the other. AI scenarios (including so-called machine learning, "deep" learning) produce - in research initially as models - adaptive solutions that are discovered "autonomously" to a certain degree. These can then be incorporated into a wide range of services and products. The potential range of technological consequences or unintended side-effects (such as social discrimination) as well as of quality deficiencies of AI solutions is therefore exceptionally large in some fields, for example in medicine, finance or security.[2]

Research ethics has some catching up to do here, especially since research ethics standards are generally not yet deeply anchored in the discipline of computer science, which is crucial for AI solutions. There is hardly any " professions culture" here that can be built on. Society is also not yet prepared for the capabilities of AI technology. Legal safety and liability regulations for products that contain AI have yet to emerge.

Many of the questions raised by AI research projects can be answered by means of "classical" research ethics. Some, however, are of a more specific nature.[3] The following paper attempts to formulate criteria for evaluating and assessing such AI-specific problems - especially for research projects from the domain of computer science. The aim is to support ethics committees in their work.[4]

Classical Criteria for the Ethical Assessment of Research

As for other research projects, the following requirements apply to computer science research projects, among others

- the prohibition of physical or psychological harm (as well as the obligation to clarify individual risks) to test subjects [TS]; if ill persons are involved, the vote of a medical ethics committee is required (the standards of the BÄK/ZEKO apply here[5]); the procedure for drug testing is regulated in the German Medicinal Products Act;
- the legal requirements for handling hazardous substances and for high-security laboratories;
- the usual rules (e.g. of the DGP) for the disclosure and consent of test subjects (including the subsequent disclosure of deception experiments);[6]
- the interdisciplinary standards regarding the representativeness of samples and, if necessary, the avoidance of discriminatory effects respectively the disadvantage of vulnerable groups;
- the documentation and archiving obligations required for the purposes of reproducibility as well as traceability of research results, which are appropriate to the selected research method/form of research;
- the provisions of the GDPR on data protection, data economy and purpose limitation of the collection or use of personal data, including the requirement to delete them once they have been used;
- the standards of good scientific practice with regard to methodological competencies of project management, authorship, academic dependencies;[7]
- the standards regarding the transparency of industry cooperation, the publication of results obtained in cooperation with industry partners on the academic side, etc.;
- the obligation to not abuse the reputation of public science for the purpose of commercial advertising;
- the rules and standards of contemporary research data management and quality assurance of data throughout the research cycle

(also from the point of view of digital sovereignty/data sovereignty of science);
- the requirement to refrain from superfluous or foreseeably useless research;
- the legal provisions of the German Criminal Code and the Foreign Trade and Payments Act.

In order to receive an ethics vote, a research project must also be described completely and with sufficient clarity regarding, among other things, methods, division of labour in the team, as well as responsibilities and, if applicable, liability issues.

Criteria for an AI-specific ethical assessment of research

The following AI-specific research ethics issues are recommended to be considered for assessment or consideration. They arise in five respects:

1. AI as the subject of a research process

1.1 Documentation. How are the data, algorithms and the dynamic interaction of the two documented?[8] The characterization of the data (type, amount, origin), the documentation of the software and training methods, the specification of the hardware, the documentation of relevant guidelines, etc., and in the case of so-called "high-risk systems"[9] also the automatically generated documentation of the computing process are required as a minimum.
If the documentation is not sufficient, a request for improvement or a negative ethics vote is appropriate.

1.2 Reproducibility. What are the inherent limits of traceability, reproducibility and predictability of the "behaviour" of the AI solution - and how does the research design or the proposed project account for the associated fuzziness/quality deficiencies? An ethics application must provide substantive answers to these questions.
If the research design does not include test strategies for the reproducibility of the problem-solving process (if necessary, with the implementation of relevant tests), an ethics vote demanding improvements or a negative ethics vote is appropriate.

1.3 Reusability, Research Data Management. How are the parts of an AI solution archived for subsequent use? The minimum requirement is that all documentation including AI-specific metadata be archived (such as risk classes, EU-conformity) as well as ethics applications and votes. In addition to compliance with the FAIR standards[10] ("Findable,

Accessible, Interoperable, Reusable"), metadata must also contain the provenance of the essential components of an AI solution or experiment.[11]
If the required metadata is not provided and its archiving is not planned, an ethics vote requiring improvements or a negative ethics vote is appropriate.

1.4 Data Protection. How is direct or indirect (re-)identification prevented when using anonymized personal data? How is compliance with the GDPR ensured when using personal data? When personal data was collected, was consent given to its use or is there an overriding public interest in the use of the data? Is the processing of personal data necessary for scientific reasons? Is the principle of data minimization (Art. 5 GDPR) followed and are sufficient security measures taken during archiving (IT security, preventive measures against sabotage, data theft, etc.)? An ethics application must address these issues conclusively.[12]
If a differentiated explanation of conformity with the GDPR, Art. 5 - 23 (including the exclusion of model inversion) is missing, also in view of Annex IV of the EU-KI Regulation and the relevant system data protection measures recommended by various parties, an ethics vote demanding improvements or a negative ethics vote is appropriate.

1.5 Effectiveness of Consent. In cases where trial subjects or data providers authorize the use of data, are the information sheets in general as well as the description of the type and objectives of the AI solution used sufficiently comprehensible? Does the information sheet contain sufficient information about the purpose of the processing, the further use and the time of data deletion?
If a sufficiently (to standard) understandable explanation and comprehensibility required for an effective authorization is missing, an ethics vote demanding improvements or a negative ethics vote is appropriate.

1.6 Data Quality. What is the type and quality of the data used? Is the dataset suitable for the "training" of algorithms? How do the quality of

the data on the one hand, the application of the data or the lack of knowledge about it on the other, affect the results of the applied methods and ultimately the quality of the research/research results?[13]
If no transparency is established about the nature of data, in particular of training data, whose characteristics have an impact on the quality of research results (e.g., as a cause of unnoticed biases), or if precautionary measures to avoid biases are not yet in place, an ethics vote demanding improvements or a negative ethics vote is appropriate.

1.7 Data purchase/"Supply Chains". Is the provenance of purchased data sets clarified and have these data sets been obtained under conditions that are justifiable in terms of research ethics (i.e., at least the European standards, ideally the standards that apply in Germany)? Proof of this may be required.
In the absence of reliable proof of provenance (or certificates) for acquired datasets (regarding compliance with European and national security standards as well as regarding the justifiable conditions of their acquisition in terms of research ethics), a request for improvement or a negative ethical vote is appropriate.

1.8 Algorithms. What are the characteristics of the algorithms used? Does the type of algorithm used involve typical risks, especially in the chosen methodological context and with regard to the data domain (e.g., high rate of not only false negatives but also false positives, easy irritability by surprising data)?
In the absence of a convincing typification of the strengths and weaknesses of the algorithms used, an ethics vote calling for improvements or a negative ethics vote is appropriate.

1.9 Non-Discrimination. Can the characteristics of training data and/or algorithms result in discrimination effects or reinforce existing discrimination (of individuals, social groups, etc.)?[14] Do the datasets represent the actual environment relevant for the research question?
In the absence of precautionary measures at the data and/or model level to prevent discrimination or reinforcement of existing

discrimination against vulnerable groups, an ethics vote requiring remedial action or a negative ethics vote is appropriate.

1.10. Sustainability. Are public research funding resources being used efficiently? How expensive is the planned computational effort/time? Are any economic or ecological gains in effectiveness possible?[15]
If, in the case of a costly computer deployment, the resources used are not accounted for or if this deployment of resources appears suboptimal, an ethics vote demanding improvements or a negative ethics vote is appropriate.

2. Using proprietary AI "tools" in the research process

2.1 Proprietary Toolboxes, "Blackboxing". Are proprietary AI components/AI "tools" used in the context of a research project and to what extent do these result in unrecoverable intransparencies for the research process (so-called blackboxing problems)? The ethics application must describe the importance of the specific parts of the research workflow that are beyond the knowledge and access of the researchers, why quality-assured alternatives developed by the scientific community cannot be used, and how the researchers ensure the validity of possible results despite the use of proprietary AI tools.
In case of missing documentation of intransparencies regarding proprietary tool-boxes as well as statements on how the validity of the results can be guaranteed despite black-box effects, an ethics vote demanding improvements or a negative ethics vote is appropriate.

2.2 Data Protection and Data Integrity on the Part of Providers of AI Tools for Science. Can the providers of analytics services based on commercial or otherwise proprietary data/software guarantee compliance with European law and European standards? The providers of research tools must also adhere to the GDPR and ensure data processing within the EU. This requires contracts, not just information in the "www".
If the providers of AI research/analysis tools do not document compliance with European standards in a legally sound form, an ethics vote demanding improvements or a negative ethics vote is indicated.

3. AI solutions in a domain predominantly characterized by research questions which are not computer science driven

3.1 Sector-Specific Research and Application. In which scientific domain[16] or in which social "sector" (e.g. medicine, mobility, finance, education, security) are the research results to be used and how does the type of application domain affect the risk balance of the planned research or the expected risk class of possible applications?[17] Do unexpected risks arise in the transfer of research results/solutions to other domains or sectors? Is there a need to limit the use of the researched/developed AI solution to certain domains or sectors as a field of application? Can this be implemented by design?
In the absence of information on the sectors in which the research results will be used and on the possible risks of their transfer to other sectors, risk assessments in the field of application-oriented AI research are incomplete. In such cases, an ethics vote demanding improvements or a negative ethics vote is appropriate.

3.2 Domain-Specific Research Competence. Domain-specific research questions may require domain-specific AI solutions. Does the research team have sufficient domain knowledge or methodological expertise tailored to the domain in question? Where researchers conducting the project (e.g., computer scientists, data scientists) are expected to have too little knowledge of the specifics of the (data) dimensions or "use cases" being researched, interdisciplinary participation of scientists with expertise in the analysed domain (e.g., social scientists, economists, medical scientists) is necessary.
If the necessary additional expertise of other disciplines (specialized knowledge, methodological competence) is lacking in view of domain-specific problems (e.g., in use cases), an ethics vote demanding improvements or a negative ethics vote is appropriate.

3.3 Interdisciplinary Teams and Interdisciplinary Publication. AI research is undergoing rapid evolution. Against this background, it must be ensured that research results in domains predominantly

characterized by non-computer science research are published in such a way that they can also be reviewed (critically, if necessary) in the research culture of the domain.[18] Interdisciplinary AI research teams should include statements on the publication strategy in ethics applications. Broadly visible publications (taken up by computer science and the domains involved) are preferable to exclusive publications in very small (sub)communities.
If a publication strategy limited to only AI-internal scholarly communication is envisaged for AI research with an expected high impact on domains other than computer science, an ethics vote demanding improvements or a negative ethics vote is appropriate.

3.4 Technology Impacts, "Society". Assessments of possible social consequences of the use of an AI solution usually require not only psychological expertise, but also the expertise of social scientists. For questions that do not concern the interface to individual users ("usage research"), but rather societal questions, the participation of social scientists needs to be envisaged.
If an assessment of the societal consequences of the use of AI systems is envisaged without involving social science or normative competence, an ethics vote demanding improvements or a negative ethics vote is appropriate.

4. Requirements to be met for future product features (e.g. "EU conformity") in application-oriented AI research.

4.1 Risk Classes/High-Risk Systems. In products of which risk class and in which areas of application will the research results be used most likely? Reference to the risk class system of the specific areas in Annex III of the draft version of the EU-AI Regulation[19] should already be made in the research process, if necessary, because additional requirements for the research process may arise here (e.g. [level 3 and 4] documentation of the names of all persons involved in the research process).
If an explicit reference to the EU risk classification (App. III) is missing in the planning/definition of expected product features ("performance

specification"), an ethics vote demanding improvements or a negative ethics vote is appropriate.

4.2 EU Conformity. The EU approval of AI systems also requires EU conformity beyond the risk categories. Controllability is one of the legal requirements for future AI products. This includes risk management (amount of damage, probability of damage, disponibility of risks) and the question of violations of fundamental rights (e.g. discrimination). The ethical justifiability of research may be diminished if an application-oriented development cannot be implemented in conformity with EU standards and thus cannot reach product maturity in the EU.[20]

4.3 Resilience. Does the research include the question of whether the AI solution can be (too) easily irritated in practice in the products in which it will be used in the future? Has it been clarified how severe the consequences would be?
Does the research design include planning about the resilience of the systems to be developed? Where this aspect is lacking in application-oriented research working towards concrete products, an ethics vote demanding improvements or a negative ethics vote is appropriate.

4.4 Explainability. Is the level of "explainability" that is required for product development/future approval ensured for the research project, its goals and in particular for the type of AI being used?
If such an explainability is not ensured in accordance with Annex IV of the EU-AI Regulation in conjunction with relevant delegated certification requirements (e.g. VDE SPEC[21]), an ethics vote demanding improvements or a negative ethics vote is appropriate.

4.5 Digital Sovereignty/Option of Non-Use. Do the AI solutions being researched contribute perspectively (possibly in indirect ways) to lock-in effects and other forms of dependency upon certain technologies or an inevitability of surveillance on the market, which can no longer be avoided in everyday life? Do development paths exclude options of non-use of emerging new technologies?[22]
If a development path taken by the research restricts or makes

impossible the option of non-use of systems ("lack of alternatives" of use), an ethics vote demanding improvements or a negative ethics vote is appropriate.

5. AI and Dual Use Constellations

5.1 Dual Use. What "dual use" problems can be expected with the planned research and the AI solutions that are developed (further) in application on the basis of this research? Dual-use constellations arise quickly in AI research because governments worldwide are currently investing in AI-based defence and security technologies. The examination of dual-use options for the results of a research project therefore concerns not only warfare-related research and research paths of research funding adjacent to military scenarios (such as disaster prevention, crime fighting), but also research for other domains.[23] Research in Germany must minimize the likelihood of using its findings for non-peaceful or otherwise unconstitutional purposes. The question is: How can undesirable uses (e.g., unpeaceful, threatening fundamental rights, compromising democracy and market stability, or ecologically dangerous) be prevented, if possible, or made less likely or less attractive?
If dual-use constellations are not reflected properly, an ethics vote demanding improvements or a negative ethics vote is appropriate.

5.2 Civil Clauses. If an institution has adopted a civil clause, the compliance of planned research with the wording of this civil clause may be subject of the ethics review.[24]
If a violation of the civil clause is found, an ethics vote demanding improvements or a negative vote is appropriate.

Further Remarks for Applicants

"Basic research"? Research on AI methods, especially in the field of algorithms, is currently often (still) of a "generic" nature; basic principles are being researched. However, due to the inherently essential interplay with concrete data or use cases, research ethics is also relevant in the realm of so-called basic research.

Timing. A research-ethical evaluation of projects/scenarios in the field of "AI" can be complex. Ethics committees should therefore be asked for a vote at an early stage. Part of good research planning is to provide the appropriate lead time for this.

Responsibilities. In principle, the ethics committees of the research institutions at which the researchers work are in charge. For medical research projects, the vote of a medical ethics committee is required. Especially in collaborative research, non-medical parts of a research project may have to be reviewed separately. A list of non-medical ethics committees ("KEF") is maintained by the German National Academy of Sciences Leopoldina.

Consultation. Advisory capacities on AI research-ethics are currently only being established in the institutions of the science system. Researchers are advised to contact the ethics committees of their respective institutions. The Center for Responsible Digitization (ZEVEDI) [office[at]zevedi.de] offers networking opportunities to chairpersons, members and administrative staff of ethics committees.

February 2023,
ZEVEDI Project Group
Regulatory Theories of Artificial Intelligence (NOKI)

Notes

[1] In case of general interest in the topic, the following is recommended to round off the picture of an ethics practice that is, however, late in adapting to the problems of digitality: Association of Medical Ethics Committees in Germany (AKEK), Journal of Academic Ethics, Jahrbuch Wissenschaft und Ethik.

[2] Cf. the Report on the Safety and Liability Implications of Artificial Intelligence, the Internet of Things and Robotics by the European Commission, Brussels 19.02.2020.

[3] A classification of AI for scientific research is provided by Gethmann et al.: Künstliche Intelligenz in der Forschung. Neue Möglichkeiten und Herausforderungen für die Wissenschaft, Berlin 2022.

[4] In its report, the Data Ethics Commission formulates recommendations for dealing with data and algorithmic systems that are not, however, specific to scientific research. Datenethikkommission der Bundesregierung: Gutachten der Datenethikkommission, Berlin 2019.

[5] Zentrale Kommission zur Wahrung ethischer Grundsätze in der Medizin und ihren Grenzgebieten (Zentrale Ethikkommission, ZEKO) at the German Medical Association: Statements.

[6] Berufsverband Deutscher Psychologinnen und Psychologen e.V. / Deutsche Gesellschaft für Psychologie e.V.: Berufsethische Richtlinien, Berlin 2016.

[7] Deutsche Forschungsgemeinschaft (DFG, German Research Foundation): Leitlinien zur Sicherung guter wissenschaftlicher Praxis. Codex, Bonn 2019.

[8] General requirements for documentation as part of good scientific practice are named in guideline 12 of the DFG Codex (loc. cit.).

[9] Like various AI ethics expert reports before it, AI Regulation of the European Union also divides "AI systems" into risk classes. The risk-based approach distinguishes applications of AI into those that pose i) unacceptable, ii) high, or iii) low or minimal risk. Cf: Proposal for a Regulation of the European Parliament and of the Council laying down harmonized rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts, Brussels, 21.04.2021, p. 15. In addition, Art. 5 para. 1 prohibits certain practices in the field of artificial intelligence, such as techniques of subliminal influence, the exploitation of weaknesses of certain groups of persons, the assessment or classification of the trustworthiness of natural persons by public authorities, or real-time remote biometric identification systems in publicly

accessible spaces for law enforcement purposes. Art. 6 refers to Annexes II and III for classification as a high-risk AI system. Annex II of the EU-AI Regulation lists the legislation on harmonised standards for products under which AI systems are classified as high-risk. Annex III of the EU CI Regulation distinguishes for the classification of high-risk AI systems approximately the following areas: 1) biometric identification and categorization of natural persons, 2) management and operation of critical infrastructures, 3) education and vocational training, 4) employment, workers management and access to self-employment, 5) access to and enjoyment of essential private services and public services and benefits, 6) law enforcement, 7) migration, asylum and border control management, 8) administration of justice and democratic processes.

For high-risk AI systems, the logging functions under Art. 12(4) must capture at least the following information: a) record of each period of use of the system, b) the reference database against which the system matches the input data, c) the input data for which the query resulted in a match, and d) the identity of the natural persons involved in the verification. Annex IV of the EU AI Regulation precisely lists the technical documentation requirements for high-risk AI systems. Annexes to the Proposal for a Regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (AI Act) and amending certain Union legislative acts, Brussels, 21.04.2021.

[10] Cf. the explanations on guideline 13 of the DFG Codex of Good Scientific Practice "Establishing Public Access to Research Results" (op. cit.) as well as the detailed description of the principles of the FAIR initiative.

[11] Cf. Guideline 5 "Nachnutzen und Reproduzieren" in the position paper of the German Science and Humanities Council: Zum Wandel in den Wissenschaften durch datenintensive Forschung, Cologne 2020, p. 42f.

[12] Cf. also the statement of the German Ethics Council: Big Data und Gesundheit – Datensouveränität als informationelle Freiheitsgestaltung, Berlin 2017.

[13] Cf. German Council for Scientific Information Infrastructures (RfII): The Challenge of Data Quality. Recommendations for Sustainable Research in the digital Turn, Göttingen 2020.

[14] In the UNESCO Recommendation on the Ethics of Artificial Intelligence, paragraphs 28-30 relate the principles of fairness and non-discrimination to the goal of a just world with respect to information, communication, culture, education, research, and socio-economic and political stability; UNESCO: Recommendation on the Ethics of Artificial Intelligence; Paris, Nov. 22, 2021, p. 9.

[15] UNESCO's Recommendation on the Ethics of Artificial Intelligence, paragraph 31, calls for a comprehensive understanding of the implications of AI technologies on the various dimensions of sustainability against the

backdrop of the goal of sustainable societies; UNESCO: Recommendation on the Ethics of Artificial Intelligence; Paris, 22.11.2021, p. 9.

[16] The EU legislation currently in the process of being updated as well as the emerging EU infrastructures (e.g. GAIA X) call for the design of sectors, domains and data spaces, which are also to be distinguished normatively. Cf. also note 9.

[17] Cf. note 9.

[18] In the UNESCO Recommendation on the Ethics of Artificial Intelligence, paragraph 110 calls for the promotion of interdisciplinary research on and with the help of AI; UNESCO: Recommendation on the Ethics of Artificial Intelligence; Paris, 22.11.2021, p. 21.

[19] Cf. note 9.

[20] Annex V of the EU-AI Regulation lists the information that must be provided for an EU declaration of conformity.

[21] Cf. the first draft for VDE SPEC 90012: VCIO based description of systems for AI trustworthiness characterisation, Offenbach am Main, 24.05.2022.

[22] UNESCO's Recommendation on the Ethics of Artificial Intelligence calls for the provision of an optional use of AI systems in para. 20; UNESCO: Recommendation on the Ethics of Artificial Intelligence; Paris, 22.11.2021, p. 7.

[23] The German Research Foundation and the German National Academy of Sciences Leopoldina have developed recommendations for dealing with security-related research against the background of the dual use problem: Scientific Freedom and Scientific Responsibility. Empfehlungen zum Umgang mit sicherheitsrelevanter Forschung, Bonn/Halle (Saale) 2014.

[24] The German National Academy of Sciences Leopoldina maintains a list of committees responsible for ethics in security-relevant research ("KEF"): Contacts and Commissions in Germany Responsible for Ethics of Security-Relevant Research. The DFG and Leopoldina also provide a model statute for committees responsible for ethics in security-relevant research, which identifies issues in the area of security-relevant research that require regulation.

ZEVEDI

ZEVEDI identifies and discusses responsibility as a crucial yet uncertain aspect of technological development and aims at making responsible digitality conceivable as well as practically feasible.

ZEVEDI engages in research projects, promotes the transfer of scientific knowledge into society and the economy and provides research-based policy advice on the topic – for a digital transformation guided by a democratic and humane orientation.

ZEVEDI is supported by the Hessian Minister for Digital Strategy and Innovation.

Contributors

AI Research Ethics Working Group

Dr. des. Andreas Brenneis
Technical University of Darmstadt, Editorial Office

Prof. Dr. Petra Gehring
Technical University of Darmstadt, Head of WG

Prof. Dr. Christoph Hubig
Technical University of Darmstadt

Annegret Lamadé
Philipps Universität Marburg

Prof. Dr. Florian Möslein, LL.M. (London)
Philipps Universität Marburg

Thanks to

The AI Research Ethics Working Group would like to thank the following experts who contributed to the creation of the handout:

Dr. Christian Geminn Mag. iur.
University of Kassel

Prof. Dr. Bert Heinrichs
Forschungszentrum Jülich / Rheinische Friedrich-Wilhelms-Universität Bonn

Prof. Dr. Sebastian Köhler
Frankfurt School of Finance and Management

Prof. Dr. Dieter Sturma
Rheinische Friedrich-Wilhelms-Universität Bonn

ZE
VE
DI

Imprint

Adopted February 2023

Centre Responsible Digitality (ZEVEDI)
Office
Technical University of Darmstadt
Residenzschloss 1
D-64283 Darmstadt

E-mail office[at]zevedi.de
Web https://zevedi.de/en

SUGGESTED CITATION

Centre Responsible Digitality (ZEVEDI): Research Ethics for AI Research Projects. Guidelines to Support the Work of Ethics Committees at Universities, Darmstadt 2023, 19 pp.

ZEVEDI emphasises gender-sensitive language. Collective terms include persons of all genders.

The German National Library lists this publication in the German National Bibliography; detailed bibliographic data are available on the Internet at http://dnb.dnb.de.

ISBN 978-3-910468-02-3

ZE
VE
DI

ISBN 978-3-910468-02-3