STEPHENSON HARWOOD

+

August 2025

AI AT WORK: NAVIGATING OPPORTUNITIES AND LEGAL RISKS IN THE EMPLOYMENT LIFECYCLE

From recruitment through to performance management and exiting employees, AI tools are being adopted by employers for use at all stages of the employment lifecycle. Whilst this heralds opportunity, it also creates potential pitfalls. In this article, we highlight key legal considerations for the use of AI in the employment journey, from both an employment law and data protection perspective. We address upcoming changes to the regulation of AI, including rules on automated decision-making ("ADM") introduced by the UK's Data (Use and Access Act) 2025 ("DUAA"), which became law on 19 June 2025, and offer practical tips for employers in navigating this evolving landscape.

AI REGULATION

There is currently no overarching statutory regulation of AI in the UK. Instead, the UK government has adopted a "pro-innovation" approach, setting out five core cross-sectoral principles for the UK's existing regulators to interpret and apply:

+ Safety, security & robustness
+ Appropriate transparency and explainability
+ Fairness
+ Accountability and governance
+ Contestability and redress

However, alongside this principles-based framework, UK organisations must still comply with domestic laws that address the development or use of AI, including the UK General Data Protection Regulation ("UK GDPR") and the DUAA. In addition, the European Union's comprehensive Artificial Intelligence Act ("EU AI Act"), with its extra-territorial application, remains highly relevant for many UK organisations. In force since 1 August 2024 and applicable in phases lasting until 2027, the EU AI Act applies to operators located outside the EU, if the AI systems or their outputs are used within the EU.

AI tools used in recruitment or for performance management are expressly categorised as high-risk AI systems under the EU AI Act, due to their potential to cause harm to individuals, unless certain conditions can be demonstrated (such as use of the AI system for a narrow procedural task only). High-risk AI systems are subject to stringent regulatory requirements both by the provider of the AI tool and by any employer seeking to deploy the AI system in its organisation; recognising that the use of such systems could have an appreciable impact on an individual's future career prospects and livelihood.

For comprehensive insights into AI fundamentals, regulatory frameworks, and practical applications, including dedicated episodes on high-risk AI systems, you can access the Stephenson Harwood AI podcast series here.

RECRUITMENT: PROMISES AND PITFALLS

Traditional recruitment processes have long been susceptible to discrimination or unconscious bias (e.g. favouring candidates who resemble the decision-maker or share similar backgrounds) and AI is unlikely to eliminate these issues. In fact, it can exacerbate them, particularly if a decision is reached by means of an automated decision-making process, with no meaningful human involvement.

Many employers use algorithmic tools to scan CVs for keywords or to assess suitability and screen out CVs at an initial stage. Such practices are not without risk. For example, it was reported that Amazon came under fire for use of a recruitment algorithm, in which AI systems trained on historical data developed a preference for male candidates, reflecting the historical trends of male dominance in the tech industry. Essentially, the algorithm taught itself that male candidates were preferable and penalised CVs that referenced women's achievements or all-women's educational institutions. The algorithm, reflecting the biases present in its training data, produced discriminatory outcomes.

In such scenarios where an AI tool is discriminating against certain groups, it's crucial to note that, in addition to possible non-compliance with the EU AI Act, liability under employment law would sit with the employer rather than the AI provider. Accordingly, when purchasing, onboarding and deploying certain tools, employers should take steps including:

- conducting thorough due diligence when purchasing AI tools, trying to gain an in-depth understanding of how the tools have been trained and how they work;
- ensuring they get adequate protections in their commercial contracts with the AI providers;
- giving careful consideration as to how they will use the AI tool in the workplace and possible unintended consequences that could arise;
- carrying out regular audits to help spot any trends or potential issues; and
- considering what level of human oversight and checks will be required.

This final point is particularly important for data protection compliance, where an employer is using AI tools in its recruitment processes. Subject to certain exceptions, the UK GDPR establishes a general prohibition on decision-making based solely on automated processing (i.e. where there is no human involvement in a decision-making process), that has a legal or similarly significant effect on individuals.

E-recruiting practices (e.g. automated screening of job applicants based on certain criteria with no human oversight of the process) would fall within the scope of this restriction, unless the automated decision could be said to be "necessary to enter into or perform a contract between the data subject and the controller". It could be argued that this would extend to a scenario where an employer makes use of an e-recruiting tool to shortlist possible candidates from a large pool of applicants and has been unable to identify a less privacy-intrusive method of achieving this.

However, when the relevant provisions of the DUAA come into force (likely at the end of this year), this will create a more permissive framework under the UK GDPR for ADM, and organisations will be able to apply ADM in wider circumstances unless special category data is involved (in which case the current restrictions remain). In theory, ADM involving personal data could be justified on the basis of legitimate interests provided that the organisation implements certain safeguards to uphold the rights, freedoms and legitimate interests of the individual, including to:

- provide individuals with information about the significant decision made about them;
- permit those individuals to contest (or make other representations about) such decision; and
- enable those individuals to obtain human intervention in respect of such decision, by someone with the authority and capability to change the decision.

Employers are also advised to:

- make frequent assessments of their personal data processing (including conducting regular quality assurance checks of any AI tools used) to check for any bias or discriminatory effects on individuals;
- record the degree of human involvement in their decision-making processes, as well as the safeguards applied, in any data protection impact assessments carried out; and
- be prepared to provide meaningful information about the logic involved in the decision-making process, and the envisaged consequences for the individual.

It is clear from the above that it is a delicate balancing act between upholding the rights and freedoms of individuals and leveraging the government's attempts to foster innovation in the UK by encouraging AI use. Unsurprisingly, the Information Commissioner's Office ("ICO") has identified ADM and the use of AI in recruitment as a key area of focus. In June 2025, the ICO launched its AI and biometrics strategy for 2025/26, which includes plans to publish a statutory code of practice on AI and ADM, as well as updated guidance on ADM and profiling in Spring 2026. A priority objective for the ICO is to set clear expectations for the responsible use of ADM in recruitment, particularly regarding transparency, discrimination and redress. To this end, the ICO is currently conducting a market study scrutinising the use of ADM in recruitment by major employers and recruitment platforms. This follows its 2024 audit engagements with developers and providers of AI-powered sourcing, screening, and selection tools used in recruitment. The ICO will publish its findings and regulatory expectations in due course.

For a deeper dive into the interrelationship of recruitment and AI, please listen to our podcast on the topic here.

PERFORMANCE: MANAGING AND MONITORING

Beyond recruitment, AI is making its mark on the management of employees throughout their tenure. AI-driven analytics can be used to monitor productivity, set targets and assess performance. Any dismissal based solely on performance metrics from an AI tool, without any human oversight, risks an unfair dismissal claim in which an Employment Tribunal would look closely at whether it was reasonable for the employer to rely only on the AI performance metrics as a sufficient reason for dismissal. There would also be data protection implications given the ADM risks outlined above, particularly if appropriate safeguards have not been implemented (such as ensuring the individual has the right to request a review and explanation for the rationale behind a decision, and having a process in place for that individual to challenge or appeal the decision).

Furthermore, AI can process vast amounts of data but may misinterpret context or reinforce existing biases – accordingly fairness, accuracy and accountability are paramount in the use of AI in this sphere. Algorithms might reward behaviours that disadvantage those with certain "protected characteristics" under the Equality Act 2010. For example, an algorithm used in an appraisal system awards high scores to those employees who demonstrate core competencies set by the company - such as active listening skills; creating rapport; using direct eye contact and confidence in public speaking. Whilst this may appear neutral as the same treatment is applied to everyone, such algorithms could be indirectly discriminating against those who, for example, have certain disabilities (e.g. hearing impairments, neurodiverse conditions or sight related issues (amongst others) which may result in them being awarded low scores). Not only would an employer be under a duty to make reasonable adjustments for a disabled employee but to avoid an indirect discrimination claim they would need to objectively justify the use of the algorithm as a proportionate means of achieving a legitimate aim. It would be tricky for them to show such objective justification unless they had audited for bias, and kept a "human in the loop" to give oversight to the appraisal outcomes.

Additionally, the use of facial recognition software in the employment sphere has created further risks. Let's take the case of Manjang v Uber Eats, in which Uber Eats introduced facial recognition software to verify drivers' identities to allow them to access work and pay. Mr Manjang, a black male driver, failed the facial recognition check and was permanently suspended from the platform. He complained that the facial recognition software was racially biased, placing people from ethnic minority groups at a disadvantage as false positive and false negative results were greater in individuals from ethnic minority groups. He requested human checks of the photograph submitted but this action was not taken and he brought claims for harassment related to race, victimisation and indirect race discrimination (which Uber Eats settled). A variety of studies have shown that facial recognition tools can be less accurate on non-white faces, such bias stemming from the datasets used to train the algorithms, which are often predominantly composed of white faces. As a result, facial recognition systems can misidentify people from certain ethnic minorities more frequently than white individuals.

It is accordingly reassuring that any ADM or profiling involving special categories of personal data (i.e. racial or ethnic origin, political opinions, religious or philosophical beliefs, trade union membership, genetic data, biometric data, health data, sex life or sexual orientation) will remain restricted when the DUAA changes come into force. Currently, conducting ADM involving special category data, with no meaningful human involvement, is only permitted under the UK GDPR if the organisation has obtained the individual's explicit consent or the processing is necessary for reasons of substantial public interest and a lawful basis applies.

With the additional HR risks presented by AI use, comes the administrative burden on organisations to ensure that individuals' rights are upheld, not least with the introduction by the DUAA of a new mandatory complaints-handling procedure. When the relevant provisions come into effect, individuals will have a new statutory right to raise a complaint to an organisation regarding their use of personal data and general UK GDPR compliance, with organisations mandated to have clear processes to facilitate complaints (including providing accessible means for individuals to submit complaints), acknowledge receipt within 30 days and respond "without undue delay".

EXITS: PREDICTING AND PREPARING FOR DEPARTURES

Using AI to monitor employee performance is increasingly common from looking at keystrokes, application usage, task tracking and even sentiment analysis to determine how engaged and happy employees are. Now, AI is even finding a role to predict the end of the employment relationship.

AI tools can be used to analyse historical HR data from previous employee departures - encompassing reasons for leaving, length of service, performance evaluations, compensation history, training records, absence rates, and results from engagement surveys - among other relevant metrics. Such tools could then identify trends and patterns within this dataset, allowing predictions for future employee attrition and identifying employees at risk of leaving. This could be of benefit to employers to get a "heads-up" before a resignation is tendered, allowing them to make a play to keep the employee if they want to retain the talent and/or investigate whether the employee has been accessing certain confidential information or speaking to more key clients than usual in recent weeks – and accordingly the employer could take necessary steps to protect the company.

However, as AI continues to automate and revolutionise employment processes, when dealing with personal data, organisations must pay close attention to two core principles of the UK GDPR: 'purpose limitation' and 'data minimisation'. To uphold the data minimisation principle, organisations must only collect and process the minimum amount of personal data necessary for a specified purpose. When conducting employee monitoring and/or using data analytics, this means ensuring that the data collected and used is adequate, relevant and limited (i.e. can the aim still be achieved with sufficient accuracy when applying data minimisation practices such as anonymisation, the removal of certain identifiers, or the use of fewer data points?).

Organisations must also ensure that any use of personal data for analytics purposes is compatible with the original purpose for which it was collected. This is particularly relevant if an organisation wishes to use a dataset to train and develop an AI model, where such purpose is more likely to differ from the original purpose for which the data was collected. The question of how broadly the "scientific research purposes" exemption (which applies a legal presumption of compatibility) can be interpreted to facilitate conducting data analytics remains to be seen, as we await further updates and guidance on the provisions of the DUAA. In the meantime, organisations must ensure that, where the new purpose is incompatible, consent to the processing is sought from the individual.

PRACTICAL TIPS FOR EMPLOYERS

It's evident that AI is already embedded in the employment lifecycle and given the direction of travel, we expect it to become even more ingrained in all ways of working. Accordingly, employers should be mindful of how they are using it and the risks it poses, this includes the following practical steps:

- Understand where and how AI is being used within your organisation.
- Conduct regular audits of AI tools to ensure they function as intended and are not inadvertently discriminating against particular groups.
- Develop clear policies and guidelines for AI use in HR processes, specifying the purposes for which AI can be used, the requirement for meaningful human oversight at key decision points and the importance of compliance with data protection principles.
- Develop clear policies and guidelines for AI use by employees in their routine daily work e.g. use of generative AI tools such as ChatGPT.
- Provide training for those using AI systems. Educate employees about the limitations of AI, the importance of data security, and the need to report concerns.
- Engage with AI providers to understand how their tools are developed and maintained, including the diversity of training data and measures taken to mitigate bias. Consider negotiating contractual protections with any AI providers.
- Get up to speed on the DUAA changes that will be coming into force and how the mandatory complaints handling procedure will impact your organisation. Consider updating your privacy notices, implementing an electronic complaints form and rolling out additional staff training. To see more information on this and a summary of other changes being introduced by the DUAA please click here.

If you would like our assistance in creating AI policies, delivering training or if you have any questions on the items in this alert please contact the authors noted below or your usual Stephenson Harwood contact.

ANNE PRITAM
Partner
+ 44 20 7809 2925
anne.pritam
@stephensonharwood.com

KATIE HEWSON
Partner
+ 44 20 7809 2374
katie.hewson
@stephensonharwood.com

ALISON LLEWELLYN
Senior Knowledge Development Lawyer
+ 44 20 7809 2278
alison.llewellyn
@stephensonharwood.com

LEANNE RAVEN
Senior Knowledge Development Lawyer
+ 44 20 7809 2560
leanne.raven
@stephensonharwood.com