

The Use of Generative AI in the UK Criminal Justice Court System

The internet lies like a post-power-cut oven clock – unfortunately, we believe it. It was Abraham Lincoln who observed, with characteristic acuity: “Don’t believe everything you read on the internet.”

This article develops a theme raised in Furnival’s previous articles regarding the criminal justice [court system](#), in which demand is becoming more complex, backlogs are at record highs, and productivity is declining. A complete examination of the merits and demerits of generative AI (“GenAI”) in the criminal justice system goes far beyond a 1,200-word article. But robing rooms have buzzed with horror stories of rogue hallucinated case citations, and this article may help readers feel more confident on the subject of AI and therefore project more credibility than a giraffe in a raincoat trying to get into the opera.

Some envisage a time when humans will be to AI and bots what dogs are to humans (loyal, slightly confused). However, for several years, [Furnival Chambers](#) has used AI for accounting function, web design conformity and now legal research – saving time, money and energy to the extent that there are now the time and resources to write this article. In short, GenAI will change the job and not steal the paycheque, and therefore (with caveats below) needs to be embraced by the profession.

The courts

The backlog and growing number of cases waiting to be heard (particularly in the Crown Court) in England and Wales reached a record high of more than 75,000 outstanding cases in 2024. This situation, which sees some serious cases waiting years for trial, is driven by more charges with more data from police and prosecutors than the system can process – compounded by shortages of barristers, judges, and courtrooms.

Simultaneously, and contrary to some populist political agendas, crime has halved in the past 20 years. The UK government’s [Crime Survey for England and Wales](#), which asks thousands of households if they have suffered crime and its nature, shows that robbery has gone down 60% in the past 20 years. Bike theft and car theft have both halved. Burglary is down by two-thirds. And all violent crime? According to the survey, it has halved since 2005. The Violence Research Group at Cardiff University corroborates this, matching surveyed crime against hospital admission. “Serious violence in England and Wales has decreased substantially,” it concludes in its latest report. NHS hospital admissions in England for “all assaults, knife assaults bodily force or firearms” peaked at 45,000 in 2005–06 and is now around 20,000.

This may be a case of reliable figures hiding unreliable facts. However, in order to increase court productivity, address the above-mentioned apparent paradox and cut through the volume and complexity of evidence, the UK government has defined its vision to embed AI across the justice system – with pilots underway for AI-powered court transcription and a “knowledge retrieval assistant” capable of searching hundreds of documents in seconds.

HM Courts and Tribunals Service (HMCTS) has trialled AI transcription for hearings and oral judgments in some non-criminal courts, with promising results – the object being to

expedite the manual transcription process, provide transcripts where none currently exist, and improve public access to proceedings. Officials say early results are encouraging, with further roll-out dependent on funding. Another pilot focused on a generative AI “knowledge retrieval assistant” that answers staff queries by interrogating more than 300 unstructured documents before producing a concise summary with a source citation. Evaluation found the tool helped court staff access and digest information more quickly, ultimately accelerating case management. Following this initial promise, options to scale the system are now being explored.

Practitioners

It is easy to urge practitioners to be brave, live like a tiger, embrace the tech, and not die like a worm. Equally, no one has heard of a wormskin rug. Most recently (6 June 2025), the High Court handed down a judgment in respect of two cases under the “Hamid Jurisdiction” (*R (Ayinde) v London Borough of Haringey* (2025) EWHC 1383 (Admin)). The High Court examined the apparent misuse of GenAI by the parties’ legal teams in both instances, reminding lawyers of their fundamental professional duty to take responsibility for their own submissions, including by taking steps to verify GenAI’s output. In reality, cited cases did not exist, and the grounds misidentified the effect of relevant legislation.

The matter was brought to the attention of the High Court on the issue of wasted costs (successfully). The judge, in a trenchant judgment, said: “It is wholly improper to put fake cases in a pleading.” It was unreasonable to say that these fake cases were “minor citation errors” or, as stated by the solicitors, “cosmetic errors”. The judge referred the barrister to the Bar Standards Board and his solicitors to the Solicitors Regulation Authority. The use of GenAI was inferred, albeit not specified in the judgment.

Privacy

Metadata and privacy are primarily governed by the [UK General Data Protection Regulation](#) and the [Data Protection Act 2018](#), which require transparency and lawful processing of personal data. Individuals have rights to access, rectify, erase, or restrict the processing of their data. Organisations must inform individuals about how their metadata is being used and provide clear privacy notices.

Since October 2023 (Google reveals), “Google Translate, while a convenient tool, presents several privacy concerns, primarily due to how it handles user-submitted text. When you use Google Translate, the text you input is transmitted to Google’s servers and stored to improve the translation engine. This raises concerns about the potential for sensitive information to be accessed by Google or even leaked in the event of a data breach. Furthermore, the accuracy of translations, especially for specialised or confidential content, can be questionable, potentially leading to misinterpretations and unintended consequences.”

Recommendations

Legal librarians – by and large beasts as understated as they are undervalued – have a nose for trouble. By way of example, Sally McLaren (Inner Temple) and Lily Rowe (Law Society)

recently ran a hallucinated citation into ChatGPT-40 (both free and paid), Claude 35, Gemini 1.5 (free and paid) and Copilot, with varying results. The optimum outcome of these queries really was a recommendation to consult legal databases or resources, with one AI search engine actually saying it did not have access to specific legal documents or case summaries. Another platform, after trying to be helpful, took some pressing to admit the case “appears to be fictional”. The librarians’ simple, key recommendation for these publicly available AI tools was to “be blunt and ask your GenAI tool: “Is this a real case or is it made up?””. Otherwise, as demonstrated earlier, the truth can be as slippery as a Teflon eel.

In order to avoid the pitfalls of hallucination or privacy breach, since July 2025, all [Furnival](#) practitioners have access to AI-assisted research on Westlaw Edge UK with CoCounsel. (Other AI-assisted legal databases are available.) With rising caseloads (eg, demand created by cloud video platforms, especially at the junior end), better productivity is key. It is important to save time and effort – and therefore increase fee-earning potential – by cutting out the background noise of research.

All of Furnival Chambers’ 92 counsel have access to more than 100,000 authoritative, contemporary, tracked, and updated legislation documents. Initial, anecdotal feedback is very positive – although the scope and manner of engagement will be monitored more formally after six months, with more concrete metrics. The author is unaware of another criminal set with the same research resource.

JB October 2025