



Step 0 – classify your system

High risk

Decisions about people: hiring, lending, healthcare, education, law enforcement

Limited risk

Chatbots, recommendation systems, content moderation tools

Minimal risk

Spam filters, inventory optimisation, internal productivity tools

Phase 1 – Design & development

GOVERNANCE

- Appoint AI Ethics Lead (or assign to CTO/CPO)
- Create cross-functional ethics review board
- Set review cadence: monthly active dev, quarterly deployed

DOCUMENTATION

- Model card: capabilities, data, limitations, metrics
- Data sheet: collection, preprocessing, known biases
- Impact assessment: who's affected, harms, mitigations

Phase 2 – Pre-deployment

BIAS & FAIRNESS TESTING

- Run fairness metrics across protected characteristics
- Test edge cases and underrepresented groups
- Compare error rates across demographics

LEGAL COMPLIANCE

- UK GDPR: lawful basis, privacy notice, DPIA
- Equality Act 2010: no indirect discrimination
- Sector rules: FCA, CQC, Ofsted as applicable

Phase 3 – Deployment

HUMAN OVERSIGHT

- Human review required before high-risk decisions
- Train reviewers on limitations and override procedures
- Log all overrides with documented reasoning

USER TRANSPARENCY

- Disclose that AI is being used
- Explain what data is processed and why
- Provide clear challenge and opt-out mechanism

Phase 4 – Ongoing Maintenance

AUDITS

- Re-run bias metrics every 3–6 months
- Review decision logs for emerging patterns
- Interview users about fairness experience

INCIDENT RESPONSE

- Playbook: bias in production, inaccurate decisions
- Playbook: data breach, regulatory inquiry
- Report data breaches to ICO within 72 hours

Sprint review ethics checklist

FAIRNESS

Equal performance across protected characteristics?

Tested on underrepresented groups?

Could this reinforce societal bias?

TRANSPARENCY

Explainable in plain language?

Users aware AI is involved?

Can users access their data?

ACCOUNTABILITY

Who is responsible if harm occurs?

Can decisions be appealed?

Decision criteria documented?

PRIVACY

Lawful basis for this processing?

Proportionate to purpose?

Could we use less data?

SAFETY

Worst-case failure mode identified?

Edge cases tested?

Human override in place?

Regulatory landscape

UK GDPR & Data Protection Act 2018

Lawful basis for processing, individual rights, automated decision-making rules

ACTIVE – ENFORCED BY ICO

Equality Act 2010

Protected characteristics, indirect discrimination, proportionality requirements

ACTIVE – ENFORCED BY EHRO

Sectoral regulations

FCA (finance), CQC (healthcare), Ofqual (education) – sector-specific AI standards

ACTIVE – CHECK YOUR REGULATOR

UK AI Act principles

Risk-based approach expected, following EU AI Act principles. Subscribe to CDEI updates

COMING – WATCH FOR CONSULTATION

Acuity Data conducts independent AI ethics audits for UK organisations. Get your systems assessed before a regulatory inquiry does it for you.

<https://www.acuitydata.io/risk-governance>

As at April 2026

