

The background illustration depicts a person with dark hair wearing a yellow sweater and a white robot sitting at a dark grey desk. The person is on the left, typing on a laptop. The robot is on the right, also with its hands on the desk. Above them, a dotted rectangular box contains several icons: a lightbulb, a mountain landscape, and two interlocking gears. To the left of the person, there are two overlapping white rectangular shapes with blue arrows pointing to the right, resembling document or code snippets. The entire scene is set against a solid blue background.

AI in audit

Illustrative example and documentation guidance

The FRC does not accept any liability to any party for any loss, damage or costs howsoever arising, whether directly or indirectly, whether in contract, tort or otherwise from any action or decision taken (or not taken) as a result of any person relying on or otherwise using this document or arising from any omission from it.

© The Financial Reporting Council Limited 2025
Financial Reporting Council
13th Floor
1 Harbour Exchange Square
London
E14 9GE

Contents

Introduction	4
Illustrative example	7
Guidance on documenting tools that use AI	12

Introduction

The use of artificial intelligence (AI) in the performance of audit procedures, and other tasks performed on an audit, has been hypothesised for many years. This prospect is fast becoming a reality - we have begun to see tools that use AI being deployed on live engagements, and many more such tools are in development.

The FRC always encouraged innovation and we believe AI, deployed responsibly and appropriately, has the potential to significantly enhance audit quality. Higher audit quality supports greater trust in UK companies' financial reporting, reducing the risk premium the market may charge them to access capital, and therefore improving their competitiveness and ability to grow.

This publication comprises two parts: an illustrative example of a potential use case of AI to enhance procedures over journals, and guidance on documenting tools that use AI.

In this publication, the term AI is used to refer to a broad range of systems. The intended scope comprises both traditional machine learning techniques and deep learning models, including generative AI.

The material in this publication is not prescriptive, and does not represent a static set of FRC expectations. The FRC recognises that this field is moving quickly and we will continue to engage across the profession, both in the UK and internationally, to ensure our standards and guidance remain appropriate.

Illustrative example

This example is intended to demonstrate a widely applicable use case of an AI enabled tool, to enhance the identification of potentially fraudulent journals. It describes some, but not all, of the judgements of the firm and the engagement team, across development and use of the tool. It is not meant to imply that these are the best or only reasonable judgements, merely to illustrate a coherent approach and emphasise some of the key considerations.

Many of the principles and judgements illustrated can translate to similar contexts with respect to other tools, with the application of professional judgement required to adapt them.

Guidance on documenting tools that use AI

ISQM (UK) 1 requires the firm to obtain or develop, implement, maintain and use appropriate technological and intellectual resources to enable the operation of the system of quality management and the performance of engagements, and document how they have addressed the risks that they might not meet these objectives.^{1 2}

This guidance seeks to clarify FRC expectations on what should be documented by firms in relation to tools that use AI, both centrally and on the audit file, in cases where these tools are deployed on an engagement. Many of these will be consistent with expectations in relation to traditional pieces of technology, but some are specific to AI, reflecting certain unique characteristics of AI models. For example, the “black box” nature of many AI models, and their ability to learn autonomously, enhances the importance of clear documentation around the governance over the tool, its development and how it was designed to be appropriately explainable. Further, our expectations are informed by the regulatory environment with respect to AI, for example the government’s 5 AI principles.

Many of the topics discussed in this guidance, for example, how a tool is described, what certification process it has been through and what training and guidance is available to teams, are covered in the FRC’s [Certification of Automated Tools and Techniques](#) thematic review. The thematic review offers insights into common and good practice with respect to these issues, which may support firms in implementing effective processes and systems.

Automated tools and techniques (ATTs) that use AI will require documentation on the audit file in accordance with the requirements of ISA (UK) 230³; this may include ATTs deployed on component audits. This need not, however, comprise all of the information that the firm has documented on the development and operation of the tool. This guidance sets out expectations on what should be found on the file, and what can reasonably be documented centrally.

¹ ISQM (UK) 1, 32f, g

² ISQM (UK) 1, 57c

³ ISA (UK) 230, 8

Tools that use AI may be developed by the firm, or obtained from a 3rd party. The overarching principles in either case are the same, but it may be impossible to obtain some information from a 3rd party that the firm would have if they had built the tool themselves. We include some considerations that the firm may wish to think about in these cases, to support them in fulfilling the requirements.

This guidance covers both automated tools and techniques, i.e. technology used to perform audit procedures, and other tools that use AI that may be deployed on an engagement. Examples of the latter may include tools that create first drafts of workpapers, tools that can review work to identify omissions or inconsistencies, or chatbots that the team can use to ask queries about the firm's methodology.

This guidance is not prescriptive, and the requirements against which firms will be assessed remain only those in the ISQMs and ISAs (UK). Indeed, professional judgement must be applied when determining what documentation is appropriate. For example, if a large language model-enabled tool is used to research an aspect of an audited entity's environment, as part of risk assessment procedures, this tool is arguably an ATT. However, if the outputs of the tool are all checked back to the cited sources, it may well be appropriate to document less about the tool than this guidance would suggest, as the corroboration limits the relevance of the fact that an AI-enabled tool has been used.

We intend for this material to support proportionate documentation on audit files, as over-documentation can divert time and resource from areas where they can better enhance audit quality.

Illustrative example

Use of an artificial intelligence enabled tool to test journals

Illustrative example

Background information

XYZ LLP have been the auditor of ABC PLC, a listed retail business, for a number of years. XYZ LLP has recently completed in-house development of a technology-enabled tool to be used as part of fraud procedures.

The tool leverages artificial intelligence to identify potentially anomalous items that are unusual relative to the population. Fraud procedures often consist of filtering the population of journals by applying rules based criteria which may indicate higher risk, for example, journals posted at certain times or with certain values. This tool allows the identification of more subtle patterns that may be indicators of risk, enhancing the quality of the procedure.

The tool has progressed through limited deployment on selected engagements, and has now been mandated on all engagements where the data quality is sufficient and where the risk of material misstatement due to fraud has not been reduced to an acceptably low level through other procedures.

Development of the tool

One of the first steps in the tool's development was the decision of what model to use at the heart of the tool, to identify potentially anomalous transactions. A wide range of artificial intelligence models can perform this function, but the firm narrowed their focus to two main options.

Firstly, they considered an **unsupervised machine learning model**. This is not trained on data – instead it applies statistical techniques that automatically adapt to the structure and patterns of each data set, to identify items that are unusual relative to the population.

Secondly, they looked at a more computationally intensive **deep learning model**. This is a neural network, trained on large quantities of current and historical data from the entity to recognise patterns in transactions. It then identifies potentially anomalous transactions by looking for those that deviate from these patterns.

Either option can be appropriate for identifying riskier journal entries, and for the purposes of this illustrative example we need not specify which is chosen as the rest of the material is compatible with both.

The development of either model would comprise numerous technical steps, and this illustrative example does not aim to cover all of these. However, the following boxes summarise some of the key aspects that this firm considers if they choose either one.

Illustrative example

The firm chooses the unsupervised machine learning model

It is important that any artificial intelligence system is appropriately explainable, but what constitutes appropriate explainability will vary widely based on context. For example, it will be impacted by the intended use case for the tool and how this fits into the wider methodological approach.

The firm judged that, in this context, users of the tool should be able to understand why a transaction has been identified as unusual relative to the population. In other words, the tool should communicate which features of a transaction most contributed to it being assessed as an outlier.

The firm felt that this was important information for the audit team to be aware of, as it would allow them to design further procedures to look at the outliers in a way that is responsive to the features that led to it being identified as higher risk.

This model does not require training, but data is utilised to test and calibrate the model. The firm uses a combination of real and synthetic data, and ensures that appropriate authorisation has been obtained from the entities whose data is used, and any data processing meets legal and regulatory requirements.

The firm chooses the deep learning model

Neural networks that employ deep learning to identify patterns are often called “black boxes” as the scale and complexity of their processing mean they resist straightforward explanation in ways comprehensible by humans. However, they can often be augmented with techniques from the field of explainable AI (XAI). These techniques, including Shapley Additive Explanations (SHAP) and Local Interpretable Model-agnostic Explanations (LIME), help explain which features of the input most influenced the decision of the model. In this context, this means they can reveal which features of a transaction contributed most to it being identified as anomalous.

The firm judged that it is important that users can understand why an item has been identified as potentially anomalous, for the same reasons as in the adjacent box. Therefore, they augmented their model with XAI techniques. The techniques they choose to implement are ones that explain input-output relationships rather than the true internal processing of the model, but nevertheless provide insight into what features of a transaction most likely contributed to its identification as an outlier.

The deep learning model that the firm chose is trained anew on the data from each entity that it is deployed on. However, the tool – i.e. the underlying architecture – can be reused again and again, without retaining learning from entities it has been previously deployed on.

Illustrative example

The firm used a combination of real and synthetic data to test the model. The firm ensures that appropriate authorisation has been obtained from the entities whose data is used, and any data processing meets legal and regulatory requirements. The firm's tests of the model include assessment of whether it displays bias.

The model was frozen prior to deployment, i.e. it no longer continues to learn after it has been rolled out on live audits.

Development of the methodology

The development of the methodology that supports use of the tool was concurrent with the development of the tool itself, and representatives from the firm's central methodology team were included on the overarching project leadership team. This promoted significant collaboration between methodology and technology experts, and ensured the final tool embedded well into the firm's methodology.

One of the decisions the firm was faced with was whether use of the tool should replace existing rules based techniques for identifying riskier journals, or whether the two approaches should be combined.

The firm felt that a combination would be most appropriate. The new tool only identifies items that are unusual relative to the population so, for example, transactions that are consistently posted to unusual accounts may not be identified.

This risk can be mitigated by running a rules based routine that identifies journals posted to account combinations deemed unusual by the firm's central technical team, alongside the artificial intelligence tool.

In fact, the firm chose to combine the artificial intelligence tool with a selection of traditional, rules based data analytics routines, into a single tool that runs them concurrently.

The rules based routines are binary pass/fail in nature, whereas the artificial intelligence tool can evaluate items along a spectrum of how unusual they appear.

Significant professional judgement was required to calibrate how much weight each routine should contribute to a transaction being identified as riskier, as well as the threshold over which a transaction is deemed high risk and must be followed up on. If a transaction were to only fail one rules based routine, or pass the rules based routines and only be deemed slightly unusual by the artificial intelligence tool, the transaction may not be high risk.

This calibration process required both theory and experimentation with data to ensure the resulting approach was robust. The firm found that, in most cases, running the final tool resulted in similar or smaller numbers of transactions being identified as high risk as under their previous approach, though this was not itself an objective of the calibration process.

Illustrative example

Use of the tool

The team begins by ascertaining whether the criteria for use of the tool to be mandatory are met. Firstly, they assess whether data quality is sufficient. This requires that the general ledger data is complete and accurate, and that it contains the requisite data fields including date, time, poster, description, value and account references. Then, they judge whether they have already obtained sufficient appropriate evidence to reduce the risk of a material misstatement due to fraud to an acceptably low level.

The team documents that the criteria are met, that they will therefore use the tool, and that it uses artificial intelligence to identify potential anomalies. The tool is integrated into the audit software, which has controls in place to only permit use of the latest approved version of the tool.

The team runs the tool, and it identifies a selection of journals deemed high risk. The team then follows these up, and obtains sufficient appropriate audit evidence that they are not fraudulent. The firm's methodology requires them to consider why the transaction was identified as potentially anomalous in order to determine what work might be appropriate to obtain this evidence.

Further, the methodology requires them to be alert for any information that indicates that the tool's assessment of items as high risk or not may be systemically flawed in the context of this engagement.

This includes a requirement to understand, for those high risk items deemed unusual by the artificial intelligence model, why it is reasonable that those items were identified as unusual.

In this instance, among the items that the tool deems high risk, there is a subset that were identified as unusual relative to the population by the artificial intelligence model, as their value is uncommon for journals with that description posted at that time. However, when the team follows this up, they conclude that the values of these journals are correct. This may indicate that the other journals with that description posted at similar times, that were not identified as unusual, may in fact be the ones that display riskier characteristics and merit further attention.

The team looks into these journals further, and finds that their description has been incorrectly recorded, but that the transactions themselves are genuine. The team concludes that each of the high risk journals, including the additional ones that were not originally selected, is not fraudulent.

The team documents their rationale for this conclusion, as well as why they judge the issue of journal descriptions not being accurate is contained to those identified.

FRC guidance on documenting tools that use AI

Document centrally, for both ATTs that use AI and other tools that use AI

Topic	Key material to be documented, where applicable	Considerations where the tool is obtained from a 3 rd party
Description of the tool and its function	Explanation of what the tool does, at a conceptual level. The objective of using the tool. For an ATT, this will include how the tool contributes to assessing risk or obtaining audit evidence, or both. The nature of the underlying technology, in broad terms.	Same expectations as for tools developed by the firm.
When it is appropriate to use the tool	The criteria that should be met for use of a tool to be appropriate. These may include, where applicable, characteristics the input data should exhibit, the categories of transaction that can be audited with the tool, any criteria the business model of the audited entity should meet, the sorts of query or task that can be asked of the tool.	Same expectations as for tools developed by the firm.
How the tool was developed	The rationale for commencing a project to develop the tool, which may include consultation with or input from methodology teams for ATTs. How use of the tool meets relevant requirements of auditing and ethical standards, and any provisions of the firm or network's methodology. The source of any data to be used to train or test the tool; why use of this data is permitted by law, regulation and relevant professional standards; and any processing it is subject to. The choice of model and why it is appropriate. Key elements of the model architecture. How the model is trained, including how bias is mitigated and what criteria performance is measured against. The history of when versions of the tool were approved and retired.	It may not be possible to obtain information on the full development history from a 3rd party. However, we would expect to see: How use of the tool meets relevant requirements of auditing and ethical standards How the firm ensured that any use of data to train or test the tool is permitted by law, regulation and relevant professional standards. This may be in the form of an independent assurance opinion. The versions of the tool currently supported by the 3rd party. If, how and when the tool will be updated.

Document centrally, for both ATTs that use AI and other tools that use AI

Topic	Key material to be documented	Considerations where the tool is obtained from a 3 rd party
Why you are confident the tool works as intended	The governance architecture around the development and operation of the tool. The key steps in any certification process that the tool has been through, which include tests of the operation of the tool.	It may not be possible to obtain all of this information from a 3rd party, so the firm may choose to rely on independent assurance that the tool is operating as intended. Where the tool is continuously updated, without identifiable version numbers, the firm documents how it ensures the tool remains appropriate for its purposes.
What training, guidance and support is available to teams	The material available on when it is appropriate to use the tool, how to use it and how to interpret its outputs, including strategies to mitigate automation bias.	Same expectations as for tools developed by the firm; the material may be obtained by the 3 rd party or developed by the firm.
How the tool is appropriately explainable	How the tool was designed so that it is appropriately explainable. The level of explainability that is appropriate may vary based on the intended use of the tool. Explainability is a measure of the extent to which the behaviour or decisions of a tool can be understood, rather than the extent to which the inner mechanics and processing of the model are transparent and understandable to humans. Appropriate explanations may, particularly in relation to tools that rely on neural networks, be approximate or post hoc explanations that seek to explain how inputs influence outputs rather than the internal features and workings of the model.	It may not be possible to obtain all of the relevant information on the design of the tool from the 3 rd party, so the firm may have to form a judgement on whether the tool is appropriately explainable based on the information they can obtain.

Document centrally, for both ATTs that use AI and other tools that use AI

Topic	Key material to be documented	Considerations where the tool is obtained from a 3 rd party
If not covered above, why use of the tool aligns with the 5 government AI principles	If not documented in relation to another topic, or at a firm-wide level, how use of the tool follows the 5 government AI principles of: 1. Safety, security and robustness 2. Appropriate transparency and explainability 3. Fairness 4. Accountability and governance 5. Contestability and redress	In relation to some of these principles, it may not be possible to obtain all of the relevant information from the 3 rd party, so the firm may have to rely on an independent assurance opinions to inform their conclusion on whether the principles have been met. This may include opinion on the 3 rd party's resilience to cyber attacks.

Document on the audit file, for ATTs that use AI

Topic	Key material it may be appropriate to document – as a guiding principle, the more widely used a tool is across engagements, the more we would expect to see the balance shift toward central documentation, and the more bespoke a tool or use case, the more we might expect to see on the audit file
Description of the tool and its function	Brief explanation of what the tool does, at a conceptual level. The objective of using the tool, including how the tool contributes to assessing risk or obtaining audit evidence, or both. The version number of the tool and model used, as applicable, which can be cross referenced to a centrally maintained list of versions approved for use at the relevant time. If controls exist that prevent access to all versions of the tool other than the latest approved one, it may be appropriate to document that control centrally, instead of documenting the version number here. Information on any configuration or modification of the tool that the team has performed, as well as what data or prompts, where applicable, were input into the tool.
Why the team considered use of the tool appropriate	The team's assessment against the centrally determined criteria that should be met for use of a tool to be appropriate. This may include record of any consultation with a relevant central function. In particular, how the team ensured any input data is complete and accurate.
Evidence of the tool's approval for use from relevant central function	Evidence that the relevant central function has approved the tool for use on the engagement. For tools that are approved for use on every engagement, this approval can be held centrally.
Consideration of the outputs of the tool	How the team used the outputs of the tool to conclude on the relevant judgement, or to inform further procedures. Where the team has run the tool a number of times, the team exercises professional judgement in determining whether to document key aspects of each, or of the final instance only.

Document on the audit file, for other tools that use AI

There may be no requirement to document the use of these tools on the audit file, if it would not be required for an experienced auditor to understand the basis for the auditor's report, or significant matters arising during the audit. However, the team may choose to do so, if they feel it would allow a reviewer to better understand the work performed.



Financial Reporting Council

**Financial
Reporting Council**

London office:
13th Floor, 1 Harbour
Exchange Square,
London, E14 9GE

Birmingham office:
5th Floor, 3 Arena
Central, Bridge Street,
Birmingham, B1 2AX
+44 (0)20 7492 2300

www.frc.org.uk

Follow us on

Linked in