White Paper #1

Efficient Advice

Let's Make it Better!

THE SIX WAYS AI OUTPUTS FAIL IN AN AUDIT WORKPAPER

Audit AI Failure Taxonomy



PRESENTED BY

EFFICIENT ADVICE, LLC

AUTHOR

ROBINSON DE JESÚS

EFFICIENT ADVICE, LLC | WHITE PAPER | NO. 1

The Six Ways AI Output Fails in an Audit Workpaper

And How to Catch Each One Before It Reaches a Finding

Robinson De Jesús | CEO, Efficient Advice, LLC

Info@efficientadvice.com | www.efficientadvice.com

Audit AI Failure Taxonomy | Efficient Advice

Version 2.0 | October 2026

Collaboration by: Claude, Gemini, Canva, and Grammarly.

WHO THIS WHITE PAPER IS FOR

Chief Audit Executives, Audit Managers, and Internal Auditors at any level who are using AI tools to draft findings, write reports, or conduct any part of their audit work — and who have not yet implemented a formal validation protocol.

Executive Summary

AI tools do not fail randomly. They fail in six specific, named, predictable patterns. Every tool built on a large language model — including ChatGPT, Microsoft Copilot, Claude, and Gemini — is subject to all six failure modes, because the modes are properties of the shared underlying architecture, not of any specific product. Internal auditors who cannot identify these patterns by name cannot catch them before they reach a workpaper, a finding, or the audit committee.

This white paper names and defines the Audit AI Failure Taxonomy — six failure modes that represent the complete set of ways AI output degrades the quality, accuracy, and professional defensibility of internal audit work. For each mode, it explains what the failure looks like in a real audit context, how to detect it, and which of the five validation steps prevents it.

The Taxonomy is not theoretical. It is grounded in how large language models (LLMs) actually generate output — through statistical prediction rather than fact retrieval — and in the specific ways that mechanism produces professional risk when applied to audit findings, workpapers, risk assessments, and reports.

KEY FINDING

Internal audit functions using AI without a named failure taxonomy cannot systematically validate AI output, because they do not have a vocabulary for what they are looking for. The Audit AI Failure Taxonomy gives auditors that vocabulary. Naming the failure is the first step to catching it.

1. Why AI Output Fails: The Mechanism Behind All Six Modes

Before naming the six failure modes, it is necessary to understand the single mechanism that produces all of them. Understanding the mechanism is what makes the Taxonomy teachable — and what separates a trained auditor from one who uses AI without knowing its failure patterns.

1.1 How Large Language Models Actually Work

A large language model does not retrieve information from a database of facts. It does not look things up. It predicts the next most statistically likely token — word or word-fragment — based on patterns learned during training on vast quantities of text. The model has no concept of truth. It has no awareness of whether a citation exists, whether a standard applies to the organization in question, or whether a conclusion is supported by evidence. It produces output that is statistically coherent and contextually plausible.

This is not a bug that will be fixed in the next version. It is the architecture. LLMs are prediction engines, not knowledge retrieval systems. Every failure mode in the Audit AI Failure Taxonomy is a direct consequence of that architecture operating in an audit context.

1.2 Why the Audit Context Creates Specific Risk

The audit context amplifies the risk of LLM prediction errors in ways that do not apply to most other professional uses of AI. Three characteristics of audit work make AI output failure particularly consequential:

- **Findings require evidence anchors.** Every factual statement in an audit finding — every condition, every cause, every effect — must trace to a specific, documented piece of tested evidence. AI output has no evidence anchor by definition. It generates plausible language, not traced conclusions.
- **Citations require source documents.** Audit findings cite standards, regulations, and policies as criteria. These citations must exist, must be accurate, and must apply to the specific entity type and jurisdiction being audited. AI produces citations with the same confident tone whether they exist or not.
- **Conclusions require auditor judgment.** The overall opinion in an audit finding — whether a control is effective, whether a risk is adequately mitigated, whether a condition is significant — must reflect the auditor's explicit professional judgment formed from tested evidence. AI generates opinions that appear to reflect judgment but do not.

These three characteristics create the conditions under which all six failure modes produce professional risk. A statistically plausible fabricated citation is indistinguishable from a real one by tone. A confabulated causal inference sounds like a logical conclusion. A certainty inflation reads like an auditor's expert assessment. The problem is not that AI output looks wrong. The problem is that it looks right.

1.3 The Audit AI Failure Taxonomy

The Audit AI Failure Taxonomy organizes AI output failure into six named, discrete modes. Each mode has a specific mechanism, a specific audit-context manifestation, and a specific catch protocol. The six modes are:

01 Fabrication	AI invents citations, standards, and regulatory provisions that do not exist.
02 Confabulation	AI fills gaps in your evidence with plausible-sounding inferences.
03 Precision Fraud	AI produces specific numbers and percentages with no calculation basis.
04 Certainty Inflation	AI presents statistical predictions as auditor conclusions.
05 Scope Blindness	AI ignores organizational and jurisdictional context and defaults to generic patterns.
06 Standard Drift	AI applies a real standard incorrectly — wrong version, wrong entity, wrong context.

These six modes are not exhaustive of every way AI can be misused. They are the six ways AI output itself degrades audit quality when applied to workpapers and findings, regardless of how carefully the auditor uses the tool. A skilled auditor using a well-written prompt can still produce output containing all six failure modes, because the failures are properties of the model's output mechanism, not properties of the prompt.

2. The Six Failure Modes — Defined

01 Fabrication

AI states something false with complete professional confidence.

WHAT IT IS

Fabrication occurs when an AI tool produces factual claims — standard citations, regulatory provisions, numerical data, publication references — that do not exist in the source material, or that do not exist at all. The critical characteristic of Fabrication is that the output reads with the same confident, professional tone as accurate content. There is no error message, no hedging, no signal that the content was invented. The model does not know the content is wrong, because the model does not know anything — it predicts statistically plausible text.

WHAT IT LOOKS LIKE IN AN AUDIT WORKPAPER

A finding cites “IIA Standard 5.2.1(c) — AI Disclosure Requirements” as the criteria for a compliance gap. The citation is formatted correctly, integrated naturally into the finding narrative, and reads authoritatively. The standard does not exist. No IIA Global Standard uses that numbering or addresses AI disclosure in that format. The citation was generated because the model predicted that a standard of that type, in that context, would be numbered and formatted that way. It was statistically plausible. It was entirely fabricated.

HOW TO CATCH IT

Open the cited source document. Find the specific section. Read the exact language. If the section does not exist — or exists but does not say what the AI claims — the citation is fabricated. No alternative interpretation is available. Fabricated citations are not recoverable through rewording. They must be replaced with citations that exist and apply. Validation Step 1 (Verify the Facts) is the systematic application of this check to every factual claim in the output before it enters the workpaper.

VALIDATION STEP

Step 1 — Verify the Facts: Confirm every citation, standard reference, and factual claim against the original source document, section by section, before use.

02 Confabulation

AI fills evidence gaps with plausible-sounding but untested inferences.

WHAT IT IS

Confabulation occurs when an AI tool takes incomplete information provided in the prompt and fills the gaps with inferences that sound logical and professionally grounded but are not derived from the auditor's tested evidence. Unlike Fabrication, which invents external references, Confabulation invents the internal logical connections that hold a finding together — the causal analysis, the risk linkages, the effect magnitude, and the root cause identification. The model infers what the causes and effects "should be" based on patterns in its training data, not on what the auditor actually found.

WHAT IT LOOKS LIKE IN AN AUDIT WORKPAPER

An auditor provides the AI with a condition statement: "The procurement system does not require dual authorization for vendor payment changes above $50,000." The AI generates: "Cause: The absence of a formal procurement governance policy has resulted in inadequate segregation of duties controls, stemming from insufficient board-level oversight of financial control frameworks." The auditor never tested the governance policy, never reviewed board oversight documentation, and never established the causal chain. The AI inferred a plausible-sounding root cause from the condition alone. None of it is in the workpaper.

HOW TO CATCH IT

Trace every Cause-and-Effect sentence in AI-generated finding content directly to a specific line, reference, or observation in your workpaper. A cause that cannot be traced to tested evidence was not established — it was inferred by the AI. The test is binary: either the workpaper contains the evidence for this specific causal claim, or it does not. If it does not, the claim is confabulated and must be removed or replaced with what the workpaper actually supports. Validation Step 3 (Link to Evidence) applies this test systematically.

VALIDATION STEP

Step 3 — Link to Evidence: Trace every Condition, Cause, and Effect statement to a specific workpaper reference before the finding is drafted.

03 Precision Fraud

AI produces specific figures with no calculation behind them.

WHAT IT IS

Precision Fraud occurs when an AI tool generates specific quantitative claims — percentages, frequencies, financial figures, rates, and comparative statistics — that appear to have been calculated from audit evidence but were statistically predicted. The term "fraud" is used deliberately: the specificity creates credibility that does not exist. A figure of "73.4%" reads as the product of calculation. It implies a sample, a denominator, a count, and a ratio. If it was generated by an AI tool, none of those elements exist. The number was predicted as a statistically plausible figure for the type of finding being generated. It is not an audit-grade figure.

WHAT IT LOOKS LIKE IN AN AUDIT WORKPAPER

A finding on vendor invoice processing states: "Approximately 34% of invoices reviewed contained at least one control exception, representing an estimated $2.3 million in unverified payment exposure." The auditor tested 40 invoices, found 14 exceptions, and noted that three involved payments over $50,000. The AI generated the percentage from the prompt context and the financial figure from a pattern-based estimate. The auditor never calculated $2.3 million. The figure entered the draft report with the same authority as a tested calculation.

HOW TO CATCH IT

Every quantitative figure in an AI-assisted finding must be independently calculated by the auditor from the source data before it is used. The auditor calculates the figure. The AI incorporates the auditor's figure into the narrative. The sequence is never reversed. When reviewing AI-generated content, flag every percentage, dollar amount, frequency, and comparative statistic and verify it has a documented calculation in the workpaper. If the calculation does not exist in the workpaper, the figure does not go in the finding. Validation Step 3 applies.

VALIDATION STEP

Step 3 — Link to Evidence: Every quantitative figure must trace to a documented calculation in the workpaper. The auditor calculates; the AI narrates.

04 Certainty Inflation

AI presents statistical prediction as professional judgment.

WHAT IT IS

Certainty Inflation occurs when an AI tool expresses conclusions, characterizations, and overall assessments with a level of confidence that the tested evidence does not support — or generates overall opinions that the auditor never formed through professional judgment. This is the most pervasive of the six modes because it is embedded in the tone of AI output rather than in specific facts or figures. Words and phrases like “clearly,” “evidently,” “no indication of,” “fully effective,” “generally strong,” and “broadly compliant” appear in AI-generated content as if they reflect tested conclusions. They reflect the AI’s pattern-based prediction of how an assessment in that domain is typically characterized.

WHAT IT LOOKS LIKE IN AN AUDIT WORKPAPER

An auditor asks the AI to draft an overall conclusion for a risk assessment of a third-party vendor. The AI produces: “Overall, the vendor’s control environment is generally sound, with no significant systemic weaknesses identified. Controls are functioning as intended across the primary risk domains, and the vendor demonstrates a strong commitment to compliance.” The auditor tested four controls, found two exceptions, and had not yet formed an overall assessment of the control environment. The AI generated one. It sounds exactly like a senior auditor’s considered professional judgment. It is the model’s prediction of what such a conclusion “sounds like.”

HOW TO CATCH IT

Read every characterization, assessment, and conclusion in AI-generated content and ask explicitly: ‘Did I form this conclusion from my tested evidence? Or did the AI generate it?’ Scan specifically for certainty language — “clearly,” “evidently,” “no indication of,” “fully,” “generally,” “broadly,” “strongly” — and interrogate each instance. If the auditor did not form the conclusion from evidence, the characterization must be removed or rewritten in language that accurately reflects what the evidence actually supports. Validation Step 4 (Interrogate the Conclusion) is the systematic application of this check.

VALIDATION STEP

Step 4 — Interrogate the Conclusion: Scan every characterization and overall opinion. If the auditor did not form it from tested evidence, remove it.

05 Scope Blindness

AI ignores your context and defaults to generic patterns.

WHAT IT IS

Scope Blindness occurs when an AI tool ignores key contextual details about the specific organization, jurisdiction, sector, regulatory framework, and entity type, and produces output calibrated for a generic default context instead. The model was trained on text from many jurisdictions, industries, and organizational types. Without explicit contextual constraints in the prompt, it defaults to the most common pattern in its training data for the type of output being requested. For internal auditors in Latin America, this often means output calibrated for US-regulated entities with US regulatory frameworks, SOX-era control language, and US sectoral norms that do not apply.

WHAT IT LOOKS LIKE IN AN AUDIT WORKPAPER

A government entity in the Dominican Republic engages an AI tool to assist in drafting a finding on financial controls. The AI-generated finding references the Sarbanes-Oxley Act Section 404, cites PCAOB Auditing Standard No. 2201, and frames the control gap in terms of material weakness disclosure obligations. None of these frameworks apply to a government entity in the Dominican Republic. The AI defaulted to the most common regulatory context in its training data for financial control findings. The output is not just incorrect — it is misleading to any reader who does not immediately recognize the jurisdictional mismatch.

HOW TO CATCH IT

Scope Blindness is the only failure mode that is primarily prevented at the prompt level rather than at the validation level. Every AI prompt for audit work must include an explicit, specific Reality Brief: the exact jurisdiction, the entity type, the applicable regulatory framework, the audit standards in use, and the relevant sector. The more specific the Reality Brief, the lower the risk of Scope Blindness. At the validation level, review every regulatory reference, every standard citation, and every sector-specific assumption in AI output and ask: does this actually apply to this specific organization in this specific jurisdiction? Validation Step 2 (Authenticate the Source) catches residual Scope Blindness after the prompt.

VALIDATION STEP

Step 2 — Authenticate the Source: Confirm every cited standard and regulatory reference applies to this specific entity type and jurisdiction, not just that it exists.

06 Standard Drift

AI applies a real standard to the wrong situation.

WHAT IT IS

Standard Drift occurs when an AI tool cites a real, existing standard but applies it incorrectly — to the wrong entity type, the wrong jurisdiction, the wrong version, or the wrong professional context. Unlike Fabrication, where the standard does not exist, Standard Drift involves a citation that is verifiable. The standard exists. The problem is in its application. The most consequential form of Standard Drift in internal audit practice is the citation of IIA professional standards as criteria in audit findings. IIA Global Standards govern how internal auditors conduct their work. They do not govern the auditee. They are not criteria for findings against the organization being audited. An AI tool trained on text that references both IIA standards and audit findings will predict that IIA standards are appropriate finding criteria — because they appear in similar contexts in the training data.

WHAT IT LOOKS LIKE IN AN AUDIT WORKPAPER

An internal audit finding on management's vendor oversight reads: "Criteria: IIA Global Standard 10.2 requires that management maintain effective oversight of third-party service providers." IIA Standard 10.2 addresses the internal audit function's Human Resources Management obligations — how the CAE manages audit staff and competency. It says nothing about management's vendor oversight obligations. Communication of results is addressed under Standard 11.3. The auditee has not violated an IIA standard — the auditee is not subject to IIA standards. The actual criteria for the finding should be the organization's vendor management policy, the applicable regulatory guidance, or the contractual terms governing the vendor relationship. The AI cited a real standard. It applied it to the wrong party in the wrong context.

HOW TO CATCH IT

For every criteria element in an AI-generated finding, ask the test question: 'Is this standard governing the auditee — or is it governing the auditor?' IIA professional standards govern auditors. The auditee's criteria must be the specific policy, regulation, contractual provision, or internal requirement the auditee violated. Additionally, confirm that the cited standard version is current and applies to the specific entity type and jurisdiction. Validation Step 2 (Authenticate the Source) systematically applies both tests to every criteria citation in AI-generated content.

VALIDATION STEP

Step 2 — Authenticate the Source: Ask 'Does this standard govern the auditee or the auditor?' and 'Is this the correct version for this entity and jurisdiction?'

3. Output Validation: A Systematic Response to All Six Modes

The Audit AI Failure Taxonomy identifies the problems. Output Validation is the systematic response — a five-step sequential validation method that addresses every failure mode before AI-assisted content enters a workpaper.

The five steps are applied in order, to every piece of AI-assisted content, before it is used. Each step has a specific target failure mode and a specific action:

STEP	NAME	WHAT YOU DO	FAILURE MODES CAUGHT
1	**Verify the Facts**	Confirm every standard citation, regulatory reference, and numerical claim against the original source document.	Fabrication
2	**Authenticate the Source**	Confirm every cited authority applies to this entity type, this jurisdiction, and this version.	Scope Blindness Standard Drift
3	**Link to Evidence**	Trace every factual statement in Condition, Cause, and Effect to a specific workpaper reference or documented calculation.	Confabulation Precision Fraud
4	**Interrogate the Conclusion**	Identify every characterization and overall opinion. Confirm each reflects your professional judgment from tested evidence.	Certainty Inflation
5	**Document the Method**	Complete the workpaper AI notation immediately adjacent to every AI-assisted item. All 8 fields required.	All six modes

Output Validation is not an aspirational standard. It is a documented, sequential checklist that produces a workpaper notation record — Step 5 — providing the CAE with evidence that validation occurred on every AI-assisted item. Without Step 5, there is no audit trail. Without an audit trail, there is no governance evidence. Without governance evidence, the function cannot demonstrate AI compliance to an external quality assessor.

4. The Professional Risk of Undetected Failure Modes

Each of the six failure modes carries specific professional risk when it reaches a completed audit deliverable. The following table maps each mode to its consequences if undetected:

Failure Mode	If It Reaches a Finding or Report	If It Reaches the Audit Committee
Fabrication	Finding criteria cannot be verified. Management can challenge and defeat the finding by showing the cited standard does not exist or does not apply. The finding is professionally indefensible.	Committee relies on non-existent regulatory authority. Any external reviewer — EQA, regulator, external auditor — who verifies the citation will identify the fabrication. CAE credibility is permanently damaged.
Confabulation	Root cause analysis is disconnected from tested evidence. Management can identify the causal chain as unjustified by the evidence and request removal. Rework is required before issuance.	Audit committee acts on a causal analysis that was never established. Management action plans may address the wrong root cause. Follow-up audits will find the original condition unresolved.
Precision Fraud	Finding quantification is not reproducible. If management requests the calculation basis, it does not exist. The figure must be removed or replaced with a calculated figure, requiring additional testing.	Committee makes resource allocation decisions based on financial exposure figures that were never calculated. The materiality assessment of the finding is unsupported.
Certainty Inflation	Overall conclusions overstate the evidence. Management may accept or reject findings based on a strength assessment the auditor never actually made. The finding misrepresents auditor professional judgment.	Committee believes the function assessed an overall condition of control effectiveness that the auditor never in fact assessed. Board reliance is based on a conclusion that does not exist in the workpaper.

Scope Blindness	Finding criteria and risk framing do not match the organization's actual regulatory environment. The finding may be technically correct for a different entity type but irrelevant or misleading for the actual auditee.	Committee receives assurance framed against standards and regulations that do not govern the organization. The assurance is not grounded in the actual compliance environment.
Standard Drift	Criteria cite standards governing auditors rather than auditees. The finding has no legal or regulatory basis for the organization being audited. The auditee cannot have violated a standard they are not subject to.	Committee believes the organization is out of compliance with a professional standard that does not apply to the organization. Corrective actions may be implemented against a non-existent obligation.

5. Why This Is Not About Banning AI

The Audit AI Failure Taxonomy is not an argument against AI adoption in internal audit. It is the exact opposite. Internal audit functions that understand these six failure modes and implement systematic validation protocols are in a significantly stronger position than functions that either avoid AI entirely or adopt it without validation infrastructure.

The functions that cannot use AI are those that adopt it without governance. A finding challenged by management because the cited standard was fabricated. A report rejected by the audit committee because the root cause analysis was not grounded in tested evidence. A quality assessment that identifies AI-generated conclusions presented as professional judgment. These are the consequences of adoption without governance, and they discredit the function in ways that take years to repair.

The goal of the Taxonomy is to give every internal auditor using AI a named, teachable vocabulary for what to look for. When an auditor knows the name “Fabrication,” they know to open the cited source and verify the specific section. When they know “Certainty Inflation,” they know to scan for characterization language and interrogate whether they formed those conclusions from their evidence. Naming the failure mode is the prerequisite to catching it.

THE GOVERNANCE IMPERATIVE

The IIA’s guidance on AI in internal auditing — including the AI Auditing Framework and the emerging Topical Requirements — consistently recommends that internal audit functions establish documented governance for AI use. This includes documented policies, output validation protocols, data governance standards, and demonstrated staff competency. The five-step method described here operationalizes the output validation component of that governance — the component most directly connected to the professional quality of the work product.

6. Monday-Morning Actions

The Audit AI Failure Taxonomy and the five-step Output Validation are actionable today. The following actions are sequenced for immediate implementation, organized by where your function currently stands.

If your function is already using AI informally

1. Pull any AI-assisted workpaper or finding produced in the last 90 days.
2. Apply the five validation steps to that content. Document what you find.
3. Identify which of the six failure modes appear in the existing content.
4. Brief your CAE on what you found. Frame it as a validation gap, not a compliance failure.
5. Implement Step 5 — workpaper AI notation — on all AI-assisted content going forward.

If your function is exploring AI adoption

6. Adopt the five-step Output Validation as the validation standard before the first AI-assisted content enters a workpaper.
7. Train every staff member on the six failure modes before they use any AI tool on engagement work.
8. Build the validation checklist into your workpaper template as a required completion step for all AI-assisted items.
9. Establish the Step 5 notation as a mandatory field in all AI-assisted workpapers from day one.

If your function has not yet adopted AI

10. Document that decision and the governance framework you will implement when adoption begins.
11. Train staff on the Audit AI Failure Taxonomy now — so they can identify AI failure modes in output produced by management and other third parties, even before the function adopts AI itself.
12. Design your AI governance infrastructure before adoption, not after. The five-step Output Validation and workpaper notation standard should exist before the first AI-assisted deliverable is produced.

Conclusion

AI output does not fail randomly. It fails predictably — in six specific modes that are direct consequences of how large language models generate text. Internal audit functions that can name these six modes, recognize them in live AI output, and apply a systematic five-step validation protocol before AI-assisted content enters a workpaper are in a fundamentally different professional position than those that cannot.

The Audit AI Failure Taxonomy is the vocabulary. Output Validation is the method. Together, they give internal audit functions the tools to use AI at a professional standard — not to avoid AI, but to use it in a way the audit committee can trust and an external quality assessor can verify.

The question is not whether your team will use AI. It is whether they will know what to watch for before it causes professional harm.

About Efficient Advice, LLC

Efficient Advice, LLC is a specialized consulting firm that works with internal audit departments on AI governance. All engagements are bilingual, available in English and Spanish.

Areas of Practice — Efficient Advice, LLC	
Audit AI Failure Taxonomy	Six common failure modes: Fabrication · Confabulation · Precision Fraud · Certainty Inflation · Scope Blindness · Standard Drift
Output Validation	**Five steps:** Verify Facts · Authenticate Source · Link to Evidence · Interrogate Conclusion · Document Method
Structured Prompting	**Five parts:** Task Authority · Reality Brief · Actionable Output · Compliance Wall · Evidence Format
AI Audit Maturity	Six levels: Absent → Informal → Governed → Capable → Optimized → Leading
Governance Evidence	Five categories: Policy · Validation · Competency · Governance · Results Evidence
Governance Lifecycle	Diagnose (P1) · Design (P2+P3) · Defend (P4+P5)

Five Products. One Complete Governance Lifecycle.

Efficient Advice offers five products covering the complete AI governance lifecycle for internal audit functions:

- **P1 — AI Audit Readiness Assessment:** Scored diagnostic baseline. Six dimensions. 25-action 90-Day Roadmap. 30 business days. 42 statement diagnostic tool.
- **P2 — AI-Powered Audit Policy Package:** Eight governance documents. 16 bilingual files. Ready to sign and distribute. 30 business days.
- **P3 — AI Audit Staff Training Intensive:** Five IIA-aligned modules. Structured prompting and output validation. Prompt library. Competency certificates.
- **P4 — Continuous AI Audit Intelligence:** Monthly advisory. Prompt library maintenance. Standards monitoring. Month-to-month.
- **P5 — AI Audit Transformation Service:** Full-service engagement. Tools, training, policy, prompts, results dashboard. 4–12 weeks.

SCHEDULE A FREE 30-MINUTE CONVERSATION
Info@efficientadvice.com | www.efficientAdvice.com

© Efficient Advice, LLC. All rights reserved. This white paper may be shared and distributed in its entirety with attribution. Reproduction of proprietary content without attribution is prohibited.

Version Notes

Version 2.0 — October 2026. Terminology and editorial updates. The six failure modes, the five validation steps, their order, and the method are unchanged.

Version 1.0 — June 2026.

EDITORIAL CHANGES IN VERSION 2.0

- Added version number and publication date.
- Updated terminology for the five validation steps and the Areas of Practice table.