

Northlight Automation

# EU AI Act Compliance Pack

Prepared for Meridian Dental Group — covering the "Mia" patient assistant chatbot and AI drafted appointment communications.

---

**Client:** Meridian Dental Group B.V., Amsterdam, Netherlands

**Prepared by:** Northlight Automation (deployment partner)

**Date:** July 5, 2026 · **Pack version:** 1.0

**Regulatory basis:** Regulation (EU) 2024/1689 (the "AI Act")

## What's in this pack

#	DOCUMENT	LEGAL BASIS	ACTION REQUIRED
1	Risk Classification Memo	Art. 5, 6, Annex III	Review with counsel; file
2	AI Disclosure Texts	Art. 50	Deploy before Aug 2, 2026
3	AI Literacy Policy	Art. 4	Adopt; brief staff; sign
4	Human Oversight & Operations Checklist	Art. 4, 50; Art. 26 as reference	Assign owner; complete; sign
5	AI System Register & Technical Documentation	Documentation practice	Keep current; review quarterly

## How to use this pack

1. Read the Risk Classification Memo first. It explains which obligations apply to your deployment and why.
2. Have your legal counsel review the pack. Every conclusion is cited to the Act so the review is fast.
3. Deploy the disclosure texts (Document 2) before August 2, 2026.
4. Adopt the literacy policy, complete the checklist, keep the register current.

**Important.** This pack contains documentation templates prepared for review by qualified legal counsel. It is not legal advice, and no attorney client relationship is created. Obligations under the AI Act depend on facts of each deployment; counsel should confirm all conclusions before reliance.

# Risk Classification Memo

## 1. SYSTEM UNDER ASSESSMENT

FIELD	DETAIL
System name	"Mia" patient assistant
Function	Website chatbot: appointment booking, opening hours, treatment FAQ, insurance questions. Also drafts appointment reminder emails sent after human review.
Underlying model	General-purpose LLM accessed via API (third party provider)
Users exposed	Patients and prospective patients in the Netherlands (EU)
Built and configured by	Northlight Automation
Operated by	Meridian Dental Group B.V.

## 2. ROLE ANALYSIS: WHO IS PROVIDER, WHO IS DEPLOYER

Meridian Dental Group operates the system in the course of its business and is a **deployer** (Art. 3(4)). Northlight Automation configured and supplied the system; whether an agency qualifies as a **provider** (Art. 3(3)) depends on the degree of development and whether the system is placed on the market under its own name. Both parties are treated here as sharing responsibility for transparency measures as a matter of good practice. *Counsel should confirm the role allocation; it determines which party carries each transparency duty.*

## 3. PROHIBITED PRACTICES SCREEN (ART. 5)

Screened against every Art. 5 category (manipulation, exploitation of vulnerabilities, social scoring, predictive policing, workplace emotion recognition, biometric categorisation). **No prohibited practice is engaged.**

## 4. HIGH RISK SCREEN (ART. 6 AND ANNEX III)

Annex III categories reviewed: biometrics; critical infrastructure; education; employment; access to essential services; law enforcement; migration; justice. The system performs **administrative scheduling and general information** functions. It does not evaluate eligibility for care, triage symptoms, or make decisions affecting access to health services.

**Boundary condition.** If Mia's scope extends to symptom assessment, triage, or prioritising access to treatment, the analysis changes — regenerate this memo before the change goes live.

## 5. CONCLUSION

**Classification: limited risk.** Applicable obligations: transparency under **Art. 50** (Doc 2), AI literacy under **Art. 4** (Doc 3), plus the operational practices in Docs 4 and 5. Art. 26 high risk duties do not currently apply (adopted as best practice in Doc 4).

## AI Disclosure Texts

Art. 50(1) requires that people interacting with an AI system are informed they are interacting with AI, unless this is obvious to a reasonably well informed person. Disclosure must be given *"in a clear and distinguishable manner at the latest at the time of the first interaction."* The texts below are ready to deploy.

### CHATBOT FIRST MESSAGE NOTICE

#### ENGLISH

Hi, I'm Mia — Meridian Dental's digital assistant. I'm an AI, not a person. I can book appointments and answer common questions. If you'd rather speak with our team, type "human" any time or call us directly.

#### NEDERLANDS

Hoi, ik ben Mia, de digitale assistent van Meridian Dental. Ik ben een AI en geen mens. Ik kan afspraken inplannen en veelgestelde vragen beantwoorden. Wilt u liever iemand van ons team spreken? Typ dan "medewerker" of bel ons direct.

### PERSISTENT INTERFACE LABEL

AI assistant — responses are generated automatically. *(Displayed permanently in the chat window header or beneath the input field.)*

### AI ASSISTED EMAIL FOOTER

This reminder was drafted with the help of an AI assistant and reviewed by our team.

### PLACEMENT RULES

- The first message notice must appear before or with the assistant's first response — never after.
- The label may not be hidden behind a menu, tooltip, or link.
- Where accessibility rules apply, notices must be available to screen readers (plain text, not image only).
- Reminder emails are reviewed by staff before sending; the footer keeps the human role documented.

**Why the email footer, when Art. 50(4) targets public interest text?** Patient emails are not "public interest" publications, so labeling is not strictly required — but disclosure costs nothing, builds patient trust, and covers edge cases. Counsel may remove it.

# AI Literacy Policy

Art. 4 has applied since February 2, 2025: "Providers and deployers of AI systems shall take measures to ensure, to their best extent, a sufficient level of AI literacy of their staff and other persons dealing with the operation and use of AI systems on their behalf."

## 1. PURPOSE AND SCOPE

This policy ensures every Meridian Dental Group staff member who works with or around the Mia assistant understands what it does, what it must not be used for, and how to intervene. It applies to reception staff, practice managers, and any employee handling patient communications.

## 2. WHAT STAFF MUST KNOW

- Mia is an AI system. It generates responses automatically and can be wrong.
- Mia handles scheduling and general information only. It must never give medical advice; conversations drifting toward symptoms or treatment decisions are handed to staff.
- Patients can always reach a human. Staff must honor "human" handover requests immediately.
- AI drafted reminder emails are reviewed by a staff member before sending — the reviewer owns the content.
- Errors, strange outputs, or patient complaints about Mia are reported to the practice manager the same day.

## 3. TRAINING MEASURES

MEASURE	AUDIENCE	FREQUENCY
30 minute onboarding briefing on this policy	All patient facing staff	At hire and at system changes
Hands on demo: handover, escalation, review flow	Reception staff	Twice yearly
Policy re acknowledgment	All in scope staff	Yearly

## 4. OWNERSHIP

The practice manager owns this policy, keeps the training log, and reviews the policy when the system materially changes or the regulation is updated.

---

Adopted by (name, role)

---

Signature

---

Date

## Human Oversight & Operations Checklist

This deployment is limited risk, so the high-risk operational duties of Art. 26 do not currently apply. The checklist below implements the duties that do apply (Art. 4, Art. 50) and adopts the core of Art. 26 as voluntary best practice – so an upgrade in risk class or a client audit finds the house already in order.

ITEM	OWNER	STATUS / DATE
<input type="checkbox"/> AI disclosure notices live in chatbot, interface label, and email footer (Doc 2)	Northlight	
<input type="checkbox"/> Named oversight owner with authority to suspend the system appointed	Meridian	
<input type="checkbox"/> Human handover ("human" / "medewerker") tested and working	Northlight	
<input type="checkbox"/> Staff briefed on AI Literacy Policy; training log started (Doc 3)	Meridian	
<input type="checkbox"/> Medical advice guardrail tested: symptom questions deflect to staff	Northlight	
<input type="checkbox"/> Conversation logs retained (6 month minimum, mirroring Art. 26(6) practice)	Northlight	
<input type="checkbox"/> Incident path defined: strange output → practice manager → Northlight → suspend if needed	Both	
<input type="checkbox"/> Reminder emails require human review before sending	Meridian	
<input type="checkbox"/> Scope change rule: any new capability triggers reclassification (Doc 1) before go live	Both	
<input type="checkbox"/> Quarterly review scheduled: register, policy, disclosures, this checklist	Both	

**Suspension rule.** If the system produces outputs suggesting risk to patient health or rights, the oversight owner suspends it first and investigates second. Downtime is cheaper than an incident.

\_\_\_\_\_  
Completed by (Meridian)

\_\_\_\_\_  
Completed by (Northlight)

\_\_\_\_\_  
Date

## AI System Register

SYSTEM	PURPOSE	MODEL / VENDOR	RISK CLASS	DISCLOSURES LIVE	NEXT REVIEW
Mia (web chatbot)	Booking, FAQ	General-purpose LLM via API	Limited	<input type="checkbox"/> pending	Oct 2026
Reminder drafting	Email drafts, human reviewed	Same	Minimal	<input type="checkbox"/> pending	Oct 2026

Add every AI system the organization adopts — shadow AI is the audit finding that hurts.

## Technical Documentation Template

Filled per system; this is what counsel or an authority will ask for first.

FIELD	ENTRY (MIA)
Intended purpose	Appointment scheduling and general practice information for patients
Out of scope uses	Medical advice, triage, treatment or eligibility decisions
Model and version	[API model + version, pinned]
System prompt / guardrails summary	Scope restriction, medical deflection, handover triggers [attach]
Data flows	Chat content → API provider (EU region endpoint); booking data → practice calendar; no training on patient data per vendor DPA
Human oversight measures	See Document 4
Logging	Full transcripts, 6 months, access restricted
Related assessments	GDPR DPIA reference: [counsel to confirm requirement]
Change log	v1.0 deployed [date]; changes require reclassification per Doc 1

**End of sample pack.** This pack was generated for a fictional client to demonstrate CompliancePack output. Your packs carry your agency's brand, your client's facts, and current regulatory text. [compliancepack.vercel.app](https://compliancepack.vercel.app)