Rules Before Tools: A Plain English AI Policy for Small and Medium Sized Firms

The internet reached the workplace before the rules for using it did, and firms spent a decade catching up. AI is repeating the pattern, faster and with higher stakes, because what leaves the building this time is the firm's own knowledge. This paper argues that a small firm does not need a compliance department to use AI safely. It needs one page of written rules, a named owner, a clear list of what never goes into a tool, and the habit of checking what comes out.

57%	48%	66%	1 page
of employees hide their use of AI at work [1]	have put company information into a free public AI tool [1]	rely on what the tool says without checking it [1]	is all the written policy most small firms need

Executive summary

Most small firms already use AI. Few have written down how, and that gap is where the risk sits. Staff reach for free tools because the tools help, and they say nothing because nobody asked. The result is customer details and client material flowing to services the firm has no contract with, and unchecked answers going out under its name. The 1990s offer a clear lesson: bans failed, vague policies were ignored, and the firms that came through well named acceptable uses plainly, taught the rules rather than filing them, and gave staff a safe approved route. This paper sets out what UK law now expects of a small firm, what the international standards are for, and a one page policy any owner can write this month.

1. The last time this happened

When the web arrived in offices in the mid 1990s, firms faced a tool that was plainly useful, plainly risky and outside their experience. Their answer was the acceptable use policy: a written agreement between employer and staff about what the connection was for. The first standard guidance came from the network engineers' own Site Security Handbook of 1991, revised in 1997 [4]. Governments mostly stood back. The United States' 1997 Framework for Global Electronic Commerce set the tone: the private sector should lead, governments should avoid undue restriction, and where rules were needed they should be predictable, minimal, consistent and simple [5]. Britain's approach to AI today follows the same instinct: no single AI law, existing regulators applying existing law [8]. The consequence is what it was in 1997: nobody is going to write your firm's rules for you.

Research from the period shows how firms got it wrong. A 1997 study of a research body with several hundred connected staff found no policy at all, an attitude of "ignorance is no excuse", and a network manager reading a daily log by hand to spot misuse [6]. The authors concluded that an effective policy needed clear lists of acceptable and unacceptable uses, named roles, sanctions that left room for discussion, active teaching rather than a document on a shelf, and a link to the firm's other policies [6].

QUESTION	UNOFFICIAL SOFTWARE, 1997	UNOFFICIAL AI, 2026
What staff did	Downloaded programs, browsed, sent personal email	Paste work into free chatbots and writing tools
What could go wrong	Viruses, lost time, embarrassing postings	Confidential material leaves the firm; wrong answers go out under its name
Could it be undone?	Usually: delete the file, revoke the access	Rarely: once a public tool has taken the text, the firm cannot get it back
What failed	Blanket bans that staff worked around	Blanket bans that push use out of sight
What worked	Clear rules, an approved route, someone in charge	Clear rules, an approved account, someone in charge

Table 1. The tools changed. The human behaviour, and the fix, did not. Sources: [2][3][6].

2. The risk now comes from inside

In the 1990s the threat was assumed to come from outside: a virus, a hacker, a hostile website. The defences were built at the boundary. AI has moved the risk inside the building. The person who causes the loss is usually a good employee trying to get the job done faster.

A 2025 global study of more than 48,000 people found that 57 per cent of employees hide their use of AI at work, 48 per cent have put company information into free public tools, and 66 per cent rely on the answers without checking them. Fewer than half had received any training [1]. The trade calls this shadow AI. In a small firm it simply looks like a colleague pasting a customer's email into a chatbot to draft the reply.

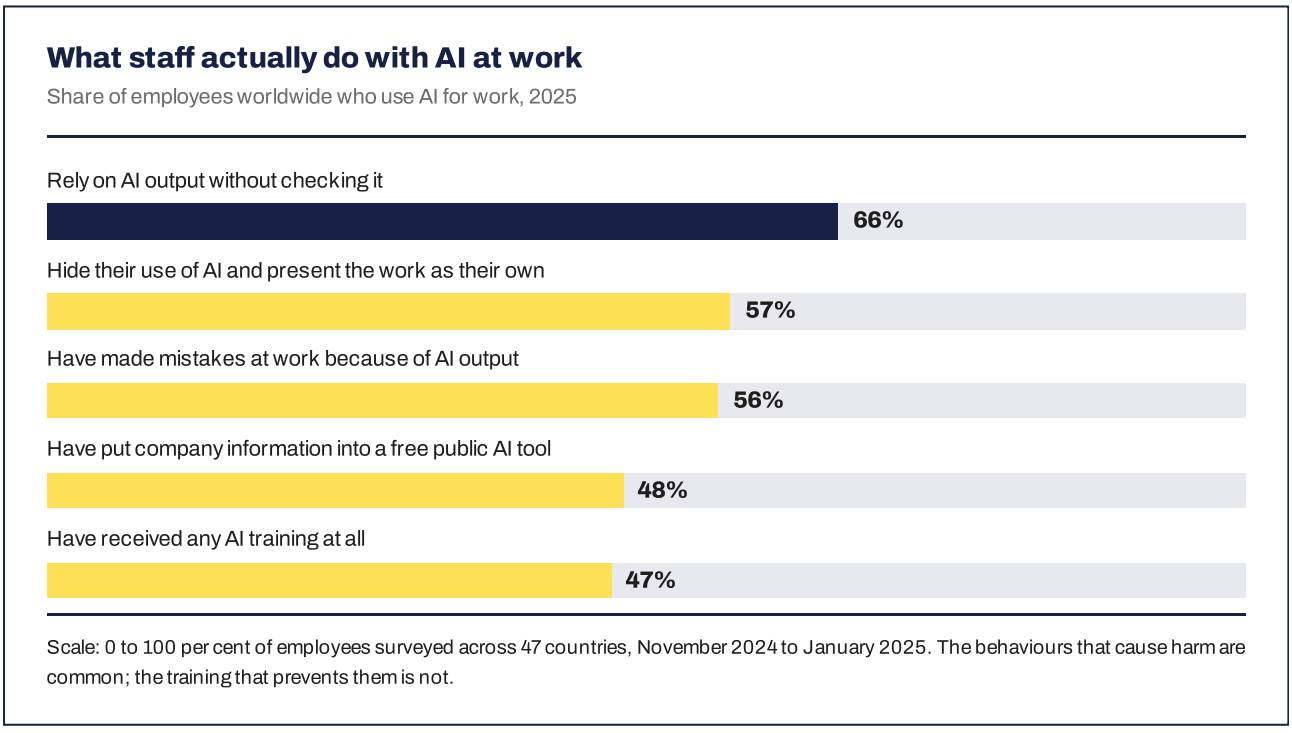

Figure 1. The gap between how AI is used and how it is governed, measured across 47 countries. Source: [1].

The best known case shows how ordinary the failure is. In 2023, over roughly twenty days, engineers at Samsung pasted confidential source code, a meeting transcript and hardware details into a public chatbot to fix bugs and write minutes [3]. There was no break-in, no stolen password and no fault in Samsung's own systems. The material left through a web browser, typed in by staff with no bad intent, and because the service could keep what it was given, there was no way to recall it. Samsung's response was to ban the tool and cap the length of what could be typed into it [3]. That is the reaction that failed against the web thirty years earlier. A ban does not stop a useful tool being used; it stops the firm seeing how it is used.

The second failure is less visible but just as costly: trusting what comes out. Amazon abandoned an experimental recruitment tool after finding it had learned, from years of historic hiring data, to mark down applications from women [7]. Microsoft withdrew its Tay chatbot within a day of launch after it began repeating abusive content picked up from users [2]. In a small firm the equivalent is a quotation sent with an invented figure, a job advert screened by a tool nobody can explain, or marketing copy that makes a claim the firm cannot stand behind. The tool will produce a confident answer whether or not it is right. Only a person can decide whether to send it.

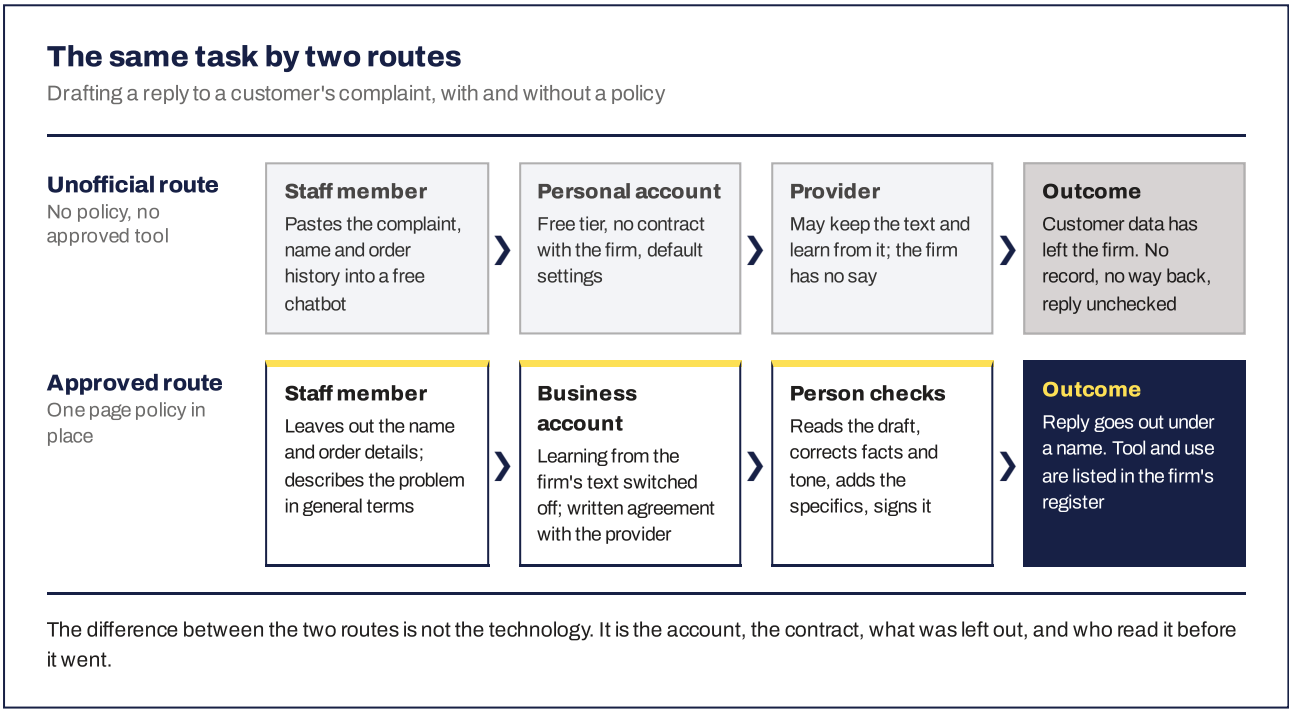

Figure 2. The unofficial route is the one most staff take today. The approved route costs a business subscription and one page of rules.

3. What the law expects of a small UK firm

The UK has chosen not to pass a single AI law. Instead, the laws a small firm already lives with apply to AI as they apply to everything else, and the regulators have been asked to say what that means in their own sectors [8]. Five points matter most.

Personal information. Customer, staff and applicant details are covered by the UK GDPR and the Data Protection Act 2018 whether they sit in a filing cabinet or a chatbot. Putting them into a public AI tool sends them to a third party the firm has no agreement with, which is a breach in itself. The Information Commissioner's Office (ICO) has published guidance on AI and data protection, and since May 2026 has been under a legal duty to produce a formal code of practice on AI and automated decisions, expected in 2027 [8][9].

Decisions about people. From 5 February 2026 the Data (Use and Access) Act 2025 rewrote the rules on automated decisions. A firm may now let a tool make or shape a decision that seriously affects a person, such as rejecting an applicant or pricing a policy, but only if it tells the person, lets them ask for a human review and lets them challenge the outcome. The ICO's draft guidance is clear that the human involvement must be real: an active review before the decision takes effect, not a signature on the way out [10].

Fairness. The Equality Act 2010 does not care whether a person or a program discriminated. If a screening tool marks down older applicants or a pricing tool disadvantages a protected group, the firm answers for it [16][17].

Directors. Under the Companies Act 2006 directors must exercise reasonable care, skill and diligence. Knowing which AI tools the firm relies on, what they touch and who is responsible is now part of that duty, as it is for any other significant risk [17]. The same expectation falls on the partners of a professional practice and the trustees of a charity.

Selling into Europe. Firms that sell AI-enabled products or services into the EU, or whose AI output is used there, also fall under the EU AI Act. The Act was amended in July 2026 and the timetable is now settled [11][12]. Since 2 August 2026, people must be told when they are dealing with a chatbot, and content produced by AI must be marked as such. Every business that uses AI must take steps to support its staff's understanding of it, a duty lightened in July 2026 but still in force. The heavier "high-risk" duties, which cover recruitment, credit and similar uses, now apply from 2 December 2027 for standalone systems and 2 August 2028 for AI built into regulated products. The Act also gives small firms specific help: fines capped at the lower of two figures rather than the higher, simpler paperwork, priority access to national testing sandboxes (due by August 2027), and a single point of contact at each national regulator [11][12]. The July 2026 changes extended these concessions to firms with up to 750 staff.

RULE	IT APPLIES WHEN	WHAT IT ASKS OF THE FIRM
UK GDPR and Data Protection Act 2018	Any AI use that touches customer, staff or applicant details	A lawful reason, a privacy notice that mentions AI, and no personal details in tools the firm has no contract with
Automated decision rules (UK GDPR Articles 22A to 22D, from February 2026)	A tool makes, or largely shapes, a decision that seriously affects a person	Tell the person, offer a real human review, let them challenge it, keep a record
Equality Act 2010	Hiring, pricing, credit and service decisions	The firm answers for bias in the tool as if a manager had shown it
Consumer protection and advertising rules	AI-drafted claims, prices, descriptions and copy	Every claim checked by a person before it goes out
Companies Act 2006, directors' duties	Every limited company	Directors know what tools the firm relies on and who owns each one
EU AI Act, as amended July 2026	Selling AI-enabled products or services into the EU, or output used there	Transparency and staff support now; high-risk duties from December 2027; SME concessions on fines, paperwork and sandboxes

Table 2. The laws a small UK firm already lives with, and what each asks of its AI use. Sources: [8] to [12], [17].

4. What the standards are for

Three international documents are quoted at small firms, usually by larger customers. None of them needs to be bought or certified to be useful.

The US NIST AI Risk Management Framework (2023) is voluntary and reduces to four verbs [13]. Govern: decide who is responsible. Map: list what you use, what it is for and what information it touches. Measure: check that it works and keeps working. Manage: fix what goes wrong and know in advance what you will do. A twelve person firm can do all four on one page. NIST's 2024 supplement adds the risks particular to writing and chat tools, above all made-up answers and the ease with which a tool can be talked into ignoring its instructions [13].

ISO/IEC 42001 (2023) is the certifiable version: 38 controls, a statement of which ones apply and why the rest do not, and internal audits at set intervals [14]. Its companion ISO/IEC 38507 (2022) tells boards what to decide rather than how to do it [14]. Most small firms will never certify, but the standard is worth knowing. A large customer may ask whether the firm's AI use is governed to it, and a one page register with a named owner is a credible answer. Its authors were also explicit about the failure they wanted to prevent: policies written for the auditor's file that change nothing on the shop floor. That warning is the whole argument of this paper.

All three documents ask the same questions of a firm of any size: who is responsible, what is in use, what does it touch, how do you know it still works, and what happens when it does not. Answer those on one page and keep it current, and the firm is already doing most of what the standards describe. The rest is paperwork that can be added if a customer or a regulator ever asks for it.

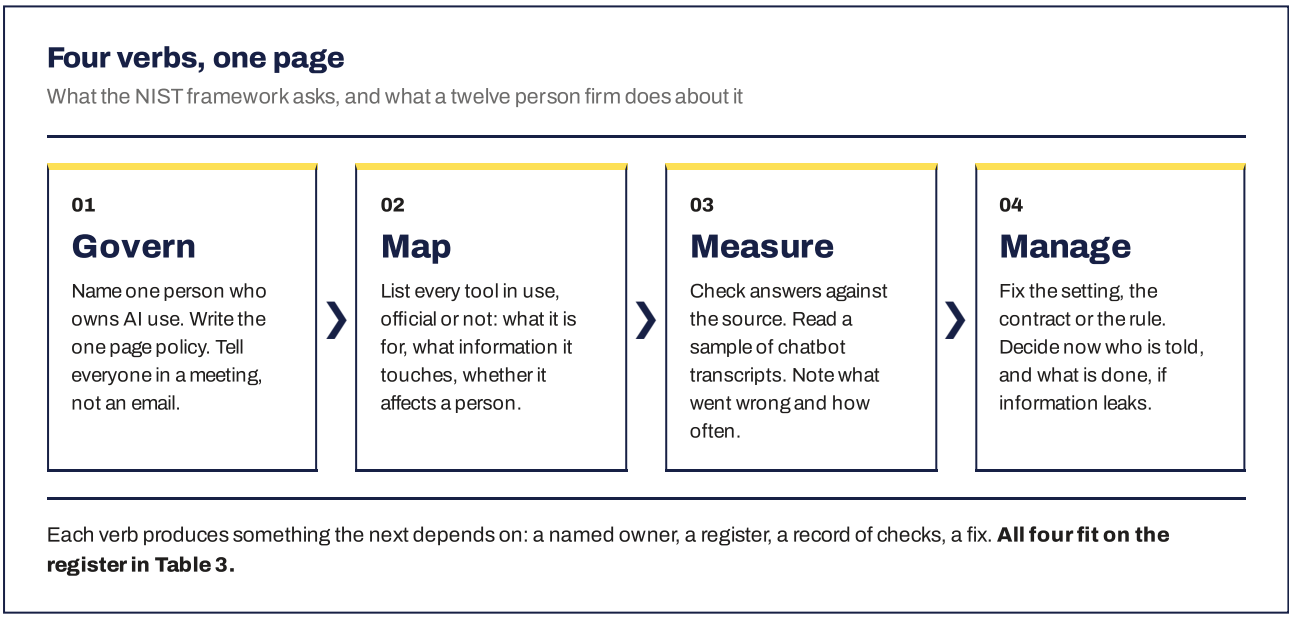

Figure 3. The four functions of the NIST framework, translated for a firm with no compliance department. Source: [13].

A ban did not stop the web in 1997 and it will not stop AI in 2026. It only stops the firm seeing how the tool is used.

5. Writing the policy: one page, eight headings

Everything above reduces to a page an owner can write in an afternoon and explain in one meeting [6] [13][14].

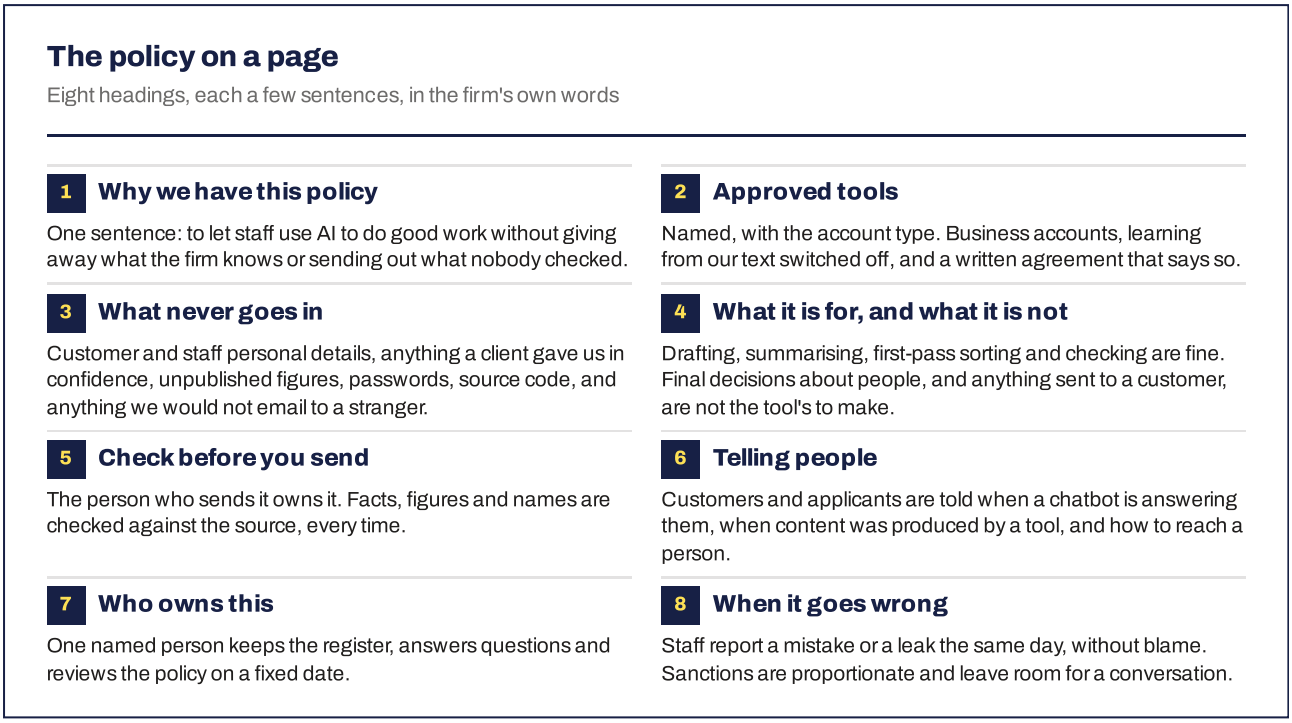

Figure 4. The eight headings of a small firm's AI policy, drawn from [6], [13] and [14].

The 1997 study also records what not to do: the organisation believed that telling staff about any rule meant telling them about every rule, so it told them nothing and relied on "ignorance is no excuse" [6]. That was untenable then and is more so now. The register below is where the policy meets the day to day.

TOOL AND ACCOUNT	WHAT WE USE IT FOR	INFORMATION IT TOUCHES	OWNER	HUMAN CHECK	REVIEW
Writing assistant, business account	Drafting quotes, letters, social posts	Product details and prices; no customer data	Office manager	Sender checks every figure and name	March 2027
CV sorting add-on in recruitment software	First sort of applications	Applicant details	Owner	Owner reads every rejected CV before a reply goes	Next hiring round
Chatbot on the website	Opening hours, stock and returns questions	Public information only	Sales lead	Weekly read of transcripts; hand-off to a person	Quarterly

Table 3. A register for a small firm: what the standards ask for, and what a regulator would want to see first.

6. An example worth copying

Large firms have run the experiment at scale. In 2023 Morgan Stanley gave an AI assistant to 16,000 financial advisers, moving over six months from a trial of 300 to the whole firm [15]. The tool answered only from the firm's own library of around 100,000 research documents, never the open web, and the adviser stayed in charge of every decision. One named head of AI kept legal, compliance and technology in step, and adoption reached 98 per cent of adviser teams [15]. A small firm copies it: a few people, one task, its own documents, one owner.

7. Recommendations

1. Write the policy on one page and teach it in one meeting. A policy that sits in a drawer protects nobody.
2. Replace the ban with an approved route: a business account, learning from your text switched off, and a contract that says so.
3. Publish the list of what never goes into a tool and put it where people work, not in the staff handbook.
4. Name one owner. Keep the register in Table 3 current and review it every quarter.
5. Keep a person truly in charge of any decision about a person. Tell them, offer a human review and record it.
6. Check every AI vendor contract for two things: the provider does not learn from your data, and you know where the data is held.
7. If the firm sells into the EU, plan for 2 December 2027, use the SME concessions and ask the national point of contact.
8. Treat the policy as a living document. When a tool changes, the register changes, and the team is told.

8. Conclusion

The firms that came through the 1990s well were not the ones with the best firewalls. They were the ones that wrote down, in language their staff used, what the new tool was for, gave people a safe way to use it, and put someone in charge. AI asks for the same thing and asks for it sooner, because what leaves through a chatbot cannot be recalled. One page of rules, one named owner, one honest list of what stays inside the firm, and the habit of checking before sending. That is the whole policy, and a small firm can have it in place by the end of the month.

REFERENCES

1. Gillespie, N., Lockey, S. et al. (2025) *Trust, attitudes and use of artificial intelligence: a global study 2025*. University of Melbourne and KPMG International.
2. BBC News (2016) 'Microsoft chatbot is taught to swear on Twitter', 24 March.
3. CIO Dive (2023) 'Samsung employees leaked corporate data in ChatGPT: report', 7 April.
4. Fraser, B. (ed.) (1997) *Site Security Handbook*, RFC 2196, Internet Engineering Task Force; superseding Holbrook, P. and Reynolds, J. (1991) RFC 1244.
5. The White House (1997) *A Framework for Global Electronic Commerce*. Washington DC, 1 July.
6. Lichtenstein, S. and Swatman, P.M.C. (1997) 'Effective Internet acceptable usage policy for organisations', *Proceedings of the 10th International Conference on Electronic Commerce*, Bled, Slovenia, 503 to 522.
7. Dastin, J. (2018) 'Amazon scraps secret AI recruiting tool that showed bias against women', Reuters, 10 October.
8. Information Commissioner's Office (2023, updated) *Guidance on AI and data protection*; Department for Science, Innovation and Technology (2023) *A pro-innovation approach to AI regulation*.
9. The Data Protection Act 2018 (Code of Practice on Artificial Intelligence and Automated Decision-Making) Regulations 2026, SI 2026/425, in force 12 May 2026.
10. Data (Use and Access) Act 2025, section 80, inserting Articles 22A to 22D into the UK GDPR, in force 5 February 2026; Information Commissioner's Office (2026) draft guidance on automated decision-making and profiling.
11. Regulation (EU) 2024/1689 (Artificial Intelligence Act), in particular Articles 4, 50, 57 and 99.
12. Regulation (EU) 2026/1744 (Digital Omnibus on AI), published 24 July 2026, in force 27 July 2026.
13. National Institute of Standards and Technology (2023) *AI Risk Management Framework 1.0*; NIST (2024) *AI 600-1: Generative AI Profile*.
14. ISO/IEC 42001:2023 *Artificial intelligence management system*; ISO/IEC 38507:2022 *Governance implications of the use of AI by organizations*.
15. IMD (2025) 'Morgan Stanley: AI maturity', entity profile, IMD Business School.
16. Chen, Z. (2023) 'Ethics and discrimination in artificial intelligence enabled recruitment practices', *Humanities and Social Sciences Communications*, 10, 567.
17. Equality Act 2010; Companies Act 2006, section 174.

REGIONAL AI READINESS

Consortium Partners & Funders

This briefing is produced by the AI Buddy Framework consortium, bringing together industrial, technology and community delivery partners with the backing of public funding.

Consortium partners

The organisations delivering the embedded adoption model described in this paper.

KERCKHOFFS Beyond Boundaries

BYDAYBREAK SOLUTIONS

SUNDERLAND COMMUNITY Soup Kitchen

Funded by

Public funding supporting the regional pilots referenced in this paper.

Funded by UK Government

TechFirst