Privacy by Design in AI context

Professor Subhajit Basu FRSA
Chair in Law and Technology
School of Law
University of Leeds
25 June 2026

What will we do in this session?

1 Meaning

Privacy by Design as a design discipline, not a slogan

2 Law

UK GDPR / Data Protection Act 2018 / DUAA duties translated for AI teams

3 Build

Lifecycle gates, minimisation, transparency, rights and security

4 Apply

Redesign an AI/data workflow and test the release checklist

Aim: by the end, the team should be able to spot privacy risk before a line of code or model fine-tuning decision locks it in.

The central proposition

Privacy is not merely secrecy, consent, or security. In AI systems, it is about power, context and human agency.

Privacy by Design is a way of asking:

What kind of relationship are we creating between the system and the person?

For AI teams, privacy is not a final audit question. It shapes what data is collected, what inferences are made, how uncertainty is handled, what people are told, and how they can contest the system.

From “can we process?” to “should this architecture exist in this form?”

The legal Basis: what current UK data protection law requires

For AI systems using personal data, privacy by design is a legal duty and an engineering requirement.

Article 25

Data protection by design and by default

Build measures into the system from the start; default processing must be limited to what is necessary.

Article 35

DPIA for high-risk processing

Assess likely high risks before processing, especially with new technologies.

Article 5

Principles + accountability

Lawfulness, fairness, transparency, purpose limitation, minimisation, accuracy, storage limitation, security.

Articles 6/9 + 22A–22D

Lawful basis, special data, ADM

Choose a lawful basis; protect special category data; for significant solely automated decisions, build in information, representations, challenge and human intervention.

Practical translation: document the purpose, prove necessity, reduce data, protect the model, explain the outcome, preserve rights, and monitor drift.

Why AI changes the privacy problem

AI often creates new personal information from old data, and new power from apparently mundane signals.

- **Scale** — more people, more attributes, more linkage
- **Inference** — systems infer traits never directly provided
- **Opacity** — people may not understand why an output appears
- **Context collapse** — data from one setting is repurposed in another
- **Feedback loops** — outputs influence future opportunities and datasets

Engineering consequence: privacy risks appear not only at collection, but also at feature engineering, model choice, prompting, logging, evaluation and deployment.

First task: map personal data across the AI system

Developers often look only at the training set. Privacy law looks at the whole processing operation.

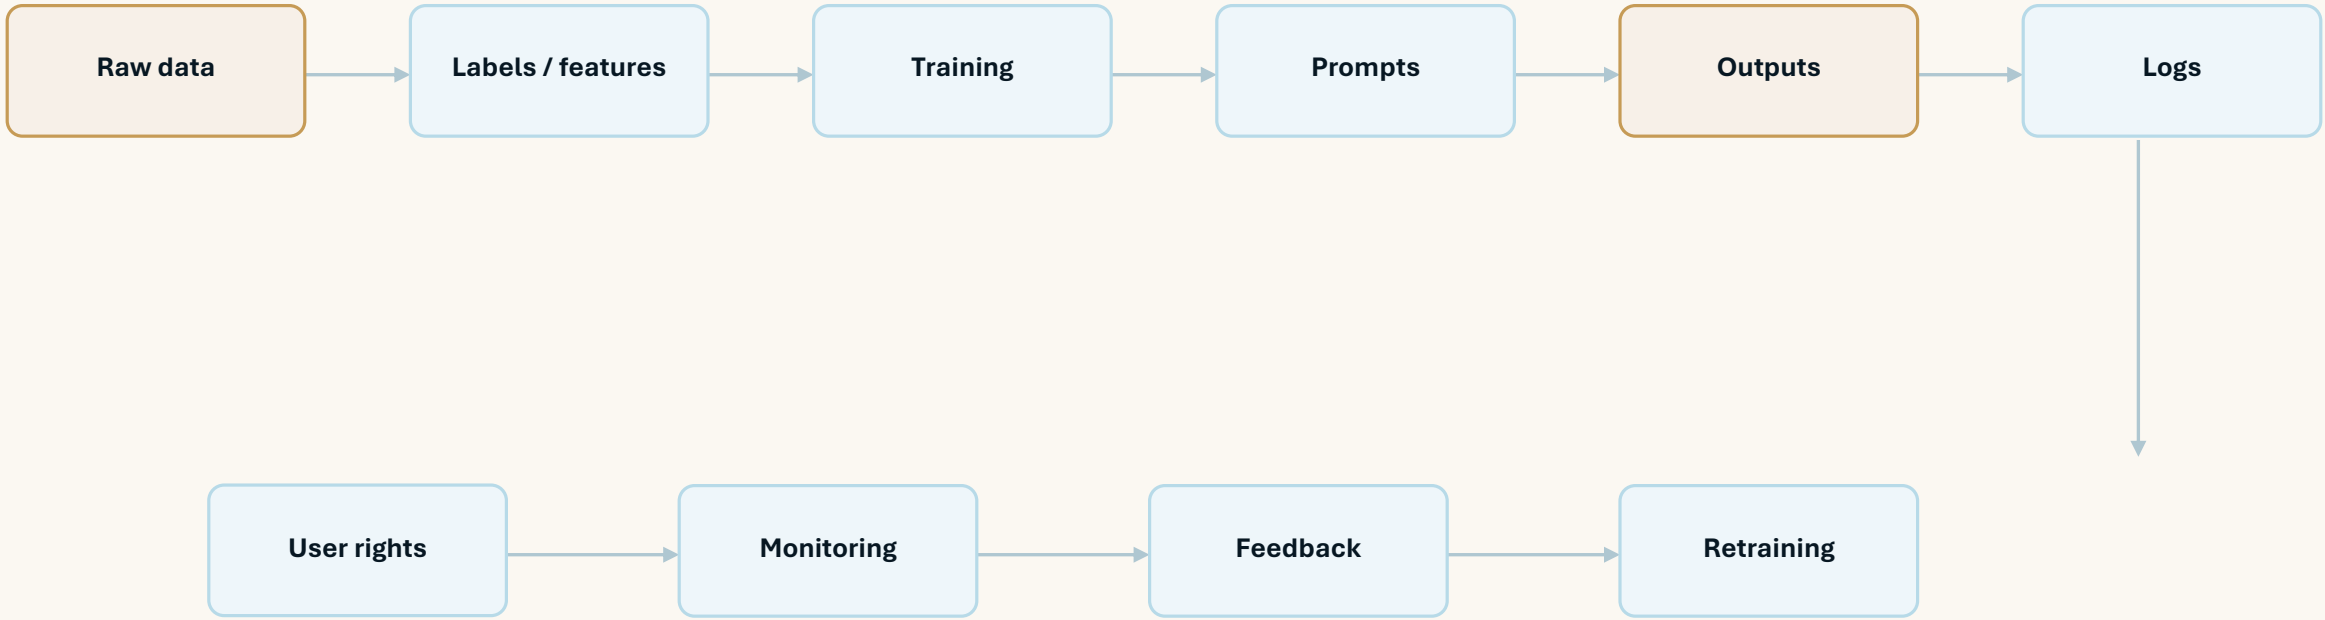

Ask at each point: whose data, whose inference, whose risk, whose control, whose remedy?

A practical PbD operating model: Five design gates

Treat privacy as a release criterion at every lifecycle stage.

Purpose gate

What problem are we solving, and is personal data necessary?

Data gate

What is the minimum data architecture that can work?

Model gate

What risks arise from training, inference and optimisation?

Deployment gate

What will people know, control and contest?

Monitoring gate

What will we log, audit, delete and retrain?

Gate rule: if the team cannot explain necessity, risk and mitigation in ordinary language, the system is not ready to progress.

Data minimisation for AI: seven engineering levers

The answer is not always “collect less”; it is often “design differently”.

Persons
fewer people in scope

Attributes
fewer fields

Granularity
less precise data

Linkage
weaker joins

MINIMUM NECESSARY

Retention
shorter logs

Locality
process near source

Substitution
synthetic / aggregate / rules

Developers’ question: what is the least invasive data architecture that can still achieve the legitimate purpose?

Lawfulness and fairness: design questions for AI teams

A lawful basis is necessary, but it does not settle fairness, necessity or proportionality.

Purpose

- What exactly is the purpose?
- Is it compatible with the context?
- Could the purpose be achieved less intrusively?

Data

- Do we need personal data?
- Is any data special category or children's data?
- Are proxy variables creating hidden sensitivity?

Model

- Does optimisation create unfair outcomes?
- Are error rates tolerable for affected groups?
- Can people understand and challenge outputs?

Fairness is not merely "no bias in the model". It includes how people are selected, represented, scored, treated, informed and given remedy.

Transparency: explain the system at the point of power

A privacy notice is not enough if people cannot understand the practical effect of the system.

Layered explanation

Short message first; detail on request.

Meaningful information

Purpose, data, logic, consequences and safeguards.

Developer artefacts

Model cards, data sheets, prompt policies, test reports.

Contestability

Clear route to challenge or ask for human review.

Test: can an affected person understand what happened well enough to decide whether to object, correct, complain or seek review?

Rights and redress must be designed into the workflow

If a person cannot exercise their rights, the system has a design failure.

Access	Can the person see the relevant personal data and derived outputs?
Correction	Can incorrect inputs, labels or profiles be corrected?
Deletion / retention	Can data be deleted without breaking the entire model pipeline?
Objection	Can a person object to processing or profiling where applicable?
Human intervention	Can significant automated decisions be reviewed meaningfully?

Redress by design means logging enough to investigate, but not so much that logs become a second surveillance system.

Security is necessary, but not sufficient

Encryption, access control and logging are essential. They do not by themselves prove privacy by design.

01

Protect data at rest, in transit and during processing

02

Protect the model against leakage, inversion and extraction

03

Protect prompts, embeddings and logs from uncontrolled reuse

04

Protect people from unfair, opaque or irreversible outcomes

Security question: “Can attackers get in?”
Privacy question: “Even if the system works as intended, does it expose people to unjustified risk?”

DPIA: the practical tool for privacy by design

A DPIA should be a design sprint with legal consequences, not a form completed after deployment.

1. Describe processing
2. Identify purposes
3. Consult people / stakeholders
4. Assess necessity
5. Assess risks
6. Choose mitigations
7. Decide residual risk
8. Record + review

Trigger: a DPIA is mandatory where processing is likely to result in high risk. AI projects often need careful screening because of profiling, inference, scale, opacity or vulnerable groups.

Exercise: redesign the workflow before it becomes harm

Group discussion with the developer/design team.

Problem scenario

A passenger/visitor manifest is visible on a shared desk or screen. Each new visitor can see previous entries. The team wants to “modernise” it using AI to pre-fill, classify and flag entries.

Prompt to group:

Where are the privacy risks? What should be redesigned before any AI feature is added?

Better design moves

- Individual entry channel
- No public display of personal data
- Role-based access + audit log
- Retention schedule
- Human review for flags
- DPIA before AI classification

Key lesson: AI does not fix a poor data process; it often scales it.

UK update: innovation space, but not a privacy holiday

The Data (Use and Access) Act 2025 modifies parts of the framework; it does not replace UK GDPR, DPA 2018 or PECR.

ADM	Articles 22A–22D create a broader framework for significant solely automated decisions, but safeguards remain information, representations, challenge and human intervention.
Children	Online services likely to be accessed by children must take their needs into account when deciding how to use personal data.
Complaints	Organisations must support data protection complaints, acknowledge them within 30 days, and respond without undue delay.
ICO direction	AI, ADM and biometrics remain priority areas: transparency, bias/discrimination, rights and redress.

Design implication: increased legal flexibility raises the importance of accountability, documentation and safeguards.

Who owns privacy by design?

No single team owns all of it. Privacy by design fails when responsibility is diffused.

Shared release decision

Product owner	purpose, scope, user impact
Developers	data flows, logs, controls
Design / UX	defaults, notices, choice
ML team	training, evaluation, drift
Security	confidentiality, integrity, access
DPO / legal	basis, DPIA, rights, records
Domain owner	context, harm, proportionality

Accountability means someone can explain why the design choices were lawful, necessary, proportionate and monitored.

Before you ship: the privacy-by-design release checklist

Ten questions every AI feature should answer before deployment.

1. **Is the purpose specific and legitimate?**
2. **Is personal data truly necessary?**
3. **Have we minimised people, fields, granularity, linkage and retention?**
4. **Have we identified lawful basis and any special category / children's data?**
5. **Have we completed a DPIA where risk is high?**
6. **Can we explain the system and its consequences to affected people?**
7. **Can users exercise rights without engineering heroics?**
8. **Are access controls, logging and deletion technically real?**
9. **Have we tested fairness, accuracy, failure modes and drift?**
10. **Is there a named owner for review, incidents and retirement?**

If one answer is "we do not know", that is not a blocker by itself — it is a design task.

The final test

Do not ask only:

“Can we collect and process this data?”

Ask:

“What is the least intrusive architecture that can achieve a legitimate, proportionate and accountable purpose?”

Privacy by Design is where law becomes system architecture.

Core sources for follow-up

These are the sources behind the legal and practical framing.

ICO	Data protection by design and by default; Privacy in the product design lifecycle; Guidance on AI and data protection; DPIA guidance; AI and biometrics strategy.
UK law	UK GDPR Articles 5, 6, 9, 22A–22D, 25 and 35; Data Protection Act 2018; Data (Use and Access) Act 2025; PECR.
EU context	EU AI Act; EDPB Opinion 28/2024 on AI models and personal data; EDPB LLM privacy risks and mitigations report.
Workshop base	AI4D Workshop PbD v03: original structure and exercise prompt supplied by workshop organisers.
Scholarly lens	Law, regulation and technology; algorithmic regulation/accountability; human-rights-centred AI governance; public law values in system design.