



How to build a judgement layer for AI use in surveying

October 2026

How to build a judgement layer for AI use in surveying

October 2026

Published by the Royal Institution of Chartered Surveyors (RICS)
Parliament Square
London
SW1P 3AD
www.rics.org



No responsibility for loss or damage caused to any person acting or refraining from action as a result of the material included in this publication can be accepted by the authors or RICS.

ISBN 978 1 78321 582 9

© Royal Institution of Chartered Surveyors (RICS) October 2026. Copyright in all or part of this publication rests with RICS. Save where and to the extent expressly permitted within this document, no part of this work may be reproduced or used in any form or by any means including graphic, electronic, or mechanical, including photocopying, recording, taping or web distribution, without the written permission of RICS or in line with the rules of an existing licence.

Acknowledgements

RICS would like to thank the following for their contributions to this insight paper:

Authors

Anil Sawhney FRICS (RICS)

Chris de Gruben FRICS (Artefact)

James Garner FRICS (Gleeds)

David Wong MRICS (The Augmented Practitioner Column)

RICS professional practice lead

Sarah Crouch

RICS editor

Sam Birch

Contents

Acknowledgements	ii
1 Introduction	1
2 Professional judgement in AI-assisted work	2
3 Evidence of professional judgement	3
4 How AI changes where judgement is applied	5
5 How does the judgement layer work?	7
6 Building the judgement layer	9
6.1 Connect	9
6.2 Codify	10
6.3 Govern	11
6.4 The three elements in practice	14
7 Using the judgement layer	15
7.1 Start with one workflow	15
7.2 Make judgement visible	16
7.3 Protect how surveyors learn	16
7.4 What this means for the profession	17

1 Introduction

The use of AI in surveying practice is rapidly increasing. Surveyors already use AI-enabled systems to support information retrieval, analysis, document review and professional decision-making. Wider adoption raises the questions of how and when professional judgement applies to AI-assisted work. RICS' [Responsible use of AI in surveying practice](#) professional standard ('the RICS AI standard') places that judgement at the heart of surveying practice.

This insight paper is written for practising surveyors, firms, educators and other professional stakeholders seeking to understand how professional judgement, accountability and professional practice are evolving in AI-enabled surveying work. It builds on the RICS AI standard and explores how judgement can be exercised and governed in practice.

RICS' [previous article](#) made the case for how AI elevates professional judgement.

This insight paper introduces the surveyor's judgement layer: a practical system that helps surveyors distinguish AI-assisted work from professional advice and decide when and how to rely on it.

The judgement layer is an organisational capability that helps firms apply professional judgement more efficiently and accountably. This paper sets out:

- 1 what that system contains
- 2 who builds each part and
- 3 how judgement is applied, evidenced and used to support professional sign-off.

To achieve this, the paper is structured as follows:

- Section 2 examines the role of professional judgement in AI-assisted work.
- Section 3 explains what counts as evidence of professional judgement.
- Section 4 explores how increasingly capable AI systems change how and where judgement is applied.
- Section 5 introduces what the judgement layer actually does.
- Section 6 explains how firms can build their own judgement layer.
- Section 7 considers the implications for firms, including sole practitioners and small firms, as well as for professional development and the profession as a whole.

2 Professional judgement in AI-assisted work

AI can produce information, analysis and draft conclusions, but it does not turn them into professional advice without professional judgement. The RICS AI standard defines professional judgement as a combination of 'knowledge, skills, experience and professional scepticism'. By applying that judgement, a surveyor decides whether they can rely on AI-assisted work for its intended purpose.

For surveyors, the value of AI therefore lies in whether it helps turn information into advice a client can act on, such as a valuation, an identified risk or a strategy for reducing carbon emissions, and **whether that advice is sound and defensible**.

An AI system can summarise a lease in minutes, but it can miss how an unusual break clause shifts the client's position. The surveyor decides whether to rely on the summary, what is missing and whether further work is needed. AI-assisted work and its output are only the starting point. The professional product is the surveyor's advice, shaped by their judgement. The practical challenge in responsible AI use is making that judgement explicit enough to rely on. The judgement layer supports this by organising the information, knowledge and controls needed to exercise and evidence that judgement.

The challenge is no longer whether AI can produce information quickly and at scale. The task for the profession is to ensure that professional judgement continues to be exercised, evidenced and relied upon when AI-produced information contributes to professional advice.

3 Evidence of professional judgement

Surveying work has always depended on sources, data, methodology and professional experience, such as the comparables considered, the cost benchmarks applied and the site conditions observed. These may explain the basis of a conclusion, but they do not always show how the surveyor exercised judgement over them, including what was:

- weighed
- challenged
- set aside or
- treated as sufficient.

In this paper, evidence of professional judgement means the working record that shows what the surveyor relied on, and how they engaged with it before accepting, changing or rejecting the result.

Until now, any gap in that record has usually been manageable because the reasoning remained with the surveyor who carried out the work. If asked why a comparable was rejected, why a rate was adjusted or why a defect was judged significant, the surveyor could explain it by retracing the reasoning through contemporaneous notes, source material and recollection.

AI-assisted work weakens that link. A system can produce an articulate conclusion that appears well-supported, even when the reasoning or evidence behind it is incomplete. This risk not only applies to the final output; it can emerge during the AI-assisted process itself – in the sources retrieved, assumptions carried forward, checks performed or intermediate steps taken. Focusing only on the conclusion can therefore obscure weaknesses in the process that led to the output.

When an AI system is used, the sources, data and method behind the conclusion are not enough on their own. A useful working record also shows how the surveyor engaged with the whole process:

- what they accepted
- what they challenged
- what they set aside and
- why they were prepared to rely on the result.

Producing such a record is not straightforward. Professional judgement rarely follows a straight line. A surveyor may return to an earlier assumption because later information changes its significance, or because something does not sit comfortably with their experience. Revisiting earlier assumptions is part of judgement. Unless the workflow captures those points of judgement as they occur, the record may show what was used, without showing how the surveyor fully engaged with it.

Professional scepticism, which the RICS AI standard identifies as part of judgement, gives the surveyor's engagement with AI-assisted work a practical form. It is not general caution or distrust of technology; it is the deliberate act of testing whether the AI-assisted process or its conclusion may be wrong.

For a surveyor, that may mean questioning a benchmark rate or a comparable, checking whether a source is current, testing whether an assumption has been carried forward properly or asking whether the process has missed something material. Verification is necessary, but it is not the same as professional scepticism. An AI system can be instructed to check an output against defined criteria. It does not – by itself – decide to look for weaknesses in a conclusion or process that appears coherent.

A completed check is therefore not evidence that the process or output has been challenged. As AI systems become more accurate and familiar, professional scepticism can erode and surveyors may accept plausible outputs too readily, risking complacency and automation bias. The AI-assisted workflow needs a clear point at which the surveyor tests the process and the output, considers whether something important has been missed and decides whether the result can be relied upon.

A properly kept working record, including an audit trail of the AI-assisted workflow, supports review by showing both the basis of the work and how professional judgement was applied. It preserves the sources, data, assumptions, prompts, checks and intermediate steps as they occurred, together with the points at which the surveyor questioned, accepted, changed, revisited or rejected them. The surveyor can then return to the record as it stood at the time, rather than relying on later recollection.

A good record does not need to prove that the judgement was correct. It shows what was considered, what was challenged and why the output was accepted, changed or rejected. That is what makes the decision defensible and open to review. Making judgement explicit means preserving enough reasoning to support the surveyor's decision as the work is carried out and leave a working basis for later review. That record also helps improve how similar work is carried out in future.

4 How AI changes where judgement is applied

As AI use moves beyond chatbots to more agentic workflows, professional judgement becomes more important and harder to exercise. It can no longer be applied only to the final output, because AI now spans the entire workflow, and a working record matters more. The distinction between a chatbot and an [AI agent](#) matters because it changes where judgement is applied:

- With a **chatbot**, judgement is usually applied at two points: framing the task at the beginning and judging the response at the end.
- An **AI agent** is different. It is given a goal, memory, tools and permission to act, then works through repeated steps to process information, act and observe until the goal is met. The path it follows is not fixed in advance, so judgement may also be needed when the workflow is set up and at one or more gates within the workflow.

With an AI agent, errors can occur long before the surveyor receives and checks the final output. For instance, a retrieval step may surface the wrong evidence, the agent may call the wrong tool, or an intermediate assumption may be unsound. Reviewing only the final output will not reveal these errors.

In surveying practice, an AI agent asked to assemble comparable evidence for a valuation might draw on two comparables from an adjoining submarket with different demand and yields. The system may not flag this and may still reach a conclusion. The draft output may look plausible, but the valuation could still be wrong, damaging client trust and exposing the surveyor and the firm to professional and regulatory risk.

Professional judgement therefore applies to both the process and the output. Deciding which tools an AI agent may use, what evidence it may draw on and where it stops and returns control is itself an act of professional judgement. The surveyor exercises it at the beginning when setting up the workflow, and at the end when they decide whether to rely on the output as professional advice. The number and placement of intermediate checkpoints may change as agent capability improves, but those beginning and end decisions remain essential.

Figure 1 compares this with the chatbot approach.




Figure 1: Application of professional judgement in chatbots and AI agents

At the outset, the surveyor decides what work the AI agent is authorised to carry out, including:

- the task the AI agent may carry out
- the evidence it may use
- the actions it is not permitted to take and
- the decisions that must return to the surveyor.

These choices define the task, evidence boundary and points at which control returns to the surveyor. If the goal is vague or the evidence boundary is wrong, a careful review at the end may only confirm a plausible answer to the wrong question.

At the end, the surveyor does more than check whether the output looks reasonable. They decide whether it can be relied upon for its intended professional purpose. This may involve:

- testing the evidence
- challenging an assumption
- asking for further work
- changing the conclusion and
- rejecting the output.

To make these decisions consistently, the surveyor needs access to the right information, the firm's relevant knowledge and clear controls across the workflow. Together, these form the judgement layer described in the next section.

5 How does the judgement layer work?

The judgement layer is not software. It is a firm's way of organising information, knowledge and controls so that AI-assisted work can be reviewed and signed off properly – from setting the task to delivering the advice the client receives. The firm builds the judgement layer, but professional judgement and accountability remain with the surveyor who uses it. In a small practice, one person may perform both roles, although the roles remain distinct.

The term 'judgement layer' is already used by software vendors, who build domain expertise and customer-specific preferences into the AI systems they sell (whether standalone systems, embedded features or autonomous agents). This insight paper uses the term differently, focusing on the firm's organisational capability rather than a feature of an AI system.

The judgement layer matters wherever AI contributes to work that may influence professional advice. The need for a judgement layer is likely to grow as AI systems operate with less direct supervision.

It has three complementary elements (see Figure 2).

- **Connect** brings together data and information spread across the organisation, making relevant evidence available and traceable.
- **Codify** captures organisational knowledge by setting out how that evidence should be treated, drawing on previous work and professional experience.
- **Govern** sets the limits, checks and records that allow AI-assisted work to be used responsibly. It defines where the work must be checked, paused or escalated.

Connect and codify together provide the context within which professional judgement is exercised, while govern provides the controls through which it is applied.




Figure 2: The three main components of the judgement layer

The judgement layer supports both quality and accountability. It is not a product a firm can buy and switch on. It is built through the way work is organised, recorded and improved over time. None of this requires a large firm's resources; it simply requires two things:

- 1 clarity about what matters and
- 2 the discipline to record it.

For a sole practitioner:

- **Connect** may begin by identifying that the condition records needed for a building survey sit in a folder the AI system cannot access, then taking a practical step to make that evidence available.
- **Codify** may start with one written rule about how those records must be checked.
- **Govern** may be as simple as knowing where the AI system stops and recording what was checked before the output is used.

6 Building the judgement layer

The connect, codify and govern elements are complementary rather than sequential because each element supports the others. The judgement layer develops through repeated use: it becomes stronger when each job produces a clearer rule, a better record or a more reliable way to handle the next task.

6.1 Connect

Connect is about ensuring the AI system can access the evidence the surveyor would normally use, and that the surveyor can see what evidence the system relied on. A good prompt and a careful final review are not enough if the system has worked from incomplete historical records. In that case, the surveyor may be judging a plausible solution without knowing what evidence was missing or incorrectly referenced.

For most firms, the starting point will be familiar: it is finding the evidence surveyors need. That evidence is often spread across shared drives, project folders, spreadsheets, email attachments, legacy systems and paper records. Its quality is uneven, and it is not always easy to find. [McKinsey's 2026 analysis](#) of the architecture, engineering and construction sector identifies the same problem in archives of drawings, specifications and close-out reports that are difficult to use because they are inconsistent, paper-based or siloed. [RICS has also highlighted](#) the risk this creates when AI is used in the built and natural environment.

Imperfect records need not delay AI adoption. The priority is to focus effort on the evidence that matters most to professional work. Agentic systems handle messy material well enough that firms '[do not need perfectly clean data to start creating value](#)', so cleaning every historic record is rarely the best first investment.

The more important question is whether the evidence the AI system relies on can be found, checked and understood before the work is signed off. A practical first step is to map where the firm's main evidence is stored and what the AI system can access.

6.2 Codify

Connect makes evidence available and traceable, but it does not provide the professional experience needed to give that information meaning and context. That leaves a knowledge gap. In the built and natural environment, the gap is acute because valuable know-how often sits with experienced professionals rather than in systems the firm can reuse. When [someone leaves or retires](#), much of that organisational memory can leave with them.

Professional judgement cannot be transferred to an AI system, but the knowledge, lessons and institutional memory that support it can be codified. This means bringing together project archives, lessons learned, technical guidance and standard procedures, so surveyors draw on organisational experience rather than individual recollection at the point of decision. Many firms have invested in knowledge-management systems for a long time, but these systems often remain separate from where professional judgement is exercised. For some firms, the codify element will therefore mean enhancing and connecting what already exists, rather than creating a new knowledge base from scratch.

Codification happens through use. An AI agent assists with a task and creates an audit trail of sources, tool calls, assumptions, intermediate outputs and review points. A surveyor reviews that audit trail, extracts the general lesson and adds it to the firm's knowledge base. This makes codify the slowest of the three elements to build. A rule can be written in one step, but a firm's body of knowledge is built case by case and improves the next job only once enough knowledge has accumulated.

As access to powerful AI models becomes more widespread, the value of organisational knowledge may increasingly lie in the record of professional decisions, corrections, exceptions and outcomes. [The World Economic Forum has argued](#) that expert judgement and institutional feedback are becoming a scarce and valuable input for AI-enabled work because competitors find them hard to acquire or reproduce. Such records therefore support both better decisions and a firm's long-term competitive advantage.

A practical starting point is a small, recurring decision. Codification captures what can be written down, and much of what makes delivery go well may be hard to capture. The judgement that matters most is often the judgement that is difficult to document, which argues for codifying recurring cases properly, rather than attempting to codify everything. Much of this unwritten judgement concerns recurring decisions about risk, exceptions, quality and escalation. For a surveying firm, these are familiar decisions, such as which comparables are admissible, when an assumption needs a second review and when a job is escalated to a senior colleague. Codifying one of these decisions produces the first entry in the firm's knowledge base.

Over time, this codified knowledge becomes part of the firm's judgement layer, although codification has limits. The AI system may read a written rule and follow it, ignore it or apply it at the wrong point. It is the third element that turns the rule into a control.

6.3 Govern

Govern turns the firm's rules into controls it can rely on. It sets:

- what an AI system is permitted to do
- where it stops
- what is checked and
- who signs off.

Some checks and controls are built into the AI system itself. Software engineers call this structure a harness (sometimes a scaffold), and this insight paper uses the term in the same sense.

The harness does not make the system safe on its own. It sets the limits within which the system can operate, outlining the AI system's permitted actions, its pause points, how its actions are recorded and reviewed, and whether the work can be halted once it has begun.

Approval gates cover the conditions the firm has anticipated. They are necessary, but they are not enough. An AI agent may still use an unexpected source, take an inappropriate sequence of actions or continue after the work should have paused. Firms therefore need a way to pause the work, withdraw permissions and decide what happens to anything the AI system has already produced. [McKinsey describes this as a management discipline](#) rather than a technical detail. Autonomous systems need a defined mandate, budget and stopping rule. For a surveying firm, the practical questions are:

- Who can stop the work?
- How quickly can this happen?
- What review process is required before any part of the work is used?

In many surveying applications, these controls will sit partly inside a vendor's tool, such as automated take-offs, suggested comparables, AI-drafted sections of a level three home survey, contract review or drafted clauses. The defaults in those tools are therefore important. They determine what the system can do, when it pauses and what it records. But the vendor has not seen the client's instruction, does not know the full professional context and does not sign the result. Vendor defaults are therefore design choices to test, not controls to accept without questioning. Responsibility remains with the surveyor and the firm that relies on the output of such AI systems.

A firm can respond in three ways:

- 1 **Specify** what the vendor's harness must do, using input from those who understand the professional work.
- 2 **Test** the vendor's harness against that specification. Not every firm can examine the system in detail, so the aim is to obtain enough evidence and transparency to judge whether the system is suitable for professional use. The RICS AI standard treats

supplier due diligence as an important requirement, but this is not just a procurement exercise. It is a professional judgement about whether the firm has enough evidence to rely on the system. A firm can ask for clear information about the system's controls and, where relevant, whether the supplier has [ISO/IEC 42001 certification](#) (the management system standard for AI).

- 3 **Add** the firm's own controls by setting limits, deciding where approval is needed and reviewing the audit trail.

These are controls the vendor's harness does not provide (see Figure 3).




Figure 3: The firm specifies and tests the vendor's harness, and adds its own controls

The firm may set its own controls, but those controls often depend on how the vendor's harness operates. If that harness does not apply the control as intended, the firm may not know until the surveyor has already relied on the output. Governance therefore needs evidence that the controls are working. That evidence could come from:

- audit logs
- testing
- exception reporting and
- periodic review of system behaviour.

The aim is to obtain enough assurance that the controls the surveyor relies on are operating as intended, without inspecting every technical component.

By specifying what the harness must do, testing it and setting its own controls, a firm puts the RICS AI standard into practice. The standard makes users and regulated firms responsible for professional work carried out with AI support, and its requirements for

a responsible use policy, an AI systems record and supplier due diligence are firm-side controls.

The RICS AI standard therefore gives firms a head start, and formal AI regulation may take time to catch up. Still, practical limits are already being set through contracts, insurance requirements and client expectations, and professional bodies are writing AI into their codes of conduct. A firm that builds these controls now is better placed to comply with the RICS AI standard, while preparing for future client and insurer expectations. Firms working towards ISO/IEC 42001 recognise similar organisational infrastructure.

Govern also produces two related but distinct records:

- 1 **The system trail:** this captures what the AI system did, including the sources used, tool calls, assumptions and intermediate steps. It supports review of how the workflow operated and feeds the codify loop.
- 2 **The record of the professional decision:** this captures the surveyor's decision about whether the output can be used for its intended professional purpose. It identifies what is being relied on, the main assumptions and reliability concerns, and where professional scepticism was applied, such as in response to a weak dataset or an unsound assumption. A named surveyor makes the decision on whether the output is used and accepts responsibility for that use. The system trail supports the professional decision, but it cannot make or replace it.

Where AI produces work at scale, the judgement may sit at the level of an approved method rather than each output. A firm may approve a defined workflow for a particular class of output, subject to stated assumptions, limits and approval conditions. Regular randomised dip sampling is required to test whether the workflow continues to operate as intended and to reveal drift, data changes, failed gates or exceptions that the approved method did not anticipate.

Sampling does not replace professional judgement. It supports the approved method and does not remove the need for separate review where an individual output is material, unusual or outside the conditions on which the method was approved.

Both records fit in the judgement layer because together they let the firm learn from a decision after the job is complete. They should not create an administrative burden or be more detailed than their purpose requires. A useful record is concise enough to be used in practice and specific enough to answer three questions:

- What was the system asked to do?
- What did the surveyor rely on?
- What caused the surveyor to accept, change or reject the result?

These records then create a feedback loop between delivery, governance and learning.

6.4 The three elements in practice

The three elements come together in ordinary practice when the AI system uses the right evidence, applies the firm's rules and stops for the surveyor to decide.

Consider a surveyor using an AI agent to assemble rates for a cost plan. Without a judgement layer, the agent might pull in rates from a different project type and produce a plausible cost estimate without flagging the mismatch. With the judgement layer in place, the workflow changes to account for this issue. The agent works from the firm's own cost data, the firm's rule identifies that the rate comes from a different project type and the workflow pauses before the cost plan is finalised.

The pause does not guarantee the estimate is right. The selected rates may still need to change and the rule may need refinement. The judgement layer provides timely visibility, bringing the issue back to the surveyor before the cost plan is completed. The working record then shows the basis on which the rate was used or changed, and makes clear that the advice to the client is the surveyor's decision, not the agent's draft.

This example also shows why connect, codify and govern are not separate compliance exercises.

- **Connect** determines whether the agent can access and use the firm's current rate library, project classifications and relevant market evidence.
- **Codify** supplies the firm's experience of which rates are comparable, which assumptions need explanation and when senior review is required.
- **Govern** sets the permitted tools, records the sources used and stops the workflow when the selected evidence falls outside the agreed conditions.

The surveyor then tests the result against the project, the available evidence and their own professional experience. If the rates are changed, the reason becomes part of the working record. If the process fails, capturing the lesson makes the same weakness less likely to reoccur.

7 Using the judgement layer

The RICS AI standard treats an output as material when it affects how the surveyor's work is 'rendered meaningful'. In practical terms, that is the point at which an AI-assisted output starts to influence professional advice. The judgement layer helps the surveyor decide whether to use that output by making the evidence, limits and controls behind the work visible to review.

The judgement layer develops through regular use. Each assignment can improve the record of where the evidence sits, clarify a rule about an exception or show where a control needs tightening. Over time, those small additions become part of how the firm works. They support compliance and improve delivery. Because they reflect the firm's own projects, clients and professional decisions, they are not easily copied or transferred beyond the firm. That is why the RICS AI standard is more than a policy-writing and compliance exercise; it gives firms a practical reason to improve how they organise AI-assisted work.

The potential gains are substantial, but applying AI without redesigning the work around it also carries significant risks. [McKinsey estimates](#) that AI and other automation technologies could unlock about \$228 billion a year in the US architecture, engineering and construction sector by 2030, with a further \$126 billion in European construction. Its analysis distinguishes between firms that redesign work and those that add AI to existing workflows. The former are likely to realise greater gains. The latter risk automating existing weaknesses, increasing the volume and speed of work without improving its reliability, and exposing the firm to errors that are harder to identify and explain. Building the judgement layer is part of redesigning the work so that firms can realise the benefits while retaining professional control.

7.1 Start with one workflow

A firm does not need to build the whole judgement layer before using AI. It can start with one recurring activity where the output could influence advice and the workflow is clear enough to map.

- 1 Identify the evidence normally used and check what the AI system can access.
- 2 Write down one or two recurring decisions and define where the system must stop for review.
- 3 Test the workflow on real examples and adjust the rule, the data or the control when exceptions arise.

This is a practical route from undocumented experience to a repeatable, defensible way of working.

The starting point will look different depending on the firm. A sole practitioner or small firm may begin with a short checklist and a properly kept working record. The judgement layer can be implemented proportionately, but applicable RICS AI standard requirements, including AI system records and supplier due diligence, still apply. More complex AI use may require additional controls.

The formality changes with scale, but the professional obligation remains the same: before relying on the work, the surveyor still needs to understand the basis on which the AI system produced it.

The judgement layer needs to handle exceptions as well as routine cases. A workflow may work well on familiar projects and still fail when the evidence is thin, the instruction is unclear or the client's need sits outside the purpose for which the workflow was designed. Firms can decide in advance which conditions will pause or escalate the work. Such conditions could include:

- encountering an unusual source
- conflicting evidence
- a material change in assumptions or
- a result outside the expected range.

The aim is to bring the surveyor back into the work before a weak result becomes bad advice.

7.2 Make judgement visible

The judgement layer makes it possible to understand and review a professional decision later. If the working record only says that an output was accepted, it is little more than a sign-off form. If it shows the evidence considered, the uncertainty identified and the reason for the conclusion, it does something more useful: it allows a client, insurer, regulator or colleague to understand, when needed, how the advice was reached.

This also benefits the individual surveyor. The firm may own the systems, records and client-specific information, but the surveyor develops professional capability through repeated use. That includes better judgement and a record of sound decisions. The AI system cannot create that capability, and the firm cannot fully retain it. For surveyors early in their careers, the RICS AI standard therefore also supports developing the judgement they will need as AI systems change.

7.3 Protect how surveyors learn

Routine work is often described as low-value because it can be automated, but it has also been one of the main ways new surveyors learn on the job. By assembling evidence, spotting exceptions, testing a first answer and seeing how experienced colleagues challenge it, they develop the judgement the profession relies on. If that work disappears without anything to replace it, firms may gain efficiency now at the expense of weakening their future professional capability.

A junior surveyor needs to see more than the finished AI-generated answer. They need supervised opportunities to take part in the work that shapes it, including:

- framing the task
- inspecting sources
- questioning assumptions and
- making proportionate decisions before a senior surveyor signs off.

Supervision is therefore most useful when it exposes the reasoning path, rather than only checking the final document for obvious errors. Senior surveyors are best placed to decide which decisions can be delegated with AI support, which require joint review and which remain with a senior professional.

Judgement is not an innate quality that professionals simply possess. It develops through repeated exposure to evidence, exceptions, uncertainty and consequences. AI changes how those experiences are acquired, but it does not remove the need for them.

Used well, the judgement layer can make apprenticeship-style learning more visible by showing where experienced surveyors intervene and why.

Used poorly, it could leave junior surveyors at the edge of AI-assisted work, reviewing conclusions without understanding the decisions behind them. Codify would also be weaker, because useful lessons come from real work, not from sign-off alone. Over time, expertise could become concentrated among fewer senior professionals, leaving firms with fewer people able to deal with AI-assisted work when it becomes uncertain, unusual or contested.

7.4 What this means for the profession

As AI changes how professional judgement develops and is exercised, the issue becomes more than an internal training concern. RICS, educators and employers will need to decide which capabilities remain essential as AI supports more routine production.

Relevant professional capabilities will increasingly include:

- defining tasks clearly
- assessing the quality and provenance of evidence
- recognising when an AI-assisted process has reached its limit and
- explaining the basis of a professional decision.

These capabilities need to be visible in training, supervised experience and continuing professional development. The key question for the profession is not whether AI can perform more work; it is whether surveyors will still be trained and assessed to exercise judgement when it does.

The RICS AI standard gives members and regulated firms a baseline for responsible AI use. The judgement layer makes that baseline practical. It links the standard's requirements and the wider implications of AI-assisted practice to the four levels of responsibility shown in Table 1.

Level	Strategic responsibility
Individual surveyor	Develop and maintain sufficient knowledge of AI, including its limitations, failure modes, bias and data risks, and apply professional judgement when deciding whether material AI-assisted outputs can be relied on.
Firm	Assess where AI is appropriate, and maintain the policy, AI systems record, risk register, data governance, supplier due diligence, controls, training, client communication and explainability needed for responsible AI use.
RICS and professional bodies	Set and maintain the standards, pathway guides, qualification requirements, continuing professional development (CPD) expectations, guidance and assurance expectations that support responsible AI practice, and help the profession identify and share emerging good practice.
Profession	Protect the practice conditions through which judgement develops, including supervised experience, exposure to evidence, delegated responsibility and review of real work.

Table 1: Applying the RICS standard through the judgement layer

The aim is not to build the entire judgement layer at once, but to strengthen it incrementally as decisions are made, recorded and reused.

AI models will keep improving, but model capability may not be the main constraint. The real constraint may be whether AI can be applied reliably to professional work. Firms do not need to wait for the next generation of AI or new regulation to start. They can decide now:

- which work is suitable for AI support
- what evidence the system may use
- where it must stop and
- how the decision to rely on its output is recorded.

The surveyor's signature has always been the point at which information becomes advice. AI changes what precedes the signature, but not what the signature means.

Delivering confidence

We are RICS. As a member-led chartered professional body working in the public interest, we uphold the highest technical and ethical standards.

We inspire professionalism, advance knowledge and support our members across global markets to make an effective contribution for the benefit of society. We independently regulate our members in the management of land, real estate, construction and infrastructure. Our work with others supports their professional practice and pioneers a natural and built environment that is sustainable, resilient and inclusive for all.

Enquiries
contactrics@rics.org



rics.org