

# AI Tools for Tax Preparers: Practical Applications and Compliance Considerations

Course Material Export – Generated for IRS Continuing Education Review

## 1.1 What Is Artificial Intelligence?

### Lesson at a Glance

To understand AI (Artificial Intelligence), it helps to first look at how traditional computers work. For decades, computers were like obedient cooks following a strict recipe. If you didn't give them exact, step-by-step instructions (code), they couldn't do anything.

**AI is different.** Instead of following a rigid recipe, AI is a computer system that has been shown millions of examples so it can **learn how to spot patterns and make decisions on its own**. It mimics human intelligence—like reasoning, solving problems, and understanding language—without a human having to program every single rule.

The Best Way to Picture It: The Fruit Sorting Machine

[Image: Fruit-Sorting-Machine.webp]

[Image: gears.svg]

### Traditional Computing:

You would have to write complex rules:

*"If the object is exactly this shade of orange, and its roundness matches this mathematical circle, it is an orange."* If a weirdly shaped, reddish orange comes along, the system breaks.

[Image: robot.svg]

### Artificial Intelligence

You simply dump 10,000 pictures of apples and 10,000 pictures of oranges into the system and say, *"Learn the difference."*

The AI analyzes the textures, shapes, and colors on its own. Soon, it can accurately sort a fruit it has never seen before.

## You likely already use AI every single day without realizing it:

### The Three Everyday Types of AI

This AI looks at data and recognizes what it is.

[Image: Recognition-AI.webp]

*Examples:* FaceID unlocking your phone, your email filtering spam, or banking apps reading the handwriting on a check you deposit via photo.

This AI studies your past behavior to guess what you will want next.

[Image: Recommendation-AI.webp]

*Examples:* Netflix suggesting a movie, Spotify building a playlist, or Amazon showing you products you might buy.

This is the newest type of AI (like Gemini or ChatGPT). It doesn't just sort or predict—it *creates* new things based on what it has learned.

[Image: Generative-AI.webp]

*Examples:* Writing a tax checklist, drafting an email, or generating a picture based on text you type in.

### How AI "Thinks" (The Human Comparison)

To explain this to beginners, think of AI like a highly trained assistant:

[Image: brain-circuit.svg]

### The Brains (Machine Learning):

You would have to write complex rules:

This is the core technology. The computer is fed massive amounts of data (books, articles, code, tax laws) and maps connections between words and concepts, similar to how human brain cells connect.

[Image: comment-dots-1.svg]

## The Output:

When you ask it a question,

it doesn't look up a pre-written answer like Google does. Instead, it predicts, word by word, the most logical and helpful response based on all the patterns it has studied.

**In short: AI isn't a "sentient robot." It is incredibly powerful software that learns from examples instead of instructions, allowing it to handle complex, messy human tasks—like summarizing notes, organizing data, or answering questions—in seconds.**

## Artificial Intelligence — A Plain-Language Definition

Artificial intelligence (AI) refers to computer systems designed to perform tasks that typically require human-like intelligence — understanding and generating language, recognizing patterns in data, making decisions based on context, and producing creative or analytical content.

The AI tools that matter most to tax practitioners today fall into a specific category called **generative AI** — systems that produce new content in response to user input. When you type a question into ChatGPT, Claude, Gemini, or Microsoft Copilot and receive a written answer, you are using generative AI.

## Large Language Models — How They Actually Work

[Image: llm-1024x683.webp]

Most generative AI tools are built on technology called a **Large Language Model (LLM)**. Understanding how an LLM works — even at a basic level — is essential for every practitioner who uses these tools professionally.

### Training: Learning from Text

An LLM is trained on an enormous volume of text — books, websites, academic papers, government publications, legal documents, news articles, and more — totaling trillions of words. During training, the model learns statistical patterns: which words tend to follow which other words, how sentences are structured, how concepts relate to each other across different contexts.

The model does not "read" or "understand" text the way a human does. It learns to recognize and reproduce patterns. The result is a system that can produce fluent, coherent, contextually appropriate language — but one whose "knowledge" is entirely derived from statistical patterns in the text it was trained on.

### Inference: Generating a Response

When you type a question into an AI tool, the model generates a response word by word — or more precisely, token by token. At each step, the model calculates the probability distribution over all possible next tokens and selects one based on that distribution. This process continues until the response is complete.

The result is text that statistically resembles what a correct, well-written response to your question would look like — based on everything the model learned during training. It does not look up your question in a database. It does not retrieve a fact from a verified source. It generates text that fits the pattern of a correct answer.

**This is the most important concept in this course:** An AI language model does not retrieve facts — it generates text that statistically resembles correct responses. This means AI can produce a response that looks, reads, and sounds exactly like an authoritative tax answer — complete with IRC citations, dollar figures, and professional language — while being entirely fabricated. There is no indicator in the response itself that anything is wrong.

### **Training Cutoff Date — What It Means for Tax Law**

Every AI model has a **training cutoff date** — the point in time at which its training data ends. The model has no knowledge of anything that happened after that date: no new legislation, no new IRS guidance, no new revenue procedures, no new court decisions.

For tax practitioners, this is a significant limitation. Tax law changes constantly:

- Annual inflation adjustments (Rev. Proc. issued each fall)
- New legislation (TCJA, SECURE 2.0, OBBBA)
- IRS notices, revenue rulings, and private letter rulings
- Court decisions that change how a provision is interpreted
- New or revised IRS forms and instructions

An AI tool operating from training data that is 12 or 18 months old may be unaware of significant changes directly relevant to a client's current-year return. Critically — the model will not know what it does not know. It will not say "I am not sure if this changed." It will answer based on the law as it existed in its training data, in the same confident tone it uses for everything else.

**Always ask yourself:** Could this answer have changed since this model's training data ended? For anything involving specific dollar amounts, thresholds, rates, or provisions that change annually — the answer is almost certainly yes. Verify against the current-year IRS guidance before relying on any AI response involving specific figures.

### **Why AI Responses Feel Authoritative**

One of the most counterintuitive aspects of LLM outputs is that incorrect responses often feel more authoritative than correct ones. This happens because the model has learned to produce text in the style of authoritative sources — formal language, precise-sounding citations, organized structure, confident tone. These stylistic features are entirely separate from factual accuracy.

A fabricated Rev. Proc. citation in a well-formatted bullet point list looks exactly like a real citation in a well-formatted bullet point list. A wrong dollar figure stated with professional precision reads exactly like a correct dollar figure stated with professional precision. The text itself provides no signal that the information is wrong.

### **Illustrative Example — Why This Matters**

A practitioner asks an AI tool: "What is the standard mileage rate for business use of a vehicle for the current tax year?" The AI responds with a specific rate, cites "Rev. Proc. 2024-XX," and formats the answer in professional language. The practitioner uses this rate on a client's Schedule C without verifying it. The actual current-year rate is different — the AI's training data predates the IRS announcement of the updated rate. The return contains an error. The preparer signed it. The \$6694 penalty analysis does not ask whether the AI was to blame. It asks whether the preparer's position had substantial authority. It did not.

[Image: Learning\_objective.svg]

### **Practitioner Tip — The Power of Starting Early**

**Know the training cutoff of every AI tool you use.** Most providers publish this information. For anything involving current-year tax figures, treat the AI's answer as a starting point for research — never as a final answer.

**The style of an AI response tells you nothing about its accuracy.** A well-formatted, confidently stated, professionally worded response may be completely fabricated. Evaluate the content, not the presentation.

**Ask the AI to identify its sources.** When using AI for tax research, ask it to identify the specific IRC section, Rev. Proc., or IRS publication it is relying on — then look up that source directly and verify it says what the AI claims.

## Common Errors

### Common Errors

#### AI Research and Verification Mistakes

##### Assuming AI Confidence = Accuracy

A confident AI response is not proof that the answer is correct or authoritative.

##### Using Unverified Dollar Figures

Always verify AI-provided rates, thresholds, limits, and phaseouts against current-year IRS sources.

##### Treating AI Like Tax Research

A well-formatted AI answer is not equivalent to authoritative tax research or primary authority analysis.

##### Ignoring Training Cutoff Dates

AI tools may not contain current-year tax law changes unless verified against updated IRS guidance.

## 1.2 Types of AI Tools Available to Tax Practitioners

### Learning at a Glance

Not all AI tools are the same. Tax practitioners encounter AI in five categories: general-purpose LLMs, tax-specific research AI, document analysis tools, AI built into tax software, and voice and meeting AI.

Matching the right tool to the right task is one of the core skills this course teaches.

#### General-Purpose LLMs

ChatGPT, Claude, Gemini, and similar tools used for drafting, brainstorming, explanations, and workflow assistance.

#### Tax Research AI

AI systems designed to search tax authorities, regulations, court cases, and IRS guidance.

## Document Analysis

Tools that review tax returns, source documents, financial statements, and PDFs.

## Tax Software AI

AI features built directly into tax preparation software, workflow systems, and practice management platforms.

## Voice & Meeting AI

AI meeting assistants, transcription tools, voice summarization systems, and note-taking platforms.

## The Five Categories at a Glance

Five categories of AI tools for tax practitioners

[Image: brain-circuit.svg]

## General-purpose LLMs

General-purpose AI tools like ChatGPT and Gemini used for drafting, summarizing, brainstorming, and explanations.

[Image: book-open-reader.svg]

## Tax Research AI

AI systems connected to tax research databases and authoritative materials designed for professional tax analysis.

[Image: page.svg]

## Document Analysis

AI tools that read PDFs, financial statements, engagement letters, and tax documents to organize and extract information.

[Image: calculator.svg]

## AI in Tax Software

AI features built directly into tax preparation software for diagnostics, workflow assistance, and automation.

[Image: microphone-lines.svg]

## Voice & Meeting AI

AI meeting assistants that record, transcribe, summarize, and organize conversations and client meetings.

**The core skill: Match the right tool to the right task — and know when no AI tool is appropriate. Category, tier, verification requirement, and data privacy implications all differ. No single tool is right for every task.**

*"AI is a tax office assistant — not a tax preparer."*

## General-Purpose LLMs — The Most Common and Most Misused

General-Purpose LLMs

General-purpose AI systems are best for tasks where a tax-specific error would likely be caught before it causes harm.

- Drafting client letters
- Creating checklists and procedures
- Explaining general tax concepts

Tax Research AI

Tax-specific research AI is generally more reliable for research questions because it is designed around tax authorities, regulations, and professional source material.

- Researching tax authorities
- Locating IRS guidance
- Analyzing technical tax questions

Practitioner Warning

Even the most advanced tax research AI is not a substitute for professional judgment.

AI-generated tax conclusions must always be verified against primary authority, IRS guidance, and the facts of the client's situation before being relied upon.

## 1.3 The Hallucination Problem

### What Is a Hallucination?

The term for AI-generated content that is confidently stated but factually wrong — often internally consistent, professionally presented, and indistinguishable from correct responses without independent verification.

[Image: Hallucination.webp]

#### What's happening in the image

- **Question:**
- “Does this taxpayer qualify for the Child Tax Credit?”
- **Hallucination:**
- “Yes, they qualify for the full credit.”
- Reality: Income is too high → not eligible

### Why Tax Information Is Especially Vulnerable

#### Why Tax Information Is Especially Vulnerable

Tax law contains thousands of highly technical rules, annual updates, and complex citation structures. These characteristics make tax information especially susceptible to AI errors and hallucinations.

#### Annual Inflation Adjustments

Tax law contains hundreds of numerical thresholds that change every year. AI systems may blend figures from different tax years and present them as current law.

#### Citation Hallucinations

Revenue Procedures, Notices, Revenue Rulings, and court cases follow predictable naming conventions. AI can generate citations that look authentic even when they do not exist.

#### Frequent Law Changes

AI may confidently describe a provision exactly as it existed under prior law while failing to recognize that Congress, Treasury, or the IRS has since changed the rule.

#### Technical Complexity

Tax rules often depend on filing status, income level, entity type, elections, phase-outs, and special exceptions. Small factual differences can produce completely different tax outcomes.

### Practitioner Takeaway

AI can be an outstanding productivity tool, but tax professionals must verify inflation-adjusted figures, citations, legal authority, and technical requirements before relying on any AI-generated answer.

## Common Hallucination Patterns in Tax AI Responses

Hallucination Type	Example	Why It Is Dangerous
Wrong dollar amount	AI states the HSA contribution limit for family coverage as \$8,300 when the correct figure is \$8,750	Client underfunds HSA; practitioner misses deduction optimization
Fabricated citation	AI cites "Rev. Proc. 2024-88" which does not exist	Practitioner uses nonexistent authority to support a return position
Wrong IRC section	AI attributes the qualified business income deduction to §199 instead of §199A	Research starts in the wrong place; prior law may not apply
Outdated rule stated as current	AI describes the bonus depreciation rate as 80% when it has since changed	Wrong depreciation calculation on the return
Wrong filing status rule	AI applies the MFJ phase-out range for the IRA deduction to a single filer	Incorrect deductibility determination; client either over- or under-deducts
Plausible but wrong exception	AI states that gambling losses are deductible on Schedule C when they are not	Improper deduction claimed on the return

## Strategies for Catching Hallucinations

### Verify Every Specific Claim

Check dollar amounts, IRC sections, form numbers, thresholds, deadlines, and Rev. Proc. citations against the actual source before relying on the answer.

### Search the Cited Source

If AI cites something like **Rev. Proc. 2025-32**, go directly to IRS.gov or your research platform and confirm both that the source exists and that it says what the AI claims.

### Cross-Check with a Second Source

For important questions, verify the answer against a professional tax research platform, official IRS guidance, or another reliable authority before giving advice.

Ask the AI to Show Its Reasoning

Asking for the reasoning or assumptions often reveals uncertainty, missing facts, outdated law, or unsupported conclusions that were hidden in the first answer.

**Practitioner Rule:** Treat every AI-generated tax answer as a draft until the numbers, citations, and legal conclusions have been verified.

## 1.4 Why AI Cannot Be a Tax Preparer of Record

### The PTIN Requirement — AI Cannot Hold One

[Image: PTIN Requirement]

Professional Responsibility

AI Cannot Sign the Return

Some practitioners ask whether AI can prepare the return while the preparer simply reviews it at the end. The answer is no. Multiple bodies of tax law and professional regulation place responsibility on the human preparer — not the software tool.

IRC §6109 — PTIN Requirement

PTIN requirements apply to human tax return preparers. AI systems cannot hold a PTIN or legally act as the signing preparer.

Circular 230 Duties

Competence, diligence, ethical conduct, and professional judgment are practitioner responsibilities that cannot be delegated to software.

§6694 Preparer Penalties

Preparer penalties attach to the human who signs the return — regardless of whether AI generated part of the analysis or content.

**Critical Professional Principle:** AI can assist in the workflow, organize information, draft explanations, and improve efficiency — but it cannot bear legal responsibility, exercise professional judgment, hold a PTIN, or sign a tax return.

The practitioner who signs the return is accountable for every line. Period.

### Circular 230 — Professional Obligations That Cannot Be Delegated

Circular 230 (31 CFR Part 10) governs the conduct of persons who practice before the IRS. Three provisions are directly relevant to AI use:

- **§10.35 — Duty of Competence**

A practitioner must possess the competence necessary for the representation or return preparation they undertake. Using an AI tool competently — understanding what it can and cannot do, knowing its limitations, and verifying its output — is itself part of the competence obligation. A practitioner who uses AI without understanding how it works, what its failure modes are, or how to verify its output is not exercising competent use of the tool.

- **§10.22 — Duty of Diligence**

Practitioners must exercise reasonable diligence in all matters. Using AI output without verification, or allowing staff to use AI without appropriate oversight, does not satisfy the diligence requirement. Diligence in an AI-assisted practice means having a documented workflow that ensures professional review of every AI-assisted output before it is used in a client matter.

- **§10.34 — Standards for Written Advice and Returns**

A practitioner may not sign a return containing a position they know or should have known is incorrect. AI-generated content that contains an error — whether a wrong dollar amount, a misapplied rule, or a fabricated citation — that the practitioner did not catch because they did not verify it creates exactly the §10.34 problem. The practitioner "should have known" the content was wrong, because a reasonable practitioner would have verified it.

## Section 6694 — The Penalty That Cannot Be Blamed on AI

IRC §6694 imposes penalties on the "tax return preparer" — the human individual who prepares the return — when the return contains an understatement of tax liability due to:

- **§6694(a) — An unreasonable position:** A position that does not meet the substantial authority standard and was not adequately disclosed. Penalty: greater of \$1,000 or 15% of income derived from the return.
- **§6694(b) — Willful or reckless conduct:** Willful understatement or reckless disregard of rules or regulations. Penalty: greater of \$5,000 or 75% of income derived from the return.

These penalties apply to the practitioner who signs the return. "The AI told me" is not a defense. "I relied on AI research without verifying it" describes reckless conduct — not a safe harbor. The §6694 analysis asks whether the position on the return had substantial authority and whether the practitioner exercised reasonable professional judgment. AI-generated content that was not verified does not establish either.

### Worked Example

#### §6694 and AI-Generated Error

A practitioner uses AI to research whether a client's home office qualifies for the deduction. The AI states the client qualifies and provides a calculation. The practitioner does not

independently verify the exclusive use requirement or the square footage calculation. The return is filed.

The IRS audits and disallows the deduction because the space does not meet the exclusive use test.

**Penalty Analysis:**

The §6694 analysis does not ask whether the AI made an error. It asks whether the practitioner took an unreasonable position — one that did not meet the substantial authority standard.

**Key Result:**

A position based on unverified AI research, without the practitioner's independent application of the exclusive use test to the client's specific facts, does not meet that standard. The §6694(a) penalty applies to the practitioner.

## The Non-Delegable Responsibility Principle

Every professional obligation a practitioner bears — competence, due diligence, accuracy, client confidentiality, professional judgment — applies in full when AI tools are used. Using AI does not transfer, reduce, or eliminate any of these obligations. It changes the workflow. It does not change who is responsible.

Professional Obligation	Does AI Use Change It?
PTIN and signature requirement	No — the human practitioner signs every return
§6694 preparer penalty liability	No — applies to the practitioner who signed, regardless of how content was generated
Circular 230 duty of competence	No — now extends to competent use of AI tools
Circular 230 duty of diligence	No — now requires verification workflow for AI-assisted outputs
Due diligence (Form 8867)	No — practitioner must conduct the actual inquiry; AI may assist with preparation
Client data confidentiality	No — fully applies to data entered into AI tools; see Module 3

## 1.5 The Guiding Principle: AI as Office Assistant

### Lesson at a Glance

The Core Principle

The Framework Behind Every AI Decision

Every concept in this course — hallucinations, PTIN requirements, preparer penalties, verification duties, and taxpayer data protection — flows from a single organizing principle.

AI may assist the workflow.

The practitioner remains responsible.

Hallucinations

AI can generate convincing but incorrect information. Verification remains the practitioner's responsibility.

PTIN Requirement

AI cannot hold a PTIN, sign a return, or legally function as the preparer of record.

§6694 Penalties

Preparer penalties attach to the human practitioner — regardless of what tool generated the content.

Data Privacy

Taxpayer data protection obligations continue to apply even when AI systems are involved.

A Way of Thinking About AI

This principle is not a single IRS regulation. It is a professional framework that keeps practitioners aligned with all the rules simultaneously.

Let AI assist.

Never let AI replace professional judgment.

[Image: ai-is-not-tp.webp]

## What the Principle Means

This course's guiding principle is not a metaphor or a marketing slogan. It is a precise description of what AI is legally permitted to do in a tax practice — and what it is not.

Think carefully about what a capable, hard-working office assistant does in a tax practice. They draft first versions of client letters. They organize incoming documents. They look up phone numbers, prepare file folders, format spreadsheets, and handle administrative tasks that free the practitioner to focus on professional judgment work. A great office assistant makes the office more efficient and the practitioner more effective at serving clients.

But a great office assistant does not sign returns. They do not provide legal or tax advice. They do not make professional judgments about a client's tax position. They do not bear responsibility for the accuracy of a return. And when they hand something to the practitioner, the practitioner reviews it — because the practitioner, not the assistant, is responsible for what goes out the door.

AI occupies exactly this role in a tax practice — when used correctly.

## What AI Can Do — The Office Assistant Role

What AI CAN Do (The Office Assistant Role)	What AI CANNOT Do (Outside the Assistant Role)
Draft first versions of client letters, tax explanations, and administrative correspondence.	Sign a tax return or hold a Preparer Tax Identification Number (PTIN).
Organize and categorize messy, unstructured incoming documents or raw notes.	Make a final determination on whether a grey-area deduction or credit legally qualifies.
Generate checklists, intake forms, and repeatable practice management templates.	Conduct due diligence—AI can help prepare frameworks, but cannot substitute for your professional inquiry.
Summarize new IRS guidance into plain, digestible language (requires manual human verification).	Provide tax advice directly to a client or act as the final consulting authority.
Identify starting points for tax law research by pointing to relevant sections to investigate further.	Represent a client before the IRS or during an audit.
Flag potential inconsistencies or missing items in a data set for the practitioner to examine.	Bear legal responsibility for an error, penalty, or omission on a signed return.
Generate tailored interview questions to use during client onboarding and intake sessions.	Be cited as the underlying authority for a tax return position (you must cite the IRC, regulations, or court cases).

## Applying the Principle — A Decision Framework

When facing any AI-related decision in your practice, apply this three-question test:

Question	If Yes	If No
Would a capable office assistant be able to do this task?	AI may assist — with practitioner review before anything goes to the client or the return	This task requires professional judgment — AI may provide information to inform the decision, but the practitioner

Question	If Yes	If No
If the AI makes an error on this task, will the practitioner catch it before it affects the return?	Proceed with appropriate verification workflow	decides Increase verification rigor or do not use AI for this task
Does this task involve client-specific identifying information that should not enter a public AI tool?	Use only an enterprise tool with a Data Processing Agreement, or anonymize the data first	Public AI tool is acceptable for this general task

### Worked Example

#### Applying the Principle

A useful test for almost every AI decision is simple:

Would I delegate this to a trusted office assistant?

If the answer is yes, AI may be appropriate with review and verification. If the answer is no, the practitioner must exercise independent professional judgment.

#### Scenario A

A practitioner wants AI to draft a letter explaining why a client's refund is lower this year than last year.

#### **Would an office assistant do this?**

Yes.

#### **Can the practitioner review it before sending?**

Yes.

#### **Data concern?**

Remove SSNs and sensitive information or use an approved enterprise AI platform.

#### Appropriate AI Use

#### Scenario B

A practitioner wants AI to determine whether a client's mixed-use rental property qualifies for the QBI deduction.

#### **Would an office assistant make this determination?**

No.

Applying §199A to specific facts requires professional judgment and legal analysis.

AI may help identify authorities to research, but it cannot make the final determination.

Research Tool Only

Scenario C

A practitioner wants AI to complete Form 8867 using analysis of a client's intake questionnaire.

**Would an office assistant complete Form 8867?**

No.

Due diligence requires practitioner inquiry, evaluation, and professional judgment.

AI may suggest interview questions, but the practitioner must conduct the interview and complete the form.

Practitioner Must Perform

Practical Rule

AI is most useful for drafting, organizing, summarizing, and researching.

The moment a task requires applying tax law to a taxpayer's specific facts, exercising due diligence, signing a return, or rendering professional judgment, the responsibility belongs to the practitioner.

Let AI assist the process.

Never delegate the judgment.

## **Why This Principle Protects Everyone**

The "office assistant" framework is not designed to limit AI use — it is designed to make AI use sustainable and professional. Practitioners who follow it:

- Get the productivity benefits of AI on appropriate tasks
- Maintain the professional judgment and verification that protects clients
- Stay within their §6694, Circular 230, and due diligence obligations
- Can confidently respond to any client, IRS examiner, or OPR inquiry about their workflow

Practitioners who do not follow it — who allow AI to make professional judgments, who rely on AI output without verification, who let AI determine due diligence eligibility — face the full weight of preparer penalties with no defensible basis for reasonable cause.

## Practitioner Tips

### Practitioner Tips

#### Keep the Guiding Principle Visible

#### Keep the Principle in Front of You

Post the guiding principle where you can see it — even if that means putting it on a sticky note next to your monitor.

“Is this an office assistant task — or am I asking AI to be the preparer?”

#### Train Your Staff on the Framework

Every employee who uses AI tools in the office should understand this framework. AI misuse by staff can create the same compliance and liability risks as misuse by the signing practitioner.

#### Due Diligence Cannot Be Delegated

A staff member who asks AI to make a due diligence determination has crossed the professional line — even if the practitioner signs the return later without carefully reviewing the work.

**Practice Management Reality:** The safest firms create office-wide AI procedures, staff training rules, verification requirements, and clear boundaries for tasks that always require direct practitioner judgment.

## Common Errors

### Worked Example

#### Applying the Principle

A useful test for almost every AI decision is simple:

Would I delegate this to a trusted office assistant?

If the answer is yes, AI may be appropriate with review and verification. If the answer is no, the practitioner must exercise independent professional judgment.

#### Scenario A

A practitioner wants AI to draft a letter explaining why a client's refund is lower this year than last year.

**Office assistant?** Yes.

**Practitioner review?** Yes.

**Data concern?** Remove SSNs and sensitive information or use an approved enterprise platform.

Appropriate AI Use

Scenario B

A practitioner wants AI to determine whether a client's mixed-use rental property qualifies for the QBI deduction.

**Office assistant?** No.

This requires applying §199A to specific taxpayer facts and exercising professional judgment.

AI may identify authorities to research, but the practitioner must make the determination.

Research Tool Only

Scenario C

A practitioner wants AI to complete Form 8867 using a client's intake questionnaire responses.

**Office assistant?** No.

Due diligence requires practitioner inquiry, evaluation, and professional judgment.

AI may suggest interview questions, but the practitioner must conduct the interview and complete Form 8867.

Practitioner Must Perform

Practical Rule

AI is most useful for drafting, organizing, summarizing, and researching.

The moment a task requires applying tax law to a taxpayer's specific facts, exercising due diligence, signing a return, or rendering professional judgment, the responsibility belongs to the practitioner.

Let AI assist.  
Never delegate the judgment.

## 2.1 The AI Task Spectrum: Low Risk to High Risk

[Image: delegate.webp]

AI Risk Spectrum

Not All AI Tasks Carry Equal Risk

The key variable in AI use is not whether AI was involved — it is the consequence of an undetected error. A mistake in a draft client letter may be corrected during review. A mistake on a signed return may not be discovered until an IRS examination, long after the return has been filed.

Low Risk

Low Risk Tasks

Drafting client letters, organizing notes, summarizing meetings, or preparing internal checklists.

Moderate Risk

Moderate Risk Tasks

Extracting figures from documents, drafting tax research summaries, or preparing preliminary analyses.

High Risk

High Risk Tasks

Research involving tax positions, eligibility determinations, deductions, credits, or filing requirements.

Never Delegate

Never-Delegate Tasks

Final due diligence determinations, signing returns, and exercising professional judgment required by law.

The Higher the Risk, the Higher the Verification

As the consequences of an AI error increase, the required level of practitioner review, verification, documentation, and direct involvement must also increase.

How much verification is required

How much practitioner oversight is necessary

Whether AI should be involved at all

## Authorized Practitioner Categories

Risk Level	Task Type	Examples	Required Practitioner Action
<b>Low Risk</b>	Administrative and drafting	Client appointment reminders; file checklists; template engagement letters; formatting spreadsheets; organizing to-do lists; drafting FAQs for the firm website	Review before sending; minimal tax-specific verification needed because content contains no specific tax claims
<b>Moderate Risk</b>	Communication and explanation	Drafting a client letter explaining why they owe this year; summarizing a new law for a client newsletter; explaining a general tax concept in plain language	Review every tax statement against a primary source; correct any inaccuracies before the letter is sent or published
<b>High Risk</b>	Research and position-related	Using AI to identify whether a deduction might qualify; using AI to summarize the rules for a specific credit; using AI to calculate a phase-out or limitation	Verify every conclusion against the IRC, IRS guidance, or a professional research platform; document the verification; practitioner makes the final determination
<b>Never Delegate</b>	Return preparation and professional judgment	Having AI "prepare" a return; using AI output as the final basis for a return position without independent review; relying on AI for due diligence determinations on EITC, CTC, AOTC, or HoH	These tasks require practitioner professional judgment — AI may assist in gathering information or drafting questions, but cannot make the determination or bear the professional responsibility

## Why the Spectrum Matters

The risk spectrum is not about whether AI is "allowed" — it is about applying the right level of verification to each task. A practitioner who applies the same verification level to every AI task — either verifying everything exhaustively (inefficient) or verifying nothing (dangerous) — is not using the spectrum effectively. The goal is calibrated verification: light

review for low-risk tasks, rigorous primary-source verification for high-risk tasks, and practitioner judgment for never-delegate tasks.

## The "never delegate" category is firm

[Image: never-delegate.webp]

**The "never delegate" category is firm. Completing Form 8867, making a final determination on due diligence eligibility, and signing a return are practitioner obligations that cannot be transferred to AI regardless of how sophisticated the tool is or how confident its output appears. If an AI tool claims to "prepare a complete, IRS-ready return," the practitioner who signs it remains fully liable for every line.**

### Practitioner Tips

Practitioner Tips

Classify the AI Task Before You Start

Classify Before You Start

Before using AI for any task, take 10 seconds to ask: **“Where does this fall on the spectrum?”** The classification determines the verification workflow before you see the AI output — not after.

Client-Facing Content Is At Least Moderate Risk

Any AI output sent to a client — a letter, email, newsletter item, or explanation — may contain tax statements that must be reviewed and verified before sending. An incorrect statement can affect a client’s planning decisions.

**Workflow Rule:** Risk classification comes first. Verification, oversight, and documentation should be based on the task’s risk level — not on how polished the AI answer looks.

## 2.2 The 3-Step AI Verification Workflow

### Lesson at a Glance

AI Verification Workflow

The 3-Step Verification Process

Any AI output used in a client matter — letters, research summaries, computations, or tax analysis — must be verified before use. This workflow is the operational expression of the practitioner’s §6694 and Circular 230 obligations.

### Step 1

#### Identify Claims

Identify every factual statement, legal conclusion, threshold, rate, limitation, and citation contained in the AI output.

### Step 2

#### Verify Sources

Confirm each claim against primary authority such as the IRC, Treasury Regulations, IRS instructions, revenue procedures, or court cases.

### Step 3

#### Correct or Remove

Correct inaccurate claims, remove unsupported statements, and document any unresolved or unverifiable items before client use.

#### Professional Responsibility Rule

AI assistance does not reduce practitioner responsibility. The practitioner signing the return or advising the client remains fully responsible under §6694 and Circular 230 for verifying the accuracy of the final work product.

## **Step 1 — Identify Every Specific Factual and Legal Claim**

Read the AI output carefully and mark every specific statement of fact or law. This includes: dollar amounts, percentage figures, IRC section references, IRS publication numbers, court case citations, filing deadlines, phase-out ranges, and descriptions of how a specific tax rule works. Do not assume any specific claim is correct — even claims in areas where you have expertise, because hallucinations frequently cluster in areas that are almost right.

The goal of Step 1 is to create an explicit list of claims that require verification — not to evaluate whether they are right or wrong yet. That evaluation happens in Step 2.

## **Step 2 — Verify Each Claim Against a Primary Source**

For each marked claim, locate the primary source — the actual IRC section, the actual Rev. Proc., the actual IRS publication, the actual court decision. Verify that: (a) the source exists, (b) it says what the AI claimed, and (c) it represents current law.

### **Acceptable Primary Sources**

- IRS.gov — Revenue procedures, notices, publications, form instructions, FAQs
- Cornell Law School Legal Information Institute ([law.cornell.edu/uscode](http://law.cornell.edu/uscode)) — IRC text
- Tax Court, Federal Circuit, and Supreme Court decisions — from official court websites or verified legal databases
- Professional tax research platforms (Thomson Reuters, Bloomberg, Wolters Kluwer) — with current content
- Current-year IRS Form instructions — for form-specific rules

### **Not Acceptable as Primary Sources**

- Another AI tool confirming the first AI's answer
- Tax blogs or practitioner websites — even reputable ones; use as research leads, not final verification
- The practitioner's memory that "the AI's answer sounds right"
- Professional tax research platforms (Thomson Reuters, Bloomberg, Wolters Kluwer) — with current content
- Prior-year materials for current-year figures that change annually

**Pasting primary source text into the AI dramatically reduces hallucination risk for that specific task.** When using AI to help understand a specific Rev. Proc. or IRS notice, paste the actual text of the document into the conversation before asking your question. This grounds the AI in the actual source and prevents it from generating text based on its (potentially outdated or inaccurate) training data about that document.

### **Step 3 — Correct, Remove, or Document Unverifiable Claims**

**For each claim reviewed in Step 2:**

- **Verified and correct:** Use the content. Note the source in your workpapers for return positions.
- **Verified but incorrect:** Correct the content using the verified information. Never forward content with a known error.
- **Cannot be verified from a primary source:** Either remove the claim from the content, or flag it as requiring additional research before any reliance. Never use unverified claims in client communications or return positions.

### **Documentation — Why It Matters**

For any return position where AI assisted in the research, the practitioner should note in the client workfile: what the AI suggested, what was verified, and what primary source confirmed the correct position. This documentation serves two purposes: it creates a discipline that prevents skipping verification, and it supports a reasonable cause and good faith defense under §6694 if the position is later challenged.

### **Practitioner Tips**

Practitioner Tips

Make AI Verification a Required Workflow

Build the Workflow First

Decide in advance that every AI output used in a client matter must go through three steps: identify, verify, and correct.

Set Aside Verification Time

AI may draft a letter in minutes, but practitioner review still requires time. The verification step cannot be skipped or guessed.

Prevent Shortcuts

Make verification a reflex, not a judgment call. A polished AI answer is not client-ready until the review is complete.

## 2.3 Preparer Penalties in an AI-Assisted Workflow

### Lesson at a Glance

#### IRC Preparer Penalties

#### AI Does Not Shift Practitioner Liability

The preparer penalty provisions were written long before artificial intelligence existed — but they apply fully to AI-assisted return preparation. The penalties attach to the human practitioner who signs the return, regardless of what software or AI tool produced the work.

#### §6694(a)

#### Unreasonable Positions

Applies when a return position lacks substantial authority or reasonable basis and causes an understatement of tax.

#### §6694(b)

#### Willful or Reckless Conduct

Applies to intentional disregard of rules or reckless conduct — including knowingly relying on unverified AI output.

#### §6695(g)

#### Due Diligence Failures

Applies when preparers fail required due diligence procedures for credits and filing status determinations.

#### Critical Rule

AI tools do not sign tax returns — practitioners do. The legal responsibility for every number, position, and statement on the return remains with the human preparer under the Internal Revenue Code and Circular 230.

### §6694(a) — Unreasonable Position

A §6694(a) penalty applies when a return contains an understatement of tax liability due to a position that does not meet the "substantial authority" standard and was not adequately disclosed. For TY2026, the penalty is the greater of **\$1,000 or 15% of income derived from the return** (per Rev. Proc. 2025-32(.54)).

In an AI-assisted workflow, §6694(a) exposure arises when a practitioner adopts an AI-generated position without independent verification and that position lacks substantial authority. The penalty analysis does not ask whether the AI made an error — it asks whether the practitioner's position on the return had substantial authority and whether the practitioner exercised reasonable professional judgment.

## **§6694(b) — Willful or Reckless Conduct**

The §6694(b) penalty — greater of **\$5,000 or 75% of income derived** — applies when a preparer willfully attempts to understate liability or acts in reckless or intentional disregard of rules or regulations. This penalty cannot be reduced by adequate disclosure.

Using AI output for a return position without any verification, when a reasonable practitioner would recognize the risk, may rise to the level of reckless disregard. A practitioner who knows AI tools frequently produce incorrect tax figures but uses them without verification is not acting with reasonable care.

## **§6695(g) — Due Diligence Penalty**

The §6695(g) penalty of **\$665 per failure** applies when a practitioner fails to meet due diligence requirements for EITC, CTC/ACTC/ODC, AOTC, or HoH. Due diligence requires: completing Form 8867, completing applicable worksheets, asking adequate questions, and not knowing or having reason to know the information is incorrect.

AI cannot satisfy any of these requirements. AI may help generate interview questions — but the practitioner must actually ask those questions and apply professional judgment to the answers. An AI determination that a client "appears to qualify" for EITC does not satisfy the due diligence requirement. The \$665 per-failure penalty applies to the practitioner regardless.

## **The penalties are per return — and they stack.**

A single return claiming EITC, CTC, and HoH where due diligence was AI-generated rather than practitioner-conducted = three §6695(g) failures = \$1,995 on one return. Multiply that across a filing season of non-compliant returns and the exposure becomes enormous. The \$665 per-failure amount is exactly why building proper due diligence into every return is non-negotiable.

## **The Reasonable Cause and Good Faith Defense**

The most common defense to §6694(a) is reasonable cause and good faith — the practitioner reasonably believed the position was supported by the law based on research,

consultation, or reasonable reliance on information provided. To establish this defense in an AI-assisted workflow, the practitioner needs documented evidence that:

### **Acceptable Primary Sources**

- The AI was used to identify relevant authorities
- Those authorities were independently verified from primary sources
- The practitioner applied professional judgment to determine the position was supported

**AI use alone — without the verification and judgment steps — does not establish reasonable cause. "The AI told me" is not a defense because reasonable cause requires the practitioner's own reasonable belief, not reliance on a tool the practitioner knows can be wrong.**

## **Practitioner Tips**

Practitioner Tips

### **Build Your §6694 Protection Workflow**

#### **Think Like a Reviewer**

Treat the verification workflow as your §6694 protection. Identifying claims, verifying authority, and documenting the review establishes the foundation for a reasonable cause defense.

#### **§6695(g) Is Mandatory**

Complete Form 8867 on every return claiming EITC, CTC, AOTC, or Head of Household status. AI may assist with preparation, but the practitioner conducts the inquiry and makes the determination.

#### **Professional Standard**

A practitioner who follows a documented verification workflow has evidence of professional judgment. A practitioner who skips verification has no meaningful defense when the AI-generated position proves incorrect.

## **2.4 What AI Can and Cannot Do in a Due Diligence Workflow**

### **Lesson at a Glance**

Due Diligence Requirements

## AI Cannot Replace Form 8867 Due Diligence

Due diligence for EITC, CTC/ACTC/ODC, AOTC, and Head of Household status is one of the most important practitioner obligations — and one of the most dangerous areas to improperly automate.

### AI CAN Assist

- Generate interview questions
- Organize client responses
- Flag inconsistencies for review
- Prepare draft explanations and notes

### AI CANNOT Replace

- Practitioner inquiry
- Professional judgment
- Form 8867 completion responsibility
- Final due diligence determination

### §6695(g) Penalty

The penalty currently applies at **\$665 per failure** and attaches to the practitioner regardless of what software or AI tool was used.

### Critical Rule

AI may help prepare the process, but the practitioner must personally conduct the due diligence inquiry, evaluate the facts, complete Form 8867, and make the final eligibility determination.

### Learning Objectives

- Identify which due diligence tasks AI may assist with and which require practitioner judgment.
- Explain why AI cannot complete Form 8867 or make eligibility determinations.
- Apply the correct due diligence workflow when AI is used in the intake process.
- Identify the documentation required to support a compliant due diligence process.

## The Five Due Diligence Triggers

Due diligence requirements under §6695(g) and Treasury Regulation §1.6695-2 apply to returns claiming:

- Earned Income Tax Credit (EITC)
- Child Tax Credit / Additional Child Tax Credit (CTC/ACTC)
- Credit for Other Dependents (ODC)
- American Opportunity Tax Credit (AOTC)
- Head of Household (HoH) filing status

For each triggered item, the practitioner must meet **four** requirements:

### Due Diligence Requirements

#### The Four Required Due Diligence Elements

For each triggered credit or filing status item, the practitioner must satisfy four separate due diligence requirements. AI may assist with preparation — but it cannot independently satisfy any requirement.

#### 1. Complete Form 8867

The preparer must complete and submit Form 8867 for every applicable return.

#### 2. Complete Worksheets

Applicable IRS worksheets must be completed to support eligibility calculations.

#### 3. Know the Law

The practitioner must understand the rules and ask adequate questions to determine eligibility.

#### 4. No Reason to Know

The practitioner cannot ignore facts or inconsistencies that suggest the information may be incorrect.

#### Critical AI Limitation

AI may help generate interview questions, organize responses, and prepare draft worksheets for Requirements 1 and 3 — but AI cannot independently satisfy any of the four due diligence requirements. The practitioner remains responsible for the inquiry, judgment, and final determination.

## What AI Can Appropriately Assist With

ask

AI May Assist?    Practitioner Must Still:

---

ask	AI May Assist?	Practitioner Must Still:
Generate a list of EITC eligibility interview questions	Yes	Review questions for completeness and accuracy; actually ask them to the client
Organize and format client intake questionnaire responses	Yes	Verify the organized information matches what the client actually stated; apply judgment to responses
Flag potential inconsistencies in client-provided data	Yes — as a review aid	Evaluate flagged items with professional judgment; ask follow-up questions
Explain the AOTC eligibility rules in plain language for a client letter	Yes — with verification	Verify the explanation against current IRS guidance before sending
Complete Form 8867	No	Complete Form 8867 personally, based on the actual client interview and professional judgment
Determine whether the qualifying child test is satisfied	No	Apply the actual test to the client's actual facts; document the conclusion
Make the eligibility determination for HoH	No	Conduct the HoH interview; verify household costs, qualifying person, and marital status

## Checkbox compliance without genuine inquiry is not due diligence

The IRS evaluates due diligence based on what the practitioner knew, what questions were asked, what documents were reviewed, and what professional judgment was exercised. An AI-generated checklist that was not used in a genuine client interview satisfies none of these requirements. Completing Form 8867 with "yes" answers based on AI-generated assumptions — without an actual client conversation — is the most common due diligence compliance failure.

### Worked Example

#### Appropriate AI Use in EITC Due Diligence

AI can help prepare the interview process, but the practitioner must personally conduct the due diligence inquiry, evaluate the client's answers, and complete Form 8867.

#### AI Prepares Draft Questions

Before the meeting, the practitioner uses AI to generate an EITC interview checklist covering earned income, qualifying children, residency, investment income, and filing status.

## Practitioner Customizes

The practitioner reviews the checklist and adds client-specific questions about the recent separation and where the child lived during the year.

## Practitioner Conducts Inquiry

During the meeting, the practitioner personally asks the questions, documents the answers, and follows up on unclear or inconsistent responses.

## Form 8867 Completed

The practitioner applies professional judgment, determines eligibility, and completes Form 8867 based on the documented inquiry.

**Appropriate AI Use:** AI generated the starting questions, but the practitioner conducted the actual due diligence. This is acceptable because the human preparer performed the inquiry, evaluated the facts, documented the responses, and made the final eligibility determination.

## The "Reason to Know" Standard

The due diligence requirement includes a critical objective element: the practitioner must not know or have reason to know that any information provided is incorrect. This standard is measured against what a reasonable, well-trained tax professional would have recognized as a red flag. If AI-generated intake analysis flags an inconsistency that the practitioner ignores — income that seems implausibly low for the lifestyle described, a child whose claimed residency conflicts with prior-year data — the practitioner "has reason to know" and must investigate further. Ignoring an AI-flagged inconsistency does not protect the practitioner; it compounds the due diligence failure.

## Practitioner Tips

### Practitioner Tips

### Use AI to Improve the Interview — Not Replace It

#### Better Questions

Use AI to help generate more complete interview questions — especially for unusual fact patterns or practice areas outside your normal experience.

#### Document the Actual Interview

The workfile should document the questions actually asked, the client's responses, and the practitioner's conclusions — not the AI-generated preparation materials.

## Important Compliance Point

The AI preparation is effectively invisible to the IRS. During an examination, what matters is the practitioner's actual inquiry, documentation, professional judgment, and completed due diligence records.

## 2.5 Practical AI Applications That Work Well

### Lesson at a Glance

Productive AI Applications

### Where AI Actually Saves Time in a Tax Practice

The earlier modules focused heavily on AI risks and limitations because safe AI use starts with understanding the risks. This lesson focuses on the other side: where AI genuinely creates productivity value when combined with proper practitioner verification.

#### Client Communication Drafting

AI can quickly draft client letters, explanation emails, appointment reminders, and tax planning summaries for practitioner review.

#### IRS Guidance Summaries

AI can summarize IRS notices, revenue procedures, publications, and announcements when the practitioner provides the actual source text.

#### Intake Checklists

AI can generate customized intake questionnaires and interview checklists tailored to specific client situations.

#### Practice Management Documents

Engagement letters, office procedures, staff policies, WISP drafts, and workflow templates can often be prepared faster with AI assistance.

#### Research Starting Points

AI can help identify relevant code sections, possible authorities, related issues, and research directions — but not replace authoritative tax research or verification.

#### Key Principle

AI delivers the most value when it accelerates drafting, organization, and preparation — while the practitioner continues to verify, analyze, and exercise professional judgment before any client use.

## Learning Objectives

- Identify five AI applications that provide genuine productivity value in a tax practice.
- Explain the appropriate verification step for each application.
- Apply the technique of pasting primary source text into AI to reduce hallucination risk.
- Estimate the time savings available from appropriate AI use in each application area.

## Application 1 — Client Communication Drafting

### Application 1

### Client Communication Drafting

AI excels at producing clear, readable first drafts of client-facing communications. A practitioner can describe a situation in general terms and receive a well-structured draft within seconds.

### Verification Step

Read every sentence carefully. Verify every tax statement against a primary source, correct inaccuracies, and personalize the language for the specific client situation.

### Time Savings

Draft with AI: approximately 2–3 minutes

Verify and personalize: approximately 8–12 minutes

Total workflow: about 10–15 minutes versus 25–35 minutes from scratch.

### Best Practice

Do not include client-identifying information such as names, Social Security numbers, or exact dollar amounts in prompts unless using an enterprise AI platform with a Data Processing Agreement (DPA). Describe the situation in general terms and personalize the document only after verification.

## Application 2 — Summarizing New IRS Guidance

### Application 2

### Summarizing New IRS Guidance

When a new revenue procedure or IRS notice is published, AI can produce a plain-language summary — but only if given the actual source text. Paste the relevant sections of the Rev. Proc. or Notice into the AI tool, then ask for a structured summary.

## Ground the AI

Use the actual IRS source text as the input. This grounds the AI in the document and reduces hallucination risk for that task.

## Verification Step

Compare the AI summary line by line against the original document. Correct misstatements and flag relevant omitted sections.

## Time Savings

AI may produce a draft summary in 3–5 minutes. Verification may take 10–15 minutes. The combined workflow is still faster than reading the full guidance cold — while preserving practitioner review.

**This is the highest-value AI application in a tax practice.** Understanding new guidance quickly is essential — and the volume of IRS guidance is high enough that AI-assisted summarization, done correctly, provides genuine time savings without meaningful accuracy risk when primary source text is provided.

## Application 3 — Client Intake Checklists and Organizers

### Application 3

## Client Intake Checklists & Organizers

AI can generate customized client intake checklists based on a short description of the client's situation. A prompt describing W-2 income, Schedule C business activity, rental property income, and dependent children can produce a comprehensive starting checklist within seconds.

## AI Creates the Starting Point

AI organizes document requests and intake questions quickly, producing a structured checklist the practitioner can build from.

## Practitioner Verification

Review the checklist against the prior-year return and add any client-specific items the AI missed. Remove items that do not apply.

## Key Principle

The AI-generated checklist is only a starting point. The practitioner's knowledge of the client's actual facts, prior-year return history, and current-year changes remains the final filter.

## **Application 4 — Practice Management Documents**

### Application 4

#### Practice Management Documents

First drafts of engagement letters, privacy notices, consent forms, office procedures, and practice management templates are well within AI's capabilities. These documents generally contain fewer technical tax law claims, making verification faster and reducing overall error risk.

#### AI Drafts the Framework

AI can quickly organize standard language, structure sections, and create readable first drafts for internal practice documents.

#### Professional Review

Review all language for completeness, accuracy, and consistency with firm procedures before using the document in practice.

#### Important Legal Protection

Have an attorney review engagement letters and legal client documents before adoption. AI-generated legal language should always be reviewed by a licensed attorney before being used with clients.

## **Application 5 — Research Starting Points**

### Application 5

#### Research Starting Points

AI is useful for identifying which areas of tax law may be relevant to a client situation — helping practitioners know where to begin their research. A prompt involving a personal injury settlement, cryptocurrency transaction, or foreign account issue can quickly produce a list of potentially relevant tax topics and authorities.

#### AI Helps Scope the Research

AI can suggest potentially relevant code sections, filing issues, terminology, and related tax concepts that may deserve further review.

#### Practitioner Performs the Research

Every item suggested by AI is only a research lead. The practitioner must verify all conclusions against the IRC, Treasury Regulations, IRS guidance, and authoritative sources.

## Key Principle

AI is helping define the scope of the research — not performing the research itself. The practitioner remains responsible for reading the authorities, analyzing the facts, and reaching the final conclusion.

## Practitioner Tips

Practitioner Tips

### Get Better AI Results with Better Inputs

#### Paste the Source Text

For tasks involving IRS guidance or technical rules, paste the actual document text into the AI conversation instead of asking broad general questions. Responses grounded in the real source material are more accurate and easier to verify.

#### Track Your Time Savings

For one week, compare the time spent on AI-assisted tasks versus your estimated time without AI. This helps identify where AI creates the most value in your practice and where verification time should be concentrated.

## Professional Workflow Principle

AI produces the best results when the practitioner provides high-quality source material, verifies the output carefully, and measures where the technology actually improves workflow efficiency.

## 3.1 What Happens to Your Data When You Use AI?

### Lesson at a Glance

Productive AI Applications

Where AI Actually Saves Time in a Tax Practice

The earlier modules focused heavily on AI risks and limitations because safe AI use starts with understanding the risks. This lesson focuses on the other side: where AI genuinely creates productivity value when combined with proper practitioner verification.

Client Communication Drafting

AI can quickly draft client letters, explanation emails, appointment reminders, and tax planning summaries for practitioner review.

IRS Guidance Summaries

AI can summarize IRS notices, revenue procedures, publications, and announcements when the practitioner provides the actual source text.

### Intake Checklists

AI can generate customized intake questionnaires and interview checklists tailored to specific client situations.

### Practice Management Documents

Engagement letters, office procedures, staff policies, WISP drafts, and workflow templates can often be prepared faster with AI assistance.

### Research Starting Points

AI can help identify relevant code sections, possible authorities, related issues, and research directions — but not replace authoritative tax research or verification.

**Key Principle:** AI delivers the most value when it accelerates drafting, organization, and preparation — while the practitioner continues to verify, analyze, and exercise professional judgment before any client use.

## The Three Tiers of AI Data Handling

### Tier 1

#### Free and Consumer Tools

- Most free AI tools — free versions of ChatGPT, Gemini, and others — reserve the right to use conversation content to improve their models. Your prompts may be reviewed by human evaluators, stored for extended periods, or used as training data for future model versions. This is disclosed in the terms of service, but most users do not read terms of service before using free tools.
- For a tax practitioner, this means: anything typed into a free AI tool may be retained and reviewed. Client SSNs, EINs, financial figures, and identifying information entered into a free tool may be stored in ways the practitioner cannot control or retrieve.

### Tier 2

#### Paid and Professional Subscriptions

- Paid subscriptions often include enhanced privacy protections — options to disable training data use, shorter data retention periods, and more explicit data handling commitments. However, "paid" does not automatically mean "private." The specific

data handling terms must be reviewed for each product. A paid subscription may still retain conversation data for weeks or months, and the specific retention and use terms vary significantly by provider.

### Tier 3

#### Enterprise Deployments

- Enterprise agreements — typically purchased by larger firms — often include a **Data Processing Agreement (DPA)** or Business Associate Agreement that explicitly prohibits the provider from using customer data for model training, commits to specific retention and deletion timelines, and provides contractual data security guarantees. For tax practices that handle significant volumes of sensitive taxpayer data and want to use AI in client-specific workflows, an enterprise tier with a formal DPA is the appropriate level.

### "Paid" does not mean "private" — read the terms.

[Image: Paid Does Not Mean Private]

#### Worked Example

##### Choosing the Right AI Tool for Client Data

A solo practitioner wants to use AI for three tasks: drafting a general client newsletter about estimated tax payments, summarizing IRS Notice 2026-11 by pasting the notice text, and drafting a letter for a specific client referencing their \$42,000 Schedule C income.

##### Task 1

Drafting a general client newsletter involves no client-specific identifying information and may be performed using a paid AI subscription after reviewing the provider's data retention terms.

##### Task 2

Summarizing IRS Notice 2026-11 by pasting the notice text also involves no client-specific identifying information and may be performed using a paid AI subscription.

##### Task 3

Drafting a client letter that references a specific taxpayer's \$42,000 Schedule C income involves client financial information and should never be entered into a free AI tool.

#### **Correct Approach:**

For Task 3, the practitioner should either use an enterprise AI platform covered by a Data

Processing Agreement (DPA) or anonymize the facts before entry. For example, enter: "My client had Schedule C income of approximately \$40,000" and personalize the draft only after the AI-generated content has been reviewed.

## What to Look For in AI Terms of Service

Question	Why It Matters
Is conversation data used to train the model?	Client-adjacent data used for training may be retained indefinitely and potentially exposed in future model outputs
How long is conversation data retained?	Longer retention = longer period of potential exposure; 30-day limits are more protective than indefinite storage
Can human reviewers access conversations?	Human review of AI conversations containing tax data is a data security risk
Is a Data Processing Agreement available?	A DPA provides contractual protections not available in standard consumer terms
Where are servers located?	Data stored outside the U.S. may be subject to different legal protections

## What Does a DPA Look Like?

A Data Processing Agreement (DPA) is usually a formal legal document attached to a vendor's terms of service or enterprise agreement. Most DPAs contain standard sections describing how customer data may be handled.

### Types of Data

Defines what information the vendor may receive, store, or process — such as uploaded documents, prompts, conversations, or client records.

### Privacy & Access

Explains who may access the data, whether human reviewers are involved, and whether subcontractors or cloud providers may also process the information.

### Retention Rules

States how long information may remain stored and whether customers can request deletion of conversations or uploaded content.

### Example DPA Language

"Vendor shall process Customer Data solely for the purpose of providing contracted services and shall not use Customer Data for model training except as expressly authorized by Customer."

### Worked Example

## Choosing the Right AI Tool for Client Data

A solo practitioner wants to use AI for three tasks: drafting a general client newsletter about estimated tax payments, summarizing IRS Notice 2026-11 by pasting the notice text, and drafting a letter for a specific client referencing their \$42,000 Schedule C income.

### Task 1

Drafting a general client newsletter involves no client-specific identifying information. It may be done on a paid subscription tier after reviewing the tool's data retention terms.

### Task 2

Summarizing IRS Notice 2026-11 by pasting the notice text also involves no client-specific identifying information. It may be done on a paid subscription tier after reviewing retention terms.

### Task 3

Drafting a letter for a specific client using their \$42,000 Schedule C income involves client-specific financial information. This should never be entered into a free AI tool.

### **Correct Approach:**

For Task 3, the practitioner should either use an enterprise AI tool with a data processing agreement, or anonymize the scenario before entering it. For example, enter: "My client had Schedule C income of approximately \$40,000," then personalize the draft after the AI-generated text is reviewed.

## 3.2 Client Data That Must Never Enter a Public AI Tool

### **NEVER enter into a public AI tool — no exceptions:**

[Image: Never Enter Client Data Into Unapproved AI Tools]

#### Prohibited Data

#### Never Enter These Items Into an Unapproved AI Tool

##### Personal Identifiers

- Social Security Numbers (SSNs) of taxpayers, spouses, or dependents
- Employer Identification Numbers (EINs)
- Driver's license numbers
- Passport numbers
- Government-issued identification numbers

##### Financial Information

- Bank account numbers
- Routing numbers
- Financial account identifiers
- Specific income figures tied to an identifiable client
- Specific asset values tied to an identifiable client

#### Protected Information

- Health information
- Medical records
- Insurance account numbers
- Information protected by privacy laws or contractual obligations

#### Taxpayer Information

- Client names combined with financial or tax data
- Client addresses combined with financial or tax data
- Prior-year tax return information of any identifiable taxpayer
- Any facts that would allow a specific client to be identified

#### Safe Practice:

If AI assistance is needed, remove names, identifying numbers, addresses, and exact financial figures before entering the scenario. Use generalized facts and reinsert client-specific information only after the draft has been reviewed.

## The Anonymization Technique

#### Instead of entering...

"John Smith, SSN 412-55-7890, had wages of \$72,000..."

"Maria Garcia at 123 Main Street, Chicago..."

"Client's EIN 82-1234567, Schedule C income \$48,000..."

"James Wilson (DOB 03/15/1985) received a 1099-R..."

#### Enter this anonymized version...

"My client had W-2 wages of approximately \$72,000..."

"My client is a sole proprietor operating from a home office..."

"My client has Schedule C income of approximately \$48,000..."

"My client, age 40, received a 1099-R showing..."

**Anonymization works for almost every AI task.** The AI does not need the client's name or SSN to help draft a letter, explain a tax rule, or analyze a general scenario. The practitioner personalizes the output after verification. The AI assists with the structure and language — not the identification.

## Incident Response

### Accidental Entry of Prohibited Data

If client information is accidentally entered into an AI tool that was not approved for that type of data, the firm should respond immediately, document the event, and update its controls.

#### 1. Document

Note what data was entered, which AI tool was used, who entered it, and when the incident occurred.

#### 2. Request Deletion

Review the tool's data deletion process and submit a deletion request immediately, keeping a copy of the request for the incident file.

#### 3. Assess Notice

Assess whether the incident triggers breach notification obligations under your state's data breach notification law.

#### 4. Update ISP

Document the incident in the firm's Information Security Plan and implement additional controls to reduce the chance of recurrence.

#### 5. Train Staff

Use the incident as a training opportunity. The goal is better controls and safer habits, not punishment.

## Common Errors

### Common Errors

Most AI-related compliance failures are not sophisticated cyberattacks — they are simple judgment mistakes made during busy workdays. These errors often begin with a practitioner trying to save time.

## SSN Disclosure

Entering a client's Social Security number into a free AI tool to obtain a "quick answer" on a tax issue.

## Uploading Tax Returns

Pasting a client's tax return or financial statements into an AI platform and asking it to "analyze" the information.

## Failure to Anonymize

Using the client's real name, business name, or exact financial figures in prompts without first removing identifying information.

## Incognito Myth

Assuming that private browsing or incognito mode protects information submitted to an AI platform. Browser privacy settings do not control what the AI provider stores.

## 3.3 IRS Publication 4557 and AI Tool Use

### Lesson at a Glance

[Image: IRS Publication 4557 Compliance]

#### WISP Requirement

##### AI Adoption Requires a WISP Update

IRS Publication 4557 requires every tax preparation business — regardless of size — to maintain a written Information Security Plan (WISP). When a practice adopts AI tools, the WISP must be updated to address the new risks those tools introduce.

#### Risk Assessment

Evaluate how AI tools collect, store, process, and retain information entered by firm personnel.

#### Acceptable Use Rules

Define what information may and may not be entered into AI systems and document approved use cases.

#### Access Controls

Restrict access to approved users and ensure firm accounts are protected with strong authentication.

## **Key Compliance Point:**

Adopting AI is not simply a technology decision. It is a security and compliance decision. The firm's WISP should document approved AI tools, acceptable uses, prohibited data, review procedures, and incident response steps before AI is used with firm workflows.

## **The Six AI-Specific Questions Every ISP Must Answer**

**1**

### **Which AI tools are authorized for use in this practice?**

- Name the specific tools and their tier (free, paid, enterprise). Any tool not on the approved list requires written approval before use on firm devices or for firm work. Update the list at least annually.

**2**

### **Which tier is authorized for which types of tasks?**

- Free tools: general research and administrative drafting only — no client-specific data. Paid/professional tools: permitted for client-adjacent tasks only if the current terms of service have been reviewed. Enterprise tools: permitted for tasks involving client-specific data only with a DPA in place.

**3**

### **What categories of data are prohibited from entry into AI tools?**

- Reference the absolute prohibition list from Lesson 3.2. The ISP should state explicitly that SSNs, EINs, account numbers, client names combined with financial data, and prior-year return data may not be entered into any tool not covered by a DPA.

**4**

### **Who in the organization is authorized to use AI tools, and for what tasks?**

- Reference the staff vs. practitioner framework from Module 4. Administrative staff may use AI for low-risk administrative tasks. Practitioners supervise all client-facing AI use. No staff member uses AI for due diligence determinations or return positions.

5

### How will the practice verify that staff are following AI use protocols?

- Describe the monitoring and training approach: annual acknowledgment of the AI Use Policy, periodic spot-checks of AI prompts used for client work, immediate reporting of any suspected protocol violation.

6

### What is the incident response procedure if protected data is accidentally entered into an AI tool?

- Reference the five-step incident response from Lesson 3.2: document, request deletion, assess breach notification, update ISP, train staff.

## Compliance Note

[Image: oops\_.svg]

An WISP that does not address AI tools is incomplete. The FTC Safeguards Rule requires the WISP to be updated when new risks are identified. AI tool adoption is a new risk. A practice that uses AI tools without updating its WISP is non-compliant with the Safeguards Rule.

## 3.4 The FTC Safeguards Rule and Tax Preparers

### Lesson at a Glance

[Image: Gramm-Leach-Bliley Act]

Legal Foundation

Gramm-Leach-Bliley Act & AI Vendor Oversight

The image above illustrates the **Gramm-Leach-Bliley Act (GLBA)**, the federal law that forms the foundation of the FTC Safeguards Rule and many of the data protection requirements that apply to tax professionals.

The FTC Safeguards Rule (16 CFR Part 314) applies to tax preparation businesses as financial institutions under the Gramm-Leach-Bliley Act. For AI tool adoption, the most important obligation is the **service provider oversight requirement**.

## Service Providers

AI vendors are third-party service providers when they receive, process, store, or access client information.

## Due Diligence

Firms must evaluate whether the provider maintains appropriate safeguards and security controls before using the service.

## Ongoing Monitoring

Oversight is not a one-time event. Firms should periodically review vendor practices, retention policies, and security commitments.

## **Practical Application:**

Before entering any client information into an AI system, a tax professional should determine whether the provider offers appropriate contractual protections, security controls, and data handling practices. Simply paying for an AI subscription does not automatically satisfy the FTC Safeguards Rule's service provider oversight requirement.

## **Why the Safeguards Rule Applies to Tax Preparers**

The Gramm-Leach-Bliley Act (GLBA) defines "financial institution" broadly — it includes any business that is significantly engaged in financial activities, including tax preparation services. Tax preparers receive, process, and store nonpublic personal financial information (NPI) from their customers. The FTC Safeguards Rule, which implements GLBA's security requirements, therefore applies to tax preparation businesses of all sizes.

## Service Provider Oversight

### AI Tools as Service Providers

Under the Safeguards Rule, a "service provider" is any entity that receives, maintains, processes, or otherwise accesses customer information in connection with services provided to the financial institution. When a tax practice uses an AI tool in any workflow involving client financial data, the AI tool vendor may qualify as a service provider.

The Rule requires the tax practice to select providers with appropriate safeguards, require those safeguards contractually, and monitor the provider for compliance.

### What Satisfies the Service Provider Assessment

#### Review Terms

Review the vendor's current privacy policy, data handling terms, retention rules, and security commitments before using the tool with client data.

#### Confirm DPA

Confirm whether the vendor offers a Data Processing Agreement or equivalent contractual commitment that addresses confidentiality and safeguards.

Document in ISP

Document the review, the approved use cases, and the vendor's security commitments in the firm's Information Security Plan.

**Do Not Use Rule:**

If the vendor cannot provide adequate security commitments, do not use that AI tool for workflows involving client financial data.

### 3.5 Practical Data Protection Protocols for AI-Using Practices

#### Six AI Security Protocols Every Tax Practice Should Adopt

[Image: six-protocol.webp]

AI Governance Framework

Six AI Security Protocols Every Tax Practice Should Adopt

Protocol 1

Anonymization Before AI Input

Before entering any client-related scenario into an AI tool, remove all identifying information. Make anonymization a reflex: before clicking "Send," ask, "**Is there any identifying information in this prompt?**"

Protocol 2

Tier Restriction Policy

**Free:** Research, templates, administrative drafting only.

**Paid:** Anonymized client-adjacent work.

**Enterprise + DPA:** Client-specific workflows involving actual client financial data.

Protocol 3

Approved Tool List

Maintain a written list of approved AI tools showing the tool name, subscription tier, authorized tasks, review date, and DPA status. Review the list annually.

Protocol 4

Staff Training Before AI Use

Train staff on approved tools, tier restrictions, prohibited data, and incident response procedures. Maintain a signed acknowledgment from every user.

#### Protocol 5

##### Annual Terms Review

AI privacy policies and terms change. Conduct a documented annual review of every approved tool and update the ISP when necessary.

#### Protocol 6

##### Incident Response

Establish a clear procedure for accidental data entry. Staff should be encouraged to report mistakes immediately. A culture that hides errors creates unreported breaches.

#### **Practice Tip:**

These six protocols should be incorporated into the firm's Information Security Plan (ISP/WISP), staff training program, and approved technology governance procedures. Together they create a practical framework for compliant AI use in a tax practice.

**These six protocols apply to a solo practitioner in exactly the same way as to a 50-person firm.** The documentation is lighter, the training is simpler (just the practitioner), and the approved tool list may have only two or three entries. But the protocols themselves — anonymize, restrict by tier, maintain an approved list, train before use, review annually, respond to incidents — apply regardless of firm size.

## 4.1 The AI-Assisted Tax Office: A Day-in-the-Life

### **Real-World AI Workflow Scenarios**

[Image: AI Assistance in Practice]

Real-World Applications

Five AI Scenarios From a Tax Professional's Workday

Client Newsletter

IRS Guidance Summary

Client Response Draft

Meeting Preparation

Research Starting Point

## Scenario 1 — Drafting a Client Letter About an IRS Notice

### Scenario #1 — CP2000 Client Letter

#### Situation

A client receives a CP2000 notice proposing additional tax related to an unreported Form 1099-DIV.

After reviewing the notice, the practitioner determines the dividend came from a **tax-exempt bond fund** and is not taxable.

The practitioner now needs to explain the situation clearly to the client.

#### AI Task

AI drafts a plain-language client letter explaining:

- What a CP2000 notice is
- What the practitioner found
- Why the dividend is not taxable
- What happens next

No client-identifying information is included in the AI prompt.

#### Verification Step

The practitioner personally reviews every sentence before sending the letter.

Verification includes:

- Confirming the CP2000 explanation against [IRS.gov/cp2000](https://www.irs.gov/cp2000)
- Confirming tax-exempt bond treatment under IRC §103
- Adding the client's notice date and response deadline
- Personalizing the final response for the client

#### Time Savings

Traditional drafting time: **~20 minutes**

AI-assisted workflow: **~3 minutes drafting + ~8 minutes verification = ~11 minutes total**

Net time savings: ~9 minutes on a single notice letter — multiplied across dozens of client notices during the year.

## Scenario 2 — Summarizing a New Revenue Procedure

### Scenario #2 — Rev. Proc. Summary Sheet

#### Situation

Rev. Proc. 2025-32 has just been published with TY2026 inflation adjustments. The practitioner needs a quick internal reference sheet of key figures most relevant to individual client returns.

#### AI Task

The practitioner pastes only the relevant Rev. Proc. sections into the AI tool and asks for a summary organized by category:

**standard deduction, EITC, retirement accounts, health accounts, and credits.**

#### Verification Step

The practitioner compares the AI summary line by line against the original Rev. Proc. text.

Two figures are slightly different and are corrected before the summary is distributed to staff.

#### Time Savings

Full manual review: **45 minutes**

AI-assisted summary + verification: **15 minutes**

Net savings: 30 minutes

Accuracy improves because the verification step forces a careful comparison against the original authority.

## Scenario 3 — Generating a Client Document Checklist

### Scenario #3 — Client Document Checklist

#### Situation

A new client is coming into the practice:

- Sole proprietor web developer
- First year in business
- Home office deduction
- One dependent child
- Owns rental property

The practitioner needs a customized document request checklist before the intake appointment.

#### AI Task

AI generates a tailored tax document checklist based on the client characteristics provided.

The prompt includes:

“Generate a tax document checklist for a client with these characteristics: sole proprietor first year in business, home office, one dependent child, rental property.”

No client-identifying information is entered into the AI system.

#### Verification Step

The practitioner reviews the generated checklist using professional judgment.

Verification actions include:

- Adding startup cost documentation
- Adding Form SS-4 EIN documentation
- Removing one irrelevant checklist item

Because this is a new client, no prior-year return exists for comparison.

#### Time Savings

Building a customized checklist manually: **~15 minutes**

AI-assisted checklist + practitioner review: **~5 minutes total**

Net time savings: ~10 minutes per new client

## **CHATGPT with paid subscription answered when I typed**

Sample AI Response

### Tax Document Checklist

#### Business Income & Expenses

- 1099-NEC, 1099-K, invoices, and business income records
- Business bank statements and payment processor reports
- Receipts for software, internet, supplies, advertising, and contractor payments
- Startup cost documentation for first-year business expenses
- Form SS-4 or EIN confirmation letter, if applicable

## Home Office

- Square footage of home and office space
- Rent or mortgage interest, utilities, insurance, and repairs
- Photos or floor plan showing dedicated business space

## Dependent Child

- Child's Social Security number and date of birth
- School or residency records, if needed
- Childcare provider statement, if claiming dependent care credit

## Rental Property

- Rental income records
- Mortgage interest statement and property tax bill
- Repairs, insurance, HOA fees, utilities, and management fees
- Purchase documents or depreciation schedule

**Practitioner Review Required:** This checklist must be reviewed and customized before sending to the client.

## Scenario 4 — Answering a Client's General Tax Question

### Scenario #4 — Routine Client Tax Question

#### Situation

A client emails the practitioner with a common question:

“I gave my old car to my college-student daughter — is that taxable to her?”

#### AI Task

The practitioner asks AI:

“What is the federal income tax treatment when a parent gives a personal-use vehicle to an adult child?”

No client-identifying information is included in the prompt.

#### Verification Step

AI correctly identifies that:

- Gifts are generally excluded from recipient income under IRC §102
- The annual exclusion and lifetime exemption apply to the donor

The practitioner then:

- Verifies IRC §102 directly
- Confirms the rule applies to personal-use property
- Checks for unusual facts such as business use of the vehicle

Verification time: **~4 minutes**

Final answer is confirmed accurate before sending the client response.

Time Savings

Researching the issue manually from scratch: **~12 minutes**

AI-assisted workflow + verification: **~5 minutes total**

Net time savings: ~7 minutes on a routine client question that appears repeatedly during filing season.

## **Gemini Responded like this - Real World Example**

[Image: gemini-response-940x1024.webp]

## **Scenario 5 — Organizing Incoming Client Documents**

### **Scenario #5 — Organizing Client Documents**

#### **Situation**

A client submits an envelope of mixed tax documents:

W-2s from two employers, Form 1099-INT, mortgage statement, property tax bill, and charitable contribution receipts.

#### **AI Task**

A staff member describes only the document types to the AI — not names, Social Security numbers, dollar amounts, addresses, or account numbers.

AI creates a categorized intake checklist showing which tax form each document connects to and what information staff should capture.

#### **Verification Step**

The staff member compares the AI categorization against the actual documents.

The practitioner reviews the completed intake form before beginning return preparation.

## Time Savings

Manual organization and intake form: **20 minutes**

AI-assisted categorization + staff review + practitioner review: **12 minutes**

Net savings: 8 minutes

The time savings come from faster organization — not from replacing practitioner review.

### **The pattern across all five scenarios:**

AI drafts or organizes. The practitioner (or supervised staff) verifies. The AI output is never the final product — it is always reviewed before any client receives it or any return uses it. This pattern, consistently applied, is what makes AI use both productive and professionally safe.

## 4.2 Staff vs. Practitioner: Who Can Use AI for What

### **Lesson at a Glance**

#### AI Task Authorization Framework

In a multi-person tax practice, clarity about who may use AI — and for which tasks — is essential for quality control and Circular 230 compliance.

The practitioner who signs the return remains responsible for every aspect of preparation, including AI-assisted work performed by staff.

#### Trained Staff May Perform

Administrative drafting, checklist creation, document organization, and non-client-specific workflow support.

#### Requires Practitioner Supervision

Research summaries, client communication drafts, issue spotting, and tax workflow assistance.

#### Practitioner Personally Required

Return positions, final eligibility determinations, due diligence conclusions, client-specific tax advice, and signing decisions.

**Core Rule:** AI can assist staff, but it cannot replace practitioner judgment. The signer remains responsible for accuracy, diligence, supervision, and compliance.

In a multi-person tax practice, clarity about who can use AI for which tasks is essential for both quality control and Circular 230 compliance. The practitioner who signs the return is responsible for every aspect of its preparation — including AI-assisted work performed by staff. This lesson provides a clear authorization framework that specifies, for each category of AI task, whether it can be done by trained staff, requires practitioner supervision, or requires the practitioner personally.

## The Authorization Framework

Task Category	Who Can Use AI	Review Required By
Administrative drafting (appointment reminders, templates, checklists)	Any trained staff member using an approved tool	Staff member — practitioner spot-check quarterly
Client-facing communication drafts (letters, emails, explanations)	Any trained staff member using an approved tool	<b>Practitioner must review every piece of correspondence before it is sent</b>
Document organization and intake categorization	Any trained staff member	Practitioner reviews completed intake before beginning return preparation
Tax research starting points (identifying relevant areas)	Trained staff may use AI to generate a research starting list	<b>Practitioner conducts the actual research, verifies from primary sources, makes the final determination</b>
Return position research and analysis	Practitioner only — AI may assist but practitioner conducts and verifies all research	Practitioner — professional judgment function; cannot be delegated to staff or AI
Due diligence (Form 8867, EITC/CTC/AOTC/HoH)	Practitioner only — AI may help generate interview questions	Practitioner — must conduct actual client interview; make eligibility determination; complete Form 8867
Return preparation and signing	Practitioner only — AI is never the preparer of record	Practitioner signs every return personally with PTIN

## Review is not rubber-stamping

The practitioner's review of AI-assisted staff work must be substantive — reading the content, verifying tax claims, and applying professional judgment. A practitioner who signs off on AI-assisted correspondence without reading it has not reviewed it. The §6694 analysis will not accept "I trusted my staff's AI tool" as a defense.

[Image: rubber-stamping.webp]

### Worked Example

#### Practitioner Review Is Not a Rubber Stamp

A staff member uses AI to determine that a new gig-worker client qualifies for the home office deduction. Based on the AI's analysis, the staff member completes the home office section of the return and forwards it for practitioner review.

The practitioner reviews the return but does not independently re-examine whether the home office qualifies under the exclusive use test. Instead, the practitioner relies on the staff member's AI-assisted conclusion and signs the return.

#### **IRS Audit Result:**

The IRS audits the return and disallows the home office deduction because the workspace fails the exclusive use test.

#### **Penalty Analysis:**

The §6694 analysis focuses on the practitioner who signed the return. A review that does not independently apply the exclusive use test to the client's facts does not establish reasonable cause. The practitioner remains responsible for the return position, even when AI and staff participated in the preparation process.

#### **Professional Responsibility Rule:**

AI can assist staff. Staff can assist the practitioner. However, the practitioner must personally evaluate material return positions before signing the return. Professional judgment cannot be delegated to AI or satisfied through a rubber-stamp review.

## 4.3 Writing Your AI Use Policy

**Every tax practice that uses AI tools needs a written AI Use Policy. It must exist, be in writing, be communicated to everyone who uses AI, and be followed**

### AI Use Policy

#### Six Required Elements of an AI Use Policy

##### Element 1: Approved Tools

List specific AI tools authorized, their tier — free, paid, or enterprise — and DPA status. Any tool not listed requires written approval before use on firm devices or for firm work. Review annually.

##### Element 2: Prohibited Data

Explicitly list categories of client data that may not be entered into any AI tool without a DPA: client SSNs and EINs; bank account numbers; client names combined with any financial or tax data; and prior-year return data of any identifiable individual.

##### Element 3: Task Authorization

Specify which tasks AI may be used for, by whom, and with what level of review. Reference the staff vs. practitioner framework from Lesson 4.2. AI tools may not be used to complete Form 8867, make due diligence determinations, or serve as the basis for a return position without independent practitioner verification from primary sources.

##### Element 4: Verification Requirement

All AI-generated content used in client communications or return preparation must be reviewed using the 3-step verification workflow: (1) identify every specific factual and legal claim, (2) verify each claim against a primary source, and (3) correct, remove, or document any claim that cannot be verified.

##### Element 5: Incident Response

Any suspected entry of prohibited client data into an unauthorized AI tool must be reported immediately. Upon report: (1) document, (2) submit a data deletion request, (3) assess state breach notification obligations, (4) update the policy if indicated, and (5) conduct additional training.

##### Element 6: Training and Acknowledgment

All firm personnel who use AI tools for firm work must complete AI Use Policy training before using any approved tool. Training should occur at onboarding and annually. All personnel must sign an annual acknowledgment.

## 4.4 Circular 230 and AI-Assisted Work Product

### **§10.35 — Duty of Competence**

Circular 230 §10.34

False Statements and AI-Generated Errors

Circular 230 §10.34 provides that a practitioner may not sign a return or other document knowing it contains a position that is not supported by a reasonable basis, or a statement that is false.

Unverified Position

A practitioner uses an AI-generated return position without independent verification and that position is incorrect.

Unverified Figure

A practitioner signs a return containing an AI-generated dollar amount, rate, threshold, or calculation that has not been verified against the current-year primary source.

Fabricated Citation

A practitioner relies on an AI-fabricated citation or authority as the basis for a return position.

#### **Important Point:**

The §10.34 analysis does not ask whether the AI made the error. It asks whether the practitioner knew or should have known that the statement was false.

#### **Professional Responsibility:**

A practitioner who knows AI tools can generate incorrect information but uses AI output without verification should have known an error was possible. The verification workflow is the operational expression of the §10.34 obligation.

### **§10.34 — Prohibition on False Statements**

Circular 230 §10.35

Duty of Competence and AI Use

Competent AI use means understanding what the tool can and cannot do, knowing its failure modes, and applying an appropriate verification workflow to every output.

Understand Capabilities

Know what the AI tool can do well and where it should be used only as an assistant rather than a decision-maker.

#### Understand Failure Modes

Recognize hallucinations, outdated training data, fabricated citations, and the difference between pattern prediction and authoritative research.

#### Verify Every Output

Apply a consistent verification workflow to every factual claim, legal conclusion, calculation, and citation.

#### **Competence Standard:**

A practitioner who uses AI without understanding how it works — or who performs no verification because the output “looks right” — is not exercising competent use of the tool.

#### **Key Principle:**

Competence does not require a practitioner to understand the mathematics or programming behind AI. It does require understanding the risks of the technology and implementing procedures that ensure its output is independently reviewed before being relied upon.

[Image: Correspondence.webp]

## **\$10.22 — Duty of Diligence**

Circular 230 §10.22

#### Duty of Diligence and AI Verification

Circular 230 §10.22 requires practitioners to exercise reasonable diligence in all matters. When AI is used in a tax practice, diligence cannot be occasional or informal.

#### **Why Occasional Review Is Not Enough:**

A practitioner who verifies AI outputs only when they remember to do so, but has no systematic workflow, is not exercising the consistent diligence §10.22 requires.

#### Written AI Use Policy

Documents the firm’s rules, approved tools, prohibited data, review standards, and accountability procedures.

#### 3-Step Verification Workflow

Requires every specific factual claim, legal conclusion, figure, threshold, and citation to be reviewed before reliance.

### **Diligence Standard:**

The written AI Use Policy and the 3-step verification workflow together document and operationalize the practitioner's diligence obligation under §10.22.

### **Key Point:**

Diligence is not satisfied by a quick glance at an AI-generated answer. It requires a repeatable process that can be followed, documented, and explained if the return position or client communication is later questioned.

## 4.5 Client Disclosure Considerations

### **Lesson at a Glance**

#### AI Disclosure to Clients

No federal law currently requires a tax practitioner to disclose to clients that AI tools were used in preparing a return or related correspondence.

However, important practical, ethical, and professional responsibility considerations still apply.

One Circular 230 principle is clear: **a practitioner may not misrepresent the nature of their workflow or professional services.**

#### Disclosure Requirements

Understand what current federal rules require — and what they do not.

#### Best Practices

Learn recommended approaches for transparency and professional communication.

#### Client Questions

Respond professionally when clients ask whether AI was involved in preparing their return.

#### Engagement Letters

Review what engagement letters should say regarding AI-assisted workflows and client confidentiality.

**Professional Responsibility Rule:** A practitioner may use AI-assisted tools — but may not shift responsibility to the AI system or falsely imply that professional judgment was not required.

**Guiding Principle:** Clients hire the practitioner for professional judgment, accuracy, confidentiality, and accountability — regardless of what technology tools assist the workflow behind the scenes.

## Engagement Letter Language

The most practical approach for AI-using practices is to include a brief, accurate description of technology use — including AI — in the engagement letter. This creates transparency at the outset of the relationship, before any work is performed. Sample language:

Taylor Tax & Advisory Group

1450 Legacy Plaza Drive · Suite 410 · Nashville, TN 37215  
(615) 555-1040 · engagement@taylortaxgroup.com

Engagement Letter — AI Use Disclosure

Dear Client,

As part of our ongoing commitment to efficient and secure professional service, Taylor Tax & Advisory Group uses technology tools — including artificial intelligence applications — to assist with limited aspects of tax research, document organization, workflow management, and drafting of client communications.

All AI-assisted work product is reviewed, verified, and approved by a licensed tax professional before delivery to you or submission to any taxing authority. The use of technology tools does not replace professional judgment, compliance review, or preparer responsibility.

You remain our client throughout the engagement, and your return remains the professional responsibility of the preparer who signs it. Confidential taxpayer information is handled in accordance with applicable federal privacy obligations, IRS Publication 4557, our Written Information Security Plan (WISP), and firm confidentiality procedures.

If you have any questions regarding our use of technology tools in the preparation process, please contact our office at any time.

Sincerely,

[Image: Anita Taylor Signature]

Anita Taylor, EA

Managing Partner · Taylor Tax & Advisory Group

## Responding to Direct Client Questions

When a client asks directly whether AI was used in their return preparation:

- **Answer honestly.** Describe how AI assists — drafting, organizing, research starting points — and how professional review and verification applies to every output.
- **Emphasize the verification workflow.** Clients are less concerned about whether AI was used than about whether the return is accurate. Explaining that every AI-assisted element is verified against primary sources before the return is filed addresses the underlying concern.
- **Affirm the practitioner's responsibility.** Remind the client that the signed return is the practitioner's professional responsibility — not the AI tool's. The signature represents the practitioner's professional judgment, not the AI's output.

## Transparency about AI use builds client trust — not the opposite.

Clients who understand how AI is used — and how verification ensures accuracy — are more confident in the practitioner's competence than clients who suspect AI is being used without disclosure. Proactive, honest communication about AI is a professional strength.

## 4.6 The Emerging Regulatory Landscape

### Lesson at a Glance

#### AI Regulation & Professional Responsibility

AI regulation in professional practice is one of the **fastest-moving areas of law and policy**.

This lesson provides a current snapshot — as of this course's publication — of how the IRS, Treasury Department, FTC, state regulatory bodies, and professional organizations are approaching AI use in tax practice.

The regulatory landscape will continue to evolve rapidly.

IRS & Treasury

Guidance, preparer standards, Circular 230, and return compliance.

FTC Safeguards Rule

Data protection, WISP requirements, and AI security controls.

State Regulators

Privacy laws, professional standards, and consumer protection rules.

Professional Associations

Ethical guidance, best practices, and practitioner education.

**Important Principle:** A practitioner’s responsibility is not to predict future AI regulation — it is to stay current through annual CE, professional publications, IRS.gov monitoring, and ongoing professional judgment.

**Regardless of how regulations evolve:** the core professional standard remains the same — the practitioner signing the return is responsible for the accuracy, confidentiality, and compliance of all work performed with AI assistance.

## IRS Position on AI in Tax Practice

Current IRS Position

Existing Rules Apply to AI-Assisted Return Preparation

As of this course's publication, the IRS has not issued a revenue procedure, notice, or regulation specifically governing tax practitioner use of AI tools. However, the absence of AI-specific guidance does not create an exception to existing law.

IRC §6694

Preparer penalty rules apply fully to AI-assisted return preparation.

IRC §6695

PTIN, due diligence, and preparer responsibilities remain unchanged when AI is used.

Circular 230

Competence, diligence, supervision, and ethical obligations apply exactly as they do for any other workflow.

### **Important Compliance Principle:**

AI-assisted preparation is not a separate legal category. Existing preparer penalty rules and Circular 230 obligations apply with full force regardless of whether the work was performed manually, with software assistance, or with AI assistance.

### **Audit Risk Consideration:**

The IRS has deployed AI and advanced analytics internally for audit selection and compliance activities. Returns that contain unusual deductions, inconsistent reporting patterns, or figures that deviate significantly from statistical norms for a taxpayer's industry and income level may attract additional scrutiny regardless of how the return was prepared.

## FTC Enforcement Activity

The FTC has been active in enforcing its Safeguards Rule and has brought enforcement actions against financial services businesses — including tax preparers — for inadequate

data security. As AI tools become standard in tax practice, FTC enforcement focus is likely to extend to AI-related data protection failures: practices that enter client SSNs into unauthorized AI tools, that fail to maintain written ISPs covering AI use, or that misrepresent their data handling practices in client communications.

## Action Plan

### What Practitioners Should Do Now

#### Implement an AI Use Policy

Adopt a written AI Use Policy today. Doing so helps satisfy current Publication 4557 and FTC Safeguards Rule obligations.

#### Monitor IRS Guidance

Subscribe to IRS e-News for Tax Professionals and monitor future IRS announcements, notices, and guidance related to AI use.

#### Add AI to Annual CE

Include at least one AI-focused continuing education course each year from organizations such as NATP, NAEA, or AICPA.

### **Future-Proof Principle:**

Specific regulations may evolve, new guidance may be issued, and AI technology will continue to change. However, the core professional framework is unlikely to change.

#### AI Is a Tax Office Assistant — Not a Tax Preparer

Regardless of how specific future AI regulations become, the practitioner remains responsible for competence, diligence, professional judgment, and the final return that bears their signature.

### **The practitioners who invest in AI literacy now will be best positioned as the regulatory landscape solidifies.**

Practitioners who understand how AI works, have implemented appropriate controls, and can demonstrate a consistent verification workflow will face no surprises from AI-specific regulation — because they are already doing what any reasonable framework will require.

## Final Lesson – Course Feedback Survey

### Course Feedback

## Share Your Experience

Your insights help us continuously improve. Please take a brief moment to share your feedback about the cybersecurity training module below.