

TUC generative AI policy toolkit

**A resource for union reps and staff in
devising organisational policies.**

Introduction

This Generative AI toolkit was produced by AWO for the Trades Union Congress.

The resource sets out an introduction to Generative AI, its use in the workplace, main issues and considerations relevant to the workplace and a discussion of legal and regulatory regimes. It also offers an initial checklist to be used by union representatives and anyone involved in devising organisational policies for the use of AI.

Generative AI is a fast-changing technology, and this guide engages with a variety of sources including academic research, investigations, media coverage and commentary, and policy analysis. The information is current as of 4 July 2024.

About AWO

AWO is the leading agency on modern technology, offering public policy and legal services. We are best known for our unique offering in rapidly developing areas of digital technology and data protection, including worked experience on legal issues arising from the use of AI. Our team have unparalleled expertise in these areas, which we use to help a diverse range of clients navigate complex areas of law in a rights-respecting, socially responsible and commercial way. We are regularly instructed on matters that shape technology for the better.

awo.agency

What is Generative AI?

Generative AI is a subset of artificial intelligence (AI) systems that produce text, audio, image and video output. The TUC's [AI draft Bill](#) sets out the following definition for AI systems:

An “artificial intelligence system” means a machine-based system that, for explicit or implicit objectives, infers from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.

As with all AI systems, developing Generative AI requires large amounts of data (known as training data) and significant computational power. A well-known example of a Generative AI tool is OpenAI's ChatGPT. As [disclosed by OpenAI](#), ChatGPT was trained with:

- information that is publicly available on the internet (regardless of copyright and potentially unlawfully)
- licensed information from third parties
- information that users or OpenAI's human trainers provide.

ChatGPT's model 'learned' patterns of text in the training data to predict the most likely combination of text based on a user's query. The same process applies to all Generative AI models. They are trained on large datasets and learn to predict output based on patterns within the training dataset, whether that's text, audio or image data. Once deployed, Generative AI models may continue to collect data for future training by capturing user interactions with the tool, such as user prompts and chat history.

Many people are familiar with Generative AI through consumer-facing products. ChatGPT had 100 million users [within two months](#) of launch, in part because it was easy-to-use and free online. However, Generative AI tools are increasingly packaged as products for business use (known as enterprise products).

Few companies train Generative AI models 'from scratch'. Instead, they leverage best-in-class Generative AI models – such as OpenAI's GPT-4 (from the same family of models as ChatGPT) and Anthropic's Claude – to create sector-specific tools. These pre-existing models can be plugged in directly or adapted based on sector and organisation-specific data.

As models are adapted for enterprise applications– especially for sector and organisation-specific needs – they will likely require specialised, high-quality datasets relevant to those sectors or specific organisations.

How is Generative AI used in the workplace?

Generative AI has wide-ranging use cases. These can appear even more expansive as similar underlying technology is applied to tasks that require distinct human skillsets. To set out an overview of the technology – and the most likely ways workers will come across it in the workplace based on user input and functionality – three broad categories are provided below.

Synthetic voice and image/video likeness generation

Digital replicas (sometimes called 'cloning') are a Generative AI task. They often require biometric data to construct convincing outputs.

In the US, background actors [have described](#) participating in face and body scanning, including on the set of the Disney+ series *WandaVision*. Workers described being caught off-guard and not informed about how these replicas would be used. In other instances, training data is collected from publicly available data (regardless of permissions) and from license-holders, such as studios.

These issues have been at the heart of Generative AI-focused labour negotiations including the SAG-AFTRA (the US film and TV actors' union) strike and [subsequent deal](#), and the recent SAG-AFTRA [deal with Replica](#), an AI voice technology company that allows professional voiceover artists to licence a digital replica of their voice.

Examples

- Resemble.AI [partnered](#) with Netflix to produce Andy Warhol's voice for *The Andy Warhol Diaries*.
- A recent lawsuit [alleges](#) that Amazon Studios used Generative AI to replicate movie actors' voices in the film *Road House* in violation of collective bargaining agreements. (Amazon Studios denies the claims).

Content generation in response to user prompts

This is the most common functionality associated with Generative AI due to ChatGPT. Such models generate text (e.g. OpenAI's ChatGPT), text-to-image (e.g. OpenAI's DALL-E; Stability AI's Stable Diffusion) and text-to-video (e.g. OpenAI's forthcoming Sora). Google's Gemini is known as a 'multimodal' large language model as it operates across text, image, video and code. For example, developers can use Gemini Code Assist, which uses Generative AI to generate and debug code.

Examples

General business function examples

- **Social media, marketing, and communication teams:** Many workplaces have begun to integrate ChatGPT in these departments. Tasks range from idea generation to content drafting.
- **Legal teams:** Companies such as [Klarna](#) use ChatGPT Enterprise within their legal departments to draft a first version of contracts.
- **Audit and general functions:** The global consultancy PricewaterhouseCoopers ([PwC](#)) recently signed a landmark deal with OpenAI to become its biggest customer and the first reseller of ChatGPT Enterprise. PwC will work directly with companies – including audit clients – to develop custom GPTs for their businesses.
- **Customer service:** ChatGPT can also be used as the interface to query specific databases (e.g. to summarise company policies for a customer service chatbot). While customer service chatbots are a common use case for Generative AI, there are substantial concerns about their accuracy.

Additional examples

- **Visual design:** Designers describe using DALL-E to create draft images that enable the discovery and ideation process.
- **Software development:** A [survey](#) by StackOverflow (the knowledge-sharing platform and community for software developers) found that over 70% of developers use or plan to use code generation tools.

Content summarisation of user-provided data.

Generative AI can also be applied to user-provided datasets to shorten, extract, and reformat data.

For example, many Generative AI tools offer summarisation features. Here, the technology is applied to user-provided data – such as a lengthy document or audio from a virtual meeting – to produce a text summary. In ChatGPT Enterprise, users can upload CSV files with numerical data to produce data visualisations, enabling no-code data analysis. It is important to note that these are still Generative AI tasks.

Users may have worked previously with software with a similar user interface and functionality – for example, creating charts from data tables in Microsoft Excel. However, such tools are not based on Generative AI. This is a critical distinction: Generative AI introduces different risks and limitations, including inaccuracy.

Examples

- **Customer Service (Call Centres):** [Genesys](#), a call centre software provider, offers customer service agents a Generative AI tool that produces customer chat summaries, company policy summaries, and conversation prompts. It also analyses agents' interactions and offers coaching. (It is unclear if performance-related analysis and advice is shared with management).
- **Human Resources, including recruiting:** The software company [Allganize](#) applies Generative AI to summarise job candidate applications, among other recruitment functions.

Companies package content generation and summarisation functionality together along with other AI models to market 'virtual assistants' or 'AI assistants'. AI assistants allow users to interact with a variety of software functionality using natural language.

A key example of this is offered by Microsoft (a main investor and OpenAI partner), which has introduced 'Copilot' in Microsoft 365. [Copilot](#) is a virtual assistant powered by OpenAI's GPT-4o, and offers Generative AI augmentation in traditional Microsoft Office products. Microsoft [advertises](#) Copilot functionality with the following examples:

- In Microsoft Outlook, users can use Copilot to summarise email threads and draft new email (after specifying length and tone).
- In Microsoft Teams, Copilot can summarise meetings in real time. Users can also enter prompts in the Copilot box (e.g. 'What questions can I ask to move the meeting forward?').
- In Microsoft Excel, users can specify (in text) for data to be transformed in a table or apply a specific function.

Issues and considerations for union representatives

Training datasets

It is likely that most, if not all, best-in-class Generative AI models have been trained on datasets with copyrighted content, without providing remuneration or acknowledgement to original creators.

Companies are reluctant to share details about training data used to develop their models. However, we know this often involves [scraping](#) sections of the internet, including custom-made tools to transcribe YouTube videos, and existing publicly available datasets. For image-generator models, this includes LAION-5B, a collection of web links to 5.85 billion images scraped from the internet. These images may be publicly accessible, but have not specifically approved this usage, potentially violating copyright, and with no remuneration or acknowledgement being given to the original creators. (The website '[Have I been trained?](#)' allows artists to identify if their images are in the LAION-5B dataset. LAION-5B is at the heart of an ongoing class action [lawsuit](#) brought by artists against Midjourney and StabilityAI, which used LAION-5B as part of their AI training datasets.)

For models trained on large text datasets, this includes Books3, a dataset with 196,000+ books in text format, many protected by copyright. Books3 is at the heart of class action [lawsuits](#) against OpenAI, Meta, Nvidia, and Databricks. Popular training datasets are supplemented by additional data collected directly by developers. For instance, OpenAI's GPT models include data scraped from the social network [Reddit](#). Such datasets may also include copyright-protected content.

In copyright complaints, a key area of focus is on the direct reproduction or similarity of generated output with copyrighted work. In The New York Times [lawsuit](#) against OpenAI and Microsoft, a part of the copyright infringement complaint is that ChatGPT produces near-verbatim output of Times' content. Generative AI tools can sometimes reproduce sections of copyrighted text or create content in the style of a specific artist.

Generative AI developers have defended the need to use copyrighted content in model training. To a House of Lords committee, [OpenAI stated](#):

Because copyright today covers virtually every sort of human expression—including blog posts, photographs, forum posts, scraps of software code, and government documents—it would be impossible to train today's leading AI models without using copyrighted materials.

While Generative AI start-ups have described copyright as an intractable obstacle, new initiatives are exploring training datasets that are completely in the public domain. These include [Common Corpus](#), a multilingual text dataset released in March 2024 and supported by the French government, and [Fairly Trained](#), a non-profit that offers certification to large language models (LLMs) that only use public domain content or copyrighted content without appropriate permissions.

These [exploratory initiatives](#) signal a shift in future approaches but –at the time of writing – these have yet to be backed by major model developers.

As media organisations enter into partnerships and licensing agreements with Generative AI developers, workers may be shut out of discussions and benefits.

Many media companies are entering into partnerships with Generative AI developers to gain lawful access to copyright-protected content. OpenAI has [struck deals](#) with Vox Media, the Financial Times, The Atlantic, [TIME](#), News Corp (parent company of The Wall Street Journal), and Axel Springer (parent company of Business Insider and Politico). These deals provide OpenAI with access to companies' archives for training data, amongst other product and licensing provisions.

Workers often [learn of](#) partnership details through public material and press releases. Members of the Atlantic Union issued a statement detailing the lack of transparency about what the deal entails and how it would affect members' work. Workers have expressed [concerns](#) about whom these deals benefit, whether they are short-term fixes, and their benefits for writers or journalism more broadly.

Other companies are finding ways to collate licensed data sets. Imaging companies Adobe, [Shutterstock](#) and Getty Images are sourcing material for AI training direct from rights holders in their existing networks. Compensation for creators is low, even compared to the low prices they already command for stock. For example, Shutterstock are estimated to pay photographers less than one US cent every 6 months on average for licensed photographs.

Generative AI's content moderation training datasets are linked to known exploitative global labour practices.

The latest GPT models, including ChatGPT, are commercially viable because they have an AI-enabled safety system. An earlier model, GPT-3, was prone to generating toxic content like hate speech. This is partially caused by large training datasets with toxic content (discussed further in a subsequent section). The latest GPT models use AI-based safety mechanisms to detect such content before it reaches the user.

However, such AI safety systems are created using additional training data with examples of toxic content. A TIME investigation found that OpenAI [outsourced](#) this work to Kenya, where workers were paid less than \$2 a day to label textual descriptions of sexual abuse, hate speech, and violence. Kenyan workers have since filed a [petition](#)

with the Kenyan government calling for an investigation into exploitative conditions for content moderators behind AI systems, describing the psychological effect of the work, low pay, and abrupt dismissals.

Generative AI has alarming environmental costs.

Generative AI systems need tremendous amounts of energy. [Research](#) estimates that such models require around 33 times more energy than machines running task-specific software. Kate Crawford, a leading scholar on the social implications of AI, [writes](#): "It's estimated that a search driven by Generative AI uses four to five times the energy of a conventional web search. Within years, large AI systems are likely to need as much energy as entire nations."

In addition, training and deploying Generative AI models requires water (including freshwater) for on-site cooling and off-site electricity generation. [Research](#) estimates that training GPT-3 required 5.4 million litres of water. Beyond training, GPT-3 would need to 'drink' a 500-millilitre bottle of water to produce 10-50 responses. These impacts are felt immediately by communities that host data centres: it is estimated that Microsoft's data centre [used 6 per cent](#) of the water supply in West Des Moines, Iowa when it hosted OpenAI's GPT-4 training, drawing from the same supply that provides drinking water to the city's residents.

Generative AI enterprise products may alleviate employers' concerns about data security and data leaks. However, there may still be linked risks related to workers' input data and iterative training of Generative AI products.

When ChatGPT was released late 2022, many organisations [banned its use](#) in the workplace due to concerns about data leaks. Such concerns focused on employees providing confidential company data, including code, to ChatGPT, and OpenAI using these inputs as training data for further iterations of ChatGPT.

At Amazon, a senior counsel [wrote](#): "I've already seen instances where its output closely matches existing [internal confidential] material. Similar concerns were one basis for the Italian Data Protection Authority's 2023 ban on ChatGPT, which [cited](#) concerns about the processing of users' personal data (through interaction with ChatGPT) for use as training data. [The ban](#) was lifted when OpenAI introduced an opt-out option to exclude chat and chat history from training data, alongside other features.

As Generative AI is increasingly packaged and deployed through existing enterprise software (such as the Microsoft Office suite), it often aligns with existing product privacy policies and data security regimes. This can alleviate employers' concerns about commercial data leakage. However, workers should continue to assess the role of user-provided input data for continuous training. For example, Adobe recently updated its Terms of Service (ToS) agreement to access users' content for machine learning techniques to improve its service and products. After substantial backlash, the company

updated [its ToS](#) to specify it does not permit the company to use users' or their customers' content to train Generative AI models.

The examples above focus on input data used by Generative AI developers to improve core models deployed across a wide range of organisations. Workers should also consider a related issue: that their use of the Generative AI product provides the specific (and potentially essential) data required to improve the product's accuracy and relevance for their specific organisation.

Deployment

Generative AI is an experimental technology with significant inaccuracy, even in enterprise products. Given limitations around accuracy testing and statistics, workers play a crucial role in determining the reliability of the technology in practice.

Generative AI produces plausible output that may or may not be accurate, including output that is completely fabricated. In the case of text generation, output is grammatically correct, which can make it more challenging to identify inaccuracies quickly. These are commonly known as 'hallucinations' although some call this a charitable term; [more accurate](#) terminology could be 'fabrication' or 'falsification'.

This has real-world consequences. In 2022, Air Canada's chatbot fabricated an airfare policy in conversation with a customer, who proceeded to act on the incorrect information. A small claims tribunal [ordered](#) Air Canada to pay compensation. "While a chatbot has an interactive component, it is still just a part of Air Canada's website. It should be obvious to Air Canada that it is responsible for all the information on its website. It makes no difference whether the information comes from a static page or a chatbot."

To improve model accuracy, AI developers use techniques such as Retrieval-Augmented Generation (RAG), which enhances the accuracy and reliability of Generative AI models. [RAG](#) does not re-train the model, which is expensive and time-consuming. Instead, it swaps out the documents used by the model when it retrieves information to promote credible and relevant source material. Such techniques are likely used in many sector-specific Generative AI applications with specialised knowledge requirements. However, such tools may still have substantial accuracy issues.

A recent Stanford [study](#) of two legal research products – WestLaw AI-Assisted Research and Lexis+AI – revealed major inaccuracies. Lexis+AI had an accuracy rate of 65 per cent while WestLaw's product was only accurate 42 per cent of the time and hallucinated more than 34 per cent of the time. Such hallucinations [included](#) fabricating references, misapplying law (citing cases that had no bearing on the query and fabricating their relevance) and misstating case law (describing the opposite decision to that which was reached). Both products [advertised](#) using techniques such as RAG.

It is challenging to understand the accuracy of Generative AI tools. Developers may share accuracy test results but, even when these figures exist, they carry substantial limitations: testing methodology is not shared and may not be reproducible; real-world testing (especially given different contexts and use-cases) may vary from developer testing conditions; and there is no external, independent verification of test results. In the case of the Stanford legal research study, the developers described their technologies as “hallucination-free” and “eliminating” or “avoiding” hallucinations. However, the research findings indicate that this was far from the case.

Many Generative AI products caveat their output. Microsoft [writes](#):

“Copilot for Microsoft 365 generates insights and formulas and explains them in a fluent, grammatically correct way, but the content it generates can be inaccurate or inappropriate. It can’t understand meaning or evaluate accuracy, so be sure to read over what it writes, and use your judgment.”

This is a severe limitation and a trade-off. Generative AI tools may offer productivity advantages, but this should factor in the need for human review. While this may be straightforward for short text such as email, it would be more challenging for complex research or lengthy Excel sheets.

Generative AI may produce biased output.

Bias is a concern across AI systems. For example, a 2023 Bloomberg investigation [revealed](#) that Stable Diffusion perpetuated racial and gender bias.

“[H]ardly any women [pictured] have lucrative jobs or occupy positions of power. Women made up a tiny fraction of the images generated for the keyword “judge” — about 3% — when in reality 34% of US judges are women.”

Such bias can be introduced through algorithm development and through continuous data collected during deployment. It is often linked to training datasets, as algorithms learn patterns during model training.

A 2021 audit of LAION-400M— a training dataset with hundreds of millions of images scraped from the internet — [found](#) that it contained troublesome and explicit images and alt-text pairs of rape, pornography, malign stereotypes, racist and ethnic slurs, among other problematic content. A more recent Stanford University study on LAION-5B identified hundreds of examples of child sexual abuse material (CSAM) in the data, indicating that these issues persist.

Generative AI model bias is a known issue, and some companies have tried to address it [clumsily](#). A previous version of Google’s Gemini overwhelmingly generated pictures of Black, Asian and minority ethnic persons, including for prompts requesting images of Nazi-era German soldiers and the US Founding Fathers.

Bias continues to be a pressing problem in Generative AI models. A 2024 UNESCO [analysis](#) of LLMs showed “unequivocal” evidence of bias against women, homophobia

and racism. In a [recent letter](#) in *Nature*, academics also expressed concerns about the limited cultural diversity in Generative AI systems, which could perpetuate Anglo-American bias.

Workers are already describing the impact of Generative AI on the labour market, especially through job loss in creative industries.

Workers are already [describing](#) experiences with deskilling and job loss. These span workers on social media and marketing teams, copyeditors, illustrators, and voice actors who have [lost](#) audiobook narration work after studios partnered with voice generator start-ups.

Descriptions of Generative AI in the workplace offer insight into how roles can change prior to redundancies. A recent BBC report described a tech company that used ChatGPT to replace most of its 60-person writing and editorial team. Initially, ChatGPT was introduced to generate an outline for writers. Eventually, the technology was used to draft all content. The editor – at this point, the only person remaining – became responsible for [editing](#) AI-generated output to make it more engaging and human-like.

Brian Merchant, a technology and labour journalist, [described](#) the process as follows:

“[This is] exactly how management can use Generative AI to degrade and replace jobs—it’s a lengthy process, with junior and precarious workers replaced first, and more senior jobs deskilled into ones tasked with overseeing AI output.”

What does Generative AI do well?

As of now, there is limited applied research on Generative AI in the workplace. In a study conducted on 758 consultants at Boston Consulting Group, the global management consultancy, researchers conclude that overall, participants benefited from AI augmentation with GPT-4. However, while Generative AI could complement and potentially displace human work on specific tasks, the technology could be inaccurate and degrade human performance on other tasks – even tasks of seemingly comparable difficulty.

For instance, on a task requiring spreadsheet data interpretation, consultants who used ChatGPT performed [more poorly](#) than those who did not. Participants using AI produced more consistent and high-quality outputs, yet had far less variability in their ideas. The authors [conclude](#):

“While some organizations might prioritize consistently high average outputs, others might value maximum exploration and innovation.”

Similar limitations of Generative AI for creative output have been echoed in other studies. For example, [research](#) with professional comedians showed that Generative AI models could generate quick content and structure for comedy but could not move away from “bland and generic” outputs, even with participant intervention. Participants

described homogenous outputs, without the point-of-view or element of surprise that characterise effective comedy.

Together, this suggests that AI may be useful with specific tasks, its utility should be assessed on a task-by-task basis. The technology may also help standardise outputs in a way that doesn't allow for creative or unexpected outputs.

Workers may be increasingly vulnerable to Generative AI-enabled scams

Earlier this year, a worker at the British engineering firm Arup fell victim to a £20 million [scam](#). The sophisticated attack used Generative AI to create digital clones (also known as 'deepfakes') of the management team on a video call, which convinced a Hong Kong-based employee in the finance department to transfer out funds. Reports of similar scams are on the rise and may be especially relevant to workers such as call centre operators.

Overview of legal and regulatory regimes

This section of the toolkit is designed to provide an overview of the legal and regulatory landscape which could be relevant to the use of Generative AI in the workplace.

This is not intended to facilitate legal or compliance analysis of any particular Generative AI tool, nor does it seek to be exhaustive. Indeed this is not possible in the abstract since:

- Generative AI can be put to almost any use: there are a wide range of legal frameworks that might be applicable for your use-case; and
- The legal issues involved are very complex and novel. Whilst this guide can help signpost you, it is quite likely that you will need some legal advice, especially if you want to challenge how a Generative AI tool is being used in your workplace.

This section provides union representatives and others with the tools necessary to identify the legal frameworks or issues related to the Generative AI use cases they are concerned about, so that appropriate further research can be undertaken and – if necessary – legal advice can be sought.

Legal rights

It is crucial to understand the legal rights of workers and unions in the context of Generative AI implementation at work. Key legal considerations include:

- Data Protection issues
- Intellectual property rights
- Equality Act considerations
- Health and Safety law
- Sector-specific regulations

Awareness of these considerations can help unions and workers to (i) hold employers accountable for complying with the law, and (ii) negotiate for collective agreements that exceed minimum legal requirements.

The sections below discuss the key legal areas to help you identify whether they are relevant to your Generative AI use-case.

Data Protection

Generative AI systems often, but not always, involve the processing of personal data. It is vital to identify *if* a specific Generative AI use-case involves personal data processing. If it does, then the UK General Data Protection Regulation (UK GDPR) will apply to that processing and could be an important tool.

Does the Generative AI system involve personal data processing?

'Processing' means *any* operation on personal data, regardless of whether it is automated. The definition is very broad and includes any creation, consultation, communication, storage or deletion of data.

'Personal data' means *any* information relating to a living, identifiable person (usually one or more employees in this context). The key question is whether a person can be 'identified' – even if only in the sense of being 'singled out' by a system. So even if employees' real names are replaced with pseudonyms or codes in a system, if they can be singled out by it, then their personal data is being processed.

Generative AI systems can process personal data in different ways, for instance:

- Input data which is personal data may be used to train the Generative AI system. For example a Generative AI system might use historic customer service chat logs between employees and customers for training purposes.
- Users may input personal data directly into the system e.g. as a query or prompt.
- Personal data can be generated as a data output, either directly or as an inference. For example, a Generative AI system might generate predictions about employees' performance in different kinds of role.

In the employment context, Generative AI will very often involve the processing of personal data. For that reason we address UK GDPR requirements in some detail.

If the Generative AI system *is* processing personal data, then the following provisions of the UK GDPR are of particular relevance.

Article 5 UK GDPR principles

These principles govern the processing of personal data and include:

- Lawfulness, fairness and transparency:** This principle¹ means that processing of personal data should be lawful and fair, and it should be transparent to workers how and why their data is being collected and used by their employer. Where Generative AI involves personal data processing, employers must communicate to workers about the processing of their data, including the

¹ Article 5(1)(a) UK GDPR.

purposes of such processing, in a concise transparent, intelligible, and easily accessible manner.²

- ii. **Purpose Limitation:** The purpose limitation principle³ says that: *“personal data shall be collected for specified, explicit, and legitimate purposes and not further processed in a manner that is incompatible with those purposes.”*⁴ This means that personal data cannot be used for purposes incompatible with the original purposes communicated to workers. This places a significant limitation upon the employer’s ability to further process data in the context of Generative AI in any way it likes, whether that takes the form of using worker data as input data in Generative AI systems or using output data from Generative AI systems. An example of an incompatible purpose might be an employer using worker shift pattern data to feed into a new product or service it is developing.
- iii. **Data Minimisation:** The data minimisation principle says that: *“personal data shall be adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed.”* Generative AI systems often consume large volumes of data for training purposes to enable it to make predictions, generate insights and create content. This can raise data minimisation concerns, although it may enhance compliance with data accuracy.
- iv. **Accuracy:** The accuracy principles says that *“personal data shall be accurate and, where necessary, kept up to date; every reasonable step must be taken to ensure that personal data that are inaccurate, having regard to the purposes for which they are processed, are erased or rectified without delay”*. Generative AI has a propensity to ‘hallucinate’ or invent information which could lead to inaccurate outputs, so this could be an area of concern.

Other principles include storage limitation and data security.

Finally, under the Accountability principle, Article 5(2) UK GDPR, the burden is on employers to *demonstrate* compliance with the Article 5 principles and with the UK

² Article 13 UK GDPR stipulates specific information that must be provided to data subjects, including workers, upon collection of their data by the employer. This includes:

- Information about the worker data that is being collected and processed by the employer,
- the purposes of the processing, including the lawful basis for processing,
- how long the data will be stored for,
- the rights that the worker can exercise over that data including the right of access; and
- whether the data is going to be used for ADM, including profiling. If it is, then the employer needs to provide meaningful information about the logic involved to the worker.

How the information is provided may vary from organisation to organisation, for instance information could be provided by a worker privacy notice or worker-facing data protection policy.

³ Article 5(1)(b) UK GDPR.

⁴ Recital 50 GDPR also sets out: *“The processing of personal data for purposes other than those for which the personal data were initially collected should be allowed only where the processing is compatible with the purposes for which the personal data were initially collected.”*

GDPR more broadly⁵. Documenting compliance with these principles should also form part of a Data Protection Impact Assessment (DPIA), if a DPIA is needed.

Articles 4, 6, 7 UK GDPR (lawfulness of processing)

Generative AI that involves the processing of worker (and of others such as job applicants and service users) personal data requires a lawful basis. Six bases are available, and they are exhaustive. Those most relevant are explored below.

- **Legitimate interests:** An often-cited legal basis for the use of Generative AI by employers is legitimate interests (LI). If an employer is relying on this to justify its processing, then it must inform workers in its privacy notice/policy. Workers can object to their data being processed, and if such an objection is raised, an employer must stop processing the data unless it can show that its own compelling interests outweigh those of the worker, which may not always be easy to demonstrate. The ICO recommends employers always conduct a “legitimate interests assessment” (LIA) and record its outcome when processing based on LI. This is a three-part test. Firstly, the employer must identify a purpose that counts as a legitimate interest. Secondly, the employer must identify that the processing is actually necessary for the purpose it has outlined. Thirdly, the employer must conduct a balancing test. This means that an employer must consider the interests and fundamental rights and freedoms of the individual and if these override the legitimate interest identified. This includes the nature of the personal data, the reasonable expectations of the individual, and the likely impact of the processing on the individual and whether any safeguards can be put in place to mitigate negative impacts. Where an employer plans to do something new or innovative with personal data (as an example, training AI models for cloning performers’ voices), then this is highly unlikely to be within individuals’ expectations for how their data will be used. An employer would need to build in a high level of safeguards to justify this kind of processing, including providing individuals with the ability to opt out of their data being used in this way, and offering appropriate levels of compensation.
- **Consent:** This is *unlikely to be a valid legal basis in the employment context*, apart from in extremely limited circumstances. Consent under the UK GDPR has strict requirements. For example, consent must be ‘freely’ given but a worker cannot freely give their consent due to the power imbalance which is inherent in the employment relationship. In non-traditional employment relationships, such as in the case of performers, consent will not be valid if the individual does not have any real choice, feels compelled to consent, or will endure negative consequences if they do not provide consent.

⁵ Note that employers as data controller are also responsible for employee use of AI and ensuring that such use complies with all aspects of the UK GDPR, although this is outside scope of this document.

- **Contract:** The employer may claim that the processing is necessary for the performance of employment contracts, but this should be interrogated: only processing that is *truly* necessary for the performance of the core of the employment contract can be justified in this way.
- **Sensitive data:** Where certain categories (like race, health or trade union membership) of personal data are being processed, then Article 9(1) UK GDPR prohibits such processing unless an employer can rely on an exemption under Article 9(2) UK GDPR. The exemption frequently relied upon is that the data is needed to fulfil the employer's obligations under employment law. To rely on this basis, the employer must have an appropriate policy document in place setting out how the employer will comply with the UK GDPR principles on processing such data, as well as applicable retention and erasure policies.

Again, an employer is unlikely to be able to rely on the "explicit consent" exemption under Article 9(2) UK GDPR to process such data of an employee.

Article 22 UK GDPR (Automated Decision Making or ADM)

Generative AI systems could potentially be used to make automated decisions about workers.

Workers have the right under Article 22 UK GDPR to not be subject to *solely* automated decisions with *legal or similarly significant* effects (such as a promotion or being dismissed).

Processing of personal data is automated where it occurs without human intervention and where it produces legal effect or significantly affects an individual. This includes profiling, which is automated processing of personal data to evaluate an individual's personality, behaviour, interests and habits to make predictions or decisions about them (such as to analyse or predict their performance at work, economic situation, health, personal preferences or interests, reliability or behaviour, location or movement).

Article 22 is an important restriction on the ability of employers to use Generative AI to make decisions about workers. In the limited circumstances where an employer may lawfully be entitled to make solely automated significant decisions about employees, there is a right to get human intervention and to contest the decision.⁶ The UK GDPR also provides workers with the right to be provided with "meaningful information about the logic involved" in such decisions. Recent case law from Europe has established that this means entitlement to model weights where AI is being used to set pay and allocate work.⁷

⁶ Article 22(3) UK GDPR.

⁷ Amsterdam Court of Appeal decisions against Uber in April and October 2023.

Article 35 UK GDPR (Data Protection Impact Assessments)

Under Article 35 UK GDPR, employers are obliged to conduct a DPIA for types of processing likely to result in a high risk to individuals' interests.

The Information Commissioner's Office (ICO), which is the UK's regulator for data protection, has also outlined certain criteria to assess if processing is high risk. Such criteria includes innovative technology which the ICO describes as *"processing involving the use of new technologies, or the novel application of existing technologies (including AI)"*. Moreover, according to the ICO, if processing involves *"vulnerable data subjects"*, then if any one of the other criteria are met, a DPIA should be conducted. Workers are deemed *"vulnerable data subjects"*⁸. Thus, the use of Generative AI in the workplace is highly likely to trigger the need to conduct a DPIA.

The DPIA duty requires an employer, in advance of processing data, to consider and identify the data protection and related risks (which could include risks of discrimination) that may arise from the processing, and to put measures in place to mitigate those risks. If the employer cannot mitigate the risk, it will need to consult with the ICO.

When conducting a DPIA, the controller, *"where appropriate, shall seek the views of data subjects or their representatives on the intended processing"*. In a workplace context, the relevant stakeholders would include workers and/or their representatives. If an employer refuses to consult with workers and / or their representative, they should [provide valid reasons](#) why such a consultation is not appropriate.

Best Practice elsewhere

There is some overlap between DPIAs and provisions of the EU AI Act which also deals with "high risk" systems. The EU AI Act, although not binding in the UK, could serve as best practice for identifying when a DPIA is needed for a Generative AI system and what it should cover. The EU AI Act's provision for Fundamental Rights Impact Assessments (FRIA) to be conducted for the first use of "high-risk" systems is also something that unions may wish to negotiate for.

Article 27 of the EU AI Act provides helpful guidance on how to conduct an FRIA⁹. A template will also be developed by the EU's AI Office to assist with conducting these assessments.

⁸ Recital 75 GDPR.

⁹ Specifically, a FRIA needs to include: a description of the deployer's processes in which the high-risk AI system will be used in line with its intended purpose; a description of the period of time within which and the frequency with which each high-risk AI system is intended to be used; the categories of natural persons and groups likely to be affected by its use in the specific context; the specific risks of harm likely to have an impact on the categories of persons or groups of persons identified pursuant to point (c) above, taking into account the information given by the provider pursuant to Article 13 of the EU AI Act; a description of the implementation of human oversight measures according to the

There are other useful best practice tools such as IFOW's [Good Work Algorithmic Impact Assessment](#). Unions could negotiate for such impact assessments to always be carried out alongside a DPIA or incorporated within the DPIA.

Articles 15 – 22 (Data Subject Rights)

As data subjects, employees have a range of rights which can help to understand your Generative AI use-case and – potentially – place constraints on its use which benefit employees:

- a right to obtain a copy of their personal data, and other supplementary information such as “meaningful information about the logic involved” in automated decisions, from their employer within a one-month period upon requesting it (known as a Subject Access Request, or SAR);
- a right to have their data corrected or erased by the employer in certain instances;
- a right to object to processing, including processing based on legitimate interests; and
- a right not to be subject to automated decision-making, including profiling.

Unions are increasingly active in developing standard forms and templates to help people exercise these rights more easily. Contact your union to find out whether they have such a tool available that could facilitate your exercise of one or more of these rights.

Action checklist

- Determine whether the Generative AI system involves the processing of personal data.
- If so, consider whether the employer has ensured compliance with the UK GDPR? In particular, consider:
 - Has the employer been transparent about the use of AI in the workplace? Consider the explainability of AI systems used in the workplace.
 - What legal basis is the employer relying on?

instructions for use; the measures to be taken where those risks materialize, including arrangements for internal governance and complaint mechanisms.

- Is Generative AI used for any automated decisions, including profiling?
- Has a DPIA been conducted and were unions and workers consulted on it?
- How are workers' rights facilitated by the employer?
- How has the employer documented compliance with legal obligations?

- Consider existing best practices/impact assessments to supplement the DPIA framework.

- If you are concerned that there is a breach of obligations/rights under the UK GDPR, or if you think that the use of data rights could help you negotiate a better outcome for employees, seek legal advice/support from a data protection lawyer.

Intellectual Property

The use of Generative AI systems may engage intellectual property (IP) rights issues, depending on the circumstances. IP rights are property rights that include the use of creative works and personal rights, including:

- Patents (inventions) and designs
- Copyrights (literary/artistic works)
- Performance rights
- Trademarks (brand-names, symbols, logos)
- Designs

Typically, traditional employment contracts contain provisions that works and inventions made in the course of employment belong to the employer, not the employee. However, workers may retain some IP rights in their work, including in the form of moral rights (although these can also be waived by agreement). Moral rights recognised in the UK include:

- The right to attribution, i.e. the right to be identified as the creator of a work
- The right to object to derogatory treatment of a work
- The right to object to false attribution

IP rights and inputs into Generative AI systems

Generative AI systems typically require training on an initial set of content of the type which is sought to be generated. Sometimes this process of training may be ongoing as the Generative AI system is used. Training inputs into a Generative AI system in a workplace context could well be IP capable of protection in IP law. This will particularly be the case in the creative industries, and may be relevant to less traditional employment context/relationships.

For example (depending on contractual provisions) a designer might have IP rights over a set of drawings which an organisation uses to train an image-focused Generative AI model. Or a series of recordings by a voice artist might be used to train a Generative AI artificial voice for future use.

A Generative AI system might also involve the use of IP-protected works as prompts into the system. For example, a system might be given a set of short stories written by a contractor and prompted to *'write a short story in the style of this author, but set it in a different country'*.

Thus workers may have bargaining power over how a Generative AI system works where both:

- Creative works are input into a Generative AI system through training or prompts; and
- Workers hold the IP in those works (or hold relevant moral rights).

This will require careful analysis of both the presence and kind of creative works or image rights, and of the contractual arrangements dictating their ownership.

Where there is leverage, unions can negotiate on behalf of their members to have an agreement around Generative AI and IP in the workplace, for instance to ensure workers are fairly compensated if their work is being used to train Generative AI systems, and ensuring there are clear policies that prohibit the use of worker's IP without their permission, as well as prohibiting the general use of IP-protected works as input data where there is no permission to use those works.

Moreover, unions may want to negotiate to have clear red lines about the use of Generative AI particularly given the thorny issues it raises about ownership and attribution. This could relate to ensuring that workers are not *obliged* to use Generative AI in the course of their work, whilst providing fair policies to ensure that where workers elect to use Generative AI this does not undermine their subsequent rights in the work.

IP rights and outputs from Generative AI systems

Whether and how IP rights attach to the outputs of Generative AI systems is a very complex question. This issue goes beyond the scope of this guidance and – given the

legal issues are untested – is unlikely to be of assistance to union reps at this juncture in constraining how Generative AI is used in the workplace.

One exception to this is in the rare/unusual case that a Generative AI system's output can be said to be imitating a specific individual worker. That individual would likely have moral rights which could prevent those outputs from being used without their consent.

Action checklist

- Determine whether the Generative AI system uses works (such as inventions, designs or literary works) which can attract IP protection (most likely copyright or moral rights, but patents, trademarks or design protections could also apply).
- If possible, identify existing contractual provisions governing IP in the workplace, for instance do IP rights reside with the employer or does the worker retain rights in their work such as moral rights?
- Where workers have leverage through owning relevant IP rights, consider how these rights need to be protected and/or can secure fair conditions. Even where rights are waived as part of the standard contract of employment, unions may still be able to negotiate on behalf of workers to secure important protections for the worker.
- Seek specialist legal advice from an IP lawyer or expert to assist with negotiations.

Checklist

Action checklist: before you start – understanding and training

- Establish a basic understanding of what Generative AI is, how it operates and the role of data
- Practise communicating about the technologies and explaining how they work to others

Action checklist: getting started – assessment of context

- Is your union recognised for collective bargaining?
- Are there any factors that are specific to your sector?
- What will the employer regard as barriers to a collective agreement?
- How has the employer met legal obligations, such as those relating to health and safety, data protection, equality and human rights law, and including statutory consultation?
- How has the employer documented compliance with legal obligations?

For Generative AI technologies that may interact with data protection:

- Determine whether the Generative AI system involves the processing of personal data. If so, consider whether the employer has ensured compliance with the UK GDPR?
- Has the employer been transparent about the use of AI and data in the workplace? Consider the explainability of AI systems used in the workplace.
- What legal basis is the employer relying on?

- Is Generative AI used for any ADM, including profiling? If so, think about the rights available to not be subject to such decisions, rights to access logic about the decision and to seek a review.
- Has a DPIA been conducted and were unions and workers consulted on it?
- How are worker's rights facilitated by the employer?
- How has the employer documented compliance with legal obligations?
- Consider existing best practices/impact assessments to supplement the DPIA framework.
- If you are concerned that there is a breach of obligations/rights under the UK GDPR, or if you think that the use of those rights could help you negotiate a better outcome for employees, seek legal advice/support from a data protection lawyer.

For Generative AI technologies that may intersect with IP rights:

- Determine whether the Generative AI system uses works (such as inventions, designs or literary works) which can attract IP protection (most likely copyright or moral rights, but patents, trademarks or design protections could also apply).
- If possible, identify existing contractual provisions governing IP in the workplace, for instance do IP rights reside with the employer or does the worker retain rights in their work such as moral rights?
- Where workers have leverage through owning relevant IP rights, consider how these rights need to be protected and/or can secure fair conditions. Even where rights are waived as part of the standard contract of employment, unions may still be able to negotiate on behalf of workers to secure important protections for the worker.
- Seek specialist legal advice from an IP lawyer or expert to assist with negotiations.

Are there other systems of accountability that interact with collective agreements? For example, algorithmic impact assessments and auditing processes.

- Are there any existing relevant collective agreements, contractual entitlements, or workplace policies?

- Can you use the Trade Union Labour Relations Consolidation Act 1992 (TULRCA) to seek information about the use of Generative AI in the workplace?

How is Generative AI being used or being planned to be used in the workplace? Carry out an investigation.

- How have workers been consulted on the use of Generative AI, and the selection of product and vendor?
- Is Generative AI being introduced through existing software or through
 - new software?
- Determine the vendor and specific software. Has the vendor provided general accuracy statistics?
- Has the Generative AI product been adapted for the employer? If so, has the employer provided company or worker-specific data?
- What assessments have been conducted on the applicability of the software to the specific task at hand?
 - What are possible risks and mitigation measures related to accuracy and bias?
 - If organisation-level accuracy testing has been conducted, what are the outcomes?
- Has the employer provided appropriate training for workers who will use the product?
 - Does this include training on how to interpret results and review output for accuracy and bias?
 - What systems exist for workers to report any incidents/concerns?
- How will the employer indemnify workers for errors resulting from the use of Generative AI, if they are being mandated to use the tools or have been following the policy set out by the employer?
- How will worker input data be collected and used by employer? By the vendor?
- [If relevant] Are there restrictions, if any, on how worker data and output could be used in Generative AI products in the future?

Additional considerations for Generative AI trials/pilots:

- How will the organisational policy on the use of Generative AI be adequately communicated to all affected workers?
- What metrics/outcomes will be documented throughout the pilot?
 - How will Generative AI-enabled 'efficiency' or 'productivity' be calculated? Does this require any additional workplace monitoring/worker-related data? Will workers be able to access these metrics and data?
 - Given substantial concerns about Generative AI accuracy, how will this be assessed during the pilot? (e.g. Has additional time been assigned for workers to review Generative AI output?)
 - How will the Generative AI product's impact on workers (e.g. wellbeing) be captured during this period?
 - How will the impact on workers' job roles be documented?
 - How can workers provide feedback to management throughout and at the conclusion of the pilot?
 - What factors determine the success/continuation of the pilot? How will workers be consulted in this process?

Where there is an absence of guidance:

- If Generative AI has not been introduced by the employer, do workers use this technology at their discretion? Have workers been informed adequately about the risks of using unauthorised tools, rather than facing liabilities they were not aware of in the case of problems arising?

Action checklist: negotiating an agreement – key issues

- Think carefully about definitions. What and who will be covered (consider employment status and third parties)?
- Draw clear red lines (these could consider factors such as the environmental cost of Generative AI system, and copyright and labour conditions linked to its development).

- ❑ Create the apparatus needed – the institutions and roles.
- ❑ Consider other accountability systems.
- ❑ Ensure workers share in the rewards from greater use of technology.
- ❑ Protect human review, contact and expertise.
- ❑ Secure full transparency and explainability.
- ❑ Introduce data power and control for workers.