SLAUGHTER AND MAY/

HOW SHOULD PRIVACY TEAMS MANAGE AI? IT TAKES A VILLAGE

JULY 2024

DIGITAL
Part of the Horizon Scanning series

A version of this briefing first appeared in the Privacy Laws & Business UK Report, Issue 134 (July 2024)

"New tech, old tricks" – those were the words of John Edwards about genAI in a recent speech. Given the unprecedented growth of ChatGPT, now hitting 100 million users, and that business uptake of AI is increasing at pace (with a nearly 40% increase in UK companies reporting using AI between 2022 and 2023 according to IBM research), it is not surprising that John Edwards was trying to reassure everyone "there are protections in place for people".

But how does that translate into practice for privacy teams that are having to move quickly to address the heightened legal and practical challenges posed by genAI without hindering business imperatives? Sometimes privacy teams will be leading the legal advice on use, training and deployment of AI, but, even where they are not, they inevitably have an important role to play. In this briefing, we look at some of the challenges posed by the current landscape, the evolving role of privacy professionals in relation to AI governance and how their existing knowledge and experience can be best leveraged as a driver for organisations' AI compliance.

Ongoing data privacy challenges

AI is not new and many of the data privacy challenges posed by AI are well recognised and documented. These include the risk of bias or discrimination, data minimisation, how to collect and use data compliantly throughout the lifecycle of AI models and issues around explainability and accuracy (as we discuss in our previous briefing). In the UK, the ICO has published a number of pieces of guidance covering these issues, many of which are helpfully listed in its response to the Department for Science, Innovation and Technology on its strategic approach to AI.

However, whilst regulatory guidance is helpful, it needs to be interpreted and operationalised on a case-by-case basis for each organisation and for each use case. In addition, whilst some issues have been litigated in the courts, such as in the Clearview AI decisions, and/or been the subject of regulatory enforcement, such as the ICO's enforcement action against Snap, for various reasons the 'meatier' issues have not yet been conclusively addressed.

AI continues to evolve and regulators have to keep up – a clear example is the launch of ChatGPT. The ICO responded quickly and pragmatically, engaging industry early on with its call for views on various issues raised by genAI. However, the ICO's final guidance is outstanding, so organisations are for now without the regulator's firm conclusions on a number of knotty issues around genAI.

Additionally, the UK's sectoral approach to AI regulation means a number of regulators are involved, which can create complexity and duplication of guidance. Having said that, the work of the Digital Regulation Cooperation Forum (DRCF), which includes the ICO, CMA, FCA and Ofcom, is demonstrating helpful collaboration between the UK digital regulators. For example, the recent launch of the DRCF's AI and Digital Hub promises to provide a way for businesses to obtain joined up advice from multiple digital regulators simultaneously.

Given the above, how can privacy teams best ensure that data privacy is not seen as a blocker to AI plans and innovation? And even better, how can they leverage GDPR compliance and existing governance processes to assist with AI deployment?

Role of the Privacy team

As organisations increasingly investigate and adopt AI solutions following the rush triggered by genAI, many organisations have turned to their privacy teams to take on legal oversight of AI given their familiarity with similar compliance principles, policies and processes (and the fact that personal data is often involved). However, it is crucial to avoid working in silos as AI (and genAI in particular) requires cross-team working, including with technical and specialist project teams within the organisation such as legal, IT, compliance, marketing business heads etc.

The ICO has echoed this view in its initial response to ChatGPT (in April 2023), where it called for privacy teams to collaborate closely with technical specialists to address the data security risks posed by genAI, such as the risks of model inversion (a machine learning security threat) and adversarial data poisoning (deliberate and malicious contamination of data to compromise the performance of AI systems). Whilst some organisations are setting up collaborative cross-functional working groups on AI, others are integrating AI working groups or councils into governance frameworks – although ownership and oversight may differ depending on the business and sector and its likely use cases for AI.

Enabling AI in a compliant manner

Many privacy teams continue to struggle with issues around visibility and being seen as 'blockers'. However, there are opportunities to leverage the business interest in AI to create a more proactive environment where privacy teams become involved early on, helping assess risk and ensuring privacy by design. Senior stakeholder support is very helpful to ensure this occurs, but processes can also assist in ensuring early engagement. For instance, supplier onboarding questionnaires and processes for data protection impact assessments can be leveraged to help ensure privacy teams are involved at the right time for AI. Going forward, it is likely that privacy teams will also play a role in considering the ethical and responsible use of AI, in collaboration with others across the business.

Information gathering and mapping

It is crucial for privacy teams to understand how AI is being, and is proposed to be, used throughout the business, certainly when personal data is involved and more generally if the privacy team has a wider remit. Although aimed at DPOs of EU institutions, the European Data Protection Supervisor's (EDPS's) orientations on genAI is useful reading for everyone, not least because the EDPS's role on the European Data Protection Board. It specifies that "comprehensive information" should be obtained including on "the origin of the datasets, the curation/tagging procedure as well as for any associated processing". The EDPS also states that DPOs should have "a proper understanding" of the life cycle of genAI models in the business and how they work, including inputs/outputs and the model's decision-making processes.

From a data privacy perspective, obtaining information such as this is important to help establish whether the organisation is acting as controller or processor (and/or a provider or deployer under the EU AI Act, if relevant), what personal data is being used and whether it is used compliantly or with risks that can/can't be mitigated (e.g. with contractual protections), how to explain the processing to individuals and generally to comply with the remaining GDPR provisions. Recognising this, the next chapter of the ICO's consultation series on genAI will focus on controllership issues.

Time spent producing an AI inventory is a good investment for many businesses so as to have a complete overview of what AI is currently used in the business. It is then of course key to keep this up to date by reflecting new uses which should be being highlighted by internal processes. Once there is an overall map of the AI in the business, existing AI use cases can be categorised according to risk thresholds. The risk categorisations in the EU AI Act may provide useful inspiration for determining what those thresholds could be.

Leveraging data privacy documentation and processes

DPIAs and supplier onboarding processes can be expanded to cover AI (as discussed above), providing a simpler process for business teams with no additional documentation to grapple with. Again, this is an area where the EU AI Act can be looked at for guidance, for instance by cross checking current processes against the requirements for a fundamental rights impact assessment to provide inspiration for additional questions that could be included.

Other data privacy policies, processes and documentation that can be leveraged, expanded or adapted to cover AI include privacy notices and audit questionnaires for third party suppliers. In some limited cases, there may even be situations where similar processes can be used for both GDPR and EU AI Act compliance – for example, when a high-risk AI system as defined under the EU AI Act involves automated decision-making as defined under the GDPR and triggers a requirement for human oversight under both Acts, the process for that oversight may be adapted to cover both.

Impact of EU AI Act

Although there is no clear indication yet that the EU AI Act will become the global gold-plated standard for AI, its effect will be felt by many UK organisations, whether directly caught or not. In addition to the risk categorisations referred to above, the transparency obligations placed on providers, including of general-purpose AI models, may filter through the UK market more generally, enabling greater information sharing and easier data privacy compliance. This may be particularly helpful in the context of 'hidden' AI in vendor offerings, which remain a difficult area for businesses to get sufficient information on to accurately assess risk.

AI literacy and training

It is likely that many DP teams will need to increase their AI literacy and/or develop a more technical understanding of how existing or proposed AI tools operate. More generally, there is a need for widespread training on AI throughout the business (even for example, to clarify what the business means by 'AI'). This can also act as a useful reminder of existing privacy policies and practices, some of which (e.g. use of personal data, house style rules etc) can be forgotten when using large language model chatbots such as Microsoft Co-pilot.

Use of standards

Privacy teams will be familiar with external standards in the context of their GDPR compliance and compliance with them is likely to be a helpful stepping-stone for achieving the requisite security levels for AI. For example, the ICO's general guidance on data security refers to resources from the UK's National Cyber Security Centre (NCSC), and the ICO's cyber/data breach enforcement actions against Tuckers and Interserve refer to industry standards such as NIST 800-53 and ISO27002.

There are specific data security concerns with genAI (see above) and compliance with external standards such as these may help. AI-specific standards are being developed (e.g. ISO42001 on AI Management (2023), US NIST Risk Management Framework for AI), although as yet there is no single standard that businesses can look to for AI, especially to satisfy the requirements of the EU AI Act. It will therefore take a little time for a market approach to develop as to the best or most helpful standards to use.

Conclusion

In our rapidly evolving AI landscape, privacy teams stand at a critical juncture. Businesses already differ in privacy maturity, and AI governance is another challenge to add to this, with AI governance being in its infancy for many.

That said, some businesses have already grappled with this, and have developed robust risk frameworks and processes, whilst others are deploying less formalised information sharing forums. However, it is clear that challenges remain across the board even for those with the most advanced AI governance.

Leveraging GDPR processes and programmes helps, but this is not a complete solution as it is crucial for privacy experts to actively seek the collaboration of others to ensure AI initiatives are tested against the business' appetite for risk across all areas, and not just privacy principles. Only then can this diverse village lawfully manage the machines, promoting fairness, transparency and accountability.

CONTACT

REBECCA COUSIN
PARTNER
T: 020 7090 4738
E: Rebecca.Cousin@slaughterandmay.com

BRYONY BACON
SENIOR PSL
T: 020 7090 3512
E: Bryony.Bacon@slaughterandmay.com

CINDY KNOTT
PSL COUNSEL AND HEAD OF DATA PRIVACY KNOWLEDGE
T: 020 7090 5168
E: Cindy.Knott@slaughterandmay.com