Journal of Ethics and Legal Technologies – Volume 8(2), November 2026

AI-Assisted Asylum Decision-Making: Source Verification and Procedural Accountability under EU Law

Khalid Ariqat, Ibrahim Ariqat***

*York University, Toronto, Canada
**KU Leuven, Leuven, Belgium

Abstract: AI-generated summaries of asylum interviews may omit important explanations and affect how evidence is interpreted in decisions. Existing scholarship examines procedural fairness and evidence interpretation in asylum administration, while UK quality checks have raised concerns about the accuracy and traceability of generated summaries. Research does not fully establish how asylum authorities should verify and document consequential AI-generated statements before relying on them in adverse assessments. Through doctrinal legal and policy analysis of EU asylum legislation, the amended AI Act and institutional guidance, this article proposes a proportionate verification procedure triggered by material reliance rather than AI use alone. The procedure connects generated statements to original sources, records human verification and the treatment of discrepancies, and links results to existing procedural safeguards. It distinguishes system-level logging and oversight from the evidentiary question of whether a particular statement was checked before it affected a person's claim. An expressly hypothetical interview-summary discrepancy illustrates how omitted context could influence a credibility assessment even when the initial interview record is accurate. The analysis considers how caseworkers would identify consequential reliance, locate and compare source material, record any correction or uncertainty, and preserve opportunities for explanation and challenge. It also identifies practical constraints, including workload, data protection and the risk of turning checks into formalities. Its operational value would require careful evaluation rather than assumption across jurisdictions. The article contributes a framework for making AI-assisted evidence assessment more reviewable and open to challenge, without claiming an empirically tested solution or a new statutory obligation.

Keywords: EU law, asylum and refugee protection, AI decision-making, source verification

1. Introduction

An asylum decision may determine whether an applicant receives international protection, which under EU law includes refugee status and subsidiary protection status. For that reason, the evidence supporting such a decision must be assessed individually and remain open to effective scrutiny under the applicable legal framework. Digital tools may assist officials with document retrieval, transcription and summarisation, but their use also raises questions about how evidence is represented after it has been recorded (Nielsen et al., 2023; Home Office, 2025). A particular concern arises when an AI-generated description omits material context and an official later relies on that description to question an applicant's account. In such circumstances, even an accurately recorded interview may be misrepresented by a later summary. The central concern is therefore not only whether the original evidence was recorded correctly, but also whether a later description can be checked before it influences a decision. This paper examines that problem without assuming that any particular national authority uses the hypothetical tool discussed below.

The importance of this distinction becomes clearer in the current EU legal context, where two regulatory timetables overlap. Regulation (EU) 2026/1744 moved the application date for principal Annex III high-risk obligations in Chapter III, Sections 1–3, of the AI Act to 2 December 2027. At the same time, Regulation (EU) 2024/1348 generally governs applications for international protection lodged from 12 June 2026 in Member States bound by it, while older applications remain within the preceding procedural regime. These dates govern different matters: postponing specified AI Act obligations does not suspend the procedural rules governing an individual asylum decision (European Parliament & Council, 2013, 2024a, 2024b, 2026). This distinction makes it important to consider how the reliability of AI-generated descriptions can be examined within the legal procedures already applicable to asylum decisions.

Against this background, the paper addresses a specific question: How can asylum authorities verify and document AI-generated information before relying on it in asylum decisions, while ensuring that errors can be identified, corrected and challenged? Formal legal rights and an official signature do not, on their own, establish what an official verified before relying on a generated statement. To address this problem, the paper proposes a proportionate case-file record connecting a consequential assertion to its source, the human verification undertaken, the treatment of discrepancies and relevant opportunities for challenge. The proposal introduces neither a new substantive credibility test nor a claim that EU law already requires a uniform

AI-specific log. The following section situates this proposal within existing legal analysis, asylum casework research and practical guidance.

2. Existing scholarship, research gap and contribution

Recent scholarship has examined how artificial intelligence may affect procedural fairness, the assessment of evidence and the protection of applicants' rights in asylum decision-making. Palmiotto (2024a) demonstrates how privacy and data protection rights can strengthen procedural fairness when automated systems provide evidence in asylum procedures. Stewart (2024a) examines the implications of AI-assisted decision-making for accuracy, efficiency and participation, while Amoroso (2024) assesses the suitability of AI for refugee status determination, including its limitations in evaluating subjective elements of protection claims. Extending these concerns, Stoyanova (2026) identifies difficulties in validating AI-supported decisions, examines procedural harm and considers how discretion may shift into technological design. Together, these contributions clarify the legal and institutional challenges of AI-supported asylum decision-making, including the importance of scrutinising automated outputs. Building on this scholarship, the present paper examines how asylum authorities can verify and document AI-generated information before relying on it in decisions, so that errors can be identified, corrected and challenged.

The importance of verification and documentation becomes clearer through studies of how asylum evidence is produced and interpreted in practice. Examining 50 Danish cases, Rask Nielsen and Holten Møller (2022) show how credibility is constructed by comparing testimony and records. Nielsen et al. (2023) analyse how case information is translated, documented and interpreted across professional settings. Palmiotto (2024b) shows how legal classifications can overlook automation that assists, rather than replaces, human judgment. Stewart (2024b), Forti (2024) and Curtin and Fia (2026) address accountability, algorithmic error and secrecy. Ozkul (2025) examines the appearance of objectivity attached to German dialect-recognition and phone-analysis technologies, not interview summaries. Alizadeh Westerling (2026) identifies procedural vulnerabilities in Finland. Therefore, this research shows how information acquires significance as it moves through administrative casework.

Research on AI governance further helps distinguish the proposed case-file record from existing system-level safeguards. Ebers (2025) examines the EU AI Act's risk-based structure and the limits of relying on system classification to protect fundamental rights. Laux and Ruschemeier (2025) analyse automation bias and the allocation of human-oversight responsibilities be-

tween providers and deployers. These analyses explain why oversight cannot be reduced to the presence of a human decision-maker. The additional issue addressed here is narrower and more concrete: whether an individual asylum file can show which source an official checked, whether an AI-generated statement was corrected and how it entered the reasons for an adverse assessment. The proposed record is designed to make that case-specific evidentiary pathway reviewable rather than to replace AI Act compliance.

The need for this kind of check is also reflected in the UK Home Office's use of AI-generated interview summaries. In its 2025 pilot evaluation, the Home Office reported that technical reviewers excluded 9% of generated summaries because they contained inaccuracies or missing information. Some users also noted that summaries did not identify the relevant passages in the original interview transcripts (Home Office, 2025). After the Asylum Case Summarisation (ACS) tool was rolled out to all asylum caseworkers in April 2026, subject-matter experts conducted spot checks on 203 summaries. Alsherif (2026), drawing on Home Office disclosures, reported that 45 failed the first-stage check and 19 were flagged for further review. Of the 56 summaries reviewed at the second stage, 30 failed; this represented 15% of the original sample of 203 summaries (Alsherif, 2026). ORG estimated that subject-matter experts checked only 203 summaries, about 0.3% of the summaries generated for approximately 70,000 cases processed between April and July 2026 (Alsherif, 2026).

The report also noted that the relevance criterion, which assessed whether a summary omitted the most relevant facts from the interview, was scored but was not one of the stated grounds for failure. It therefore raised concern that a summary could pass even if it omitted material facts, provided that the remaining content was accurate (Alsherif, 2026). These findings come from a limited set of checks and do not establish the overall accuracy of the tool or show that a particular asylum decision was affected. The 2025 pilot and the later checks were conducted under different review arrangements and should not be treated as directly comparable measures. Although this evidence comes from the UK, outside the European Union, it illustrates why authorities should be able to compare generated summaries with original interview records before relying on them.

Any additional check must also work with existing procedures. The EUAA's Practical Guide on Evidence and Risk Assessment already addresses information gathering, credibility assessment and risk assessment through guidance and checklists (European Union Agency for Asylum [EUAA], 2026). The task is therefore to incorporate verification of consequential AI-generated assertions into that established process.

Taken together, these sources reveal a specific implementation gap. Palmiotto, Stewart, Amoroso and Stoyanova address legal protections and risks; Rask Nielsen and colleagues, Ozkul and Alizadeh Westerling examine information practices; and the Home Office and EUAA provide operational findings and evidence-assessment guidance. Together, these sources leave a narrower question unanswered: how can an individual asylum file show that a potentially adverse AI-generated statement was checked against its original source before it influenced an assessment? Technical logging, human oversight and interview documentation serve important purposes, but they do not necessarily show which generated statement an official checked, whether a discrepancy was identified or how that discrepancy was addressed in an individual case. This article makes three contributions. First, it proposes a material-reliance trigger rather than treating all AI use alike. Second, it specifies a source-to-output record identifying what was checked, any discrepancy, its disposition and relevant opportunities for explanation or challenge. Third, it situates the procedure within EU legal timetables and proposes criteria for future assessment of reconstruction, correction, workload and privacy. The proposal integrates existing safeguards; it creates no new legal right and has not been empirically tested.

This article adopts a focused doctrinal legal and policy analysis to examine how existing EU safeguards apply when AI-generated information influences asylum evidence assessment. The legal analysis first identifies the provisions governing when AI requirements apply and how individual asylum claims are recorded, assessed, explained and challenged. It then compares those provisions with the information an authorised reviewer would need to establish whether a particular AI-generated statement was checked against its original source. Institutional guidance and the cited scholarship are used to assess how the proposed verification record could fit existing casework, rather than to suggest that EU law already requires it. The article does not present a systematic literature review. The UK Home Office's 2025 evaluation, the later sample-check findings reported by Alsherif (2026), and the parliamentary response by Turley (2026) are used as non-EU operational comparators, not as evidence of the prevalence of errors or their effects on decisions within the European Union. The analysis then applies the relevant legal and practical concerns to an expressly hypothetical interview-summary error to develop a proportionate case-file verification procedure. Because no interviews, case-file audits, technical tests or deployment datasets were collected, the article cannot establish the procedure's effectiveness or estimate AI-related error rates. The following section sets out the applicable legal framework.

3. Regulatory timing and applicable safeguards

The AI Act obligations relevant to asylum decision-making depend first on the tool's intended function, because different uses of AI do not attract the same legal requirements. Annex III, point 7(c), of the AI Act covers certain systems assisting authorities with asylum applications, including the assessment of evidence reliability. Whether a system is classified as high risk depends on Article 6 and its intended purpose. Article 6(3) provides qualified exclusions for certain Annex III systems that do not pose a significant risk of harm, subject to specific conditions and limitations, including those concerning profiling. Scheduling, transcription, summarisation and automated credibility assessment should therefore not be treated as legally interchangeable (European Parliament & Council, 2024a).

Classification alone, however, does not establish when the relevant obligations take effect. Under the 2026 amendment, the principal obligations in Chapter III, Sections 1–3, apply from 2 December 2027 to Annex III systems classified as high risk under Article 6(2). Article 111 establishes separate transitional arrangements, including a 2 August 2030 deadline for necessary compliance steps concerning certain high-risk systems intended for use by public authorities. Although the AI Act generally applies from 2 August 2026, individual provisions have different application dates. Article 86, for example, is not among the provisions postponed until December 2027 and remains subject to its own conditions. Importantly, none of these dates suspends the applicable obligations under EU asylum law (European Parliament & Council, 2024a, 2026).

Beyond classification and timing, it is important to distinguish system-level compliance from the review of evidence in individual asylum cases. For qualifying high-risk systems, Article 12 of the AI Act addresses technical logging, Article 14 human oversight, Article 26 deployer obligations and Article 27 fundamental-rights impact assessment in specified circumstances (European Parliament & Council, 2024a, Arts. 12, 14, 26, 27). Although these requirements may strengthen AI governance, they do not automatically produce an individual asylum-file record showing whether a generated assertion was compared with its original source and later retained, corrected or excluded. The proposed material-reliance check addresses this more specific issue by documenting the verification of an individual evidentiary proposition before it may influence an applicant's case. It is intended to complement the applicable AI Act requirements, rather than reinterpret Article 12 logging obligations or replace mandatory human oversight.

While the AI Act establishes obligations for qualifying AI systems, the assessment of an individual asylum claim remains subject to its governing

procedural rules. For applications lodged from 12 June 2026 within the scope of Regulation (EU) 2024/1348, Article 34(2) requires objective, impartial and individual examination. Article 34(2)(a) refers specifically to the applicant's statements and documentation under Article 4(1) and (2) of Regulation (EU) 2024/1347. Article 4(3) of that Regulation requires the relevant elements to be assessed in accordance with Article 34 of the Procedures Regulation. Article 4(5) establishes cumulative conditions under which aspects of an applicant's statements that lack supporting evidence do not require additional evidence. It does not require an authority to accept every inconsistency or establish an unconditional presumption of credibility. Accordingly, an AI-generated description of an alleged inconsistency cannot replace the assessment of the applicant's statements, explanations and other relevant evidence as a whole (European Parliament & Council, 2024b, 2024c).

These requirements for individual assessment are reinforced by procedural safeguards operating at different stages of the asylum process. Article 12(2) of Regulation (EU) 2024/1348 provides applicants with an opportunity to explain missing elements, inconsistencies or contradictions during the substantive interview. Article 14 governs the interview report or transcript and audio recording, including opportunities to comment on or clarify factual mistakes. Where there is doubt about statements made during the interview, the audio recording prevails under Article 14(4). Article 14(6) further provides for access to the report or transcript after the interview and in due time before a decision, as well as access to the recording during an appeal, subject to the provision's terms. Article 18 establishes qualified access for a legal adviser representing the applicant to information in the file on which the decision is or will be based. Where an application is rejected in circumstances covered by Article 36(2), the written decision must state the factual and legal reasons for rejection, while Article 36(3) governs information about how to challenge specified decisions. Article 67 provides for an effective remedy within its scope, supplemented by Article 47 of the Charter. Together, these provisions establish safeguards concerning interview documentation, evidence assessment, access to information, reasons for decisions and appeals. However, they do not expressly require the proposed AI-specific verification record or provide an unrestricted right to confidential information (European Parliament & Council, 2024b; European Union, 2012).

The AI Act may provide additional protections where its relevant conditions are satisfied. Article 86 concerns the right to obtain an explanation of the role of certain high-risk Annex III systems in decisions producing legal or similarly significant effects, subject to its conditions, exclusions and other applicable Union-law rights. This provision does not apply automatically to all software or establish an unrestricted right of access to source code. Rights

under the General Data Protection Regulation may also be relevant where their statutory conditions are met. These provisions address related but distinct questions concerning AI system classification, the fair assessment of asylum evidence and an applicant's ability to understand or challenge a decision. Table 1 summarises how these safeguards relate to the proposed verification procedure. It presents the the authors' synthesis rather than an additional legal requirement. The following section illustrates why these distinctions matter when an AI-generated statement enters an individual asylum assessment (European Parliament & Council, 2016, 2024a).

Table 1. The interaction of legal safeguards relevant to a proposed AI-source verification record (the authors' synthesis).

Review stage	Relevant existing rules	Question addressed
Tool and timing	AI Act, Art 6, Annex III and transitional rules	What is the tool's intended function, how is it classified, and when do the relevant obligations apply?
Interview and source record	Asylum Procedures Regulation, Articles 12 and 14	What did the applicant say, and how can the original interview record be checked?
Evidence assessment and reasons	Asylum Procedures Regulation, Articles 34 and 36; Qualification Regulation, Article 4	Was the original evidence accurately represented and assessed individually before informing the decision?
Access and challenge	Asylum Procedures Regulation, Articles 18 and 67; AI Act, Article 86, where applicable	Can a consequential AI-generated statement be identified, explained and challenged under the applicable procedures?

4. When an AI-generated statement enters the evidence assessment

The practical importance of these safeguards can be illustrated through a hypothetical personal interview. An applicant explains that dates recorded using different calendars refer to the same sequence of events, and the interview report accurately records this explanation. A later AI-generated summary, however, omits the explanation and describes the dates as inconsistent. If the case officer relies on that description without checking the original interview record, the misleading summary may influence an adverse credibility finding.

This example highlights two distinct stages of evidence handling: documenting the original interview and later relying on an AI-generated description of it. Article 14 of the Procedures Regulation governs the documenta-

tion of personal interviews, but an accurately recorded interview does not prevent omissions or inaccuracies in a later summary. Article 4 of the Qualification Regulation, read alongside Article 34 of the Procedures Regulation, concerns the assessment of the applicant's statements and other relevant circumstances. Article 36 separately governs the factual and legal reasons for specified adverse decisions. The relevant concern is therefore not simply whether an official signed the decision, but whether the alleged inconsistency was checked against the original interview record before influencing the assessment and the reasons given for an adverse decision (European Parliament & Council, 2024b, 2024c).

The hypothetical example illustrates a potential evidentiary problem rather than alleging misconduct or describing an observed practice within a particular national authority. Palmiotto (2024a) and Stewart (2024a) examine rights and procedural fairness in automated asylum procedures, while Rask Nielsen and Holten Møller (2022) and Nielsen et al. (2023) demonstrate how information and interpretations are constructed through asylum casework. Stoyanova (2026) further highlights difficulties in validating AI-supported asylum outcomes and the movement of discretion into technological design. Together, these studies help explain why the relationship between original evidence, generated information and later administrative decisions requires scrutiny. The specific concern addressed here is whether a reviewer can establish which AI-generated assertion an official encountered, which original source it represented and how any discrepancy was handled.

Establishing this connection does not, however, replace the ordinary assessment of evidence. The European Union Agency for Asylum's 2026 Practical Guide on Evidence and Risk Assessment treats information gathering, credibility assessment and risk assessment as connected stages of the asylum procedure. The proposed verification record is intended to complement this guidance, not replace it or determine whether an applicant's statements are credible. Verifying that an AI-generated summary accurately represents its source is necessary for responsible reliance on that summary, but source accuracy does not itself establish credibility. An accurate summary may still contain disputed evidence, while an inaccurate summary may introduce an inconsistency that was absent from the original interview (EUAA, 2026).

The UK Home Office evidence discussed earlier provides a limited practical illustration of why this distinction matters (Home Office, 2025; Alsherif, 2026). The reported omissions and missing source references highlight the importance of checking AI-generated summaries against their original sources. However, as non-EU evidence, these findings do not establish the frequency of such problems in EU asylum systems or their effects on asylum decisions. Together, the hypothetical example and this limited operational

evidence demonstrate the need to consider what an authority should document when relying on an AI-generated statement that may disadvantage an applicant. The following section develops a proposed verification procedure to address this concern.

5. A proposed standard of reviewable evidence

This article proposes a short, reviewable case-file record triggered by material reliance on AI-generated information. Material reliance occurs when an identifiable AI-generated assertion is used to form, support or justify a potentially adverse factual or credibility assessment, and correcting or removing that assertion could reasonably change the evidentiary basis of the assessment. A final rejection need not have been made, and the assertion need not appear verbatim in the written decision. This threshold is proposed as a risk-based administrative measure rather than an existing statutory requirement. An AI-generated observation that prompts further questioning but is later checked and discarded should be treated differently from an assertion retained in an official's reasoning. However, tools that influence the interpretation of evidence indirectly may also require broader process-level supervision. Table 2 presents the proposed five-stage verification procedure, which is explained in the following paragraphs.

Table 2. Proposed case-level verification sequence (the authors' synthesis; not a statutory requirement).

Proposed stage	Question or action
1. Identify reliance	Was a specific AI-generated assertion used, or proposed for use, to support a potentially adverse assessment? If not, routine supervision and process-level checks remain relevant.
2. Locate the source	Where material reliance is possible, identify the relevant output version and the underlying interview passage or document.
3. Verify fidelity	Compare the AI-generated assertion with the original source and identify missing context, altered wording or other discrepancies.
4. Record disposition	Document whether the assertion was confirmed, corrected, disregarded or left unresolved. Do not treat an unverified assertion as independent proof.
5. Connect to procedure	Record how the assertion influenced the assessment and identify applicable opportunities for explanation, file access, reasons and remedy, subject to governing law.

The proposed sequence distinguishes verification of an AI-generated statement against its source from the later legal assessment of the applicant's claim. It is intended to support, rather than replace, that assessment.

The first stage is to establish the AI system's role and determine whether its output has been, or may be, relied upon in a potentially adverse assessment. The authority should identify the system's function and the point at which its output reached the caseworker. This helps distinguish routine retrieval or formatting from substantive summarisation and apparent credibility assessment. At this stage, the official does not need to know whether an AI-generated statement is inaccurate. The initial question is whether it could affect an adverse assessment. Where that possibility cannot reasonably be ruled out, the original source should be located and checked before the statement is used. Where a tool may influence the questions asked or the interpretation of evidence without producing an identifiable final assertion, authorities should consider sampling and broader process-level supervision. The absence of an explicit AI-generated finding does not necessarily mean that the system had no influence on the assessment. Any classification under the AI Act remains a separate legal determination (European Parliament & Council, 2024a).

Once potential reliance has been identified, the second stage is to establish a clear and proportionate link between the AI-generated assertion and its original source. The record should identify the relevant interview passage or document and the specific version of the generated statement considered by the official. This link should allow later comparison without necessarily retaining every intermediate prompt or disclosing protected information. In the hypothetical calendar example, a concise record could identify the relevant interview segment, the AI-generated statement describing inconsistent dates, the omitted calendar explanation and the action later taken. This is an illustrative record format, not evidence from an actual asylum case. Applicable requirements concerning confidentiality, retention and data minimisation must continue to be respected.

Locating the original source, however, does not establish whether the generated statement accurately represents it. The third stage therefore requires the official to compare the AI-generated assertion with the underlying material and identify any missing context, altered wording or other discrepancies. This examination is particularly important where an apparent inconsistency may result from an omission in the generated summary rather than from the applicant's original statements. The purpose is to establish whether the generated information faithfully represents its source, not to determine the credibility of the asylum claim.

The fourth stage is to document the result of that verification. The official should record whether the assertion was confirmed, corrected, disregarded or remained unresolved, together with the action taken. If the original source cannot lawfully or practically be retrieved, the generated statement should not be treated as independent proof supporting an adverse conclusion. Instead, the official should seek lawful corroboration, document the limitation and assess the remaining evidence without silently adopting the unsupported assertion. This distinction is essential because verifying the accuracy of a summary does not establish the truth of the underlying statement. The assessment of credibility and entitlement to protection remains governed by the applicable substantive and procedural rules (EUAA, 2026; European Parliament & Council, 2024b, 2024c).

The fifth stage connects the verification record to the relevant procedural safeguards so that discrepancies can be addressed rather than merely documented. Before a decision, applicants should retain the opportunities provided under the applicable law to explain inconsistencies during the substantive interview and to address mistakes in the interview record. Where an AI-derived discrepancy emerges only after the interview, separate notification or discussion is proposed as good practice, subject to applicable legal constraints, rather than as an automatic statutory disclosure right. If an application is rejected, Article 36 governs the factual and legal reasons required for specified decisions, and the verification record may help establish whether a generated assertion influenced those reasons. Access to file information under Article 18 and remedies under Article 67 remain subject to their respective legal conditions. These safeguards are particularly important where an applicant needs to identify and challenge a disputed factual proposition rather than merely receive general notification that AI was used (European Parliament & Council, 2024b).

These five stages are intended to operate together rather than as separate administrative formalities. Recording an AI-generated statement does not establish that it was accurate, just as marking a verification field as completed does not demonstrate that the original evidence was meaningfully examined. Similarly, opportunities for appeal may be difficult to exercise when the disputed assertion cannot be identified or traced to its source. The effectiveness of the proposed procedure should therefore be assessed by whether an authorised reviewer can reconstruct how an AI-generated statement entered the assessment, identify any consequential discrepancy and determine how the authority addressed it. This provides the basis for the practical recommendations developed in the following section.

6. Recommendations for asylum authorities and policymakers

Making the proposed record useful in practice requires clear instructions, not simply another administrative form. Authorities considering AI-assisted evidence work could pilot a short case-file entry covering the tool's function, source location, relied-upon assertion, verification result and disposition. Written instructions should define the material-reliance threshold, require escalation where the source is unavailable and distinguish errors found before a decision from those detected only after a decision. In the latter situation, the responsible authority would need to examine existing correction, reconsideration and remedy routes under the applicable law; the proposal does not create a new power to reopen a concluded claim. Training should focus on reviewing the underlying material, not simply completing a field in software.

A parliamentary response published on 3 November 2026 provides a practical example of why these procedures matter. Asked what happens when an AI-generated summary used in an asylum claim later fails expert review, the Home Office stated that ACS is a decision-making aid and that an asylum decision cannot rely on a summary alone. It added that decision-makers retain authority, have access to all relevant evidence and are reminded to check summaries for accuracy. The response stated that the Home Office's general privacy notice informs applicants about its use of AI, but that applicants are not individually told when the tools are used in their cases (Turley, 2026). Alsherif (2026), however, reported that the version of the notice updated on 6 August 2026 did not specifically mention either tool. The parliamentary response did not explain what action follows a failed expert review or whether the decision-maker is informed of that failure. This leaves a practical question for the proposed procedure: can an authorised reviewer establish whether a failed summary influenced the assessment and what action was taken in response?

Putting this record into practice also depends on what authorities can retrieve from the systems they procure. Procurement and case-management specifications could therefore require retrieval of relevant source passages, identification of output versions and lawful preservation of a review trail. The requirements would have to accommodate confidentiality, national rules on file access, third-party information and data security, and should not presume that proprietary code must be disclosed. These measures respond to documented concerns about procurement secrecy and institutional responsibility (Curtin & Fia, 2026; Stewart, 2024b). They remain recommendations in addition to, rather than substitutes for, any applicable AI Act duties such as risk assessment, record-keeping or fundamental-rights impact assessment

when the relevant statutory provisions apply (European Parliament & Council, 2024a).

Those implementation measures do not remove the proposal's costs or risks. A verification field can become formulaic, particularly under heavy caseloads; recording an unsupported assertion could give it undeserved authority; and retrieving sources may impose costs without improving decisions. A narrow focus on quoted assertions can also miss earlier effects on caseworker attention. Authorities should therefore test actual reconstruction and correction by an independent reviewer, assess time and privacy costs, and monitor indirect influence through sampling and practitioner feedback. Where these procedures fail to improve meaningful scrutiny, the recording requirement should be revised rather than expanded by default. No evidence from the present analysis establishes that such a pilot would succeed. The next section sets out the research needed to assess these limitations rather than assume the procedure will work.

7. Future research and limitations

The first research need is to establish the practical baseline. Field research should map which AI functions are procured, piloted or used in asylum procedures across participating jurisdictions, when their outputs enter a file and whether they influence questions, fact selection or decision reasons. Comparative analysis should distinguish automated administrative support from systems affecting the interpretation of claims, and account for differences in the applicable legal regime. Such research is needed before generalising from the hypothetical example; the FRA's related project description is not evidence of measured deployment or error rates (European Union Agency for Fundamental Rights [FRA], 2026).

Once the systems and points of reliance are mapped, controlled evaluations could use synthetic or lawfully de-identified interview passages to compare review with and without the proposed source-to-output record. Independent assessors could evaluate whether reviewers detect omitted qualifications, altered chronology or mistranslations and whether correction changes a proposed factual finding. Studies should predefine error categories, use suitable language expertise, distinguish fidelity from the ultimate assessment of credibility and measure time spent. Without an appropriate comparison design, improved outcomes should not be attributed causally to the record.

Controlled tasks alone, however, cannot show whether an applicant can meaningfully understand or challenge an error. Ethically approved practitioner and applicant-centred studies could therefore test whether disputed

assertions are intelligible and challengeable at three different points: before a decision where lawful engagement is available; when reasons for refusal are given; and in later remedies. Researchers should examine legal assistance, interpretation, disability-related accessibility, confidentiality restrictions and national time limits, as well as how unresolved or indirectly influential outputs are handled. The central outcome would be whether applicants can identify and address a material discrepancy, rather than merely receive notice about AI use.

Taken together, these lines of inquiry could provide evidence about both administrative accuracy and practical opportunities for challenge. Such evaluations could later be repeated as AI Act requirements take effect for different system categories. They should report the proportion of consequential outputs with retrievable sources, the proportion of material discrepancies identified and resolved, the rate of unresolved assertions excluded from reliance, and administrative costs. This article cannot establish whether any measure reduces unlawful decisions, since it reports no controlled tests, interviews or EU deployment data. This limitation is central to the recommendation that authorities evaluate the procedure before wider adoption, and to the qualified conclusion that follows.

8. Conclusion

This article has examined how AI-generated information can influence asylum evidence assessment when a consequential statement is relied upon without adequate verification against its original source. Building on existing scholarship concerning automated evidence, procedural fairness and asylum administration, the analysis identifies a specific implementation gap between established legal safeguards and the practical documentation of AI-assisted evidentiary assessments. To address this gap, the article proposes a material-reliance threshold and a source-to-output verification procedure that records what was examined, how discrepancies were addressed and how consequential statements relate to existing opportunities for explanation and challenge.

The proposed framework does not replace the legal obligations governing individual assessment, interview documentation, written reasons and effective remedies. Instead, it offers a way to make those protections more operational when AI-generated information enters asylum decision-making. Its effectiveness, administrative costs and privacy implications require empirical examination. Ultimately, the value of the proposed approach will depend on whether it enables applicants and authorised reviewers to identify and

challenge consequential errors, rather than merely adding another layer of administrative documentation.

Research ethics and AI-use disclosure

This desk-based analysis used publicly available legislation, institutional guidance and published research. No human participants were recruited and no non-public case files or technical tests were used. ChatGPT (OpenAI) assisted with language editing, structure, drafting and bibliographic verification. The authors reviewed and edited the AI-assisted output and retain responsibility for checking the legal analysis, references and final text. AI assistance does not constitute authorship.

Funding: No external funding was received for this study.

Competing interests: The authors declare no competing interests.

References

Alizadeh Westerling, F. (2026). Digitalisation of the Finnish asylum procedure: From efficiency to procedural vulnerabilities. *Journal of Refugee Studies, 39*(1), 80–97. https://doi.org/10.1093/jrs/feaf071

Alsherif, S. (2026, October 12). *The use of ChatGPT in asylum decisions: Errors, inaccuracies & Home Office incompetence* [Policy paper]. Open Rights Group. https://www.openrightsgroup.org/publications/the-use-of-chatgpt-in-asylum-decisions/

Amoroso, M. P. (2024). Intelligent borders: Exploring the suitability of artificial intelligence systems in refugee status determination under international law. *Refugee Survey Quarterly, 43*(4), 410–426. https://doi.org/10.1093/rsq/hdae021

Curtin, D., & Fia, T. (2026). Cracking AI secrecy: (Beyond) EU migration law and governance. *European Law Open, 5*(2), 305–334. https://doi.org/10.1017/elo.2025.10064

Ebers, M. (2025). Truly risk-based regulation of artificial intelligence: How to implement the EU's AI Act. *European Journal of Risk Regulation, 16*(2), 684–703. https://doi.org/10.1017/err.2024.78

European Parliament & Council. (2013). *Directive 2013/32/EU on common procedures for granting and withdrawing international protection.* https://eur-lex.europa.eu/eli/dir/2013/32/oj

European Parliament & Council. (2016). *Regulation (EU) 2016/679 (General Data Protection Regulation).* https://eur-lex.europa.eu/eli/reg/2016/679/oj

European Parliament & Council. (2024a). *Artificial Intelligence Act: Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (as amended).* https://eur-lex.europa.eu/eli/reg/2024/1689

European Parliament & Council. (2024b). *Asylum Procedures Regulation: Regulation (EU) 2024/1348 establishing a common procedure for international protection in the Union (as amended).* https://eur-lex.europa.eu/eli/reg/2024/1348

European Parliament & Council. (2024c). *Qualification Regulation: Regulation (EU) 2024/1347 on standards for the qualification of beneficiaries of international protection.* https://eur-lex.europa.eu/eli/reg/2024/1347/oj

European Parliament & Council. (2026). *Regulation (EU) 2026/1744 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI).* https://eur-lex.europa.eu/eli/reg/2026/1744/oj

European Union. (2012). *Charter of Fundamental Rights of the European Union (2012/C 326/02).* https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:12012P/TXT

European Union Agency for Asylum. (2026). *Practical guide on evidence and risk assessment.* https://www.euaa.europa.eu/publications/evidence-and-risk-assessment-practical-guide

European Union Agency for Fundamental Rights. (2026, February 20). *Use of artificial intelligence in asylum and immigration procedures – Fundamental rights implications* [Project description]. https://fra.europa.eu/en/project/2026/use-artificial-intelligence-asylum-and-immigration-procedures-fundamental-rights

Forti, M. (2024). Addressing algorithmic errors in data-driven border control procedures. *German Law Journal, 25*(4), 635–645. https://doi.org/10.1017/glj.2023.102

Home Office. (2025, April 29). *Evaluation of AI trials in the asylum decision making process.* 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

Laux, J., & Ruschemeier, H. (2025). Automation bias in the AI Act: On the legal implications of attempting to de-bias human oversight of AI.

European Journal of Risk Regulation, 16(4), 1519–1534. https://doi.org/10.1017/err.2025.10033

Nielsen, T. R., Menendez-Blanco, M., & Møller, N. H. (2023). Who cares about data? Ambivalence, translation, and attentiveness in asylum casework. *Computer Supported Cooperative Work, 32*, 861–910. https://doi.org/10.1007/s10606-023-09474-7

Ozkul, D. (2025). Constructed objectivity in asylum decision-making through new technologies. *Journal of Ethnic and Migration Studies, 51*(14), 3629–3648. https://doi.org/10.1080/1369183X.2025.2513161

Palmiotto, F. (2024a). Procedural fairness in automated asylum procedures: Fundamental rights for fundamental challenges. *Computer Law & Security Review, 55*, Article 106065. https://doi.org/10.1016/j.clsr.2024.106065

Palmiotto, F. (2024b). When is a decision automated? A taxonomy for a fundamental rights analysis. *German Law Journal, 25*(2), 210–236. https://doi.org/10.1017/glj.2023.112

Rask Nielsen, T., & Holten Møller, N. (2022). Data as a lens for understanding what constitutes credibility in asylum decision-making. Proceedings of the ACM on Human-Computer Interaction, 6(GROUP), Article 6, 1–23. https://doi.org/10.1145/3492825

Stewart, L. S. (2024a). Fair and efficient asylum procedures and artificial intelligence: Quo vadis due process? *Computer Law & Security Review, 55*, Article 106050. https://doi.org/10.1016/j.clsr.2024.106050

Stewart, L. S. (2024b). The regulation of AI-based migration technologies under the EU AI Act: (Still) operating in the shadows? *European Law Journal, 30*(1–2), 122–135. https://doi.org/10.1111/eulj.12516

Stoyanova, V. (2026). *Artificial intelligence and asylum decision-making: Any role for human rights law?* (Delmi Report 2026:3). Delegation for Migration Studies. https://www.delmi.se/en/publications/report-2026-3-artificial-intelligence-and-asylum-decision-making-any-role-for-human-rights-law/

Turley, A. (2026, November 3). *Asylum: Artificial intelligence (UIN 37730)* [Written answer to parliamentary question]. UK Parliament. https://questions-statements.parliament.uk/written-questions/detail/2026-10-23/37730/