AI USE AND AI ETHICS POLICY (AI POLICY)

Document ID: AI-CFDD-v2.0

Owner: Center for Digital Dannelse ApS (CfDD) / MapUs

Address: Bådehavnsgade 42P, 2450 Copenhagen SV, Denmark | CVR No.: 39607581

Compliance Officer & DPO: Anders Skov

Executive Management (CEO): Jacob Brøndum

Email: info@mapus.io

Phone: +45 71 74 38 76 (Mainline)

Version: 2.0

Effective Date: November 2026 | Review Cycle: Annual

Revision & Version History

Version	Date	Approver	Summary of Changes
1.0	2024 / 2025	CfDD	Initial approved version
2.0	November 2026	CfDD Management / Lasse Embøl	Annual revision: Standardizing CfDD, updating CVR 39607581 and address (Bådehavnsgade 42P), and legal harmonization (EU AI Act, EU Data Act, NIS2, opt-out model for AI meeting summaries, and internal AI driving license)

1. Purpose, Vision, and Scope

At Center for Digital Dannelse (CfDD), we work at the intersection of technology, digital literacy, and human learning. We view artificial intelligence (AI) as a valuable tool that can support idea generation, optimize administrative processes, and enhance professional analyses.

Our use and development of technology must always be lawful, ethically sound, and human-centric. This policy sets out the overarching framework for responsible AI usage in compliance with Regulation (EU) 2024/1689 of the European Parliament and of the Council (the EU AI Act), the General Data Protection Regulation (GDPR), and the EU Data Act.

As a knowledge and technology company, this policy addresses the organization's two complementary roles:

- The Role of Deployer: Guidelines for employees' daily use of generative AI tools within the company for sparring, analysis, advisory services, educational development, and communication.
- The Role of Provider: Principles and technical safeguards for our proprietary SaaS platform, MapUs, used for mapping digital competence and AI readiness (e.g., based on the DigComp framework).

2. Fundamental Ethical Principles and Prohibited Practices

All use of AI in the organization—whether regarding internal workflows or digital platforms—is based on three non-negotiable core principles:

- Human-in-the-Loop: AI must always support human decision-making, never replace it. We practice a strict principle of human oversight and professional supervision. Critical decisions and external deliverables are never left to autonomous algorithms.
- Fairness, Objectivity, and Bias Minimization: We actively work to prevent technological tools from reproducing or amplifying unintended biases or discrimination based on gender, age, ethnicity, disability, or other protected characteristics.
- Transparency and Data Integrity: We maintain transparency regarding our use of technology, protect confidential data against unauthorized model training, and ensure that sources and facts can always be verified.

Prohibited AI Practices (EU AI Act Article 5)

In accordance with Article 5 of the EU AI Act, we fundamentally distance ourselves from and under no circumstances use, develop, or integrate AI systems that pose an unacceptable risk to fundamental human rights:

- Social Scoring: Systems that evaluate or classify the trustworthiness or personality traits of natural persons based on their social behavior or socioeconomic status.
- Emotion Recognition: The use of AI systems to infer, record, or assess the emotions of employees in the workplace or students in educational settings.
- Manipulative Systems: Systems that deploy subliminal techniques or exploit specific vulnerabilities (such as age, disability, or social conditions) to materially distort behavior in a manner that causes or is likely to cause harm.
- Sensitive Biometric Categorization: Systems that use biometric data to categorize individuals based on race, political opinions, trade union membership, religious or philosophical beliefs, or sexual orientation.

3. Internal Use of AI in the Company (The Role of Deployer)

This section sets out the guidelines for employees, consultants, and student assistants in their daily use of AI tools (including language models such as Gemini and ChatGPT).

3.1 Data Confidentiality and Protection of Information

Protecting personal data and confidential client information is our highest priority:

- Public Consumer Models: It is strictly prohibited to input personally identifiable information (GDPR), confidential client data, unpublished analytical material, or internal trade secrets into public AI services (e.g., free consumer versions of chatbots) where data may be used by the provider for model training.
- Closed Business and API Solutions (Google Workspace Business Standard & Gemini): When AI is used to process internal working documents or business data, we primarily utilize Gemini integrated into Google Workspace Business Standard (or approved paid API connections) backed by a formal Data Processing Agreement (DPA) that contractually guarantees "Zero Data Retention" (ZDR) and prohibits the provider from storing prompts or using data for model training.
- Pseudonymization as Standard Practice: Even when using approved closed environments, best data practices require de-identifying or pseudonymizing names, institutions, or individual cases prior to prompting.

3.2 Specific Guidelines for Internal Use of AI Meeting Assistance

The use of AI tools for audio/video recording, live transcription, and automated meeting summaries is subject to strict confidentiality and privacy safeguards:

- Approved Tools and Prevention of Shadow IT: Employees may only use approved business meeting assistants operating within the organization’s secure, governed perimeter (specifically Gemini within Google Workspace Business Standard under an active Data Processing Agreement (DPA) with Zero Data Retention). The use of unauthorized third-party tools, browser plugins, or external consumer bots (e.g., free transcription bots) is strictly prohibited to prevent confidential discussions from being routed through untrusted environments or used for model training.
- Advance Notice and Opt-Out Model for Transcription: Written AI transcription and action-item generation are conducted on the basis of legitimate interest (GDPR Art. 6(1)(f)). The meeting host must explicitly inform participants in advance (via meeting invites) and verbally reiterate at the meeting's immediate start that AI note-taking is active. Silent or unannounced AI transcription is strictly prohibited.
- Immediate Respect for Opt-Out: If any participant (internal or external) objects or requests not to be transcribed or recorded, the host must immediately disable the AI assistant and cease recording without hesitation. In such cases, note-taking must proceed manually.
- Strict Deletion of Raw Recordings: Raw audio and video files used to generate the written transcript must be deleted immediately once the transcription is complete (and no later than within 14 days), and recordings must never be used for employee surveillance, performance tracking, or behavioral monitoring. Only the final written summary may be retained. Storing raw video or audio recordings (e.g., for internal staff training) requires prior active consent (opt-in) from all participants.
- Confidential Matters and Absolute Prohibitions: Certain meetings may under no circumstances be recorded or transcribed using AI. This strict prohibition applies categorically to employee performance reviews (MUS), HR or personnel matters, disciplinary actions, salary negotiations, and meetings subject to strict non-disclosure agreements (NDAs) or involving special categories of personal data (GDPR Article 9).

3.3 Human Control and the GOAL Framework

Generative AI serves as a sparring partner for idea development, linguistic refinement, and structuring. To ensure consistently high professional quality and mitigate the risk of hallucinations (factually incorrect information), we internally apply CfDD's GOAL framework:

- G – Goal: Formulate a clear purpose for the task and define the precise role of the AI tool.
- O – Output (Format/Result): Specify detailed requirements for format, structure, target audience, and tone of voice.
- A – Assurance (Quality Control & Human-in-the-Loop): All AI-generated output undergoes systematic human review, source verification, and fact-checking. No material is delivered to a client or published without explicit employee approval.
- L – Learn (Learning & Reflection): The interactive process is used reflectively to strengthen the employee's own professional insight, task understanding, and digital literacy.

The employee always bears full professional and legal responsibility for the final work product. AI can never be held responsible for errors, omissions, or inaccuracies.

3.4 Copyright and Intellectual Property

We respect copyright principles. AI tools must not be used for the purpose of deliberately imitating or plagiarizing named living authors, researchers, or artists without permission. Cited sources, quotations, and research references must always be verified directly against the original primary sources.

3.5 AI Literacy (EU AI Act Article 4)

In accordance with Article 4 of the EU AI Act (applicable from February 2025), Center for Digital Dannelse (CfDD) ensures that all employees possess the necessary skills to use AI systems responsibly and critically.

Operational Implementation and Competence Frameworks:

- Internal AI Literacy & Micro-Assessments (AI Driving License): CfDD maintains an ongoing internal training program in safe prompting, data protection, and technical risk mitigation (hallucinations, bias, automation bias). All employees complete targeted micro-assessments and mini-questionnaires hosted on MapUs (our internal 'AI Driving License') covering our approved tool stack and this AI Policy, documenting verifiable organizational literacy in compliance with Article 4.
- External AI Competency Frameworks: For external clients, educational institutions, and enterprises, CfDD provides structured AI-readiness mapping frameworks and customized survey modules hosted on the MapUs platform, enabling organizations to evaluate, benchmark, and strengthen digital and AI competencies across diverse target groups.

Employees are encouraged to maintain professional curiosity and experiment safely within the governed boundaries established in this policy.

4. Transparency and Labeling (EU AI Act Article 50)

Center for Digital Dannelse (CfDD) strives to serve as a role model for digital transparency. Our labeling and disclosure practices adhere to Article 50 of the EU AI Act and European Commission guidelines:

4.1 Professional Articles, Reports, and Educational Material

- Legal Basis: Article 50(4) of the EU AI Act stipulates as a baseline that AI-generated or manipulated text published to inform the public on matters of public interest must be labeled as artificially generated.
- Statutory Exemption (Article 50(4), second subparagraph): The labeling requirement does not apply where the AI-assisted content has undergone a process of human review or editorial control and a natural or legal person holds editorial responsibility for the publication.
- Our Practice: Because articles on digitaldannelse.org, professional analyses, course presentations, and advisory briefs are always thoroughly developed, edited, fact-checked, and approved by our staff (who carry editorial responsibility), such materials are exempt from mandatory AI labeling.
- Full Disclosure for Substantial Machine Generation: Should we exceptionally publish content where text is directly machine-generated without substantial human editing, this will be explicitly declared to the reader.

4.2 Chatbots and Interactive Systems (Article 50(1))

If we provide AI-driven chatbots or conversational features on our websites or within our solutions, we ensure that users are clearly and promptly informed that they are interacting with an artificial intelligence system, unless this is obvious from the context of use.

4.3 Images, Graphics, and Synthetic Content (Article 50(4))

As a general rule, we use authentic photos and licensed illustrations. Where synthetically generated images, graphics, or video are exceptionally used in our communication, they will be accompanied by a clear source attribution (e.g., "Illustration: AI-generated / edited by CfDD"). Synthetic media must never be used for misleading purposes.

5. The MapUs Platform (The Role of Provider)

MapUs is a specialized software platform (SaaS) used by schools, municipalities, and companies to map digital competence and AI readiness based on recognized professional frameworks (including DigComp).

5.1 Technological Status and Scope (Non-High-Risk AI)

- MapUs mappings, respondent scores, and aggregated reports are based exclusively on deterministic algorithms, validated survey frameworks, and standardized scoring rules (pursuant to EU AI Act 3(1), 6(3), and 26).
- The platform does not use autonomous machine learning models or AI to analyze respondent data, conduct automated profiling, or make decisions regarding evaluation, recruitment, or promotion. Specifically, unlike qualitative AI synthesis platforms (such as Sopact), MapUs does not execute autonomous LLM text analysis on respondent qualitative data in production. Should qualitative AI analysis features be introduced in the future, they will strictly be subject to a comprehensive prior AI Impact Assessment (pursuant to EU AI Act 6(3) and GDPR) and mandatory Human-in-the-Loop verification.
- Because MapUs does not utilize AI components to make or materially influence critical decisions regarding educational outcomes or employment relationships, the platform is classified as a Non-High-Risk AI system under the EU AI Act (Annex III and Article 6(3)).

5.2 Administrative Support Tools

Internal use of AI in relation to MapUs (e.g., as administrative assistance for technical tasks, inspiration for or import and conversion of content for the platform and survey frameworks, or translation aid within closed environments) functions exclusively as administrative support for our internal specialist team. No content is activated on the platform without prior manual review and approval by our staff, and it does not affect clients' existing result data.

5.3 Client Data Protection (No Training Policy)

We guarantee complete confidentiality and integrity for data stored in MapUs:

- No Model Training on Client Data: Responses, personal data, organizational data, and reports in MapUs are never used to train, further develop, or fine-tune external or internal AI models. This does not preclude internal, fully anonymized process learning and the capture of general methodological best practices (such as developing internal workflows, SOPs, and AI harness skills free of client-specific or personal data).
- Zero Data Retention (ZDR): Technical integrations with language models for administrative purposes are executed exclusively via closed business APIs with a Data Processing Agreement guaranteeing Zero Data Retention (ZDR).
- Data Ownership and Export (EU Data Act & GDPR): The client retains full ownership of their data. In compliance with the EU Data Act and GDPR, MapUs provides standardized functionality for exporting or importing data using common,, machine-readable formats (CSV and JSON).

5.4 Future Development and Prior Impact Assessment

Should future plans arise to integrate features based on machine learning or AI directly into MapUs core functionality, Center for Digital Dannelse (CfDD) commits to conducting a thorough prior risk assessment (AI Impact Assessment pursuant to the EU AI Act and GDPR). The policy and technical measures will accordingly be updated prior to deployment.

6. Responsibility, Enforcement, and Revision

- Management at Center for Digital Dannelse (CfDD) ApS bears overall responsibility for compliance, communication, and oversight regarding this policy.
- Compliance with this policy is mandatory for all employees. Intentional violations—particularly uploading confidential client data or personal data to unsecure, public AI services—are considered a material breach of company security regulations and may lead to employment sanctions.
- This policy is reviewed and updated at least once annually as part of the company's annual compliance cycle, or immediately in the event of significant legislative changes (e.g., entry into force of new parts of the EU AI Act), updated guidance from authorities, or changes in the company's technology setup.
- In addition to the annual review, management and the DPO evaluate the practical implementation of this policy annually. Effectiveness is measured against concrete metrics, including a 100% completion rate and passing score on employees' annual MapUs AI Driving License, zero reported data breaches associated with AI tools, as well as screening of the company's API and software usage to prevent unauthorized Shadow IT.

- Questions regarding this policy, data ethics, or the handling of AI-related data may be directed to management or our DPO (Anders Skov) via info@mapus.io.

TECHNICAL APPENDIX: ARCHITECTURE AND HIGH-RISK BOUNDARY SAFEGUARDS

MapUs as a Non-High-Risk AI System (pursuant to EU AI Act Art. 6(3)):

Where MapUs provides prompt-based analytical or presentation assistance, the platform remains strictly Non-High-Risk under the EU AI Act (Regulation (EU) 2024/1689) by adhering to the "Deterministic-First" principle and the statutory exemption under Article 6(3).

A. Comparison Matrix: MapUs vs. High-Risk Criteria

Parameter	High-Risk Criteria (EU AI Act Annex III)	MapUs Safe Architecture (Art. 6(3))
Decision-Making Power	Autonomous or binding evaluations of individuals' future prospects (educational admission, grading, dismissal, promotion, cf. Annex III, points 3 & 4).	Purely advisory and formative. Data is presented to a human decision-maker (school leader, educator, HR), who always makes the final determination.
Data Granularity & Scope	Individual rankings, personal scorecards, or performance evaluations of named individuals.	Aggregated group level only (N >= 5) for cross-cutting syntheses. No individual AI profiling or ranking.
Free-Text & Qualitative Feedback	Raw personally identifiable survey responses tied directly to an individual's identity.	Anonymized/pseudonymized synthesis prior to processing.
Computational Engine	AI ("black box" machine learning) calculates scores, percentages, and assessment levels directly.	Deterministic-First. All mathematics and scoring are computed deterministically in the database. Any AI component merely structures linguistic presentation.

B. 4 Guardrails for Compliance & Security

1. Aggregation and Ban on Individual Profiling (N >= 5): The AI model is never fed data from a single named respondent to evaluate work performance, reliability, or personality (preventing GDPR Art. 4(4) profiling). Group syntheses require a minimum threshold of 5 responses.

2. Deterministic-First Calculation (Zero LLM Math): Averages, standard deviations, response rates, and benchmarks are always computed deterministically by the server backend/database (zero numerical hallucinations).

3. Google Workspace Gemini Perimeter (ZDR + EU Processing): Data is processed within the Google Workspace Business Standard / Vertex AI API EU region with Zero Data Retention (ZDR) under an active Data Processing Agreement (DPA).

4. Article 50 Labeling & GOAL Assurance: Any AI-assisted summaries are visibly labeled in the user interface, and administrators always retain access to underlying raw data and full editing control.