Appendix 4: Our Don'ts for AI

It's important that we are clear about what we do not do with AI. This section outlines the boundaries we commit to respecting. These 'don'ts' are guardrails for our teams, helping us to reduce risk, protect people, and uphold trust in how we use AI.

Don't use AI without clear human oversights

- AI is never left to operate without designated persons or teams responsible for its outcomes. We keep humans in the loop.

Don't use AI in ways that could cause harm

- We avoid using AI that could harm people, organisations, communities, or the environment.

Don't use biased or non-representative data

- We do not train AI on data that excludes or unfairly represents certain groups, and act quickly to correct imbalances.

Don't process personal or sensitive data carelessly

- Never use identifiable data in AI systems without privacy safeguards.

Don't hide when AI is being used

- We are always open and transparent, so users know when AI is in use, what it does, and how it impacts them.

Don't use AI to replace critical human judgement

- AI is a support tool; we avoid full automation in sensitive or safety-critical domains.

Don't ignore long-term risks

- We do not implement AI without plans for long-term monitoring, maintenance, and responsible management.

Don't allow AI to be a 'black box'

- If we cannot explain how AI works, or support us, we do not use it.

Don't exclude people from the benefits of AI

- We don't use AI which is inaccessible due to language, ability, or other equitable barriers.

Don't ignore ethical, legal, or regulatory boundaries

- We follow applicable laws, and do not use or build AI systems that are prohibited.