

AUTOMATING THE HOSTILE ENVIRONMENT

AI in the asylum decision-making process

January 2026

In April 2025, the Home Office quietly published a summary note evaluating the pilots for the Asylum Case Summarisation (ACS) and Asylum Policy Search (APS) tools.¹ These tools were developed with the company Methods Analytics.²

Through this pilot, the Home Office plans to use Azaru and ChatGPT-4 to summarise asylum-seeker interview transcripts and to build an AI search assistant that locates and summarises country policy information from Country Policy and Information Notes (CPINs) and Country of Origin Information Requests (COIRs). While the Asylum Policy Search (APS) tool is already being rolled out and is now available to all asylum decision-makers, the Asylum Case Summarisation (ACS) tool remains in advanced development and is planned for full rollout in January 2026.³

The Home Office has stated that individuals affected by these tools will **NOT** be informed about the use of AI in their cases.⁴

The use of ChatGPT-based tools in such a sensitive domain risks undermining basic principles of transparency, accountability, equality, and impartiality. The model's training data – and the biases embedded within it – shape every output. Even with feedback and additional constraints, LLM behaviour remains unpredictable. Much of the training data comes from English-language, Western-dominant sources such as the internet, Wikipedia, and social media. Applying such a model to summarise interviews from people whose first language is not English introduces additional risks of inaccuracy, partiality, and misinterpretation.

The Home Office highlights time savings as a key benefit – 23 minutes per case for

¹ <https://www.gov.uk/government/publications/evaluation-of-ai-trials-in-the-asylum-decision-making-process/evaluation-of-ai-trials-in-the-asylum-decision-making-process>

² https://www.linkedin.com/posts/methods-analytics_evaluation-of-ai-trials-in-the-asylum-decision-activity-7326259630030864388-iLLa?utm_source=chatgpt.com

³ <https://questions-statements.parliament.uk/written-questions/detail/2025-11-19/92650>

⁴ https://www.whatdotheyknow.com/request/asylum_case_summarisation_tool

ACS and 37 minutes per case for APS.⁵ However, these savings may come at the cost of accuracy, fairness, accountability, and legal safeguards. There is no evidence-based research demonstrating that LLMs improve productivity in this context; on the contrary, many examples show that the need for human verification increases workload.⁶ Errors produced by these systems may also contribute to backlogs in the judicial system, and certainly avoids the real challenge the system faces, which is a need to process claims both with speed and accuracy, and avoid claims being resolved through expensive court processes.

What we know so far about these tools:

- Asylum Case Summarisation (ACS) uses OpenAI's GPT-4, accessed through the Azaru interface and built on Azure AI Foundry. This means that the large language model (LLM) runs on Microsoft's AI platform.
- The tool will automatically convert lengthy asylum interview transcripts into summary documents. The ACS tool will be accessed via Azure rather than Home Office servers, and data will be stored in Amazon cloud infrastructure⁷.

Using ChatGPT-4, a pre-trained transformer neural network trained on extensive but undisclosed datasets, creates risks of bias and discrimination, particularly because the training data is not publicly known.

- **The Home Office claims that ACS “does not learn and change based on user interaction.”:**
 - This is a significant claim, made without clarification of how this is guaranteed, or how ChatGPT-4's pre-trained model has been fine-tuned for this specific purpose. In an FOIA response,⁸ the Home Office refused to disclose the prompts used for summarisation. LLMs such as ChatGPT match patterns in language and generate responses based on statistical likelihood; they do not ‘read’ or ‘understand’ content. Summarising asylum cases necessarily requires omission, simplification, and selection, all of which pose serious risks. These omissions may affect the outcome of cases in ways that are neither transparent nor predictable. LLMs depend on instructions, weights, and preferences, raising questions about how prompts can avoid

⁵ <https://www.gov.uk/government/publications/evaluation-of-ai-trials-in-the-asylum-decision-making-process/evaluation-of-ai-trials-in-the-asylum-decision-making-process>

⁶ <https://pmc.ncbi.nlm.nih.gov/articles/PMC12237199/>

⁷ ibid

⁸ ibid

introducing exclusion, bias, or discrimination. There is also no guarantee that prompts cannot be altered case by case.

- 9% of ACS-generated summaries were so flawed that the Home Office removed them from the pilot. This means nearly one in ten summaries were too inaccurate or incomplete to be retained. However, the distribution of these errors by nationality, language, or case complexity remains unknown.
- The repeated Home Office assertion that “the ACS tool does not learn and change based on user interaction”⁹ is unverified. It implies that GPT-4 is used solely in inference-only mode, with no additional machine-learning components or feedback loops. However, Azure regularly updates and replaces model versions, meaning their behaviour may shift without the Home Office’s knowledge or control. Prompt changes can also significantly affect outputs. Caseworker edits can indirectly influence future prompts, creating a form of ‘shadow training’ even without formal retraining. Even if ACS does not learn directly from transcripts or user interactions, it can still change due to model updates, prompt adjustments, infrastructure changes, or hidden system drift.
- The evaluation summary showed that **23% of caseworkers lacked full confidence** in the tools’ outputs. Research indicates that humans are more likely to accept AI-generated advice when it confirms their existing biases.¹⁰
- **Asylum Policy Search (APS)** is designed as an AI assistant that finds and summarises country policy information from CPINs, COI reports, and other guidance. Beyond the broader concerns of bias and reliability, APS relies on CPINs, which are themselves summaries and already contain known limitations. APS therefore produces summaries of summaries. Country guidance documents are slow to update, leaving APS dependent on incomplete or outdated information. Asylum claims often involve unique or highly specific circumstances; reliance on narrow sources risks disadvantaging people whose cases fall outside common patterns.
- Although the Home Office describes these tools as “aids,” they directly shape the information base on which decisions are made. APS amplifies the structural biases within CPINs because it draws almost exclusively on Home Office documents. This can lead to inaccuracies, opacity, and potential over-reliance by

⁹ ibid

¹⁰ <https://www.sciencedirect.com/science/article/pii/S2949882124000264>

overstretched caseworkers. Without evidence of robust and enforced human oversight, it is unclear whether APS will genuinely support, rather than replace, human judgment.

- APS operates in an area where interpretation of policy directly affects outcomes. It can misquote, invent citations, over-generalise, or distort nuanced information. APS does not simply retrieve text; it interprets it, meaning it plays an active role in shaping the inputs to decision-making.
- **During the APS pilot, one caseworker reported that the tool failed to provide source references**, creating additional verification work, and calling into question whether these tools will actually save caseworkers' time. This highlights the risk of hallucination,¹¹ a known issue in LLMs. There are well-documented incidents of ChatGPT generating false legal citations.¹² Without reliable referencing, outputs may be misleading and significantly increase the burden of manual fact-checking.

Using large language models (LLMs) in the immigration and asylum system does not improve fairness, legality, or respect for human rights. Achieving fairness would require a fundamental shift in policy, as the Home Office continues to treat migration as a threat rather than acknowledging people's dignity and rights under the Human Rights Act and the ECHR.

Although the Home Office presents LLMs as a way to reduce the asylum backlog, the practical effect is to reinforce high refusal rates and removal policies. LLMs follow user instructions and system design; they are not built to ensure impartiality or justice, and they can easily reproduce or amplify bias. They cannot safeguard rights to a fair process, non-discrimination, or an adequate remedy.

AI can serve the public only when deployed with a genuine intent to protect rights, ensure transparency, reduce bias, and uphold data protection. This requires strong safeguards and accountability, which are currently lacking.

Using ChatGPT-4 in asylum decision-making will not improve accuracy or efficiency. The resources spent on these systems would be better invested in caseworker training, additional staffing, and strengthening fair decision-making.

LLMs risk creating only the appearance of speed while increasing errors, placing further strain on appeals and judicial review, and undermining trust in the system.

¹¹ <https://openai.com/index/why-language-models-hallucinate/>

¹² <https://www.theguardian.com/law/2025/sep/03/lawyer-caught-using-ai-generated-false-citations-in-court-case-penalised-in-australian-first>

Such risks are incompatible with the principles of justice, accountability, and human dignity at the core of the Human Rights Act and the ECHR.

AI tools such as LLMs are not neutral; they can be easily tweaked to produce preferred results. Implementing such tools in a very vulnerable population in very critical situations, without safeguards, transparency, or a chain of accountability, can have fatal or severely damaging consequences.