

Artificial Intelligence in Accounting

Catherina Wong

De Anza College

Business

CC BY 4.0





Conceptual illustration of artificial intelligence in accounting, featuring a digital human-like profile integrated with circuit patterns, data visualizations, and financial chart elements. The image highlights the intersection of AI technologies, data analytics, and modern accounting practices.

Artificial Intelligence in Accounting

Catherina Wong

De Anza College

This text is disseminated via the Open Education Resource (OER) LibreTexts Project (<https://LibreTexts.org>) and like the thousands of other texts available within this powerful platform, it is freely available for reading, printing, and "consuming."

The LibreTexts mission is to bring together students, faculty, and scholars in a collaborative effort to provide an accessible, and comprehensive platform that empowers our community to develop, curate, adapt, and adopt openly licensed resources and technologies; through these efforts we can reduce the financial burden borne from traditional educational resource costs, ensuring education is more accessible for students and communities worldwide.

Most, but not all, pages in the library have licenses that may allow individuals to make changes, save, and print this book. Carefully consult the applicable license(s) before pursuing such effects. Instructors can adopt existing LibreTexts texts or Remix them to quickly build course-specific resources to meet the needs of their students. Unlike traditional textbooks, LibreTexts' web based origins allow powerful integration of advanced features and new technologies to support learning.



LibreTexts is the adaptable, user-friendly non-profit open education resource platform that educators trust for creating, customizing, and sharing accessible, interactive textbooks, adaptive homework, and ancillary materials. We collaborate with individuals and organizations to champion open education initiatives, support institutional publishing programs, drive curriculum development projects, and more.

The LibreTexts libraries are Powered by [NICE CXone Expert](#) and were supported by the Department of Education Open Textbook Pilot Project, the California Education Learning Lab, the UC Davis Office of the Provost, the UC Davis Library, the California State University Affordable Learning Solutions Program, and Merlot. This material is based upon work supported by the National Science Foundation under Grant No. 1246120, 1525057, and 1413739.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation nor the US Department of Education.

Have questions or comments? For information about adoptions or adaptations contact info@LibreTexts.org or visit our main website at <https://LibreTexts.org>.

The LibreTexts name, logo, and book covers are trademarked by LibreTexts, Inc. (a 501(c)(3) not-for-profit organization) and protected against unauthorized use.

This text was compiled on 09/16/2026

Program Page



De Anza College sits in the heart of Silicon Valley, where innovation and opportunity are part of daily life, and also has one of the highest costs of living in the country. For our students, that makes every dollar count, and the price of a textbook can be the difference between staying enrolled or dropping a class. Our Zero Textbook Cost (ZTC) program closes that gap by giving students full access to high-quality, faculty-created course materials at no cost from the first day of class.

De Anza is proud of its diverse student community, and we hold our ZTC materials to the same standard of excellence that defines the college. Faculty across departments write, adapt, and review these resources themselves, keeping them accurate, current, and grounded in the perspectives of the students who use them, with accessibility and inclusive design built in from the start.

This program reflects the time and care our faculty invest in their students. If you have questions, feedback, or come across any issues with a ZTC resource, contact De Anza College's ZTC Grants Coordinator, [Shagun Kaur](#).

Learn more about [De Anza ZTC Program](#) [www.deanza.edu/academic-services/oer/]

Note: The De Anza logo is a copyrighted trademark of De Anza College and may not be reproduced without permission.

TABLE OF CONTENTS

Title Page

Program Page

Licensing

Accessibility Statement

About the Book

About the Author

1: The roles of intelligence and artificial intelligence (AI) in accounting

- 1.1: Introduction and Technology in Accounting Today
- 1.2: The Concept of Intelligence
- 1.3: Different types of Artificial Intelligence
- 1.4: The Historical Development of Intelligence in Accounting
- 1.5: Analyzing Artificial Intelligence Strengths and Weaknesses in Accounting
- 1.6: Detailed Text Descriptions for Chapter 1

2: AI in Financial and Managerial Accounting

- 2.1: Distinguish the roles of financial and managerial accounting purposes
- 2.2: AI-Driven Tools in Financial Accounting
- 2.3: AI's Role in Managerial Accounting
- 2.4: Opportunities and Challenges in AI Applications

3: Applications of AI in Audit and Tax

- 3.1: AI's Evolution in Audit and Tax
- 3.2: AI as a Tax Preparation Assistant
- 3.3: AI's Role in Auditing

4: Ethics of AI in Accounting

- 4.1: Why Ethics Matters?
- 4.2: Individual Professional Responsibilities
- 4.3: Professional Accounting Ethical Standards
- 4.4: Societal Consequences
- 4.5: AI Ethics Policy for an Accounting Firm

5: Emerging Trends in AI and Accounting Software

- 5.1: Emerging AI Trends in Accounting
- 5.2: Evaluation of the Impact of AI in the Accounting Profession
- 5.3: Creating and Implementing Guardrails

Index

Glossary

Detailed Licensing

Licensing

A detailed breakdown of this resource's licensing can be found in **Back Matter/Detailed Licensing**.

This page titled [Licensing](#) is shared under a [Public Domain](#) license and was authored, remixed, and/or curated by Catherina Wong.

Accessibility Statement

Accessibility

De Anza College is dedicated to creating learning materials that are accessible and benefit all students. Our materials follow the Web Content Accessibility Guidelines (WCAG) 2.1 Level AA. Our authors and contracted accessibility specialists checked all resources using this ASCCC OERI checklist, the accessibility tool in LibreTexts, WAVE, JAWS, and other more detailed checks for complex materials such as interactive components.

We selected LibreTexts as our preferred platform because of its [commitment to accessibility](#). In addition to the platform's inherent accessibility, we collaborated with the LibreTexts team to fix any accessibility issues we found.

Our Accessibility Review Process

Accessibility was built into the resource as it was written. This resource went through the following stages of review:

1. **Template check.** The project coordinator and instructional designers checked the first chapter for accessibility as a model to ensure later chapters meet accessibility standards.
2. **Progress checks.** The manuscript was checked at approximately 40%, 60%, and 80% completion to catch and correct accessibility issues in production.
3. **Specialist review.** Once the manuscript was complete, an accessibility expert conducted an accessibility check.
4. **Remediation.** The author and the project manager who conducted the earlier progress checks worked together to remediate the issues identified by the accessibility expert.
5. **Final check.** The resource underwent another full check by the accessibility expert before the book was deemed ready for publication.

This process combines ongoing checks during writing with a final specialist review, so issues are caught early and confirmed fixed before publication.

Report an Accessibility Problem

We are always looking for ways to make our resources more accessible. If you run into a barrier while using this resource, please let us know so we can look into it. When you [contact us](#), it helps to include:

- The name of the resource
- The location of the problem (a web address or page/chapter description)
- A description of the problem you encountered
- The device, browser, and any assistive technology you were using (e.g., Windows 11, Google Chrome, JAWS screen reader)

Contact us: [De Anza ZTC team](#)

This statement was last updated August 2026

This resource was project-managed by **Shagun Kaur**, overseeing scope, budget, and schedule; stakeholder coordination via monthly meetings; training and platform support (LibreTexts, troubleshooting, technical support); QA/accessibility check scheduling; and licensing and copyright governance.

A Note on Known Limitations

No accessibility check is perfect. We've made a genuine, sustained effort to catch and fix issues throughout the writing and review process described above, but some barriers may remain, especially in complex or interactive content. We'd rather be upfront about that than claim a standard of perfection we can't guarantee.

About the Book

About the Book

Artificial intelligence is no longer a distant concept for the accounting profession, rather, it is reshaping daily workflows and redefining professional roles. Additionally, AI is raising new questions about ethics, responsibility, and trust. This textbook offers a comprehensive and accessible introduction to AI in accounting, written for business students and emerging professionals who will work with AI technologies throughout their careers. The book is organized around five interconnected themes that progress from foundational understanding to applied practice and forward-looking awareness. Beginning with the nature of intelligence itself and the historical development of AI in accounting, the text moves through hands-on applications in financial and managerial accounting, the transformation of audit and tax practice, the ethical imperatives facing individual professionals and firms, and the emerging technologies that will continue to shape the field.

The Approach

This textbook is designed with concept understanding at its core. Each chapter opens with clear learning objectives and a chapter overview that frames the “why” before the “what.” Real-world examples, case scenarios, and practical exercises are woven throughout to connect theory to professional practice. Students are not simply asked to learn what AI tools do but they are challenged to evaluate how those tools work, where they fall short, and what responsibilities professionals carry when deploying them.

This book was written based on developments in AI for accounting applications up to Summer 2026. Some applications discussed, such as the Metaya TaxBot, are still in active development at the time of publication. Given the fast-paced nature of AI advancements, readers should expect frequent updates to this text as the field continues to evolve.

About the Author

Catherina Wong



Catherina Wong is an Accounting Instructor at De Anza College, where she is passionate about making accounting concepts accessible and engaging, particularly for English Learners. She holds an MBA from the University of Toronto and an active Certified Public Accountant (CPA) license, bringing both academic depth and real-world professional experience to her teaching. She is equally passionate about seeing her students through to finding meaningful employment, believing that education is most powerful when it opens real doors.

With a deep commitment to educational equity and concern over the rising cost of learning materials in higher education, Catherina is especially proud to contribute to the world of Open Educational Resources (OER). In addition to authoring OER materials, she has championed student-centered initiatives such as the LinC program, which works to dismantle barriers that prevent students from thriving in higher education.

Outside the classroom, Catherina dedicates her time to animal welfare as a volunteer at the Foster Animal Vaccination Station (FAVS) at the Humane Society of Silicon Valley, where volunteers support foster families by ensuring animals in their care receive timely vaccines, treatments, and essential supplies. She also opens her home to orphaned kittens as a foster caregiver.

Catherina is deeply grateful to her family members, Alex and Dominic, who are lifelong learners themselves. Their curiosity, encouragement, and love of learning inspired her to write this book.

1: THE ROLES OF INTELLIGENCE AND ARTIFICIAL INTELLIGENCE (AI) IN ACCOUNTING

Chapter Overview

Imagine reviewing thousands of transactions, identifying unusual patterns, and forecasting future outcomes in seconds rather than days. This is the promise of artificial intelligence. Just as ledgers gave way to spreadsheets and spreadsheets evolved into enterprise software, AI is transforming what accountants can accomplish. Yet technology alone is not enough—the quality of accounting decisions still depends on human judgment, ethical reasoning, and professional expertise.

The five sections in the chapter move from context to concept to critique. Section 1.1 sets the stage with the current technology landscape in accounting. Section 1.2 defines intelligence itself before applying the term to machines. Section 1.3 maps the different types of AI you will encounter in practice. Section 1.4 traces how AI entered and evolved within the accounting profession. Section 1.5 brings it together by evaluating AI's real strengths and weaknesses for accounting work — the foundation for everything that follows in this course.

Learning Objectives

By the end of this chapter, you should be able to examine the roles of intelligence and artificial intelligence (AI) in accounting.

- Explain the concept of intelligence.
- Identify the different types of artificial intelligence (AI).
- Outline the historical development of AI in accounting.
- Analyze the key components of AI to evaluate their strengths and weaknesses in accounting.

1.1: Introduction and Technology in Accounting Today

1.2: The Concept of Intelligence

1.3: Different types of Artificial Intelligence

1.4: The Historical Development of Intelligence in Accounting

1.5: Analyzing Artificial Intelligence Strengths and Weaknesses in Accounting

1.6: Detailed Text Descriptions for Chapter 1

1.1: Introduction and Technology in Accounting Today

A Revolution in the Books

Accounting is as old as trade itself. For as long as businesses have existed, someone has kept records to tell a financial story. What has changed, through every era, is not the purpose of that work but the tools used to do it.



Figure 1.1.1: Accounting. (CC BY-SA 3.0; Nick Youngson via Pix4free). Detailed Description of Figure 1.1.1.

It was the same way for centuries – ledgers were written down in leather books using ink, closing books was about literally closing the book and opening a new one. Accountant was just a keeper of records who used such tools as pen, ten-key calculator and lots of time and patience for performing calculations. The work was not only precise but demanded knowledge about which records were important and how they were related to another. However, there appeared a computer. It did not substitute an accountant but made his life easier by shifting time spent on some routine work.

The transition from a paper ledger to electronic spreadsheet was the first step in the evolution of the field. At first, accountants performed their routine work but instead of paper and filing cabinet they used computer and floppy disk. Thus, this kind of work became faster but tasks were still the same – the accountant entered transactions, balance accounts and generated reports.

But when the Internet and cloud computing emerged, the evolution of the field reached its peak. Accounting information did not reside on a single computer anymore. Such popular applications as QuickBooks, Xero, NetSuite and lots of others allowed transferring general ledger into the cloud. Now everything can be done in real-time mode and from any place. Bank feed records transactions, point-of-sale system records them right after purchase, automatic payroll integration, photo of a receipt is transformed into the record of expenses and everything is recorded in the general ledger in the same time.

Influence of Technology in Accounting

However, it is not only about automation but also about artificial intelligence which can interpret invoices, find cash flow problems before their occurrence, discover the patterns which can be suspicious in the context of fraud and draft preliminary financial statements. Moreover, blockchain ensures the tamper-proof audit trail while robotic process automation performs procedures of month-end close which took whole weekends. There appear lots of new technologies and with them new requirements and possibilities.

Evolution of accounting is going to continue but not from compliance to strategy since good accountants already perform both of them. Instead, technologies are just going to rebalance the equation and reduce the time spent on records and increase the amount of work with analysis and forecasting. The question is not whether technologies will change the field – they already did

that. The question is what accountants will do with this time.

From Midnight to Dinner Case

It's 4:30 PM on the last day of the month, and Sarah needs to close the books for her consulting firm. Closing the books means verifying that all the company's financial records for the month are accurate and complete. Ten years ago, this process meant staying at the office until midnight to match bills with payments, find missing receipts, and get approval signatures from managers. Today, she will be home for dinner.

Sarah begins by reviewing the company's bills. Her accounting software has already processed the vendor invoices throughout the month. As each bill arrived as a PDF in the company inbox, artificial intelligence read the important information: the vendor name, what was purchased, the amount owed, and the payment due date. The software then compared these bills to the company's purchase records, identified any discrepancies, and sent them to the appropriate managers for approval. Bills that matched the purchase records perfectly were queued for payment. The three bills with discrepancies appear in Sarah's review queue with the specific problems highlighted. She resolves these issues in minutes.

Bank reconciliation was once Sarah's least favorite task. Now she clicks a single button and the software completes most of the work. The system pulls transaction data directly from the bank and compares it to the records in the company's accounting system. It identifies only the few transactions that need her attention. One transaction for \$247 appears as unmatched. The software suggests this might be the office supply purchase from last week that was recorded under a different description. The suggestion is correct. Sarah confirms the match with one click.

Sarah must categorize the month's expenses, but the software has already sorted most of them. The system has learned patterns over time. Payments to certain vendors are always categorized as consulting services. Monthly software fees belong in the technology category. Purchases from the deli during lunch hours are categorized as meals. The software even determined that a new annual software license should be recorded as a long-term asset rather than a regular monthly expense by analyzing the cost and reading the contract terms.

Sarah generates preliminary financial reports. These reports show how much money the company earned, how much it spent, and what assets and obligations it has. Calculations that previously required hours now complete instantly. The software has drafted explanatory notes for significant transactions. Sarah reviews the numbers, makes two manual corrections for unusual items, and the reports update immediately.

Before finishing for the day, Sarah schedules payments for the approved invoices. There are no checks to write or envelopes to mail. The following day, the system will transfer funds electronically to vendors, record each payment in the accounting system, and send confirmation emails. The software will also update the company's cash projections based on these outgoing payments.

Sarah's work on the month-end close took half a day. This same work previously required three full days. The complete closing process will take several more days because department managers must review budget allocations, Sarah's supervisor will verify the financial statements, and the chief financial officer will review unusual transactions before finalizing everything. However, the repetitive tasks of matching transactions and checking numbers that once kept Sarah working past midnight are now complete.

Behind Sarah's work, multiple systems operate invisibly. When the office supply invoice arrived, optical character recognition technology scanned it, natural language processing interpreted the content, and machine learning algorithms classified it automatically. When Sarah approved payments, encryption protocols secured the data, application programming interfaces connected the accounting system to the bank's system, and automated processes sent notification emails. When Sarah generated financial reports, algorithms applied accounting standards, calculated financial ratios, and formatted the reports according to professional templates.

These systems were active throughout the entire month. Each time an employee used the company credit card, the transaction was recorded immediately and categorized based on the merchant type. Transactions outside normal patterns were flagged for review. Customer invoices were generated automatically from project management software and monitored for payment receipt. Payroll calculations for taxes and benefits occurred through systems that update automatically when tax laws change. Financial activity for the month was captured continuously in real time by interconnected systems.

Sarah's role has evolved rather than disappeared. She no longer spends her time entering data and performing manual calculations. Instead, she reviews the exceptions that automated systems flag, makes informed decisions about complex transactions, explains financial results to company managers, and ensures the automated systems function correctly. Technology handles routine tasks. Sarah applies professional judgment and expertise.

Your Connection

Think about the last time you paid for something online. The transaction recorded itself instantly. Now scale that up to a business level.

If technology cuts accounting work from three days to half a day, what happens to the extra time? Accountants can shift from recording what happened to analyzing what it means and advising on what to do next. They focus on questions about profitability, growth opportunities, and strategic decisions that shape the business.

Technology handles the data entry. Accountants interpret the results and guide decision making.

What skills do you think accountants need today that go beyond just being good with numbers?

1.2: The Concept of Intelligence

What is Intelligence?

One of the best ways to understand intelligence is to view it as a process that moves from raw data to meaningful action. Consider a doctor diagnosing a patient. Test results, blood pressure readings, and symptoms are all just pieces of data. When the doctor interprets them, looks for patterns, and compares them with previous cases, that data becomes information, such as a possible diagnosis. Intelligence comes into play when the doctor uses that information to recommend treatment, suggest lifestyle changes, or adjust medication based on what the data reveals. That is when knowledge truly becomes useful.

The same process applies in accounting. Financial data begins as individual transactions: a sale to a customer, a payment to a supplier, an employee paycheck, or a utility bill. When accountants organize these transactions and examine them in context, comparing current month spending to previous months or analyzing which products generate the most revenue, that data becomes information. Financial reports showing profit margins, cash flow trends, or expense patterns represent this information stage. Intelligence emerges when accountants use this information to provide recommendations and insights. They might advise renegotiating a supplier contract because costs have risen significantly, recommend discontinuing an unprofitable product line, or confirm the company has enough cash to expand while cautioning about expense trends that need attention. Data records what happened. Information explains the patterns. Intelligence guides decisions and actions.

When Intelligence Fails

Data can be very misleading if the interpretation process lacks intelligent insight. For instance, consider the case of a doctor interpreting the results of a diagnosis procedure. A doctor might use data and interpret it in terms of the disease or condition that is most likely and ignore other diagnoses that are rarer. The outcome would be the prescription of medicine and treatments for the wrong ailment. Moreover, the application of medical data from one population group to another in medical research could lead to ineffective treatment or even cause harm to those from different groups since the data was not interpreted correctly because it was selective and thus misleading. The data is there, but poor interpretation as a result of **selective analysis** turns it into a liability instead of an asset.

Similarly, in accounting, an accountant can analyze data related to products in terms of sales performance and recommend investments in products that yield the most sales without analyzing the margin of profits. The accountant could find out later that the sales leader does not have any profit at all. Similarly, if the accountant uses expenses data collected in summer time, when many expenses are low due to seasons, he or she may find out after making cuts to the budget on the basis of such limited data that expenses are higher at other times of the year. Interpretation of data selectively, incorrectly, or incompletely means the failure of intelligence process. Instead of leading to action, such data will lead to losses for the company.

Artificial intelligence is also affected by the same problem because of similar nature of its functioning. Artificial intelligence systems use the data they were provided during training and make recommendations and predictions based on them. However, if the training data contains biases, the system will reflect those biases when making predictions. For example, if an artificial intelligence system was trained using data from large corporations, it would make recommendations that will be good for large corporations but not for small businesses. Moreover, if it used a historical data set with misclassified expenses, it will misclassify even more expenses. The model trained to detect fraud in one industry will detect fraud in other industries even if no fraud exists. Such cases illustrate the **amplification effect**, which occurs due to the lack of correct training. This means that the system is not biased intentionally but reflects the biases of the training set. In conclusion, intelligence – human or artificial – is only as accurate as the data it is based on.

Definition: selective analysis

A human error in the interpretation process where someone examines only part of the available data while ignoring other relevant information, either intentionally or unintentionally, leading to incomplete or flawed intelligence.

Definition: amplification effect

When AI systems reproduce and scale up errors, biases, or limitations from their training data, automatically spreading flawed patterns across thousands or millions of decisions, transactions, or cases.

Exercise 1.2.1

Where should Bean & Brew Coffee spend their \$5,000 marketing budget next month?

Scenario: Bean & Brew sells two signature drinks. Their caramel latte generated \$50,000 in sales last month with a profit of \$1.50 per drink. Their specialty cold brew generated \$25,000 in sales with a profit of \$4 per drink. The owner sees that caramel lattes brought in twice as much sales and wants to spend the entire marketing budget promoting lattes.

1. If the owner only looks at total sales, what would they decide?
2. What information does the profit-per-drink data reveal?
3. Test your intelligence (what is your recommendation in one sentence?)

Answer available at the end of this section.

Answers

Exercise 1.2.1

1. They would spend all \$5,000 marketing caramel lattes since they have higher sales.
2. Cold brew is significantly more profitable per sale, even though it sells less overall.
3. Invest the marketing budget in cold brew because each new customer generates much higher profit for the business.

1.3: Different types of Artificial Intelligence

What sets AI apart from rule-based software

To understand what makes AI different from traditional accounting tools, imagine two approaches to learning a new skill. Traditional accounting software operates like someone following a detailed instruction manual, which means that every action requires a specific, pre-written rule. If you encounter a situation not covered in the manual, the system can't help. AI, by contrast, operates more like an apprentice who learns by observing patterns and applying those insights to new situations, even ones they've never encountered before.

This distinction becomes immediately practical in daily accounting work. Traditional software might have a rule stating "If expense type = office supplies, then post to Account 6100." That works perfectly until you encounter something the rule doesn't cover.

The difference in how AI systems work lies in how they learn based on the examples you give them. In other words, your software recognizes the patterns in your past decision-making and applies them to the new situations when trying to suggest how the situation should be categorized. Thus, for example, when your firm starts cooperation with a certain vendor, you may categorize the first few transactions yourself, and then the AI will analyze these examples based on parameters such as the vendor's name, the amount of transaction, the descriptive text and the categorization. Thus, AI will identify the patterns in your decision-making process, and based on these patterns, it will calculate the probability of how the next transactions with this particular vendor should be categorized.

Instead of asking "Does this transaction comply with the rule I programmed?", an AI will ask "Is this transaction similar to the patterns identified in my past examples?". Such approach enables the AI to categorize even the situations it never met before by comparing them with the patterns it has already learned from the previous experience.

However, the apprentice-like learning is not only done in one way. As well as human apprentices who acquire various sets of skills – the ability to notice exceptional cases, the ability to understand complex written texts and the ability to deal with full projects without any supervision, for instance – the AI has developed different skills as well. Instead of considering the AI as the only technology, it may be better to realize that there exist different AI systems each of which has different characteristics, and thus can be recognized as such in the accounting processes.

The History of AI Explained: Crash Course Futures of AI #1

Video Length: 12 minutes

Context: In the Crash Course video, "The History of AI Explained: Crash Course Futures of AI #1" builds the foundation needed to understand how artificial intelligence evolved into what it is today before exploring the six types of AI used in accounting. Starting from early computing milestones through Moore's Law, it traces the key breakthroughs that led to modern AI systems, including the benchmarks used to measure AI progress and the scaling principles that help predict where AI is headed. Chapters include Introduction, How did AI get here?, How do we measure AI progress?, How can we predict the future of AI?, and Conclusion.

This interactive element is not available in this format. Visit the link or scan the QR code to access it online.

[View this YouTube Video online](#)



Six Core Types of AI

Accountants today encounter AI in audit, tax, and financial reporting, but not all AI operates the same way. Knowing what each

type actually does, and what it cannot do, is what separates informed use from unverified reliance.

Machine Learning (ML)

Machine learning represents AI's ability to learn from historical data and apply those learnings to new situations. The key to identifying ML in action is watching for systems that improve their accuracy over time based on the feedback you provide. Unlike traditional software that performs exactly the same way on day one as it does on day one hundred, ML systems get smarter with use.

Consider the following example. After you've categorized 100 coffee shop transactions as "Meals & Entertainment," an ML system recognizes this pattern and begins automatically suggesting the same category for future transactions from similar vendors. When you encounter a transaction from an unfamiliar vendor, the system draws on these historical patterns to suggest a classification. If you correct that suggestion, the system incorporates your feedback and handles similar vendors more accurately in the future. The crucial point is that no one needs to write new rules. The system learns from experience, much like a human would.

Definition: machine learning

A type of AI that learns to recognize patterns from historical data and applies those patterns to new situations. In accounting, machine learning categorizes transactions by learning from past examples, such as automatically classifying a coffee shop purchase as a business meal expense.

Natural Language Processing (NLP)

While machine learning identifies patterns in data, **natural language processing** enables computers to read and comprehend human language in financial documents. The hallmark of NLP is its ability to extract meaningful information from unstructured text, the kind of documents accountants handle every day.

NLP systems can scan an invoice and understand that "Net 30" refers to payment terms, "Acme Corp" is the vendor name, and "\$1,250" represents the amount due, all without requiring this information to be entered into structured data fields. The technology goes beyond simple keyword matching; it understands context and can distinguish between similar-looking documents, recognizing whether a transaction represents a sale, an expense, or a payroll entry.

Definition: natural language processing (NLP)

AI technology that reads and extracts information from text-based documents. In accounting, NLP scans invoices, receipts, and tax forms to pull out key data like vendor names, dates, and amounts without manual typing.

Large Language Models (LLMs)

Large language models, such as GPT-4, represent a more recent evolution in AI technology. These systems have been trained on vast amounts of text, enabling them to both understand complex language and generate human-like written responses. You can identify LLMs by their ability to engage in conversational interactions and produce coherent written content.

In accounting contexts, this means you can pose questions in plain English and receive a written analysis in return. For example, you might ask "Summarize last quarter's expense trends by department" and the system generates a comprehensive response. LLMs power systems that can draft routine correspondence, interpret complex financial documents, and assist with various communication tasks that traditionally required human writing.

Definition: large language models (LLMs)

AI systems trained on vast amounts of text that can understand context, answer questions, and generate human-like responses. In accounting, LLMs like ChatGPT or Claude can explain complex financial reports, summarize audit findings, or help draft client communications.

Did you know?

Key Findings of AI Impact on Accounting Productivity

Choi and Xie's 2025 study of 79 firms and 277 accountants reveals significant productivity gains from AI adoption in accounting. Accountants using AI tools that incorporate natural language processing, optical character recognition, and large language models to process unstructured documents like invoices and receipts saved 3.5 hours per week previously spent on routine data entry.

Survey results show data entry and administrative work are the most common AI applications (44.40% each). AI adoption delivered measurable improvements of an 18% increase in weekly client support capacity, a 12% increase in general ledger detail, and completion of monthly financial close 7.5 days faster. Critically, accountants redirected time saved from manual tasks toward higher-value activities including client communication and quality assurance, demonstrating AI's role in augmenting rather than replacing professional expertise.

Anomaly Detection

The fourth kind of AI works differently, however, still performing a vital task, which is to detect anomalies within financial data. The **anomaly detection** system knows how "normally" your financial data looks and marks the outliers which need further investigation. Anomaly detection systems can be spotted through their ability to identify outliers regardless of there being any threshold rule programmed into them.

Thus, if in your organization the amount of expense claims usually lies between \$200 and \$500, but there appeared an expense claim of \$5,000, then it will be detected by the anomaly detection system as an outlier. It will happen not due to anyone setting a threshold rule "\$5,000," but because the system recognizes what is normal and what is an exception. Such technologies are revolutionizing both fraud detection and audit process making accountants able to pay attention to those outliers.

Credit card fraud detection is one of the most common examples of anomaly detection technology that one deals with daily. When traveling to a new city and using your credit card in a restaurant, the system of the company issuing your credit card compares the transaction with your previous pattern of spending money, such as merchants you spend at, amount you usually spend and the geographical locations where you make your transactions. If you have not done your purchase in that city before, then the system flags it or even rejects your transaction and calls you up to check whether you are indeed using your card. Nobody at your bank has set a rule that states "reject all purchases made in Denver." Instead, the system has learned your spending patterns and detected an anomaly.

The professional society recognizes the importance of that technology as well. According to the results of a 2024 survey conducted by the Association of Certified Fraud Examiners among nearly 1,200 fraud professionals, in two years from now the number of companies using AI in anti-fraud programs will triple and exceed 50%.

Definition: anomaly detection

AI that identifies unusual patterns or outliers in data that don't match expected behavior. In accounting, anomaly detection flags potentially fraudulent transactions, such as duplicate payments or expenses that fall outside normal ranges for a particular vendor or category.

Conversational AI Chatbots

Conversational **AI chatbots** represent specialized AI assistants designed specifically for accounting tasks and knowledge. While

these chatbots often utilize LLM technology under the hood, what distinguishes them is their focused design for accounting-specific interactions and their integration into accounting platforms. You can identify chatbots by their interface. The telling sign is that they present as an on-demand assistant within your software, and are ready to answer questions and guide you through processes.

Unlike general-purpose AI tools, accounting chatbots are trained on accounting terminology, regulations, and best practices. When you ask "What's the depreciation method for office furniture?" the chatbot provides guidance aligned with accounting standards. Depreciation refers to how businesses gradually write off the cost of long-lasting items like furniture over their useful life. The chatbot understands this concept and can explain it according to GAAP (Generally Accepted Accounting Principles, the rules used in the United States) or IFRS (International Financial Reporting Standards, rules used in many other countries), often citing specific regulations. When you're confused about a feature in your accounting software, the chat bot can walk you through the exact steps needed, understanding the context of where you are in the system.

What sets accounting chatbots apart is their specialized design for the accounting field. An accounting chatbot understands the specific terminology that accountants use in their daily work. For example, it knows that when accountants say "reconcile the account," they mean comparing two sets of records to make sure they match, such as verifying that the company's internal records align with the bank statement. The chatbot won't confuse this with the everyday meaning of reconcile, which means to restore a friendly relationship. This specialized knowledge, combined with conversational ease, makes chatbots particularly valuable for training new staff, providing instant answers to routine questions, and offering guidance on unfamiliar accounting scenarios without requiring you to search through manuals or help documentation.

Definition: AI chatbots

Conversational AI assistants designed to answer questions within a specific domain or platform. In accounting, chatbots like QuickBooks Assistant help users navigate software features, explain accounting terminology, or troubleshoot common bookkeeping issues.

Agentic AI

Agentic AI represents the most autonomous type of AI in accounting. These are systems that can independently execute complete processes involving multiple steps and finish entire accounting functions from start to finish. The defining characteristic of agentic AI is its ability to take action across multiple systems, make decisions at various points in a process, and deliver completed work rather than just recommendations. You can recognize agentic AI when a system doesn't just tell you what to do but actually does it.

Consider bank reconciliation, a critical accounting task where you verify that the company's internal records match what the bank says happened. Traditionally, this process requires an accountant to compare the bank statement against the company's own records, examining each line, investigating any discrepancies (differences between what the bank shows and what the company recorded), and making corrections. An agentic AI system can autonomously access bank statements, retrieve corresponding entries from the accounting software, match transactions using sophisticated pattern recognition, identify discrepancies, propose correcting entries, and even execute those corrections after applying learned rules about the organization's policies. The system handles the entire process, alerting you only when it encounters genuinely ambiguous situations requiring human judgment. Organizations must establish clear boundaries for what actions the system can take independently versus what requires human approval, ensuring that efficiency gains don't come at the expense of control and accuracy.

Definition: Agentic AI

AI systems that can independently complete multi-step tasks from start to finish with minimal human intervention. In accounting, agentic AI performs complex workflows like bank reconciliation by automatically matching transactions, identifying discrepancies, and generating reconciliation reports.

Understanding Open-Source vs. Closed-Source AI

Closed source models of AI are proprietary systems created and managed by particular companies that do not disclose their

source codes and internal structure. Examples of these models include GPT-4 (OpenAI), Claude (Anthropic), and Gemini (Google). When using tools based on these models, you are connecting to the connection points of these companies, where software products interact with each other. This is similar to using subscription services such as Netflix or Spotify. You get access to media content, but you do not own the technology and know nothing about its inner workings. The benefits of these models include their state-of-the-art capabilities due to huge investments in the development and improvement of AI models. On the other hand, you rely on the company providing the service, your data pass through its infrastructure, and thus, there is no confidentiality (this issue may be especially important in working with sensitive financial data), and you do not know exactly how the model functions.

As for open source models of AI, these models make their code and structure accessible. Examples of these models include Llama (Meta), Mistral, and many open source models from the community. These models are supposed to be downloaded and installed on the organization's infrastructure; however, in most cases, it is impossible due to the requirements of powerful computing hardware, special technical skills, and high costs of maintenance of infrastructure. Thus, most medium-sized and small companies cannot really download these models and use them. In practice, "open source" models are used via third-party services hosting them.

In spite of that, the main advantage of open source models is transparency and choice of whom to host these models. Large organizations with the presence of the department of technologies may install these models on their infrastructure and control the whole process. Small practices will find it easier to use open source models that are hosted by the third-party services specializing in providing the necessary infrastructure. These services will provide more transparency and control comparing with closed source models.

Many organizations use both types of models according to the sensitivity of the data they need to analyze. Thus, closed source models are used for doing research about certain accounting issues (learning about new legislation or the best practices), as well as preparing internal materials, while open source models are used for working with the real data of the clients (creating tax returns or analyzing financial statements). The distinction between these two situations is not the complexity or routine nature of the task.

Table 1.3.1: AI Types, Applications, and Software Tools in Accounting

AI Type	Characteristics	Applications	Where They Are Found
Machine Learning (ML)	Learns from examples and improves accuracy over time; identifies patterns in historical data; adapts based on user feedback.	Transaction categorization; vendor classification; expense sorting.	QuickBooks, Xero
Natural Language Processing (NLP)	Reads and understands human language in documents; extracts information from unstructured text; recognizes context and meaning.	Invoice data extraction; receipt processing; tax form data extraction (Tax forms such as, W-2s, 1099s).	Dext, Hubdoc, Intuit ProConnect, CCH Axxess Scan
Large Language Models (LLMs)	Generates human-like written responses; engages in conversations; trained on vast amounts of text; can draft and analyze content.	Summarizing financial reports; drafting emails; explaining complex concepts; tax research assistance.	ChatGPT, Claude, Microsoft Copilot
Anomaly Detection	Identifies unusual patterns without pre-programmed rules; learns what "normal" looks like; flags outliers automatically.	Fraud detection; unusual expense identification; audit risk assessment.	MindBridge
Conversational AI Chatbots	Specialized for accounting domain; integrated into accounting platforms; understands accounting terminology; provides guided assistance.	Answering software questions; explaining accounting and tax rules; walking through processes.	QuickBooks Assistant, CCH Answer Connect
Agentic AI Systems	Executes complete multi-step processes independently; makes decisions at multiple points; takes action across systems; delivers finished work.	Full bank reconciliation; end-to-end invoice processing; automated tax return preparation; sales and use tax compliance.	FloQast, Black Line, Thomson Reuters ONESOURCE.

Exercise 1.3.1

Instructions

Match each accounting scenario (1-6) with the most appropriate AI type (A-F). Each AI type should be used exactly once.

Scenarios

1. An accounting firm needs to read supplier invoices in various formats (PDF, email, scanned documents) and extract key information like invoice numbers, dates, amounts, and line items into their accounting system. Staff will verify the extracted data before posting.
2. A tax accountant wants an AI assistant that can have a back-and-forth conversation to help explain new tax code changes, answer questions about specific deductions, and guide users through filing requirements. The accountant will confirm all advice before providing it to clients.
3. A company wants to analyze three years of expense reports to learn which expenses typically belong in which general ledger accounts, then apply those patterns to categorize new expenses. The accounting team will review categorizations before finalizing entries.
4. A CFO needs the system to identify when expense claims or financial transactions don't fit normal company patterns, such as an employee suddenly submitting unusually high travel expenses or payments to unfamiliar vendors. Management will investigate all flagged items.
5. A CPA firm wants junior staff to use AI to help research complex accounting standards (like ASC 606), generate first drafts of technical memos, and get explanations of difficult concepts while working on client engagements. Senior staff will review all outputs before use.
6. A multinational corporation wants a system that can complete their month-end close process by pulling data from multiple subsidiaries, performing intercompany eliminations, generating consolidated statements, and creating management reports. Controllers will review the final package before distribution.

AI Types

- A. Machine Learning (ML)
- B. Natural Language Processing (NLP)
- C. Large Language Models (LLMs)
- D. Anomaly Detection
- E. Conversational AI Chatbots
- F. Agentic AI Systems

Answer available at the end of this section.

Want To Dig Deeper?

The productivity and fraud detection findings discussed in this section draw from two industry studies, a 2025 Stanford research paper on [AI adoption in accounting firms](#) and a [2024 benchmarking report](#) from the Association of Certified Fraud Examiners.

Answers

Exercise 1.3.1

Table 1.3.2: AI Types and Explanations for Accounting Scenarios

Scenario	Answer	Explanation
1	B (NLP)	Reading and extracting information from unstructured documents (invoices in various formats) requires natural language processing capabilities to understand document structure and extract specific data points.

Scenario	Answer	Explanation
2	E (Conversational AI Chatbots)	The interactive, question-and-answer format with domain-specific tax knowledge indicates a conversational AI chatbot specialized for tax.
3	A (ML)	Learning from historical patterns in the data and applying those learned patterns to categorize new transactions is classic machine learning functionality.
4	D (Anomaly Detection)	Identifying unusual patterns that don't fit the norm without being explicitly programmed what to look for is the core function of anomaly detection systems.
5	C (LLMs)	Researching standards, generating written drafts, and explaining complex concepts are characteristic applications of large language models. This supports junior staff development and professional work.
6	F (Agentic AI Systems)	Executing a complete multi-step process independently across multiple systems and delivering finished work describes agentic AI capabilities.

1.4: The Historical Development of Intelligence in Accounting

The Evolution of Intelligence in Accounting

To understand the partnership between human and artificial intelligence in accounting today, we need to trace how "intelligence" in accounting has evolved over time. Understanding where we are requires understanding where we've been. The evolution of accounting technology reveals a consistent pattern: each advancement didn't replace human intelligence but rather augmented it, enabling accountants to apply their expertise to increasingly sophisticated problems.

The Manual Ledger Era (pre-1970s)

In the age of manual ledgers, intelligence in accounting meant meticulous arithmetic accuracy and organizational skill. Accountants spent hours performing calculations by hand, double-checking their work, and maintaining orderly records. The "intelligence" was purely human: the ability to add, subtract, and organize information without error. A skilled bookkeeper was valued for their computational accuracy and neat handwriting, literally counting every transaction. This gave the profession the nickname "bean counters." Human intelligence was consumed by mechanical tasks, leaving little time for analysis or strategic thinking.

The Calculator Era (1970s-1980s)

Electronic calculators freed accountants from manual arithmetic. For the first time, intelligence could shift from computation to verification and analysis. An accountant could now quickly test different scenarios or verify complex calculations in seconds rather than hours. The intelligence evolved from "Can I calculate this correctly?" to "What does this number mean, and is it reasonable?" Calculations were still isolated, though. Each computation was independent, requiring manual transfer to written records.

The Spreadsheet Era (1980s-2000s)

Spreadsheets introduced a revolutionary concept involving formulas that automatically recalculated when inputs changed. This created what we might call "formula intelligence," the ability to model relationships between numbers and instantly see the impact of changes. An accountant could build a financial model where changing one assumption, such as a 5% sales increase, automatically flowed through dozens of calculations. Intelligence now focused on building logic, testing scenarios, and asking "what-if" questions. For the first time, accountants could spend their time thinking about business problems rather than performing calculations.

The Cloud Era (2000s-2020s)

Cloud-based systems introduced real-time data feeds and collaborative workflows. Intelligence became distributed across teams and systems. Bank transactions flowed automatically into accounting software. Multiple people could work on the same data simultaneously. The "intelligence" evolved to include data integration, workflow design, and real-time decision-making. Accountants shifted focus from data entry to data interpretation, asking "What story does this real-time information tell about our business?"

The AI Era (2020s-present)

AI is a major step in the development of intelligence available to accountants. As Accenture defines AI as "a suite of advanced technologies enabling machines to perceive, understand, act and learn," today's AI is capable of reading invoices, processing and categorizing thousands of transactions, detecting anomalies, and predicting future cash flow trends. It is something that would take days or weeks to be done by a person.

KPMG's 2024 global AI in Finance report ¹ claims that according to a survey among 2,900 companies in 23 countries, almost

two-thirds of companies pilot or use AI in accounting. But that does not mean that three-quarters of accounting positions would be lost. Accounting professionals have been relieved from mundane data processing tasks to perform those which require unique human skills, such as professional judgments about complex transactions, business context, ethical considerations in case of pressure on financial reporting, and relationship-building with clients.

Today accounting is performed by a collaboration of artificial and human intelligence. Whereas AI can detect patterns in big amounts of data and process information rapidly, human accountants offer judgment in uncertain situations and business context. The evolution of intelligence in accounting has never stopped; it continues to change the very essence of accounting through the emergence of new technologies.

To understand why AI represents such a significant shift in accounting, it helps to trace how the profession's relationship with intelligence and technology has changed over time. Figure 1.4.1 illustrates this progression, from hand-recorded ledgers to AI-powered automation.




Figure 1.4.1: Evolution of Intelligence in Accounting. (CC BY 4.0; Catherina Wong). [Detailed Description of Figure 1.4.1.](#)

Reference

1. KPMG. *AI in Finance: US Report*. KPMG LLP, Dec. 2024.

1.5: Analyzing Artificial Intelligence Strengths and Weaknesses in Accounting

AI Strengths

AI brings distinct advantages to accounting because it thrives on precision, repetition, and data-driven analysis. Its foremost strength lies in processing speed and capacity. Tasks that would occupy human accountants for days, such as scanning invoices, reconciling transactions, or matching expense data, can be completed by AI in seconds. This scalability makes AI particularly effective for document processing, data extraction, and real-time reporting, where volume and consistency matter most.

Beyond speed, AI demonstrates exceptional pattern recognition. When trained on historical data, it identifies similarities and trends that guide classification and prediction. For instance, after analyzing prior expense reports, an AI system can consistently categorize new transactions based on vendor names, payment descriptions, and typical amounts. This capability not only increases efficiency but also promotes standardization across large data sets, creating uniformity that would be difficult to achieve manually.

Consistency represents another defining strength of AI systems. Unlike human accountants who may tire, lose focus, or experience variations in judgment throughout the workday, AI applies the same criteria uniformly to every transaction. This ensures reproducible accuracy, an essential factor in audit reliability. The consistency extends beyond individual transactions to create predictable, dependable results across entire accounting periods. Moreover, machine learning models demonstrate the ability to improve through feedback. When accountants correct classification errors, the AI updates its internal model, reducing future mistakes and becoming progressively more accurate over time.

AI also enhances risk management through anomaly detection. It can quickly identify unusual activity such as duplicate payments, atypical expense levels, or transactions outside normal business patterns. While human accountants might overlook subtle irregularities under time pressure or among thousands of routine transactions, AI flags them in real time. This enables faster review and more proactive fraud prevention, transforming risk management from a reactive process into a continuous monitoring function.

AI Weaknesses

Can AI truly understand the context behind financial decisions, or does it only replicate patterns it has seen before?

This question reveals one of AI's most pressing weaknesses in accounting: its inability to grasp meaning beyond data. While AI can detect recurring patterns, it lacks awareness of the unique circumstances that shape business transactions.

Consider a company that regularly pays two thousand dollars each month to a vendor called ABC Equipment Rental. After learning this pattern, the AI system confidently classifies these payments as rental expenses. When the company later makes a single payment of twenty-five thousand dollars to purchase the equipment, the AI still categorizes it as an expense rather than an asset acquisition. A human accountant, noticing the change in payment amount or reading the description, would immediately question the transaction. The difference lies in **context**, something AI does not perceive.

Definition: Context

In this example, context means understanding that the transaction's nature changed from renting to purchasing, which requires different accounting treatment despite involving the same vendor. It encompasses recognizing the business meaning behind the payment rather than just matching familiar patterns in the data.

Another major limitation involves the problem of **hallucination**. Large language models can produce convincing but incorrect explanations, sometimes inventing details that appear authoritative. For example, an AI might propose an accounting treatment that sounds valid but violates **generally accepted accounting principles**. Because of this, professionals must verify every AI-generated explanation and treat the system's confidence as a starting point, not a conclusion.

Definition: Hallucination

Hallucinations are instances when AI generates plausible-sounding information that is actually incorrect or fabricated. Large language models can produce convincing but wrong explanations, sometimes inventing details that appear authoritative.

Definition: Generally accepted accounting principles

Generally Accepted Accounting Principles (GAAP) are the standardized set of accounting rules, standards, and procedures that companies in the United States must follow when preparing financial statements. These principles ensure consistency, comparability, and transparency in financial reporting across different organizations.

AI also struggles with **professional judgment** and critical thinking. It cannot assess materiality, balance competing factors, or apply ethics to ambiguous scenarios. These abilities rely on human reasoning, experience, and regulatory understanding. When complex or novel issues arise, the absence of judgment exposes a critical boundary between automation and expertise.

Definition: Professional Judgement

Professional judgment in accounting refers to the ability to make informed decisions in ambiguous or complex situations where rules don't provide clear answers. It involves weighing competing considerations, assessing whether an issue is material enough to matter in the context of the overall financial statements and considering the ethical implications of different accounting treatment options.

Adaptability presents another weakness. Since AI models learn from historical data, they often lag behind new accounting standards or regulatory updates until retrained. This delayed adjustment can lead to noncompliance and inconsistency during transition periods. Over-reliance amplifies the risk further, especially among junior professionals who may accept AI outputs without adequate scrutiny. When accountants fail to question the system, the technology's blind spots become their own.

Definition: Adaptability

Adaptability in the context of AI in accounting refers to the ability to recognize and respond to changes in accounting standards, regulations, tax laws, or business practices without needing to be completely retrained.

Understanding these limitations underscores why AI should serve as an assistant, not a replacement. It excels at speed, accuracy, and detection, but human professionals remain essential for interpretation, ethical reasoning, and final decision-making. The value of accounting lies not only in processing information but in understanding what that information truly means.

The Human and AI Partnership

From this comparison of strengths and weaknesses, one can clearly see why an effective accounting process involves the use of a partnership approach. First, the AI does the initial processing of the data by collecting the information from the documents, categorizing transactions, and highlighting the unusual transactions. Second, human accountants conduct the review of those flagged transactions, correct the mistakes, and make decisions in ambiguous cases. Third, the AI learns from this interaction and increases the accuracy of its work. As a result, human accountants can concentrate on doing the value-added activities. However, this process works only when professionals have a good understanding of the abilities and limitations of AI.

Ethical and Professional Considerations

The knowledge about the abilities and limitations of AI technologies helps to identify ethical issues in the area of accounting. One should pay attention to such things as data security and privacy. These issues become important because AI needs to access sensitive financial information. Moreover, one should consider where cloud-based solution stores the data, and what is the security level provided by the vendor of the technology.

The other issue relates to bias and fairness. As mentioned earlier, the AI learns from the data. Thus, if the data contains some biases, the AI uses them to make decisions based on the patterns. Finally, professional accountability always remains the key requirement to the professionals regardless of the use of AI. When there are any mistakes made by AI, accountants are responsible for them because they must be accountable for the work done according to the professional standards.

1.6: Detailed Text Descriptions for Chapter 1

Figure 1.1.1: Accounting

The image is a top-down flatlay photograph of items commonly associated with accounting and financial work.

The items depicted:

- A Clipboard holds a notecard on which the word "accounting" is handwritten in casual, rounded lettering.
- A pair of glasses rests at the base of the clipboard.
- Several sheets of paper printed with columns of numbers and financial data are partially visible beneath other items.
- A thick bundle of \$100 bills secured with a paper band labeled "\$10,000" is prominently displayed. The serial number "LB685309271" is visible on the top bill.
- A handheld calculator with function keys including OFF, MRC, M-, M+, √, and % is positioned to the right.
- A pen rests diagonally in the upper area of the image.
- The edge of a lined notepad is visible along the right edge.

Figure 1.4.1: Evolution of Intelligence in Accounting

The image illustrates the evolution of intelligence in accounting through a winding timeline that moves from top to bottom, marking six key technological eras. Beginning with Pre-1970s Manual Ledgers, where intelligence meant hand-counted arithmetic accuracy, the timeline progresses through the 1970s Calculator, which shifted the accountant's role toward verification and analysis, and the 1980s Spreadsheet, which emphasized logic-building and financial modeling. The 2000s introduced Cloud computing, enabling real-time collaboration across teams, while the 2020s brought AI-driven pattern recognition and task automation. The timeline concludes with a look Beyond the present, envisioning a partnership in which AI efficiency is guided by human judgment and strategy.

The stops progress as follows:

1. **Pre-1970s, labeled Manual Ledger.** Before modern technology entered the accounting office, intelligence was entirely human. Accountants were valued for their arithmetic accuracy and meticulous attention to detail, recording every transaction by hand in physical ledgers. A single error could cascade through an entire set of books, making methodical thinking the most essential professional skill.
2. **The 1970s, labeled Calculator.** The introduction of the calculator did not replace the accountant — it freed them. With computation no longer the primary burden, intelligence in accounting shifted upward toward verification and analysis. Accountants could now spend less time adding columns and more time asking whether the numbers made sense.
3. **The 1980s, labeled Spreadsheet.** The spreadsheet era transformed accounting intelligence once again. Tools like Microsoft Excel allowed accountants to build dynamic financial models, test assumptions, and simulate scenarios that would have taken days to calculate manually. Intelligence now meant the ability to think logically, structure data, and anticipate outcomes.
4. **The 2000s, labeled Cloud.** Cloud computing broke accounting out of the individual workstation and into shared, real-time environments. Intelligence became collaborative where multiple team members, clients, and advisors could access and act on the same data simultaneously. This era rewarded accountants who could coordinate information across people, systems, and time zones.
5. **The 2020s, labeled AI.** Artificial intelligence marks the most significant shift yet. For the first time, the technology itself can recognize patterns, flag anomalies, and automate routine tasks without being explicitly programmed for each situation. Intelligence in this era belongs both to the machine and to the accountant who knows how to interpret, direct, and validate what the AI produces.
6. **The final era, labeled Beyond.** The next frontier is not about replacing human judgment but combining it with AI efficiency. The accountants of tomorrow will be valued not for what they can compute or even what they can automate, but for the strategic thinking, ethical reasoning, and contextual insight they bring to decisions that AI cannot make alone. Intelligence, in its most advanced form, will be a partnership.

This page titled [1.6: Detailed Text Descriptions for Chapter 1](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by Catherina Wong.

2: AI IN FINANCIAL AND MANAGERIAL ACCOUNTING

Chapter Overview

Accounting has always been about using numbers to tell the financial story of a business. Artificial intelligence is accelerating that process by automating repetitive tasks, uncovering patterns in large datasets, and generating forward-looking forecasts with greater speed and precision.

This chapter builds a practical understanding of where AI enters the accounting workflow and what it changes. You will first distinguish between financial and managerial accounting as distinct disciplines with different audiences and rules. From there, you will examine specific AI-driven tools in each field, weigh the opportunities against the challenges, and practice using AI in accounting simulations.

Learning Objectives

By the end of this chapter, you should be able to evaluate the impact of artificial intelligence (AI) on automating processes, enhancing decision-making, and addressing challenges in financial accounting and managerial accounting.

- Distinguish the roles of financial accounting and managerial accounting purposes, users, reporting requirements, and impact on business decision-making.
- Demonstrate how AI-driven tools improve financial accounting.
- Evaluate AI's role in managerial accounting.
- Assess the opportunities and challenges AI presents in financial and managerial accounting.
- Create financial and managerial reports using AI.

2.1: Distinguish the roles of financial and managerial accounting purposes

2.2: AI-Driven Tools in Financial Accounting

2.3: AI's Role in Managerial Accounting

2.4: Opportunities and Challenges in AI Applications

2.1: Distinguish the roles of financial and managerial accounting purposes

Why do we need financial information?

What is Accounting?

Context: Before diving deeper into the ledger entries, watch this short, animated primer explaining how daily business transactions build the baseline of accounting data. This video gives a brief explanation of what accounting is through telling a story.

Video Length: 9 minutes and 44 seconds

This interactive element is not available in this format. Visit the link or scan the QR code to access it online.

[View this YouTube Video online](#)



Exercise 2.1.1

Instructions

After watching the video on what accounting is, answer the following questions.

1. What is Accounting?
2. What is Business?
3. What is the Income Statement, and what does it show?
4. What is the Balance Sheet, and what are its three components?
5. What is the fundamental Accounting Equation?
6. How do the Income Statement and the Balance Sheet connect?

Answer available at the end of this section.

Accounting is the Language of Business

All organizations need information from accounting for their operation. Without financial information, it would be impossible to manage cash, equipment, buildings. That is why accounting is considered to be the language of business. Nowadays, there are numerous innovations in this field, and artificial intelligence contributes greatly to this process. However, people do not use financial information equally. It means that all aspects, including types of data and principles of its management depend on who needs it and for what purposes.

External Users and Financial Accounting

An external user refers to people who are neither part of nor work for the company but have an economic interest in the business entity. This includes shareholders assessing whether to invest or sell their shares in the company, banks trying to determine if it would be safe to lend money, government bodies looking at how much to charge the company in taxes, regulators monitoring compliance, and analysts looking at the performance of the firm among others. Since such users do not have access to the company's records, financial statements are prepared. Financial statements have pre-set formats that try to ensure uniformity,

compatibility and objectivity. This field is referred to as **financial accounting**. The use of AI tools has significantly revolutionized the process of analyzing such reports more accurately and quickly.

Internal Users and Managerial Accounting

Internal users are the managers of the business entity themselves, including managers responsible for the allocation of departmental budgets and executives responsible for setting the organization's strategy. Their requirements are quite different. First, they need detailed information concerning particular products, departments or projects. Secondly, they need forecasting information rather than historical information. Finally, they need quick and flexible analysis that is easily adaptable to each particular decision. The branch of accounting discussed above is known as **managerial accounting**. Unlike financial accounting, it is not constrained by any standardized external principles and it aims at being as relevant and useful as possible.

Did you know?

Financial accounting data shared with external users is like a common language spoken across a country, broadly understood and used by many. The data prepared for internal users, often referred to as **managerial accounting** information, however, is more like a mother tongue, intimate, detailed, and tailored specifically to those who live within it.

Tax Accounting

Beyond financial and managerial accounting, there is a third distinct branch worth noting. **Tax accounting** governs how organizations measure and report income to government authorities, operating under its own set of rules that often differ significantly from standard financial reporting. We will explore tax accounting and how AI is reshaping it in a later chapter.

Financial and Managerial Accounting Information

Financial accounting information appears in financial statements that are intended primarily for external use (although management also uses them for certain internal decisions). Stockholders and creditors are two of the outside parties who need financial accounting information. These outside parties decide on matters pertaining to the entire company, such as whether to increase or decrease their investment in a company or to extend credit to a company. Consequently, financial accounting information relates to the company as a whole, while managerial accounting focuses on the parts or segments of the company.

Managerial accounting information is for internal use and provides special information for the managers of a company. The information managers use may range from broad, long-range planning data to detailed explanations of why actual costs varied from cost estimates. Managerial accountants in a company prepare the financial statements. Thus, managerial accountants must be knowledgeable concerning financial accounting and reporting. The financial statements are the representations of management, not the public accounting firm that performs the audit.

Understanding the Difference between Financial and Managerial Accounting Case 2.1.1

Imagine you are the accountant for a growing technology company. On Monday morning, you are preparing financial statements that will be reviewed by the company's bank to determine whether to approve a loan. By Wednesday afternoon, you are analyzing which of three new product lines the company should invest in next quarter. These two tasks may seem similar since they both involve numbers and financial analysis, but they represent fundamentally different branches of accounting with distinct purposes, audiences, and rules.

Financial accounting focuses on reporting what has already happened to people outside the company such as, investors who are considering buying stock, creditors deciding whether to lend money, and regulators ensuring the company follows the law. **Managerial accounting**, in contrast, helps managers inside the company make decisions about the future, for example, which products to develop, whether to expand into new markets, or how to reduce costs without sacrificing quality.

Understanding the difference between these two types of accounting is crucial because artificial intelligence is transforming both fields, but in very different ways.

Your Connection

Think about the last time you applied for a job, internship, or scholarship. You produced two very different kinds of information during that process, and you probably did not even realize it.

The resume or personal statement you submitted was carefully formatted and designed for someone outside your life to read. It followed conventions, invited comparison with other candidates, and had to make sense to a stranger. The notes you kept for yourself looked nothing like that. How much were you hoping to earn? How would the hours affect your classes? Did the role actually match where you wanted to go? That information was flexible, detailed, and shaped entirely around what you needed in order to decide, never intended for anyone else to see.

One type of information was built for an external audience making a judgment. The other was built for you, making a decision. Financial accounting and managerial accounting follow exactly the same logic.

Discussion Questions

1. Why does financial accounting require strict, standardized rules (e.g. US generally accepted accounting principles), while managerial accounting does not? What problems would arise if they swapped rules?
2. The case states that AI is transforming both types of accounting "in very different ways." Based on what you know about each type, predict one specific way AI might be applied in financial accounting and one in managerial accounting.
3. A startup founder says: "I only need one type of accounting and I will report the same numbers to my bank and to my team." Evaluate this statement. What risks does this approach create?

Solution

1. Financial accounting requires standardized rules because its audience who are external users such as, investors, banks, and regulators, cannot verify the information themselves. They depend on consistent, comparable reports across companies to make fair decisions. Without generally accepted accounting principles or international financial reporting standards, every company could present numbers differently, making comparison impossible and enabling manipulation. Managerial accounting has no such requirement because its audience is internal. Management already has access to the underlying data and understands the company's context. Flexibility allows them to design reports that are actually useful whereas a compliance format might hide the exact insight a manager needs. If the two swapped rules, external reporting would become chaotic and untrustworthy, while internal reports would become unnecessarily constrained, potentially slowing down good decision-making.
2. In financial accounting, AI is increasingly used for automated report generation, real-time compliance checking, and fraud detection. For example, AI can scan thousands of transactions to flag anomalies that might indicate errors or misrepresentation before a report is filed. This is work that once required entire audit teams. In managerial accounting, AI's role shifts toward prediction and scenario modeling. A managerial AI tool might analyze sales data, cost trends, and market conditions to forecast which product line will be most profitable next quarter.
3. The founder's approach is risky and conceptually flawed. Financial accounting reports are prepared under strict rules that prioritize comparability and legal compliance which are not necessarily optimized for internal decision-making. A bank reviewing a loan application needs standardized, and verifiable figures. A management team planning next quarter's investment needs granular cost breakdowns, projections, and scenario analyses that would never appear in a standard financial report. Using the same report for both purposes means one of two things: either the external report reveals sensitive strategic information (a competitive risk), or the internal team is making decisions based on oversimplified data (an operational risk). In practice, growing companies separate these functions precisely because each type of accounting becomes more valuable when it is optimized for its intended audience.

Table 2.1.1: Summarizes key differences between financial and managerial accounting

Category	Financial Accounting	Managerial Accounting
Primary Users	External (Investors, banks, regulators)	Internal (Managers, executives)
Time Focus	Past Performance	Present and Future Decisions
Governance	US Generally Accepted Accounting Principles	No Mandatory Standards
Report Format	Standardized Financial Statements	Customized Reports as needed

Category	Financial Accounting	Managerial Accounting
AI Application	Automated reporting, fraud detection, compliance	Predictive analysis, scenario modeling and budget forecasting

Reporting Requirements

Reporting requirements reveal the fundamental difference between the two disciplines. Financial accounting is governed by strict external standards, while managerial accounting gives companies the freedom to report what is most useful internally.

Financial Accounting Reporting Requirements

Financial accounting operates under mandatory frameworks that specify exactly how transactions must be recorded and reported. In the United States, companies follow Generally Accepted Accounting Principles, commonly known as GAAP. International companies typically follow International Financial Reporting Standards, or IFRS. These aren't suggestions or guidelines, they're mandatory rules enforced by regulatory bodies like the Financial Accounting Standards Board and the Securities and Exchange Commission.

These standards dictate everything from how companies must recognize revenue to how they report their assets and liabilities. The core financial statements required include the Balance Sheet, which shows what the company owns and owes at a specific point in time; the Income Statement, which reports revenues and expenses over a period to show profit or loss; the Statement of Cash Flows, which tracks where cash came from and where it went; and the Statement of Stockholders' Equity, which explains changes in ownership value. Companies must also provide extensive footnotes explaining their accounting policies and providing additional details that help users understand the numbers.

Publicly traded companies must have their financial statements audited by independent certified public accountants who verify that the statements fairly present the company's financial position in accordance with GAAP or IFRS. This audit requirement adds another layer of oversight and accountability. The reports are based on historical results, typically published quarterly and annually, and they cannot be changed or customized based on who's reading them. Everyone gets the same information at the same time.

Managerial Accounting Reporting Requirements

Unlike financial accounting, managerial accounting has no mandatory reporting standards. There is no managerial accounting equivalent of GAAP or IFRS. No regulatory agency requires specific reports or formats. No external auditors verify the information. This freedom exists because managerial accounting reports stay inside the company, they're not published to investors or submitted to regulators. Managers can design whatever reports and analyses are most useful for their decisions without worrying about whether they conform to external standards.

This flexibility manifests in several ways. Managers can define costs and revenues differently than GAAP requires if those definitions better support decision-making. For example, while GAAP requires that certain overhead costs be allocated to inventory, a manager might prefer to analyze contribution margin by excluding those allocations to better understand which products are most profitable. Managers can create reports at any time interval that's useful, daily, weekly, or even hourly rather than waiting for quarter-end. They can segment information by product line, geographic region, customer type, or any other dimension that helps them understand the business.

Common managerial accounting reports include budgets showing planned revenues and expenses for the coming period, variance analysis comparing actual results to budgets and explaining differences, cost-volume-profit analysis examining how changes in sales volume affect profitability, and segment reporting showing profitability by product, customer, or region. The emphasis is on future-oriented information that supports planning and control rather than backward-looking historical reporting.

Impact on Business Decision-Making

Accounting information only has value if it leads to better decisions. How that happens differs significantly between financial and managerial accounting.

Financial Accounting Impact

The impact of financial accounting on business decisions is profound but indirect. While financial statements don't tell managers what decisions to make, they powerfully influence what decisions external stakeholders make about the company. Strong financial statements showing consistent profitability and healthy cash flows can help a company secure loans at favorable interest rates, attract new investors willing to pay premium prices for stock, and build credibility with suppliers who may offer better payment terms.

Financial statements that reveal losses, declining sales, or cash flow problems can trigger serious consequences. Banks may demand immediate loan repayment, investors may sell their shares causing the stock price to drop, and suppliers may require cash payment upfront rather than extending credit. This external pressure then forces internal decisions. Management may need to cut costs, sell assets, or restructure operations to restore financial health.

The quarterly and annual reporting cycle also creates its own decision-making pressures. Managers know that investors and analysts closely scrutinize every earnings report, comparing actual results to expectations. This can lead to short-term thinking where managers prioritize decisions that will improve the next quarter's numbers, even if longer-term strategies might create more value. The transparency required by financial accounting also means that companies must disclose competitive information that rivals can use to their advantage.

Managerial Accounting Impact

Managerial accounting information directly drives internal business decisions in ways that financial accounting does not. When a product manager receives a managerial accounting report showing that Product A generates a 45% contribution margin while Product B generates only 18%, she can immediately decide to shift marketing resources toward Product A. When a budget variance analysis reveals that manufacturing costs are running 12% over budget due to overtime wages, the operations manager can take corrective action by hiring additional workers or adjusting production schedules.

Consider a company deciding whether to accept a special order at a discounted price. Financial accounting will eventually record whatever decision is made and report the results to investors, but it provides no guidance about whether to accept the order. Managerial accounting, however, can analyze the relevant costs and revenues, calculating the incremental cost of producing the additional units, comparing that to the offered price, and considering whether accepting the order might damage relationships with regular customers who pay full price. This analysis directly informs whether management should accept or reject the order.

The same pattern holds for virtually every internal decision. Should we make a component in-house or buy it from a supplier? Managerial accounting compares the relevant costs. Should we drop an unprofitable product line? Managerial accounting analyzes whether the product covers its variable costs and contributes to fixed costs, or whether eliminating it would hurt sales of other products. Should we expand capacity? Managerial accounting projects the cash flows and calculates the expected return on investment. The key difference from financial accounting is that these analyses are specifically designed to support the decision at hand rather than to report results according to external standards.

How AI Applications Differ by Accounting Type

Understanding the distinction between financial and managerial accounting becomes especially important when considering how artificial intelligence is applied in each context. Machine Learning for pattern recognition, Natural Language Processing for extracting information from documents, and Robotic Process Automation for routine tasks are designed primarily to improve accuracy, ensure compliance with Generally Accepted Accounting Principles (GAAP), and enhance fraud detection. These objectives directly serve financial accounting's purpose of providing reliable information to external stakeholders.

AI applications in managerial accounting, while using similar underlying technologies, serve very different purposes. Machine Learning in managerial accounting might be used to forecast future sales, predict which customers are likely to stop buying, or identify which production processes are most likely to experience quality problems. These are forward-looking, decision-oriented applications rather than backward-looking compliance checks. Natural Language Processing might analyze customer feedback to identify product improvement opportunities or scan supplier contracts to optimize purchasing decisions. Robotic Process Automation might automate the generation of customized reports for different managers rather than standardized financial statements for external users.

The same AI detecting fraudulent transactions for financial reporting purposes might be adapted in managerial accounting to

detect operational inefficiencies, patterns where certain suppliers consistently deliver late, or where specific products experience higher-than-expected returns. The key difference is that managerial accounting AI focuses on supporting internal decisions and improving operations rather than ensuring compliance with external reporting standards. When we study specific AI applications throughout this course, it's essential to recognize whether they're designed for financial accounting's external reporting purpose or managerial accounting's internal decision-making purpose, because that fundamentally shapes how the AI is designed and what value it provides.

Answers

Exercise 2.1.1

1. Accounting is the systematic recording, analyzing, and reporting of financial information. It communicates the financial health of a business to stakeholders, and is often called the "language of business" because it is essential for informed decision-making.
2. The video illustrates business through the story of Claudio, who buys colorful glass plates at €1 each and sells them to tourists for €5 each. Business is the activity of exchanging goods or services for profit which might involve buying, selling, and creating value in the process.
3. The Income Statement is also called the Profit & Loss Statement, the Income Statement shows a business's sales (revenue), costs, and resulting profit or loss over a period of time.
4. The Balance Sheet is a snapshot of a business's financial position at a specific point in time. Its three components are: Assets (resources owned, like cash and inventory), Liabilities (debts owed to others), and Equity (the owner's remaining interest after subtracting liabilities from assets).
5. The fundamental accounting equation is: $\text{Assets} = \text{Liabilities} + \text{Equity}$. Both sides must always balance.
6. Net income from the Income Statement flows into the Equity section of the Balance Sheet. This link reflects how a business grows over time.

This page titled [2.1: Distinguish the roles of financial and managerial accounting purposes](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by Catherina Wong.

2.2: AI-Driven Tools in Financial Accounting

How AI-Driven Tools Improve Financial Accounting

Financial statements exist to inform people outside the company, including investors deciding where to put their money, lenders deciding whether to extend credit, and regulators deciding whether a company is operating within the rules. A single mistake, revenue recorded in the wrong period or a payment posted to the wrong account, can force a restatement, trigger regulatory fines, and erode the trust of everyone who depends on those reports to make consequential decisions.

Accounting software transformed the profession decades ago, and the advances came in waves where spreadsheets replaced ledger books, enterprise systems automated transaction posting, and digital workflows made it possible to close the books in days rather than weeks. Each wave made accountants faster and more consistent, yet none of them changed the fundamental structure of the work itself. The higher the volume of transactions, the greater the demand on human attention, and sustained attention across thousands of repetitive entries is precisely where small errors tend to slip through, not from carelessness, but from the limits of human concentration.

AI in the Financial Reporting Process

AI technologies are now helping to enhance this process at three key stages in the financial reporting process. From 45 peer-reviewed studies that have been reviewed systematically for the purpose of this paper between 2019 and 2024, Awad et al. suggest that the three key factors enhancing accuracy, compliance, and detection of fraud in the process of external financial reporting are Machine Learning (ML), which is trained on historical data and helps to identify the patterns and anomalies in the data; Natural Language Processing (NLP), which helps in reading and understanding the unstructured text found in contracts, statements, and footnotes; and Robotic Process Automation (RPA), which automates the process of performing the rule-based tasks that have to be performed repetitively. This section matches each technology with a particular AI-based tool and a practical example from financial accounting. This section also includes an activity, where free AI-based tools will be used by the readers to see how they function and evaluate any new tool they come across in practice.

Three AI Technologies in Financial Accounting

Financial accounting has been transformed by three distinct AI technologies, each solving a different type of problem. Understanding what each one does and how it works makes it easier to see where AI adds the most value in the accounting process. The Table 2.2.1 introduces all three technologies side-by-side.

Table 2.2.1: Comparison of three AI technologies in accounting and a representative tool for each

AI Technology	How it Works	Primary Use in Financial Accounting	Featured Tool
Machine Learning (ML)	Trains on historical data to recognize patterns. Flags transactions that deviate from what is normal.	Automates balance sheet account reconciliations using ML transaction matching; clears matched items automatically and reports exceptions for human review.	Learn more about BlackLine's use of ML to automate account reconciliations and manage the financial close process.
Natural Language Processing (NLP)	Reads and interprets unstructured text in contracts, filings, and reports.	Extracts financial obligations from lease and revenue contracts; checks that footnote narratives match reported figures.	Learn more about Workiva's use of NLP to streamline financial reporting and ensure consistency across filings

AI Technology	How it Works	Primary Use in Financial Accounting	Featured Tool
Robotic Process Automation (RPA)	Executes rule-based, repetitive tasks exactly as programmed, at high volume.	Processes invoices, posts journal entries, runs bank reconciliations without fatigue or transposition errors.	Learn more about UiPath's RPA platform and how it automates high-volume accounting tasks like invoice processing and journal entries.

Did you know?

Each technology represents a distinct mode of AI operation. Machine Learning is adaptive, it improves as it encounters more data. NLP is interpretive, it reads and understands language. RPA is executional, it follows rules precisely at scale. Together they address the full range of tasks in a financial close cycle.

Machine Learning

The machine learning algorithm (**ML**) learns from historical transactional data how normal transactions look, such as how an average journal entry looks like, how much money vendors invoices have on average, and which accounts are often posted together. When the algorithm understands how normal transactions look, it will detect all transactions and patterns deviating from normality beyond the range of acceptable deviation. The process is known as anomaly detection or identification of abnormal transactions and patterns deviating from statistical normality in a dataset.

This approach differs from the rule-based system which can only detect whatever has been programmed into it. The ML algorithm detects abnormal transactions and patterns even when they have not been specifically programmed in the system as rules. This ability is most useful for detecting errors and frauds in the whole general ledger, the master list of all transactions in the accounting system sorted by accounts where it is impossible for any human team to manually audit every transaction.

Featured Tool - BlackLine 2.2.1

Account reconciliations of the entire balance sheet of a company are automated by the Transaction Matching engine offered by BlackLine using machine learning-based business rules. Accountants no longer need to reconcile individual transactions by comparing each line of entries in the sub-ledger to the ledger manually because BlackLine's matching engine uses business rules to match and clear all such transactions without any manual work. The unmatched items in the balance sheet that could be due to timing difference, missing transaction, or posting error are reported as exceptions that require an accountant's attention. As far as 2025, BlackLine provides its artificial intelligence abilities via the Verity AI platform. It offers a dedicated layer of generative AI called Verity Prepare that streamlines the account reconciliation process by analyzing risky accounts, sourcing documents, and explaining discrepancies in the account before finalizing the process. BlackLine is specifically designed for the financial close process and is employed by many firms to automate close cycles.

Quick Check

A company has 300 balance sheet accounts. BlackLine's ML-enhanced matching engine clears 255 of them automatically. What does this tell you about the other 45?

Solution

They contain items that deviate from patterns the model has learned to expect, for example, timing differences, missing transactions, or posting errors. These are the accounts that need human judgment. The ML doesn't flag them because something is necessarily wrong; it flags them because it can't confirm they're right.

Hartwell Industrial Case 2.2.1

Hartwell Industrial is a hypothetical company created for illustrative purposes. The scenario is representative of how BlackLine is applied in practice. Hartwell Industrial is a mid-size manufacturer that closes its books monthly. At each month-end, the accounting team must reconcile 240 balance sheet accounts which involves comparing each account's general ledger balance to supporting sub-ledger detail or external documentation. Manually, two staff accountants spend approximately three weeks each period preparing, reviewing, and certifying these reconciliations, leaving little time for analysis or error correction before the financial statements are finalized.

When BlackLine is deployed for the month-end close, its rules-driven transaction matching engine (enhanced by machine learning for complex scenarios) processes all 240 accounts and produces the following results in the first period:

Table 2.2.2: Results of BlackLine AI-driven tool deployment in the Hartnell Industrial case

What the System Found	How Machine Learning (ML) Resolved It	Financial Reporting Impact
196 of 240 accounts auto-cleared by Machine Learning (ML) matching (82%)	Transaction amounts, dates, and references in the sub-ledger matched the GL balances exactly, following patterns the model recognized from prior periods.	These accounts require no accountant time; balances are confirmed accurate and the reconciliation is automatically certified in BlackLine's audit trail.
31 accounts flagged as exceptions requiring human review (13%)	Sub-ledger amounts do not match GL balances which may be driven by timing differences, missing invoices, or transactions posted to the wrong account.	Each exception routes to the responsible accountant with the specific unmatched item highlighted; unresolved differences could cause a balance sheet account to be misstated in the financial statements.
13 accounts identified as high-risk due to large unexplained variance (5%)	Variance between GL balance and supporting documentation exceeds materiality threshold; ML model flags these as requiring supervisor sign-off before the account can be certified.	A balance sheet account with a large unexplained variance that is certified without investigation could result in a material misstatement; these items get escalated before financial statements are finalized.

Further Results

The two accountants now focus their time on the 44 accounts that require human judgment which include exceptions and high-risk variances rather than manually preparing all 240 reconciliations. Investigation of the high-risk accounts reveals a \$47,000 accrued liability that had been omitted from the sub-ledger, an error that would have understated liabilities on the balance sheet if not caught before the financial statements were issued.

Automated vs. Manual Reconciliation

What makes this different from manual reconciliation?

BlackLine's Transaction Matching applies two mechanisms. First, accountants configure matching rules that define what a correct match looks like, for example, exact amounts, date ranges, and reference codes. These rules automatically clear accounts where transactions align perfectly, which is the source of most auto-certifications. Second, machine learning handles the complex cases rules cannot resolve: multiple invoices matched to a single payment, small variances from bank fees or currency exchange, or transactions where reference numbers differ across systems. ML learns from how accountants resolved those complex cases manually in prior periods and applies that logic at scale.

The value of rule-based matching is consistency because the same rule runs identically every period. The value of ML is judgment at scale which involves handling exceptions too complex to pre-program. Both are operating in this scenario, and the 82% auto-clear rate reflects primarily the rules engine, with ML extending automation into the harder cases.

Exercise 2.2.1

Background

This activity uses a simple business expense list, the kind any small company might keep, to observe how an AI tool identifies entries that look out of place. The same core skill that a free publicly available AI tool such as, ChatGPT applies here is what enterprise tools like BlackLine apply at scale across thousands of transactions per month.

Table 2.2.3: A sample expense list to be used in this exercise

Date	Day	Description	Amount	Vendor
March 1	Saturday	Office Supplies	\$47.23	Staples
March 3	Monday	Client Lunch	\$89.50	The Capital Grille
March 3	Monday	Client Lunch	\$89.50	The Capital Grille
March 5	Wednesday	Software subscription	\$299.00	Adobe
March 7	Friday	Taxi to Airport	\$38.75	Uber
March 10	Monday	Team Lunch	\$1,200.00	The Capital Grille
March 11	Tuesday	Postage	\$12.40	USPS
March 12	Wednesday	Office Supplies	\$2,000.00	Amazon
March 14	Friday	Client dinner	\$88.90	The Capital Grille
March 15	Saturday	Equipment repair	\$500.00	Unknown Vendor
March 17	Monday	Software subscription	\$299.00	Adobe
March 18	Tuesday	Travel reimbursement	\$47.23	J. Smith
March 20	Thursday	Office Coffee	\$8.45	Starbucks
March 21	Friday	Office Supplies	\$47.23	Staples
March 24	Monday	Consulting fee	\$5000.00	Cash
March 26	Wednesday	Team Lunch	\$67.80	Panera Bread
March 28	Friday	Software subscription	\$299.00	Adobe
March 29	Saturday	Client entertainment	\$450.00	Unknown
March 31	Monday	Office supplies	\$47.23	Staples

Instructions

1. **Look at the data with fresh eyes.** Read through the 20 expense entries above. Using only common sense to spot 3 entries that seem odd and a one-sentence explanation of why each caught your attention. Think along the lines like: Does this amount seem too large or small for what was purchased? Does anything repeat in a way that seems suspicious? Does the day of the week seem unusual for this type of expense?
2. **Ask ChatGPT or any AI tool of your choice.** Go to chat.openai.com. Copy the entire dataset above, then paste it into ChatGPT along with this prompt exactly as written: *"You are reviewing a small business expense list. Identify every entry that seems unusual or worth investigating. For each one, explain in plain language why it stands out, consider things like duplicate entries, amounts that seem out of proportion to the description, payments to unnamed vendors, transactions on weekends, and any amounts that repeat across different entries. Do not assume any accounting knowledge on my part."*
3. **Compare the results.** Compare what ChatGPT flagged to your own list from Step 1. Record: (a) entries the AI flagged that you missed, (b) entries you noticed that the AI did not mention, and (c) any entry where the AI's explanation surprised you or differed from your reasoning.

4. **Ask a follow-up question.** Choose one entry ChatGPT flagged and send this follow-up: *"If this turned out to be a legitimate expense, what information would you need to see to confirm it?"* Record what the AI says it would need. This is the AI describing what a human reviewer must verify which is similar to the same function a BlackLine exception workflow performs automatically at scale.

Answer available at the end of this section.

Natural Language Processing

First of all, Natural Language Processing (**NLP**) allows an AI system to work with text rather than numbers. This feature is especially critical for financial accounting since a huge amount of the information that has to be included in financial statements is not stored in a database but rather in various forms: contracts, leases, suppliers' agreements, SEC forms, etc.

There are two tasks that NLP can perform in financial accounting. First, there is an extraction task which means that the tool is used to retrieve the necessary information from the text. It can be thought of as an automatic highlighting mechanism. For instance, when a company signs a lease agreement, this document may be tens of pages long. However, it is necessary only to extract some pieces of information from it: What are the terms of payment? What is the length of the contract? Is the contract renewable? What additional expenses are unpredictable? All of this is critical since according to the accounting standard ASC 842, it became necessary for the companies to disclose their lease obligations in the financial statements: basically, they had to say that "we have this obligation". Before the appearance of ASC 842, the majority of firms did not mention such obligations at all. To sum up: the NLP tool performs the reading, takes necessary numbers and passes them to the accounting system.

Featured Tool - Workiva 2.2.2

Workiva connects financial data directly to the language in SEC filings, annual reports, and footnote disclosures. Its NLP capability checks consistency between numbers in financial tables and the surrounding narrative, catching the common but consequential error where a figure is updated in the table but the corresponding footnote description is not revised to match. Workiva also handles XBRL tagging, where XBRL (eXtensible Business Reporting Language) is the standardized digital format the SEC requires for financial data in public company filings, ensuring that tagged numbers match the narrative disclosures that describe them. This is a specific, high-stakes financial accounting problem that Workiva is purpose-built to address, and it is used by thousands of public companies for SEC reporting.

Quick Check

Why would a company use Workiva instead of just carefully proofreading their own reports?

Solution

Imagine a company's annual report as a living document that hundreds of people are editing at the same time, for example, finance teams updating numbers, lawyers revising language, executives approving sections. Every time a single figure changes, it might appear in a table, a footnote, a summary paragraph, and a digital tag all at once.

A human proofreader can miss the footnote that still says "\$4.2 million" after the table was quietly updated to "\$4.8 million." To the SEC, that inconsistency isn't just sloppy but it's a compliance problem.

Workiva acts as the connective tissue between all those moving pieces. When a number changes in one place, it flags everywhere else that number lives and checks that the surrounding language still matches. It also ensures the digital tags the SEC requires are accurate because regulators don't just read filings, they also process them as machine-readable data. So the answer isn't that companies can't proofread but it's that at the scale and stakes of public SEC reporting, manual proofreading alone isn't reliable enough.

Meridian Logistics Case 2.2.2

Meridian Logistics, a delivery company with trucks and warehouses signed 23 new lease agreements in a single year. Under accounting rules, each one has to be carefully analyzed before it can be recorded in the books.

For a staff accountant doing this by hand, each lease takes 45 minutes to an hour and a half to process. Across 23 leases, that adds up to roughly two full weeks of work which means just reading contracts and filling out forms, before any actual accounting even begins.

What Has to Happen for Each Lease

Think of processing a lease like a five-stop assembly line.

Table 2.2.4: The 5 steps and describes each step to process a lease

Step	What it means in plain English
Read the contract	Find the key details — payment amounts, how long the lease lasts, any renewal options
Classify the lease	Decide which "type" it is — this affects how it shows up in the financial statements
Calculate present value	Figure out what all future payments are worth <i>today</i>
Record it	Write the formal accounting entry that puts it on the books
Check for consistency	Make sure every summary and description of the lease says the same thing

Consider what it takes to process a single lease contract by hand. An accountant must read every page, check each rule, run the calculations, write the journal entry, and then compare documents for consistency. NLP handles each of those steps automatically, leaving the accountant to review and approve rather than build from scratch.

Table 2.2.5: Improvements to lease processes using NLP-based tools such as Workiva

Step	By Hand	With NLP
Read the contract	Accountant reads every page and highlights key terms	Tool automatically pulls out the important details
Classify the lease	Accountant checks each rule one by one	System flags which rules are met — and which aren't
Calculate present value	Accountant runs the math manually	System builds the payment schedule automatically
Record it	Accountant writes the entry from scratch	System drafts it; accountant reviews and approves
Check consistency	Accountant compares documents side by side	System flags any mismatch between numbers and descriptions

The Bigger Win Isn't Just Speed

When 23 leases are processed manually possibly by different accountants, small differences in judgment quietly add up. One accountant might interpret a renewal clause one way; another reads it differently. Multiply that across dozens of contracts and the financial statements start to reflect *who did the work* as much as *what the contracts actually say*.

Exercise 2.2.2

Background

Businesses sign lease agreements when they rent office space, equipment, or vehicles. The contract spells out how much they pay, for how long, and what the rules are. Accountants have to read these contracts carefully to record them correctly in the company's financial books. AI tools are increasingly being used to speed up this kind of reading. In this activity, you will put that to the test: read a short real-world lease yourself, then hand it to an AI and see how the results compare.

The contract you are using is real. It was filed with the U.S. Securities and Exchange Commission (SEC) in 2023 by Co-Diagnostics, Inc., a biotech company in Salt Lake City. In the U.S., public companies must disclose their contracts publicly. You can look up the original filing at sec.gov if you want to see the full document.

Step 1

Read the excerpted lease contract below. You do not need to understand every word, just get a feel for what it says before you answer the questions in step 2.

Commercial Lease Agreement

Filed with the SEC: September 18, 2023

Landlord: Ge Estate, LLC — Sandy, Utah

Tenant: Co-Diagnostics, Inc. — Salt Lake City, Utah

Property: 3215 S 300 W, Salt Lake City, UT (3,200 sq ft warehouse + office)

Section 2 Length of lease.

This lease begins October 1, 2023 and runs for thirty-nine (39) months, ending December 31, 2026.

Section 3 Option to extend.

The tenant may extend the lease for one additional 2-year period by notifying the landlord in writing no later than May 1, 2024. If extended, rent increases by 3% for each year of the extension.

Section 5 Monthly rent.

Year 1 (months 1–12): \$2,500 per month

Year 2 (months 13–24): \$2,575 per month

Year 3 (months 25–39): \$2,652 per month

Rent is due on the first day of each month, paid by direct bank deposit.

Section 6 Who pays what.

This is a gross lease. The landlord pays property tax, building insurance, and utilities. The tenant pays only the monthly rent listed above.

Section 8 Landlord concession.

Provided that Tenant has fully performed all of its obligations through the first six (6) months of the Term and is not then in default, Tenant shall be entitled to an abatement of one (1) monthly installment of Base Rent (the “Abated Rent”). The Abated Rent shall apply to the sixth (6th) full calendar month of the Term. The parties acknowledge that the Abated Rent is a conditional concession and is not a reduction of the Base Rent schedule set forth in § 5. If Tenant is in default and such default results in early termination of this Lease, any Abated Rent shall, at Landlord’s option, be deemed earned and immediately due and payable as additional rent.

Section 17 Ownership.

The tenant may sublet or assign the lease only with the landlord’s written permission. Assigning the lease to another party does not remove the tenant’s obligations. There is no option to purchase the property.

Step 2

Find the key information by yourself. Before asking an AI tool anything, extract the relevant information from the contract and fill in the Table 2.2.6 provided to print. Write your answer in the column on the right. Time how long it takes you to finish the task.

Table 2.2.6: Key lease information for manual extraction with corresponding response fields

What to find	Your Answer
How long is the lease? Count the months from start to end dates.	
What are the monthly payment amounts? List all 3 amounts.	
Is there an option to extend the lease? If yes, how long?	
Does the tenant ever get to own the building? Quote the relevant sentence from the contract.	
What is the total amount the tenant will actually pay out of pocket over the 39-month term? Add up all payments across all three year. Read every clause before calculating	

Step 3

Ask any free publicly available AI tool to perform the same task in step 2. Once you have an AI tool open, paste the lease text, and then paste the prompt below on a new line. Time how long the tool takes to finish the task.

Copy this prompt:

Read this lease agreement and answer four questions: (1) How long is the lease in months? (2) What are the total rent payments over the full lease term? (3) Does the tenant ever get ownership of the property? (4) Is there an option to extend, and if so, on what terms? (5) What is the total rent paid over the full 39 months? Show your work for each answer.

Read the full AI response before moving to Step 4.

Step 4

Compare and reflect on the AI response by answering the following questions.

1. Did the AI tool find the same answers you did? Note any differences.
2. The extension options says the tenant may extend the lease term, it does not say they will. Should the AI tool include those extra 24 months in the lease length? Why or why not?
3. Based on this exercise, where do you think AI is most useful in reading a contract? Where would you still want a human to double check the work?
4. How long did the task take you? How long did the AI take to respond? Now consider: if you had 50 leases to review instead of one, how would that change the case for using AI?

Answer available at the end of this section.

Robotic Process Automation

The concept of Robotic Process Automation (**RPA**) involves the use of software robots (programs) that run through the sequence of programmed actions to perform the tasks just like humans do but consistently regardless of transaction size. It is different from machine learning as it does not have a feedback loop or a process that allows improvement of the program based on experience. RPA runs the program exactly as it was set and produces the same output when executing it 10,000 times as it did with the first transaction. This is exactly why it is useful in financial accounting. As it is known, whenever there is clear rule-based process and high transaction volume, the risk of human mistake is not in the mistake itself but in the accumulation of the mistakes that happen because of a human's inconsistency while performing the process again and again in a time-constrained environment. RPA makes people get rid of this inconsistency, which is why organizations entrust it with execution of the processes where small human mistakes can make a financial statement completely wrong.

Speaking of financial accounting, RPA is most frequently used in accounts payable processing as in this case there is high transaction volume, well defined steps and great importance of the task. Accounts payable is the sum of money that a company

owes outside parties including suppliers, service providers and even landlords for the products and services that were already provided. Every transaction results in an invoice that is to be checked by the accounting team before payment approval. This check is done through the three-way match – a comparison of the three documents that have to be consistent with each other. Purchase order is the first document that contains information about the company's initial permission to buy something with indicated quantity and price. Receiving report confirms delivery of the goods or services. Vendor invoice is the document that indicates the demand of the supplier to pay him/her for the products/services. In case the documents are consistent, the payment process goes further. In the case of inconsistencies someone should investigate them.

In accounts payable, traditionally, the process of comparison of the three documents has been performed manually by an accounts payable clerk. However, in order to automate it, a UiPath robot can be used. It performs the same comparison, extracts necessary information from the three documents, compares them, posts the invoice to an appropriate expense account and cost center in the accounting system, and routes only inconsistencies to a human. Expense account is the account that shows the reason for expense such as utilities, office supplies, professional services, etc. Cost center is the account that shows which department made this expense.

Featured Tool - UiPath 2.2.3

UiPath automates the repetitive, rule-based steps in accounts payable processing that generate the most errors when performed manually at high volume. Its core function in this context is executing the three-way match, the verification step that compares a purchase order, a receiving report, and a vendor invoice before any payment is approved. A UiPath robot extracts the relevant figures from each document, checks them against one another in the accounting system, posts matched invoices to the correct expense account and cost center, and routes any discrepancies to a human reviewer for investigation. This is a specific, high-volume financial accounting problem where RPA delivers consistent results because the rules are fixed and the steps are identical for every invoice processed. UiPath is used by organizations worldwide to reduce processing time and error rates in accounts payable, and its Community Edition makes the same workflow-building environment available to students at no cost.

Quick Check

A mid-sized company processes 600 invoices per month. After deploying a UiPath robot, 510 invoices are processed automatically and 90 are flagged and routed to a human reviewer every month. A manager argues that the robot should be reconfigured to process all 600 automatically so the AP team can be fully reassigned. Why is that a problem?

Solution

Think about what those 90 flagged invoices actually represent. They are not random failures. They are the invoices where something did not add up, where the amount on the vendor invoice does not match the purchase order the company originally approved, where the receiving report shows a partial delivery but the invoice charges for the full order, or where the same invoice number has already been submitted once before.

A human reviewing those 90 cases is not doing leftover work. They are doing the only work in the entire process that actually requires judgment, because each flagged invoice represents a situation where money might leave the organization incorrectly.

If the robot were reconfigured to push all 600 invoices through without stopping, the company would lose its last checkpoint before payment is issued. A duplicate invoice would be paid twice. A vendor who delivered half an order would be paid in full. An invoice that was never approved through a purchase order would be processed without anyone noticing.

The 90 exceptions are not a sign that the robot is underperforming. They are the robot doing exactly what RPA is designed to do, which is handle the routine work at speed and reveal the non-routine work for human decision-making. Removing that final review step does not improve the process. It removes the control that makes the automation trustworthy in the first place.

Pinnacle Retail Case 2.2.3

Pinnacle Retail operates 28 store locations and receives around 600 vendor invoices every month from suppliers, utility companies, and service providers. Each one has to be checked, coded, and recorded before any payment can go out the door.

For an AP clerk doing this by hand, each invoice takes several minutes to process when everything goes smoothly, and significantly longer when something does not match. Across 600 invoices, that adds up to a full week of work every month, and that is before investigating the ones with errors, which means the accounts payable process alone becomes the main bottleneck in Pinnacle's monthly financial close.

What Has to Happen for Each Invoice

Think of processing an invoice like a four-stop checkpoint system where the payment cannot move forward until each stop is cleared.

Table 2.2.7: Illustrates the 4 accounts payable processing steps.

Step	What it means in plain English
Receive and collect	Download the invoice PDF from email or a supplier portal and bring it into the system.
Verify the details	Check that what the vendor is charging matches what the company actually ordered and received.
Code the expense	Assign the invoice to the right spending category so it appears correctly in the financial reports.
Record and log it	Post the entry to the books and create a record that an auditor can follow later.

Table 2.2.8: Contrasts manual and RPA accounts payable processing

Step	By Hand	With UiPath RPA
Receive and collect	AP clerk checks email and the supplier portal each day and downloads PDF files manually.	Robot monitors both continuously and downloads new invoices automatically as they arrive.
Verify the details	Clerk pulls up the purchase order and compares it to the invoice line by line.	Robot performs the three-way match across the invoice, purchase order, and receiving report simultaneously, and flags any amount that does not align.
Code the expense	Clerk selects the spending category from a dropdown based on the vendor type and invoice description.	Robot applies predefined rules consistently, for example routing utility invoices to one account and supply invoices to another, every time without variation.
Record and log it	Clerk enters the journal entry manually and saves a note in a tracking spreadsheet.	Robot posts the approved entry automatically and creates a time-stamped log with the invoice image and all matching results, giving auditors a complete record of every action taken.

Processing time drops from a full week to several hours. Coding errors decrease substantially because the robot applies the same rules to every invoice. The AP clerk's role shifts from data entry to reviewing exceptions, the roughly 8% of invoices each month where the three-way match fails or the coding rule does not clearly apply.

The Bigger Win Is Not Just Speed

When 600 invoices are processed manually, small inconsistencies accumulate quietly. One clerk might code a mixed-use invoice one way while another codes it differently. One might flag a five dollar discrepancy while another lets it through. Multiply those small judgment calls across hundreds of invoices and the financial records start to reflect individual habits as much as company policy.

RPA does not just make the process faster. It makes the process the same every time, which means the financial records reflect one consistent set of rules rather than the accumulated variation of whoever processed the invoices that week.

Exercise 2.2.3

Background

Businesses receive invoices from vendors every time they purchase goods or services. Before any invoice can be paid, an accounts payable clerk must verify that the amount being charged matches what the company originally agreed to pay, check that the same invoice has not already been submitted, and assign the invoice to the correct spending category so it appears in the right place on the company's financial reports. At a company like Pinnacle Retail, this happens for hundreds of invoices every month.

In this activity you will process a small batch of invoices by hand exactly the way an accounts payable clerk would, then use a spreadsheet with built-in formulas to run the same process automatically. The formulas you will use represent the same decision logic that an RPA robot executes in an enterprise accounts payable system. By doing the work both ways, you will be able to see not just what automation does, but where it helps, where it fails, and what it cannot replace.

The invoices below are fictional but structurally accurate. Each one reflects the kind of document a real AP clerk processes daily.

Step 1

Read the invoice batch below. You do not need to understand every detail yet, just get a feel for what each document contains before you answer the questions in Step 2.

Table 2.2.9: Invoice 1001 information

Vendor	Acme Supplies
Invoice Number	INV-1001
Invoice Date	March 1, 2024
Bill To	Pinnacle Retail, Accounts Payable Department
Description	Office supplies, bulk order
Amount Due	\$450.00
Payment Terms	Net 30

Table 2.2.10: Invoice 1002 information

Vendor	City Power Co
Invoice Number	INV-1002
Invoice Date	March 1, 2024
Bill To	Pinnacle Retail, Accounts Payable Department
Description	Monthly electricity, Store 7
Amount Due	\$312.50
Payment Terms	Net 15

Table 2.2.11: Invoice 1003 information

Vendor	BuildRight Inc
Invoice Number	INV-1003
Invoice Date	March 3, 2024
Bill To	Pinnacle Retail, Accounts Payable Department
Description	Emergency HVAC repair, Store 12
Amount Due	\$2,800.00
Payment Terms	Net 30

Table 2.2.12: Invoice 1004 information

Vendor	Metro Cleaning
Invoice Number	INV-1004
Invoice Date	March 5, 2024
Bill To	Pinnacle Retail, Accounts Payable Department
Description	Weekly cleaning services, Store 3
Amount Due	\$175.00
Payment Terms	Net 30

Table 2.2.13: Invoice 1005 information

Vendor	FastFreight LLC
Invoice Number	INV-1005
Invoice Date	March 6, 2024
Bill To	Pinnacle Retail, Accounts Payable Department
Description	Freight delivery, February shipment
Amount Due	\$990.00
Payment Terms	Net 15

Table 2.2.14: Invoice 1006 information

Vendor	TechRent Corp
Invoice Number	INV-1006
Invoice Date	March 7, 2024
Bill To	Pinnacle Retail, Accounts Payable Department
Description	Monthly equipment rental, store tablets
Amount Due	\$640.00
Payment Terms	Net 30

Table 2.2.15: Invoice 1007 information

Vendor	City Power Co
Invoice Number	INV-1007
Invoice Date	March 8, 2024

Vendor	City Power Co
Bill To	Pinnacle Retail, Accounts Payable Department
Description	Monthly electricity, Store 14
Amount Due	\$289.00
Payment Terms	Net 15

Table 2.2.16: Illustrates Pinnacle Retail purchase order records for March 2024

Purchase Order Number	Vendor	Authorized Amount	Expense Category
PO-2201	Acme Supplies	\$450.00	Supplies
PO-2202	City Power Co	\$312.50	Utilities
PO-2203	BuildRight Inc	\$2,500.00	Repairs
PO-2204	Metro Cleaning	\$175.00	Services
PO-2205	FastFreight LLC	\$1,100.00	Shipping
PO-2206	TechRent Corp	\$640.00	Equipment
PO-2207	City Power Co	\$289.00	Utilities

Step 2

Process the invoices by hand. Before using any formulas or tools, work through the invoice batch yourself and fill in the table below. Time how long it takes you to finish.

A purchase order is the company's original written authorization to spend money on something. Before paying any invoice, an AP clerk checks that a matching purchase order exists and that the invoice amount does not exceed it. This is called the three-way match. A general ledger account code is the number the accounting system uses to categorize what the spending was for, the same way a personal budget separates groceries from rent.

Table 2.2.17: Accounts payable template showing columns and rows for invoice data (Invoices 1001-1007)

Invoice Number	Vendor Name	Invoice Amount	Matching PO Amount	Status: Approved, Exceeds PO, or Review Required	Duplicate Invoice? Yes or No	GL Account Code: 6100 for Utilities, 6200 for Repairs, 6300 for Supplies, 6400 for all others
INV-1001						
INV-1002						
INV-1003						
INV-1001						
INV-1004						
INV-1005						
INV-1006						
INV-1007						

Once you have filled in the table, answer these 2 questions before moving on. What is the total dollar amount of invoices you would approve for payment, and how many invoices would you flag for human review?

Step 3

Now run the same process using spreadsheet formulas that replicate RPA logic.

Opening the Application

Go to sheets.google.com and sign in with your Google account. Click the large + (Blank) button to create a new spreadsheet. If using Excel, open Microsoft 365 through your institution portal, click New, and select Blank Workbook.

Open Google Sheets at sheets.google.com or Microsoft Excel through your institution's student license, Microsoft 365 Office. Using Table 2.2.17 as a guide, set up your spreadsheet with the following eight column headers in row 1: Vendor Name, Invoice Number, Invoice Amount, Purchase Order Amount, Expense Category, Status, Duplicate Flag, and GL Account Code.

Understanding the Grid

A spreadsheet is organized like a map. Columns run vertically and are labeled with letters (A, B, C...). Rows run horizontally and are labeled with numbers (1, 2, 3...). A cell is the intersection of a column and row. Cell A1 is the top-left corner, and B3 means column B, row 3. You click a cell to select it, then type to enter data.

Entering the Column Headers

Click cell A1 and type Vendor Name, then press Tab to move right. Continue typing each header and pressing Tab: Invoice Number, Invoice Amount, Purchase Order Amount, Expense Category, Status, Duplicate Flag, GL Account Code. When you finish, click on cell A2 to begin entering data.

Entering a Formula

A formula always begins with an equals sign (=). When the spreadsheet sees =, it knows to calculate rather than display text. Click cell F2, type the formula exactly as shown in the sections below, and press Enter. The cell will immediately display the result of the calculation.

Filling a Formula Down

After entering the formula in F2, click F2 again to select it. Move your mouse to the bottom-right corner of the cell. A small square appears, called the fill handle. Click and drag it down to F9. Release the mouse and each cell will contain its own adjusted copy of the formula. Alternatively, select F2, hold Shift, click F9, then press Ctrl+D on Windows or Cmd+D on Mac to fill down the entire selection.

Invoice and Purchase Order Data Entry

Continue using table 2.2.18 as a template, enter the given invoice data from Tables 2.2.9 to Table 2.2.15 into rows 2 through 9, and the purchase order records for March 2024 found in Table 2.2.16 to fill in the Purchase Order Amount and Expense Category columns. Leave the Status, Duplicate Flag, and GL Account Code columns empty for now.

Three-Way Match Check

For the Status column, click on cell F2 and enter the formula, **=IFS(D2="", "Review Required", C2=D2, "Approved", C2>D2, "Exceeds PO", C2<D2, "approved")**. **<="" strong="">It would be the easiest to copy and paste from this text. Then, fill it down to F9.</d2, "approved">**

This formula checks the purchase order amount first. If no purchase order exists, it returns Review Required. If the invoice matches the purchase order exactly, it returns Approved. If the invoice exceeds the purchase order, it returns Exceeds PO. If the invoice is below the purchase order amount, it returns Approved, though the reflection questions will ask you to think carefully about whether that is always the right decision.

Duplicate Flag

For the Duplicate Flag column, click on cell G2 and enter the formula, **=IF(COUNTIF(\$B\$2:\$B\$9, B2)>1, "Duplicate", "Clear")**, then fill it down to G9.

This formula counts how many times each invoice number appears in the full list. If the same invoice number appears more than once, it returns Duplicate. If it appears only once, it returns Clear.

General Ledger Account Code Lookup

For the GL Account Code column, click on cell H2 and enter the formula, `=IFS(E2="Supplies",6300,E2="Utilities",6100,E2="Repairs",6200,E2="Services",6400,E2="Shipping",6400,E2="Equipment",6400)`, then fill it down to H9.

This formula reads the expense category and assigns the corresponding account code. The final TRUE condition is a catch-all that assigns 6400 to any category not on the predefined list.

Once all three formulas are filled in, filter column F to show only rows that are not marked Approved. Note how many exceptions appear and which invoices they are. Time how long it takes you to set up the spreadsheet and run all three formulas.

Step 4

Compare your manual results to the spreadsheet results and answer the following questions.

1. Did the spreadsheet produce the same Status, Duplicate Flag, and GL Account Code results that you did by hand? Note any differences.
2. INV-1005 from FastFreight LLC shows an invoice amount of \$990 against a purchase order of \$1,100. The formula marked this invoice Approved because the amount does not exceed the authorized limit. Is that the right outcome? What could go wrong if a company automatically approves every invoice that comes in below the purchase order amount?
3. TechRent Corp changes the label on their next invoice from Equipment to Technology Rental. What will the GL Account Code formula return for that invoice, and what does that tell you about who is responsible for maintaining an automated system after it is deployed?
4. How long did the manual process take you? How long did the spreadsheet formulas take once the data was entered? Pinnacle Retail processes 600 invoices per month, not 8. How does that change the case for automation?

Answer available at the end of this section.

What AI Does Better and What it Does Not

The three tools covered in this section each solve a different problem in financial accounting. UiPath handles the volume problem, processing hundreds of invoices through the same decision rules without variation. Workiva handles the consistency problem, connecting numbers across tables, footnotes, and digital tags so that a change in one place is reflected everywhere it needs to be. Machine learning handles the detection problem, identifying patterns and anomalies across datasets too large for any human team to review manually. Together they address the tasks that have historically consumed the most staff time and introduced the most preventable errors in financial reporting.

What none of these systems can do is exercise judgment. An RPA robot cannot decide whether a flagged invoice reflects a vendor error or deliberate fraud. Workiva cannot determine whether an inconsistency in a filing is a drafting mistake or a disclosure that requires legal review. A machine learning model cannot evaluate whether an anomaly it flagged reflects a genuine business event or a problem that needs escalation. The table below shows where that boundary sits in practice.

Table 2.2.19: AI-suitable tasks versus tasks requiring professional judgment

AI Performs Consistently Well	Human Judgment Remains Essential
Applies the same coding rule to every invoice regardless of volume.	Recognizes when a business situation falls outside any existing rule.
Reconciles hundreds of accounts by matching items against defined criteria.	Interprets contracts with ambiguous or contradictory terms.

AI Performs Consistently Well	Human Judgment Remains Essential
Checks numerical consistency across hundreds of footnotes and tables simultaneously.	Understands the business context behind an unusual-looking transaction.
Detects statistical anomalies that fall outside expected patterns.	Determines whether an anomaly represents an error, a fraud, or a legitimate event.
Maintains consistent performance across a full close cycle.	Takes professional and legal responsibility for the financial statements.

The accountant's role in this environment is not diminishing. It is shifting toward the work that requires training, context, and accountability, which is the work that automation reveals but cannot resolve. The practical implication for anyone entering the profession is that familiarity with these tools is no longer optional. Understanding what they do, where their rules break down, and when to override their output is part of what it means to be a competent financial accountant.

Want To Dig Deeper

The [study](#) by Awad et al. exploring the three branches of AI used in external financial reporting is available for further reading.

Answers

Exercise 2.2.1

1. An example of something unusual is the \$1,200 spent for the Team lunch on March 10. Most food entries in this dataset are under \$90, the next closest is an \$89.50 client lunch.
2. ChatGPT most likely would catch the duplicate entry of the client lunch. The identical \$89.50 client lunch at The Capital Grille appears twice on the same date. This is either a data entry error or a duplicate reimbursement submission — both represent a real financial risk.
3. (a) An item that might be missed by a Human Reader can be the exact duplicate entry (Rows 2–3): Reading the list linearly, a first-time reader can overlook an identical row. The AI compares every entry against every other entry simultaneously, making this kind of pattern-match automatic. (b) What an AI would miss is the concentration of Capital Grille visits: Four separate visits to the same fine-dining restaurant totaling nearly \$1,500 may raise questions about a vendor preference relationship that the AI may not highlight unless specifically prompted. (c) The answer to this part depends on the result that was generated by the prompt in question 2.
4. If the consultant fee of \$5,000 turns out to be legitimate, the name of the consultant, scope of services, and agreed fee should be verified so that the \$5,000 can be matched to a specific, documented deliverable or invoice.

Exercise 2.2.2

1. The AI matched my answers for A, B, C, and D. It found the 39-month term, listed all three rent amounts correctly, identified the extension option, and confirmed there is no ownership transfer. For E, however, it got the grand total wrong. The AI calculated $\$30,000 + \$30,900 + \$39,780 = \$100,680$. It read Section 5 carefully but never connected Section 8 to the payment calculation. It found the free-month concession but treated it as a separate note rather than subtracting it from the total.
2. The AI should not automatically include the 24-month extension. The contract says the tenant "may" extend — that word means it is optional, not guaranteed. Just because the option exists does not mean the tenant will use it. To decide whether to include it, you would need information the contract does not provide, for example, has the company said publicly that they plan to stay? Have they made expensive improvements to the building that would make moving costly? Is the rent below market rate, so renewing would be a good deal? None of that is in the lease. The AI included the extension in its answer without explaining why and it simply added 24 months to the lease term. That is a guess, not a judgment. A real accountant would need to speak with company management before making that call.
3. AI is most useful for straightforward reading tasks, for example, finding dates, amounts, and specific clauses. It does this quickly and does not get tired or distracted the way a person might when working through a long document. If a company has 200 leases to review, having AI pull the basic facts from each one would save enormous time. But I would still want a human to check whether one clause changes the meaning of another as we saw with Sections 5 and 8. And I would want a human to make judgment calls: whether the tenant is likely to renew, whether the contract is missing anything important, and whether the numbers make sense given what the company actually does. Those questions require understanding the business, not just reading the document.
4. Steps 1 and 2 took me about 12 minutes. The AI responded in roughly 15 seconds. That speed difference becomes dramatic at scale. If a company had 50 leases to review, my manual reading would take around 10 hours. The AI would finish in under 15 minutes. That is not a small efficiency gain, it is the difference between a task that takes one person a week and a task that takes one person an afternoon to check. But the Section 8 mistake complicates the picture. The AI was fast, but it was also wrong in a way that is not obvious unless you already know the right answer. If I were reviewing 50 leases with AI, I would need a system for spot-checking the totals, otherwise I might accept a \$2,500 error across every single one of them and never notice. Therefore, AI changes the job but certainly does not eliminate it. The accountant's time shifts from reading every word to designing good checks and knowing where errors tend to hide.

Exercise 2.2.3

These are not the only correct answers, especially for the reflection questions, which depend on your reasoning. Use these as a guide for the level of detail and thinking expected.

Table 2.2.18: A sample accounts payable processing chart prepared by a human

Invoice Number	Vendor Name	Invoice Amount	Matching PO Amount	Status	Duplicate?	GL Code
INV-1001	Acme Supplies	\$450.00	\$450.00	Approved	Yes	6300
INV-1002	City Power Co	\$312.50	\$312.50	Approved	No	6100
INV-1003	BuildRight Inc	\$2,800.00	\$2,500.00	Exceeds PO	No	6200
INV-1001	Acme Supplies	\$450.00	\$450.00	Approved	Yes	6300
INV-1004	Metro Cleaning	\$175.00	\$175.00	Approved	No	6400
INV-1005	FastFreight LLC	\$990.00	\$1,100.00	Approved	No	6400
INV-1006	TechRent Corp	\$640.00	\$640.00	Approved	No	6400
INV-1007	City Power Co	\$289.00	\$289.00	Approved	No	6100

The total approved for payment before exceptions is \$2,406.50 across five invoices. Three invoices require human review: the two duplicate submissions of INV-1001 and the BuildRight Inc invoice that exceeded its purchase order.

1. The spreadsheet should match on every row. If there are differences, the most likely cause is a data entry error in the spreadsheet or a capitalization mismatch in the expense category that caused the GL code formula to return 6400 instead of the correct code. For example, if you typed "supplies" in lowercase instead of "Supplies," the formula would not recognize it and would assign the default code instead of 6300.
2. The formula marks it Approved because \$990 is less than \$1,100, and the rule only stops invoices that exceed the purchase order. But approving it without review could still be a problem. If FastFreight delivered the full shipment and only charged \$990, the company might owe them \$1,100 and be underpaying, creating a dispute later. If FastFreight intentionally billed less because part of the shipment never arrived, the receiving report should reflect that and a human needs to verify it. An invoice that comes in lower than expected is not automatically a clean transaction. It may indicate a delivery discrepancy that the three-way match is supposed to catch. A more complete rule would flag invoices that deviate from the purchase order in either direction, not just those that exceed it.
3. The formula returns 6400, the default catch-all code, because the text "Technology Rental" does not appear anywhere in the formula's predefined list. The invoice would be coded to the wrong account and would appear under the wrong spending category in Pinnacle's financial reports. No error message appears and nothing flags the problem unless a human reviews the output. This is one of the most important limitations of rule-based automation: the system does exactly what it was told, and when the real world changes, the rules do not update themselves. Someone at Pinnacle would need to notice the miscoding, trace it back to the formula, and update the rule manually. At scale, across hundreds of vendors who occasionally change their invoice formats, this kind of maintenance becomes a significant ongoing responsibility.
4. Processing 8 invoices manually takes most students between 8 and 15 minutes depending on how carefully they read each document. The spreadsheet formulas, once the data is entered, run instantly. At 600 invoices per month, the manual process would take the equivalent of a full work week. The formula-based process would still take time because someone has to enter the invoice data, but the decision logic runs in seconds across the entire batch. More importantly, the formulas apply the same rules to every single invoice without variation. A human working through invoice 580 at the end of a long day is not making the same quality of decision as they were on invoice 1. The formulas are consistent regardless of volume. But the FastFreight and TechRent examples show that speed without accuracy is not a gain. The real value of automation in accounts payable is not that it eliminates human work entirely. It is that it shifts human attention away from repetitive data processing and toward the exceptions that actually require judgment, which is exactly what the sample data in this activity was designed to surface.

2.3: AI's Role in Managerial Accounting

How AI-Driven Tools Improve Managerial Accounting

Managerial accounting involves internal decision-making by managers, who are responsible for planning, resource allocation, and monitoring performance. The importance of these functions can be distinguished by comparing managerial accounting to financial accounting – the first contact people have with accounting. Financial accounting is a highly regulated field with a well-defined process for preparing reports for third parties, including investors who decide where to invest and lenders who determine whether credit is warranted. Managerial accounting has nothing to do with reporting information to external entities. It is meant for internal use, and it seeks answers to different questions, such as what the firm needs to do next.

The need for accuracy is even greater in managerial accounting compared to financial accounting. While a financial report describes the past for an outsider, a managerial report gives the leaders of a business an overview of its current state and is used as a basis for making important decisions related to future development. Thus, information used in managerial accounting not only describes a business. It directs the business's activities.

The main instrument that managers use for making sense of the provided information and translating it into a strategy is an operating budget. In accounting terminology, a budget is a detailed quantitative plan. It states the expected sales, production costs, number of employees required, cash flows, and other aspects of a business activity during a certain period of time. Each aspect is interrelated, which means that the whole budget is only as good as its base.

That starting point is the sales forecast, which is management's best estimate of how many units the company expects to sell and at what price during the budget period. All other components of the operating budget are built on top of it.

1. **Sales Budget** estimates how much revenue the company expects to generate, based on projected units sold and price per unit.
2. **Production Budget** uses the sales estimate to determine how many units need to be manufactured to meet demand.
3. **Materials Budget** uses the production target to calculate how much raw material needs to be ordered and at what cost.
4. **Direct Labor Budget** determines how many workers need to be scheduled and for how many hours to meet the production plan.
5. **Manufacturing Overhead Budget** estimates the indirect costs of running the production process, such as utilities, equipment maintenance, and supervisory salaries.
6. **Budgeted Income Statement** pulls all of the above together to project whether the company will be profitable during the budget period.

Because each budget depends on the one before it, an inaccurate sales forecast does not create just one problem. It creates a chain of compounding miscalculations that works its way through the entire operating plan before a single unit has been produced.

AI in the Budgeting Process

Artificial intelligence is now improving the budgeting process at three specific points. The same three technologies examined in [Chapter 2.2](#) for financial accounting apply here as well. Machine Learning, Natural Language Processing, and Robotic Process Automation each address a distinct weakness in how organizations traditionally build their operating budgets.

Table 2.3.1: AI technologies, how they work, and applications in managerial accounting budgeting.

AI Technology	How it Works	Primary Use in Managerial Accounting	Featured Tool
Machine Learning (ML)	Trains on historical data to recognize patterns and generate forecasts. Improves over time as more data is processed.	Generates sales forecasts by analyzing dozens of demand drivers simultaneously and flags budget variances automatically.	Explore Datarails ' forecasting and budgeting features to see how AI analyzes variance drivers and supports scenario planning across departments.
Natural Language Processing (NLP)	Reads and interprets unstructured text in reports, articles, and documents.	Analyzes market reports, economic data, and operational narratives to surface budget-relevant signals.	See Planful 's planning and budgeting platform to learn how it uses predictive analytics and scenario modeling to highlight budget-relevant insights from financial data.
Robotic Process Automation (RPA)	Executes rule-based, repetitive tasks exactly as programmed, at high volume.	Collects and consolidates budget submissions from multiple departments without manual re-entry or formatting errors.	Discover UiPath 's finance and accounting automation solutions to see how RPA bots consolidate data from multiple systems and reduce manual work across financial processes.

Machine Learning

A machine learning model is designed to analyze many factors that affect sales at the same time. It studies years of past sales data and looks for patterns between those sales and the external factors listed above. Over time, it learns which combinations of those factors have most reliably predicted demand in the past, and it applies that knowledge to generate a forecast for the future.

The other issue which is solved by using machine learning is one of the most frequent problems related to the budgeting process. The quality of a sales forecast depends on the data used as input and the traditional sales forecast is built based exclusively on the historical sales data. In other words, the sales figures for certain products for previous months and years are used as inputs. Even though the historical sales data is very helpful, there is a lot more than that which impacts the future demand.

In order to better understand why this method does not provide the best forecasts possible, let us examine the factors which determine how many items of each product would be sold in a particular period of time. Those factors may include

- Seasonality, which means the changes in customers' buying behavior based on the time of the year, like more purchases of heating supplies in winter or pre-holiday sales
- Promotion, which means special offers such as price reductions or discounts.
- Various economic indicators, which measure the condition of the entire economy. Those could be unemployment levels, consumers' indexes and so on.
- Changes in pricing of competitors, which mean whether their prices are growing or declining, thus encouraging or discouraging customers to buy from a particular competitor.

Machine learning allows analyzing all of these factors simultaneously. It analyzes the historical sales data and tries to find any correlation between the historical sales figures and the listed external factors. Based on those correlations, it makes predictions of the future demand.

Featured Tool - Datarails 2.3.1

Datarails is a financial planning and analysis platform that embeds machine learning directly into the budgeting and forecasting process. It connects to a company's accounting systems and enterprise resource planning systems, which are the large software platforms companies use to manage day-to-day operations including purchasing, inventory, payroll, and order processing. By pulling data from all of these sources, Datarails consolidates all financial information into a single model

that updates automatically. Its AI layer analyzes historical sales patterns by category, incorporates economic indicators and promotional schedules, generates rolling forecasts, and flags when actual results deviate from budgeted figures. That variance is then routed to the responsible manager for review.

Rather than replacing the spreadsheet workflow most finance teams already use, Datarails operates within Excel, making it accessible without requiring teams to abandon familiar tools.

Tool Information

Visit [Datarails website](#) for product information and demo resources.

Quick Check

A company deploys Datarails to generate its annual sales forecast. The AI model processes the company's full sales history and flags 4 product categories as requiring manager review before the budget can be finalized. A manager argues that the flags mean the tool is not working correctly and that the forecast should be accepted as is for all categories. Why is that reasoning a problem?

Solution

The flags are not a sign that the tool is underperforming. Datarails generates forecasts automatically for the categories where the data is consistent and the patterns are clear. It flags the remaining categories because the historical data alone is not sufficient to produce a confident estimate, whether because the sales trend is inconsistent, or an economic indicator is pointing in a different direction than the internal sales history suggests. Accepting those forecasts without review means building the budget on figures the model itself indicated were uncertain. The manager's job is not to approve everything the AI produces. It is to investigate the cases the AI cannot resolve, which is what the flags are identifying.

Thornton Goods Case 2.3.1

Thornton Goods is a hypothetical mid-size consumer goods company that sells seasonal home products such as storage organizers, cleaning supplies, and kitchen accessories through retail partners like department stores and home goods chains. This scenario is representative of how Datarails and similar financial planning platforms are applied in practice.

The Problem with the Traditional Approach

At the start of each year, Thornton's finance team builds a sales budget the way many companies do. They take last year's actual sales figures and apply an estimated growth rate across all product categories. For example, if the company sold 10,000 units of a storage organizer last year and management believes the business will grow by 5 percent, the forecast for this year becomes 10,500 units. That same 5 percent assumption is applied to every product category regardless of whether the evidence supports it.

The process takes approximately three weeks and relies heavily on input from regional sales managers who estimate how much they think they can sell based on experience and current customer relationships. In practice, the forecast consistently misses for two reasons.

First, the sales team anchors its estimates to the prior year. If storage organizers sold well last year because the company ran a major promotion, that increase in sales due to the promotional discount is baked into the baseline. When no promotion is planned for next year, the forecast still reflects last year's inflated number, leading to overproduction and excess inventory.

Second, the finance team applies the same growth assumption to all twelve product categories, even though some categories grow faster, some are declining, and some fluctuate heavily by season. Treating them all the same produces a budget that is accurate on average and off for individual category.

How Datarails Changes the Process

When Datarails is deployed, its machine learning features examine four years of the company's own sales records across all twelve product categories. Rather than applying a single growth rate, the system looks for patterns in the data across five specific factors that have historically influenced Thornton's sales results.

- **Product pricing changes**, meaning whether prices were raised or lowered and how that affected unit sales in each category.
- **Promotional schedules**, meaning which months had discount campaigns and how much those campaigns increased purchases.
- **Category-level sales trends**, meaning whether each product category has been growing, declining, or staying flat over multiple years, not just the most recent one.
- **Economic indicators**, specifically consumer confidence scores, which are publicly reported measurements of how willing people feel about spending money on non-essential purchases.
- **Historical seasonal patterns** already present in Thornton's own sales records, meaning which months or quarters consistently produce higher or lower sales for each category.

The system does not replace the finance team's judgment. It processes all five factors simultaneously across four years of data and produces a category-by-category forecast that the finance team then reviews. Where the system finds patterns the team had not accounted for, it flags them for discussion. The table below shows three examples of what the system found in Thornton's first budget cycle with Datarails.

Table 2.3.2: Findings, Datarails resolutions, and impact on Thornton Goods' first budget cycle.

What the System Found	How Datarails Resolved it	Impact on the Operating Budget
3 of 12 product categories consistently generate their highest sales in Q3, a pattern the flat growth rate method had spread evenly across all four quarters.	The system identified the concentration by comparing four years of quarterly sales data for each category separately, rather than averaging across the whole year.	The production budget for those 3 categories is revised to concentrate manufacturing activity before Q3. Raw materials are ordered 6 weeks earlier, avoiding last-minute purchases at premium prices.
Promotional campaigns drive a significantly larger share of sales in 4 categories compared to the other 8, but the manual forecast treated all categories as equally affected by promotions.	The system compared historical sales figures during promotional periods to non-promotional periods for each category and identified which ones show the strongest response.	Quarterly sales estimates for those 4 categories now reflect planned promotional periods rather than treating every quarter as equivalent.
One category shows a consistent three-year decline in annual sales. At the same time, the consumer confidence score signals that shoppers are reducing spending on non-essential items.	The manual forecast had projected flat growth for this category based on the most recent year alone, which happened to be a slightly stronger year than the two before it.	The direct labor budget for that category is reduced for Q2. Rather than scheduling workers against a demand level the data does not support, staff are reassigned to a category the system is forecasting to grow.

What Changes for the Finance Team

The finance team still spends three weeks on the budget. What changes is how that time is used. Previously, the three weeks were spent gathering data, building the forecast from scratch, and following up with sales managers to collect their estimates. With Datarails, the initial forecast is generated automatically. The three weeks are now spent reviewing the results, challenging the assumptions behind them, and adding the business knowledge the system cannot access on its own.

Consider what that means in practice. Datarails can detect that a product category has been declining for three consecutive years based on the sales data. The AI tool cannot distinguish that Thornton's leadership has decided to relaunch that category next quarter with new packaging and a revised price point. A forecast that treats the category as continuing to decline would lead to reduced production, fewer staff scheduled, and less material ordered, at the same time the business is preparing to grow it. Only a manager who understands what the company is planning can catch that gap and adjust the budget accordingly. The data shows the pattern, but the manager needs to bring the context that determines what the pattern actually means for the business going forward.

Quick Check

Thornton's finance team used to spend three weeks building the annual sales forecast manually. After deploying Datarails, the initial model runs in two days. A manager proposes cutting the forecast review period from five days to one day, reasoning that the AI is doing the work and human review time can be reduced. Why might that be a problem?

Solution

The AI produces the forecast faster, but the review period is when managers apply business context the model cannot replicate. A model trained on historical data does not know that Thornton just lost a major retail partner, that a competitor is launching a directly competing product next quarter, or that the company's largest region has a new sales manager still building client relationships. The review period is not a formality. It is where human judgment is added to the model's output. Compressing it to one day does not save time. It removes the step that makes the AI forecast reliable enough to build the rest of the operating budget on.

Note: Adaptive vs. Traditional sales Forecasting

What makes this different from traditional sales forecasting?

What makes this different from a traditional spreadsheet forecast is how the prediction is made. A standard spreadsheet forecast applies a fixed growth percentage to last year's sales figure, for example adding 5 percent across the board because that is what growth looked like the year before. A machine learning model does not use a single fixed formula. Instead, it continuously tests the relationship between dozens of variables and actual sales outcomes and updates its predictions as new information arrives. This approach is called adaptive forecasting because the model adapts its output as conditions change, rather than locking in an assumption at the start of the budget period and holding it fixed for the rest of the year.

Exercise 2.3.1

Quarterly sales data for Thornton Goods have been collected for the four previous years. The data is expressed as sales data in units sold. Management is preparing the sales budget for the following year and wants to identify any unusual patterns before committing to a forecast.

Table 2.3.3: shows Thornton Goods quarterly sales data, Years 1 through 4, expressed in units sold.

Quarter	Year 1	Year 2	Year 3	Year 4
Q1	4,200	4,410	4,630	4,860
Q2	3,800	3,990	4,190	4,400
Q3	5,100	5,360	5,630	8,900
Q4	6,200	6,510	6,840	6,920

1. Look at the data. Using your own observation, identify two quarters that appear unusual compared to the surrounding pattern that you can summarize. Write one sentence for each explaining what caught your attention. Consider whether any quarter breaks a consistent trend or appears disproportionate to the quarters around it.
2. Ask ChatGPT or any free AI tool. Open chat.openai.com. Paste the entire dataset above, then paste this prompt exactly as written: "You are reviewing quarterly sales data for a consumer goods company. Identify any quarters where the figures seem unusual compared to the overall trend. For each one, explain in plain language why it stands out and what might have caused it. Do not assume any accounting knowledge on my part."
3. Compare your results to the AI output. Record: (a) anything the AI flagged that you missed, (b) anything you noticed that the AI did not mention, and (c) any case where the AI's explanation differed from your own reasoning.
4. Ask a follow-up question. Choose one quarter the AI flagged and ask this follow-up question: "If management wanted to include this quarter in next year's sales forecast, what additional information would help determine whether this was a one-time event or a new trend?" Record what the AI says it would need. This is the AI describing what a human budget analyst must verify, the same judgment step that no forecasting model performs automatically.

Answer available at the end of this section.

Natural Language Processing

When the manager begins to compile the sales forecast spreadsheet, the figures on the sheet constitute only part of the story. The second part of the narrative exists in written form and consists of industry reports, press releases, economic forecasts, and trade publications describing the state of the industry outside the enterprise. Such data do not appear in rows and columns. Instead, it must be carefully read in order to find relevant pieces of information which take much time.

Natural Language Processing is the field of artificial intelligence technology that allows a computer to read and comprehend written text in the way humans do. The system possessing such capabilities will process a large number of text documents, analyze their content and provide summarized answers to questions about relevance in minutes instead of hours.

In the case of developing the sales forecast, such technology may prove helpful in practice. Before committing to a forecast, there should be certain answers obtained to the following questions: will the construction industry slowdown in the next year, will people spend less money on non-essential items, did the rival firm reduce its prices recently, etc. The answers to these questions cannot be found within the enterprise's dataset. However, the answers exist in written documents, reports, and briefing which need to be reviewed by somebody. In the vast majority of firms, there are one or two analysts handling all this.

Featured Tool - ChatGPT 2.3.2

ChatGPT is a general-purpose AI assistant developed by OpenAI that can read, summarize, and draw connections across large volumes of written text. In the budgeting context, finance teams use it to process external documents such as industry reports, economic briefings, and trade publications that contain information relevant to sales forecast assumptions. A manager pastes the text of a document or a series of documents into the tool and asks it to identify findings relevant to a specific question, such as whether conditions in a particular sector are expected to improve or weaken in the coming year.

Unlike specialized financial software, ChatGPT does not connect directly to a company's financial data. Its value in this context is reading and synthesizing external written information quickly, a task that would otherwise require an analyst to read each document manually. ChatGPT is free to access at chat.openai.com. No account is required to begin using it, though a free account allows for longer conversations and access to more capable versions of the tool.

One important consideration for organizations is data confidentiality. When text is pasted into a publicly available tool like ChatGPT, that information is transmitted outside the company's internal systems to servers operated by a third party. For documents containing sensitive financial projections, internal strategy, or customer data, many companies choose not to use public AI tools for that reason.

A growing number of organizations address this by deploying AI models privately within their own technology environment. Rather than sending information to an external service, the company installs and runs the AI model on its own servers, so the data never leaves the building. One well-known example is Meta's Llama 3, an AI model that Meta has made available for organizations to download and run on their own systems under a license that permits most business uses. Financial services companies and other regulated industries have adopted this approach specifically because it provides the same text-reading capability as a public tool while keeping all data within the organization's own infrastructure. Meta's official deployment documentation even includes guidance written specifically for regulated industries covering financial services compliance requirements. Setting up a private AI environment of this kind requires technical resources and expertise, but for organizations handling confidential budget information it is a practical and increasingly common solution.

Quick Check

A budget analyst argues that using ChatGPT to read industry reports is unnecessary because the team already reads those same reports. What does this argument overlook about how AI creates value?

Solution

The argument assumes the analyst reads every available report with the same consistency and without any tendency to favor sources that confirm existing assumptions. In practice, an analyst working under time pressure reads a limited number of documents and naturally gravitates toward sources that align with what the team already believes. ChatGPT processes all the documents the analyst provides and returns findings from sources that both support and contradict the current forecast. Its consistency is greater than a rushed human reader, though results can still shift depending on how a question is phrased.

or how certain language appears in the source. The value is not that AI eliminates bias entirely. It is that AI reduces the tendency to skip inconvenient sources, broadens coverage, and slows down the pull toward confirmation without removing it completely.

Ridgeway Manufacturing Case 2.3.2

Ridgeway Manufacturing is a hypothetical company that produces industrial components for the construction sector. Its annual sales budget is built in August, before the construction industry's busy season, using a combination of prior year sales and qualitative input from account managers.

The Problem with the Traditional Approach

In a recent budget cycle, account managers projected 8 percent growth based on current backlog and relationship strength. That estimate came from experience and current customer conversations, which are valuable inputs. What it did not include was any systematic review of what external sources were saying about conditions in the construction sector for the coming year. The finance team accepted the estimate and moved forward with budget preparation.

This is a common pattern in manual forecasting, particularly in organizations where the finance team is small and the planning cycle is compressed. Qualitative input from the sales team is collected, a growth assumption is set, and the budget is built from there. The sales team's judgment is not the issue. What gets missed is the external picture, not because analysts are careless, but because most planning cycles run four to eight weeks, department deadlines stack on top of each other, and reading through industry reports and analyst commentary falls to the bottom of the list when consolidation work is still unfinished.

How ChatGPT Changes the Process

Before accepting the 8 percent estimate, the senior financial analyst gathered four external documents. These included two industry association reports on commercial construction activity, one economic outlook published by a regional bank, and one publicly available summary of a competitor's recent investor presentation. Reading all four manually and extracting the relevant findings would have taken most of an afternoon. Instead, the analyst pasted the text of each document into ChatGPT one at a time and submitted the same prompt for each one.

Across the four documents, ChatGPT identified a consistent pattern the account managers had not raised. Three of the four sources referenced a projected slowdown in commercial construction starts driven by rising financing costs. One report cited a specific estimate of 12 percent fewer new project starts compared to the prior year. A fourth source noted that a direct competitor had recently reduced pricing on a key product line, a signal that demand pressure in that segment was already being felt in the market.

The analyst brought these findings to the budget review meeting with the specific source sentences attached so the team could read them in context. The finance team used the findings to stress-test the 8 percent growth forecast before it was submitted. The production budget, materials budget, and direct labor budget were all modeled at both 8 percent growth and 4 percent growth. Management entered the fiscal year with a contingency plan already in place rather than discovering the need for one mid-year.

What Changes for the Analyst

The analyst's role did not disappear. Choosing which documents to gather, deciding which findings were relevant to Ridgeway's specific situation, and judging how much weight to assign each signal all required human judgment. Processing the written material, however, took the AI seconds per document rather than an afternoon.

This reallocation of effort matters in a practical sense. When reading and note-taking are handled by the AI, the analyst spends the planning cycle interpreting findings and making decisions rather than working through a stack of reports. The work becomes less about processing information and more about acting on it.

Quick Check

Ridgeway's analyst used ChatGPT to process four external documents before the budget review meeting. A colleague suggests the analyst should have processed ten documents instead of four to get a more complete picture. What does this suggestion get right, and what does it miss?

Solution

More sources generally produce a more complete picture of market conditions, and one genuine advantage of an AI tool is that it processes additional documents quickly without adding much time to the process. The more important point, however, is that the constraint in this scenario is not how many documents the AI can read. It is which documents the analyst chose to include in the first place.

If the analyst gathered only four sources, the AI will only reflect those four sources. Selecting the right documents requires knowing which industry associations publish reliable data, understanding which competitor disclosures are meaningful, and recognizing when a source carries a potential bias. These are all judgment calls that happen before the AI is involved at all. The tool is only as useful as the inputs the analyst brings to it.

Note: Automated vs. Traditional Document Review

What makes this different from traditional document review?

What makes this different from traditional document review is how the information is extracted. A standard approach relies on an analyst reading through each source, deciding what is relevant, and manually recording key findings. Natural language processing does not follow a fixed reading checklist. Instead, it allows the analyst to query unstructured text quickly and consistently across a large number of documents. The practical advantage in budgeting is speed and coverage. An analyst can review more sources before the deadline arrives and apply the same question to each one without the fatigue or inconsistency that comes with reading them manually. The tool responds to what the analyst asks, so the quality of the findings still depends on how well the question is framed and which documents were selected to begin with.

Robotic Process Automation

The problem that Robotic Process Automation aims to solve appears even before the budget analysis commences. Entering data into the model is easy, but it can be very laborious. It involves collecting spreadsheets, identifying any omissions, and keying in the data manually, which could take several days before the analysis commences.

Normally, in companies, departmental heads send their financial projections for both income and expenditure in spreadsheets. In turn, a financial management team compiles all the spreadsheets and consolidates them to make one master budget. In an organization that has ten or more departments, it results in several different spreadsheets sent in different formats and at different times.

This entire procedure is undertaken by a Robotic Process Automation robot. It monitors the submission portals and collects all the completed spreadsheet templates. It then identifies missing entries and enters the data into the master budget.

Featured Tool - UiPath 2.3.3

UiPath is a Robotic Process Automation platform used across accounting and finance operations worldwide. In the managerial accounting context, a configured UiPath robot monitors department budget submission portals, retrieves completed templates as they are filed, validates that required fields are present, and writes approved submissions into the master budget model. Any submission that fails the completeness check is automatically flagged and routed back to the submitting manager with a specific note on what is missing. The robot applies the same validation rules to every submission regardless of which department filed it or what time it arrived, which is the core advantage over manual review.

Quick Check

A finance manager asks why the company should invest in setting up a UiPath robot for budget consolidation when the current manual process, though slow, produces acceptable results. What is the strongest argument for automation in this specific context?

Solutions

The argument for automation is not primarily about speed, though time savings are real. The strongest argument is consistency. In a manual process, the person entering data on submission 14 at the end of a long day is not making decisions with the same attention as the person who entered submission 1. Small inconsistencies in how fields are formatted, how missing data is handled, and how categories are coded accumulate across dozens of submissions. Those inconsistencies do not appear as visible errors. They appear as subtle inaccuracies in the consolidated model that are difficult to trace back to their source. An UiPath robot applies the same rules to every submission in identical fashion regardless of volume or timing. The budget reflects one consistent set of standards rather than the accumulated variation of whoever processed each file.

Cascade Foods Case 2.3.3

Cascade Foods is a hypothetical food manufacturing company with 16 operating departments, each of which submits a quarterly budget update to the central finance team. This scenario is representative of how UiPath and similar Robotic Process Automation platforms are applied in budget consolidation.

The Problem with the Traditional Approach

Manually, the budget consolidation process takes the finance team four days. The process begins with collecting templates by email, a task that involves monitoring a shared inbox, downloading attachments, and logging which departments have submitted and which have not. Incomplete submissions require follow-up emails and waiting for responses before consolidation can proceed.

Once all submissions are collected, each one must be reformatted to match the master budget template because different departments use slightly different layouts, labels, and category names. An analyst then re-enters the figures from each submission into the master model by hand. Errors introduced during re-entry are typically discovered during the review period, requiring the process to be partially repeated.

How UiPath Changes the Process

A configured UiPath robot takes over the collection and consolidation steps automatically. The table below contrasts the manual and automated approaches across four steps in the process.

Table 2.3.4: Manual and RPA-assisted approaches across four steps in budget consolidation

Step	Manual Process	RPA-Assisted Process
Collect submissions	Finance team monitors email daily, downloads files manually, and logs receipt in a tracking spreadsheet.	Robot monitors the submission portal continuously and retrieves new files automatically as they are filed.
Check completeness	Analyst opens each submission and visually checks for missing fields, a process that takes 5 to 10 minutes per department.	Robot applies a completeness rule to each submission immediately upon receipt and flags any missing field, routing the incomplete file back to the submitting manager with a specific note.
Reformat and consolidate	Analyst reformats each submission to match the master budget template and manually re-enters figures, approximately 30 minutes per department.	Robot maps each submission to the master template automatically and writes all figures directly into the model. No manual re-entry occurs.
Deliver to finance team	Finance team receives a consolidated model after 4 days, with re-entry errors that surface during review.	Finance team receives a consolidated model within hours of the submission deadline, with a complete log of all inputs and validation results.

What Changes for the Finance Team

The four-day consolidation process is reduced to several hours. More importantly, the errors introduced by manual re-entry are eliminated because no human touches the data between the department submission and the master model. The finance team's time shifts from assembling the budget to analyzing it. That shift changes the nature of the finance team's contribution. In the manual process, the team's skill is applied to a task that a robot can perform more reliably. In the automated process, that skill is applied to reviewing whether the consolidated numbers make sense, identifying departments whose assumptions may be unrealistic, and preparing for the management conversation that follows. The robot handles the mechanics. The finance team handles the interpretation.

Quick Check

Cascade Foods deploys UiPath to collect budget submissions from 16 departments. After the first run, the robot flags 4 departments for incomplete submissions. A manager argues this proves the robot is creating more work, not less. Why is that interpretation incorrect?

Solution

The robot identified 4 incomplete submissions that would have entered the master budget model undetected if processed manually. In a manual process, a missing field or placeholder value often passes through data entry without being caught, producing a budget built on incomplete assumptions. The 4 flags are not extra work the robot created. They are errors the robot caught before they compounded into a flawed consolidated model. The finance team is now correcting 4 known, specific problems rather than troubleshooting an unknown number of silent errors embedded across 16 department submissions.

Note: RPA vs. Manual Budget Consolidation

What makes this different from traditional budget consolidation process?

Traditional budget consolidation is a manual assembly process. A finance team collects spreadsheet submissions, reformats each one to match the master template, and re-enters figures by hand. The process is time-consuming, but the more significant concern is consistency. When dozens of submissions are processed over several days by one or two analysts, small differences in how fields are formatted, how missing data is handled, and how categories are recorded accumulate across the model. These inconsistencies rarely appear as obvious errors. They show up as subtle inaccuracies that are difficult to trace once the consolidated model is built.

Robotic process automation removes human handling from the mechanical steps entirely. A configured robot collects submissions as they arrive, applies the same validation rules to every file, and writes approved data directly into the master model. No reformatting decisions are made differently on the last submission than on the first. The finance team receives a consolidated model built to a single consistent standard, and their time shifts from assembling the budget to evaluating whether the numbers inside it make sense.

Putting it Together

In this section, three technologies that use artificial intelligence and their functions in the management accounting budgeting process were discussed. Machine learning provides better estimates for forecasting the sales because it can analyze multiple demand drivers at the same time rather than making one growth assumption for all products. NLP enables faster review of external documents due to querying written information and using a uniform set of questions on a greater number of documents than in case of manual document review. RPA automates the mechanical processes of gathering, validation, and consolidation of departmental budgets from departments. In all three cases, people play an important role. AI is responsible for data processing while the managers and analysts make decisions based on the results.

Answers

Exercise 2.3.1

1. The most visible anomaly is Year 4 Q3: 8,900 units against a three-year pattern of 5,100, 5,360, and 5,630. That is a 58 percent jump in a quarter that had been growing at roughly 5 percent annually. Year 4 Q4 is also worth noting at 6,920, immediately after the spike, suggesting growth slowed sharply following the Q3 surge.
2. ChatGPT will most likely flag the Year 4 Q3 figure. It should also note the Q4 deceleration and may observe that the spike and subsequent slowdown together

suggest a one-time demand event rather than a structural trend shift.

3. The important insight here is that these two anomalies are likely connected. When one quarter is unusually high and the following quarter is unusually low, a common explanation is demand pull-forward, which means sales that would normally have been recorded in Q4 were instead counted in Q3.
4. To determine how to treat Year 4 Q3 in the Year 5 forecast, a manager would need to know whether a bulk order, early promotion, or recording change occurred in Q3 of that year, whether the conditions that caused the spike are expected to repeat, and whether Q4 came in low because demand was genuinely weaker or because it had already been captured in Q3.

2.4: Opportunities and Challenges in AI Applications

Where AI Tools Perform Best

Across financial and managerial accounting, AI tools perform best where the work is high in volume, governed by fixed rules, and structurally consistent. The errors that accumulate in manual accounting processes do not usually arise from a lack of skill. They arise from the limits of human concentration when the same task must be repeated hundreds or thousands of times without variation. AI eliminates that source of error because consistency is precisely what these systems are built to deliver.

Managerial Accounting

In managerial accounting, three tools address the stages where traditional budgeting consumes the most time and generates the most preventable mistakes. Datarails targets the forecasting problem by generating sales estimates from more variables than any manual model can handle. ChatGPT targets the assumptions problem by synthesizing external documents so that the assumptions behind a forecast reflect what outside sources actually say rather than what an analyst expected to find. UiPath targets the data collection problem by assembling the budget model from department submissions without the errors that manual re-entry introduces. Each tool solves a different problem, and together they cover the full span of the budgeting process.

Financial Accounting

In financial accounting, a parallel set of problems is addressed by a different set of tools. UiPath handles the volume of invoice processing that breaks down under manual repetition by executing three-way match verification across hundreds of invoices using the same decision rules every time. Workiva handles the consistency problem by connecting figures across tables, footnotes, and digital tags in complex filings so that a change made in one place is automatically reflected everywhere it appears. BlackLine handles the detection problem by applying machine learning to account reconciliations, automatically clearing matched transactions and routing exceptions to human reviewers.

The Pinnacle Retail case from this chapter demonstrates the principle directly. A UiPath robot processing 600 invoices a month applies identical matching logic to every one of them. An accounts payable clerk working through invoice 580 at the end of a long day is not making decisions at the same quality as on invoice one. The formulas apply the same rules to every invoice without variation, and that consistency is precisely where the value of automation lies.

What AI Cannot Do

What none of these systems can do is exercise judgment. A BlackLine exception report identifies that a reconciliation needs attention. It does not explain whether the discrepancy reflects a timing difference, a posting error, or something more serious. A UiPath robot flags an invoice that fails the three-way match. It cannot determine whether that mismatch is a vendor mistake or deliberate fraud. Workiva catches a figure that appears in a table but not in the corresponding footnote. It cannot determine whether the inconsistency is a drafting oversight or a disclosure that requires legal review. These determinations require context, professional training, and accountability. No algorithm substitutes for them.

AI Tools Performance Boundaries in Accounting Practice

Every task AI handles consistently is one that can be defined in advance and applied without variation. Every task that remains with a human is one where the right answer depends on context that no rule can fully anticipate.

Table 2.4.1: AI-handled tasks versus tasks requiring professional judgment

AI Performs Consistently Well	Human Judgement Remains Essential
Generates sales forecasts by analyzing multiple variables across years of historical data.	Evaluates whether a forecast reflects a real market shift or a data quality problem.
Flags budget variances and routes them automatically to the responsible manager.	Interprets what a variance means in the context of current business strategy and conditions.
Reads and synthesizes external documents systematically, treating all sources equally regardless of whether they confirm or contradict existing assumptions.	Determines which sources are relevant and how much weight to assign each finding.
Consolidates budget submissions from multiple departments without re-entry errors.	Assesses whether department assumptions are realistic against the strategic plan.
Applies the same coding rule to every invoice regardless of volume.	Recognizes when a business situation falls outside any existing rule.
Reconciles accounts by matching transactions against defined criteria at scale.	Interprets contracts with ambiguous or contradictory terms.
Checks numerical consistency across footnotes, tables, and digital tags simultaneously.	Understands the business context behind an unusual transaction.
Detects statistical anomalies across datasets too large for manual review.	Determines whether an anomaly represents an error, a fraud, or a legitimate event.
Performs repetitive tasks at full volume without fatigue or variation.	Takes professional and legal responsibility for financial statements and operational decisions.

What this Means for the Accounting Profession

The accountant's role and the manager's role are not diminishing in this environment. They are being redirected toward work that requires training, context, and judgment. As routine processing is automated, what remains is the harder and more consequential part of the job. A professional must interpret what an anomaly means, evaluate whether an assumption holds, and decide when an automated result should be questioned or overridden. For students entering either profession, developing that capacity is not an optional skill. It is what competent practice looks like when AI is part of the process.

3: APPLICATIONS OF AI IN AUDIT AND TAX

Chapter Overview

Taxation and auditing have traditionally depended on careful human judgment, including reviewing documents, tracing transactions, and interpreting complex and evolving regulations. Today, artificial intelligence is transforming these practices by automating data intensive tasks, identifying patterns and anomalies that may escape human attention, and delivering research and compliance insights with unprecedented speed.

This chapter develops a practical understanding of how AI is integrated into tax and audit workflows and the implications of that integration. It begins by examining the evolution of technology in both fields, from early spreadsheet tools to current AI enabled platforms, incorporating perspectives from accounting and information systems professionals. The chapter then explores AI applications in tax preparation, including automated data collection, deduction identification, and faster tax research. Next, it analyzes how AI is reshaping the audit process, from improving risk assessment and sampling to supporting continuous monitoring.

Learning Objectives

After studying this chapter, you will be able to:

- Analyze how AI is applied across tax and audit practice
- Explain how automation and optimization tools reshape tax preparation
- Assess how AI transforms auditing from risk evaluation to continuous monitoring
- Explain how AI enhances tax preparation through automated data collection and optimization strategies
- Assess the role of AI in transforming auditing processes

3.1: AI's Evolution in Audit and Tax

3.2: AI as a Tax Preparation Assistant

3.3: AI's Role in Auditing

3.1: AI's Evolution in Audit and Tax

A Changing Profession

Most people have an idea about taxes and audits associated with paperwork: piles of receipts, rows of numbers and people who analyze all of this carefully. Until several decades ago, this perception used to be absolutely accurate. However, with the development of technologies, most of routine activities associated with accounting have become fully automated, allowing people to use their skills for judgment and interpretation. This knowledge should be considered while exploring the role of artificial intelligence in the field.

From Paper Ledgers to Connected Systems

The first methods of accounting were completely manual. Accountants calculated the figures manually, referred to various documents, and submitted papers to appropriate government bodies. It was difficult work and it took a lot of time, but there was a track and it could be seen easily.

The introduction of the accounting information system marked a radical change in everything. This kind of a system is an application software used to gather, store and process the financial data of the entire company. A single individual could do what an entire department used to do. The audit trail, automatic payroll and the financial statements were produced through the computerized system.

Table 3.1.1: Evolution of technology in tax and audit practice from the 1980s to the present, with key tools and impacts

Era	Technology	Impact on Tax and Audit Practice
1980s to 1990s	Spreadsheets and desktop software	Programs such as Lotus 1-2-3 and early QuickBooks replaced paper ledgers. Tax preparers moved calculations off paper and onto screens, reducing manual arithmetic errors and speeding up basic recordkeeping.
Late 1990s	Tax preparation software and e-filing	Programs such as TurboTax guided users through returns using question-and-answer formats, making filing accessible to non-professionals. The IRS introduced e-filing, enabling direct digital submission and faster processing of returns.
2000s	Enterprise systems and big data	Enterprise resource planning systems integrated every business function into a single platform, generating large volumes of structured financial data. This information became available for both tax compliance and audit review, reducing the need for manual data gathering.
2010s	Data analytics and cloud platforms	Audit software tools gave professionals the ability to examine entire transaction populations rather than relying on selected samples. Cloud storage enabled real-time collaboration across firms and clients, expanding both the speed and reach of audit work.
Today	Artificial intelligence	AI now handles work once assigned to junior professionals, recognizing patterns across millions of transactions, identifying compliance risks, and producing analytical insights in seconds. Large language models assist with drafting client communications, interpreting regulatory guidance, and answering complex tax queries. Professionals concentrate on judgment-intensive work requiring interpretation, strategy, and ethical accountability.

Expert Systems as a Bridge to AI

One important transition technology was the expert system. Expert systems were computer programs that allowed for decisions making in the manner of an expert person. Questions were asked and the answers helped determine the solution or rule. For example, tax preparation programs would ask whether a person is single or married, whether he owns property or has any dependents, and would make use of these answers to decide which tax rules applied in such cases. The early users of these programs realized that by answering these questions they could file their tax returns without hiring an expert.

Similarly, in audit, the exception approach was used. This involved looking out for items that were out of the ordinary. Transactions that were out of normal range or were entered into the book at abnormal time, or had unusual patterns of activity.

Concepts in Practice

Accounting information systems made it possible for one person to manage tasks that once required an entire team. But that efficiency came with new risks. When a single employee has access to an entire company's financial records, including purchasing, payroll, and accounts receivable, the potential for undetected error or fraud increases. A transaction entered incorrectly looks identical to a correct one to the machine processing it. This is why accountants use the phrase "garbage in, garbage out." The system is only as reliable as the data fed into it. Building controls into computerized systems, rules that catch bad data before it enters, became one of the most important skills in modern accounting and remains essential in an AI-enabled environment.

How AI is Transforming Tax Preparation and the Audit Process

Artificial intelligence is reshaping both tax preparation and auditing by taking over tasks that were once among the most time-consuming parts of a professional's day. While the two functions serve different purposes, they share a common thread: both have historically depended on professionals manually gathering, checking, and organizing large volumes of financial data before any meaningful analysis could begin. AI is changing that starting point in both fields at the same time.

Automation in Tax Preparation

The most remarkable breakthrough in preparing taxes in recent time is automated processing of tax slips. People receive various documents annually related to their employment income, investment income, pensions, tuitions, and many others. In the past, it was necessary for the preparer to go through each document, find the needed numbers, and put them into the respective fields of the return form. This process is time-consuming, monotonous and error-prone.

AI solutions perform this task in an efficient way using techniques of optical character recognition and machine learning. As soon as the client uploads all of his documents, the system automatically recognizes the type of the document, reads the slip, finds the data fields and puts them into the corresponding lines of the return form. For instance, the T4 slip will be recognized based on its format and issuer data, and the information about employment income, tax withheld and benefits will be automatically entered into the proper lines of the return form without any effort of a preparer.

What is even more remarkable, the software performs automatic checks and error detections. After entering all the data from slips, AI tools compare the figures from different slips and raise warnings if something is wrong. For example, it happens when there is inconsistency in income from last years' returns, or the amount of the deduction is greater than the contribution limit allowed. All these checks which were previously done mentally or on a separate spreadsheet are now performed automatically by AI solutions.

AI-Driven Continuous Audit

The traditional audit controls involved the process of periodic human inspection. Managers would inspect the transactions, approve the payments, and verify the records on a timely basis. This approach was effective in small organizations, but as the number of transactions increased, this process started becoming more difficult. Quarterly inspection allows for much time for something to go wrong and remain unnoticed.

The AI resolves the issue of periodic human inspection by changing it into continuous audit procedures. With machine learning technology, the transaction information gets inspected continuously and compared to the already available patterns. Since the process happens continuously, any anomaly will be revealed right away and allows the organization to resolve an issue before the damage grows.

In addition to that, another key feature of the audit process is the principle of segregation of duties to prevent any one person from handling the entire process of financial transactions from beginning till their end. This process is made possible due to automation of the process through computers.

Real World Example 3.1.1

KPMG Clara

KPMG built Clara as an AI-supported audit platform designed to handle the data-intensive work that precedes human judgment. The platform applies machine learning to sort and score large volumes of transactions, flagging those that deviate from expected norms and warrant closer examination. Auditors can then direct their attention to the highest-risk items rather than working through entire ledgers manually. Clara also uses predictive modeling to help identify financial risks before they appear in the formal record, supporting a more proactive approach to audit work.

How AI is Changing the Roles of Audit and Tax Professionals

Each new generation of technology changes the nature of professional responsibilities. While the advent of software changed tax professionals' work to focus on client counseling on effective planning, timing, and deduction, and audits changed into the analysis of control environment and communication of results to the management. AI takes this tendency even further, allowing professionals to delegate tasks of analysis to automated systems to be able to concentrate on other responsibilities and making use of the opportunities which AI gives them.

AI is not an individual technology, but a set of them, including machine learning, predictive analytics, natural language processing, and generative AI. Using these systems, companies are able to perform analytical works of high volume, identify patterns within huge volumes of information, and get compliance reports extremely quickly comparing to manual works.

Instead of spending time collecting and structuring information, these professionals use this time to interpret the obtained results and give advice to the clients, thus making use of the unique ability of humans, which artificial intelligence does not have yet. More and more companies require from the candidates not only the knowledge of basics of accounting, but also familiarity with AI technologies.

However, the development of AI brings up some questions related to the responsibility for these systems. Automated systems may ignore some contexts, make incorrect assumptions, or provide information based on false results. Professional skepticism is one of the most useful characteristics which a tax or audit professional should possess.

3.2: AI as a Tax Preparation Assistant

How AI-Driven Tools Improve Tax Preparation

If you have prepared your own tax returns in the past, you probably agree that sorting through documents, reading forms, and inputting numbers correctly is quite a challenging experience. Now try doing that for hundreds of people every year. This is what public accounting firm employees do when it comes to tax preparation.

Tax preparation has always been one of the most resource-intensive activities in the field of accounting. For each customer, a preparer needs to gather and verify source documents (official financial documents containing information about income, interest paid, donations, and much more), extract the required data from each of the source documents, input it in a correct place on a tax form, use tax laws and regulations in order to determine all the available deductions or credits that may help the customer pay less, and provide explanations to the customer. At each step, the human element is critically important, but even minor mistakes may lead to costly consequences: a review of a filed document by the IRS – an administrative authority responsible for enforcing the tax code, the payment of penalties, and not getting the refundable tax that was due to the client.

In the process of tax preparation, artificial intelligence helps to improve the efficiency of the two activities mentioned above. First, the collection and organization of source data is done with the help of AI tools. Second, the development of strategies for reducing the amount of tax payable becomes more efficient. Here, the same AI technologies discussed in [Chapter 2](#) will be applied. Machine Learning and Robotic Process Automation will assist in the first case; Natural Language Processing and Generative AI will help in the second one.

Before going further, here is a plain-language reminder of what those terms mean.

- **Machine Learning (ML):** A type of AI that learns patterns from examples. In tax preparation, ML learns to recognize different tax documents and pull the right numbers from them.
- **Robotic Process Automation (RPA):** Software that performs repetitive computer tasks automatically, such as moving data from one program into another without a human doing it by hand.
- **Natural Language Processing (NLP):** AI that reads and understands written human language, including legal documents and tax codes written in complex formal English.
- **Generative AI (GenAI):** AI that produces new text, summaries, or recommendations based on what it has read and learned, much like a knowledgeable colleague drafting an explanation.
- **Certified Public Accountant (CPA):** A licensed accounting professional who has passed a national exam and met state requirements. CPAs are the professionals most commonly responsible for preparing and signing tax returns at accounting firms.

Table 3.2.1: AI technologies, how they work, and applications in tax preparation

AI Technology	How it Works	Use in Tax Preparation	Featured Tool
Machine Learning (ML) / Robotic Process Automation (RPA)	Reads, classifies, and extracts data from documents; executes structured tasks at high volume without manual input.	Collects and organizes client source documents; auto-populates tax return fields for professional review.	Metaya TaxBot
Natural Language Processing (NLP) / Generative AI (GenAI)	Reads tax law in plain language; generates synthesized answers with citations; drafts client-ready documents.	Researches applicable deductions, credits, and planning strategies; generates advisory recommendations tied to a client's specific return.	Thomas Reuters CoCounsel

Automated Data Collection and Return Preparation

The first step in preparing a tax return involves collecting client documents, extracting relevant figures, and entering them into the return form. This is by far the most time-consuming of all preparatory stages. The traditional process of preparing tax returns involves receiving from a client tax documents and analyzing each one, extracting relevant figures and entering them into tax

software.

Source documents used most often include the W-2 (tax document issued by an employer indicating salary and taxes withheld), the 1099 (report issued by a third party reporting income received not from an employer, e.g., from freelance activities or investments), and 1098 (mortgage statement from lenders reporting interest on a home loan). In a company preparing thousands of tax returns every year, the time spent on analyzing and manually entering information into these documents can take up the majority of preparation time.

With the use of machine learning and robotic process automation, this bottleneck can be addressed. An adequately trained ML algorithm can analyze and read tax source documents as well as a tax preparer would do it. It will recognize the type of a document, find relevant fields in it, extract figures and map them to the appropriate line on a tax return. With RPA, the data will be automatically entered into the preparation program.

This shift is similar to what AI tools accomplish in budgeting, as described in [Chapter 2](#). Just as a configured automation tool can collect budget submissions from multiple departments and enter them into a master model, a tax automation platform reads client documents, checks them for completeness, and populates the return to a review-ready state. In both cases, the AI handles the mechanical steps. The professional handles the interpretation.

Featured Tool - Metaya TaxBot 3.2.1

Metaya Tax Bot is an AI tax preparation platform designed by Metaya Consulting specifically for CPA firms. Rather than replacing a firm's existing tax software, such as CCH Axcess, UltraTax, Lacerte, or Drake, Metaya integrates directly with those systems as an intake and preparation layer.

How it Works

There are four steps of the workflow. First, the company uploads clients' source documents to the platform. Then, the artificial intelligence reads and classifies each document, extracts all financial data necessary for tax preparation, considers applicable tax laws, and brings prepared return to the stage when it needs only a review and signature. Third step consists of delivering of prepared returns to the company's existing software for further review by the CPA. Finally, in the last stage, the tax optimization strategy is proposed by the AI tax tool.

Metaya provides services for different types of accounting companies, starting from medium-sized firms having dozens of employees, who prepare hundreds of tax returns every tax season, down to sole practitioners, who serve smaller number of clients in specialized industries. For accounting firms preparing many tax returns every year, the company provides its AI-enabled intake workflow to read and classify financial documents, which are submitted every year by their clients such as salary statements provided by employers, documents related to income from investments, as well as financial information from business partnerships. The platform extracts relevant numbers and provides ready return which needs only a review and signature by the CPA, not requiring any manual typing of numbers by people working in the firm at the stages of data processing.

The AI engine works using the secured cloud infrastructure. At the same time, as it is not a completely cloud-based solution, client data stay on the company's own network, which means that sensitive taxpayer documents are not uploaded to some third party servers, but are processed locally. All sensitive and personal data are anonymized and encrypted before being processed by AI model. This solution is SOC 2 compliant, meaning that it was independently audited according to recognized standard of data protection. In addition, accuracy of processing is guaranteed by independent CPAs who work on the platform.

Note: SOC 2 (Service Organization Control 2) is a certification that proves technology company's compliance with the established standards of protection of handled data. It is a common requirement of businesses before providing such data to vendor.

Westbrook Tax Group Case 3.2.1

The Problem with the Traditional Approach

Westbrook Tax Group is a hypothetical CPA firm with eight full time staff and a client base of approximately 600 individual tax clients. This scenario is representative of how AI preparation platforms are applied in small to mid sized tax practices.

Each filing season, Westbrook's preparers spend the first two weeks of February processing incoming client documents. As W-2s, 1099s, and mortgage statements arrive by email, client portal, and physical mail, staff open each document, manually verify its type, extract the relevant figures, and enter them into Lacerte, the firm's tax software. For a straightforward individual return, this intake and entry process takes an average of 90 minutes. For returns involving multiple investment accounts, self-employment income, or rental properties, it regularly takes three hours or more.

Two problems repeat every season. First, errors introduced during data entry are rarely caught until the review stage, requiring preparers to trace the mistake back to the source document and correct the return. Second, the volume of incoming documents in February creates a backlog that delays completion of returns until late March, even when the underlying tax situations are simple. Clients who file in March receive the same turnaround time as clients with genuinely complex returns, which frustrates both groups.

How Metaya Changes the Process

When the Riveras' documents are submitted to the firm's client portal, the intake and preparation stages are handled automatically.

Table 3.2.2: Manual vs. AI-assisted approaches across key steps in Westbrook's preparation process

Step	Manual Process	AI Assisted Process
Document Intake and Classification	Staff downloads files from email and portal, manually identifies each document type, and routes it to the correct tax schedule by hand. Missing documents are caught only if the preparer thinks to check.	Metaya monitors the client portal, classifies each document automatically, and checks the incoming set against the prior-year return. Missing or inconsistent items are flagged before preparation begins.
Data Extraction and Return Population	Preparer reads each document, identifies the relevant figures, and enters them by hand into tax software. Transcription errors surface during review.	The AI reads each document, extracts figures from the correct fields, and writes them directly into the return. No human touches the data between document receipt and a completed, review-ready return.
Flags and Review	Preparer identifies issues during review, rebuilding context from scratch and verifying each figure against source documents. Items requiring tax law knowledge may be missed depending on the preparer's experience.	All items that are highlighted appear together in a single view before any review takes place by the preparer, organized by risk rating along with the basis for law that applies to each issue.
Tax Optimization	Senior staff research applicable deductions and planning strategies case by case. Junior preparers may not recognize when an opportunity exists or know what questions to ask.	The AI scans the completed return against current tax law and suggests optimization opportunities with estimated dollar impact and priority ratings. The preparer evaluates the findings and determines what to recommend to the client.

Tracing a Complex Return Stage by Stage

Marcus and Elena Rivera live in Sacramento, California, and file their federal taxes jointly. They have two dependent children. Their return is more complex than most, because they have income from four different sources: wages from two employers, a small business Marcus owns, and a rental property. This is a great example of Westbrook's typical more complex tax client and where Metaya TaxBot adds the most value.

The Four Stages of Metaya TaxBot

When the Rivera documents arrive at Westbrook's client portal, Metaya begins processing automatically. No preparer touches the data until Stage 3 and even then, they are reviewing a completed return rather than building one from scratch.

Stage 1 Document Intake and Classification

Six different documents are uploaded by Marcus and Elena, including two W-2 statements, the QuickBooks income statement for Creature Comforts, Form 1098 mortgage interest, and the documents for rental condo plus a set of receipts that are bundled together. Under a manual system, a preparer opens each document, recognizes its contents and routes it to the proper tax schedules.

Metaya reads the document and routes the document to the appropriate tax schedule automatically. Two of the W-2 statements are routed to Form 1040 Line 1a. The QuickBooks income statement is identified as a business Profit & Loss and is routed to Schedule C. The 1098 is sent to Schedule A (Itemized Deductions). The rental documents are routed to Schedule E. The bundle of receipts is identified as needing additional work since it is not clear whether it is made up of deductible or non-deductible items without more context.

Metaya reviews the new documents against the prior-year tax return. The lack of a rental depreciation schedule is detected in the prior return. A preparer might have noticed this omission in reviewing the documents against the prior return but only if he thought to compare the new documents against the old return, actually had the prior-year return available, and took the time to do so in the middle of a busy intake process.

Table 3.2.3: Main components of the Stage 1 document intake tab, including labels and their purpose

Interface Label	What it shows	Why it matters
Document Intake	The active tab selected at the top of the Tax Review Center; clicking it reveals all uploaded documents and their classification results.	Tells the preparer they are in Stage 1. The other tabs (Extracted Data, Tax Forms, Flags & Review, etc.) represent later workflow stages — this tab anchors where intake results live.
Uploaded Documents (6)	Lists every file uploaded, with columns for Document Type (AI), Mapped To (tax form/schedule), Confidence score, and Status.	Shows the AI's classification decision for each file in one place — preparer can immediately see whether the AI correctly identified a W-2, Schedule C source, etc., without opening each document.
Review Needed	A Yes/No flag next to each document; "Yes" entries are highlighted with an orange dot, signaling that the AI cannot complete classification alone.	Directs preparer attention only to documents that require human judgment. Documents the AI is uncertain about (e.g., mixed receipts, business P&L), therefore saving time on clear-cut files like W-2s.
Intake Summary	A donut chart and counts showing how many documents were Auto Mapped, Needs Review, or Unmapped, plus a breakdown by document type (W-2, Receipts, etc.).	Gives the preparer an at-a-glance status of the entire intake batch before reading any individual row; it quickly reveals if most documents are clean or if a large portion needs attention.

Stage 2 Data Extraction and Return Population

Once categorization is done, Metaya pulls out the relevant numbers and inputs them in the tax return. With respect to the Riveras, the Schedule C page demonstrates how the income statement of QuickBooks is dissected line by line to find out revenue (\$191,400) and each type of expenses associated with it. The business income form also identifies the net profit at \$62,000. For each page of the tax return, Metaya gives a confidence score. For instance, 96% score implies that the AI reviews the extracted data and judges how confident it is with regard to accuracy of those numbers and categorizations. High score indicates that preparers should pay attention to some other area while a low score highlights that particular area should be carefully reviewed.

The preparer receives a return that already looks like a finished document, assembled and populated and ready for professional review before a single number has been typed by hand.

Table 3.2.4: Main components of the Stage 2 data extraction tab, including labels and their purpose

Interface Label	What it shows	Why it matters
Extracted Data	The active tab at the top of the Tax Review Center, now advanced from Document Intake to show the results of AI data extraction across all schedules.	Signals the preparer has moved into Stage 2. Sub-tabs (Schedule C, Schedule A, Schedule D, Schedule E) let the preparer navigate between each section of the return that was auto-populated from source documents.
Schedule C (Business)	A line-by-line breakdown of every revenue and expense figure extracted from the QuickBooks P&L statement, mapped to the correct Schedule C lines, with net profit calculated automatically at \$62,000.	How AI translates a source document into a tax form, each row corresponds to a real Schedule C line, so no manual data entry or formula work is required from the preparer.
Confidence Score	A percentage rating the AI assigns to its own extraction, along with a plain-language label (High Confidence, Needs Review, etc.) displayed alongside the populated schedule.	The AI model assesses its own certainty. A high score (90%+) tells the preparer to focus attention elsewhere; a low score flags the section for closer examination before accepting the numbers.
AI Notes	A plain-language explanation of where the data came from and which specific line items (Meals & Entertainment, Travel) are flagged with a "Review" tag because they require human judgment on business purpose or allocation.	AI-assisted tools don't just produce numbers randomly, they explain their reasoning and highlight ambiguous items. This reveals the boundary between what AI can resolve automatically and what requires professional judgment.

Stage 3 Preparer Review

All issues in the tax return cannot be addressed through analysis of the documents alone; some issues need to be assessed professionally, taking into account the knowledge of tax laws, business environment, and making some decisions that cannot be made by the artificial intelligence. Metaya highlights eight such issues for the Riveras at four different risk levels; all of which are found in the Flags and Review tab before the preparer works on the return.

The AI categorizes the flags in terms of risk level so that the preparer will know immediately where to focus. There are four high risk items, three medium risk items and one low risk item. Some of the most enlightening issues are three flags in which the AI has already reached a conclusion rather than just posing a question. The fine amounting to \$300 paid to the city was entered under miscellaneous business expenses in QuickBooks, but according to the AI, it is a penalty paid to the government and is therefore non-deductible expense under IRC §162(f). The political contribution of \$250 is listed with the charities on the Schedule A, but since political contributions do not satisfy the requirements of a charity, they should not be included. The most financially valuable flag is the rental property depreciation flag. The condo owned by the Riveras satisfies the requirements for depreciation of 27.5 years, but there is no depreciation claimed in the previous years.

Rather than leaving the preparer to figure out what additional information is needed, the AI also generates a follow-up checklist automatically. By the time the preparer begins review, the list of items to request from the Riveras is already ready to send.

Table 3.2.5: Components of the Stage 3 preparer review tab with labels and their functions

Interface Label	What it shows	Why it matters
Flags and Review	The flags and review tab consolidates every item the AI could not resolve on its own, the items are organized into a single review queue before the preparer touches the return.	Stage 3 signals the transition from AI automation to professional judgment. Everything completed before this stage, the document intake, data extraction, form population has already been handled. This tab is where the preparer's expertise is actually required.

Interface Label	What it shows	Why it matters
Flags and Review Items	A prioritized list of flagged items, each with the tax area affected, risk level, AI finding, CPA action required, and the source document that triggered the flag. Items are sorted high to low risk so the preparer works through the most consequential issues first.	AI does not just flag uncertainty, it has already reached a conclusion ready for human review. For example, a city fine may be identified as non-deductible under IRC §162(f), a political contribution excluded from Schedule A, or missing rental depreciation calculated across multiple unclaimed years. In each case, the preparer must verify the AI's conclusion, not discover it independently.
Risk Summary	A color-coded count of flags by risk level. Each level is represented by a distinct color bar matching the risk badges in the table.	The risk summary gives the preparer an immediate sense of the return's overall complexity before reading a single flag. A return with four high-risk items demands a different level of attention than one with four low-risk items, and this panel communicates that at a glance.
Need from Client	An automatically generated checklist of documents and information the preparer must request from the client.	This illustrates one of the more practically valuable AI output. Rather than leaving the preparer to reconstruct what is missing, the AI compiles the follow-up list directly from the flags. By the time the preparer begins review, the client request is already ready to send.

Stage 4 Tax Optimization

Upon assembling the return and resolving the flags, the preparer executes an optimization query. The AI scans the Rivera return against the existing tax laws and comes up with five areas in which an associate might overlook possible optimizations.

First, Marcus makes a contribution of \$7,000 to his traditional IRA. This contribution is subject to full or partial elimination due to Marcus being covered by a pension plan at work and combined income of the Riveras' exceeding approximately \$162,660. The flagging system alerts the preparer of the need to calculate the exact deduction elimination value based on the Riveras' Modified Adjusted Gross Income that cannot be calculated yet due to the incomplete return.

Second, there is a question related to self-employment tax deduction. Being the owner of his own business, Marcus is required to pay both portions of Social Security and Medicare taxes – altogether around 15.3% of his net business income. He is able to reduce his income by the portion of that tax (approximately \$4,700 after calculating Schedule C). No further documentation from Marcus is needed to execute this deduction.

Third, there is a question of rental property depreciation detected in Stage 1. The potential recovery of this opportunity in Stage 4 is calculated to be between \$8,000 and \$12,000, depending on the depreciable basis of the property that is to be determined by the CPA.

Estimated total potential benefits in all five optimization areas are calculated to be between \$18,700 and \$26,700. The range is used since not all of the opportunities can be quantified without the addition of further information.

Adjusted Gross Income (AGI) is the total income of the taxpayer minus specific deductions. It is used to determine the taxpayer's eligibility for various tax advantages, including the right to apply certain deduction. The AGI is shown on Line 11 of Form 1040. Phase-out is a gradual decrease of the benefit in case if the income exceeds a certain level rather than sudden elimination. In case of Marcus, the IRA deduction begins to disappear if the Riveras' income exceeds a certain limit for married people with workplace pension plans.

Self-employment tax applies to taxpayers who own their businesses. Employees have their employer contributing half of Social Security and Medicare taxes while the self-employed person pays both halves himself. In order to offset the additional expense of the latter, the self-employed individual is allowed to deduct half of this tax from his income before calculating any other deductions.

Table 3.2.6: Main components of the Stage 4 tax optimization tab, including labels and their purpose

Interface Label	What it shows	Why it matters
Optimization Opportunities Identified	A numbered list of five planning opportunities the AI has discovered from the return data, each with a plain-language explanation, an estimated dollar impact, and a priority rating of high or medium.	The list distinguishes stage 4 from the earlier stages. Stages 1 through 3 focused on accuracy while this stage shifts to strategy, showing that AI can also scan a completed return for deductions and planning strategies.
Annual Tax Impact	An estimated range of the total potential tax benefit if all identified opportunities are acted upon, accompanied by a note that the figures require CPA validation and additional documentation before they can be applied.	The impact gives the preparer and client a concrete sense of what is at stake before a single strategy is implemented. The range format, rather than a single number, reflects the AI's acknowledgment that outcomes depend on verification.
Confidence in Analysis	A plain-language confidence rating with a visual bar, indicating that the analysis is grounded in the extracted return data and current tax law. The label distinguishes this as an assessment of the analysis quality, not of the dollar estimates themselves.	The AI model rates its own certainty. High confidence in the analysis does not mean the dollar figures are guaranteed, since those still depend on documentation and Tax Preparer judgment.
Ask follow-up: Ask TaxBot a question about these opportunities	A chat input at the bottom of the panel that allows the preparer to ask follow-up questions about any of the identified opportunities directly within the workflow, without leaving the interface.	This introduces a different mode of AI interaction. The earlier stages used AI to process and classify; this bar enables a conversational back-and-forth where the preparer can probe specific strategies, ask about edge cases, or request clarification on tax law, illustrating how generative AI functions as an on-demand research assistant within a professional tool.

What Changes for Westbrook's Staff

Rivera case exemplifies the way in which tax preparation has been transformed by the assistance provided by the technology. While Metaya does not substitute for the services of a CPA, it alters the process through which the CPA's time is allocated. In the absence of AI assistance, a staff preparer of Westbrook accounting will spend the first part of the process opening and registering seven documents, as well as reading through all of them manually and manually entering all figures in the tax software. Preparing the tax return such as that of the Riveras, who have both a business and a rental property, as well as capital losses, along with an array of mixed receipts, is labor-intensive at this stage.

Metaya prepares the return and populates it before it comes to the attention of the staff preparer of Westbrook. W-2 figures are entered. The Schedule C corresponds to the income statement in QuickBooks. City fine is recognized as non-deductible. The missing depreciation is highlighted. The political contribution is accounted for. The preparer starts analyzing the return from the very beginning instead of spending time entering the data. What is more, the nature of the work has been altered. At this point, the preparer does not have to construct the return. She/he has to review it. The professional judgments are made here as well. The meals should be allocated by their purposes, the vacation to Las Vegas split by business and personal days, and the depreciable basis of the rental property verified. The Riveras' decision regarding their IRAs should also be recommended.

What has not changed is the professional accountability. All tax returns prepared by Westbrook are analyzed and signed by a licensed accountant before filing. AI processes the information, and the professional makes decisions based upon the processed data. The lack of some of the documents due to the client's failure to upload them is a responsibility of the firm.

Reflection Questions

1. The \$300 city fine was listed under miscellaneous business expenses in the QuickBooks report. A Westbrook preparer doing manual data entry might have entered it on Schedule C without a second thought. Why did Metaya flag it instead, and what does that tell you about how the AI reads documents differently from how a person does?

2. The rental property had no depreciation claimed in the prior two years. This was caught in Stage 1 through a prior-year comparison. A Westbrook preparer could catch the same issue by checking the prior return, but this depends on the preparer thinking to do it. What is the cost to the Riveras if it goes uncorrected, and what does this tell you about the difference between what a preparer can do and what they consistently do under time pressure?
3. The optimization query estimated a total tax benefit of \$18,700 to \$26,700 as a range rather than a single number. Why does the AI present a range, and what would have to happen before a Westbrook CPA could give the Riveras a precise figure?
4. Metaya flags items and assembles the return, but a junior Westbrook staff member could still submit the return before a licensed CPA has reviewed and approved it. What kinds of safeguards do you think a system like Metaya would need to prevent that from happening, and why is controlling who can take which action just as important as what the AI can do?

Tax Optimization and Advisory

The second stage of tax preparation where artificial intelligence generates value is completely different from the first one, as the latter involves data entry and producing a return without errors. At this point, tax planning involves finding applicable deductions, credits and planning strategies based on the analysis of relevant tax law. Deduction is a reduction of gross income from taxation. Tax credit is a direct reduction of tax liability; hence, credits are always better than deductions of the same value. Tax planning strategy is a legal approach to arranging your financial life in order to pay less taxes both at the moment of filing and in the following years. This function was presented in stage 4 of the Metaya Taxbot operation.

Traditional process of finding opportunities for tax savings is time consuming and unevenly distributed throughout a firm. Experienced professionals who are aware of where to search for necessary information can perform this task easily. Junior employees responsible for the majority of clients' returns usually lack experience and cannot conduct deep enough research to find potential optimization possibilities. A client's return prepared by junior employee will be less thoroughly analyzed than a return reviewed by senior partner.

With help of natural language processing and generative AI technologies, optimization research becomes a much simpler task. Tax AI research platform analyzes client's return and financial information, searches for applicable provisions in its authoritative database and presents the professional with optimization strategies based on those provisions. The professional then reviews the reasoning of AI and makes a decision based on his or her understanding of client's particular circumstances.

Featured Tool - CoCounsel 3.2.2

CoCounsel is an AI platform developed by Thomson Reuters, a global information and technology company. It is designed specifically for tax, audit, and accounting professionals. Unlike a general-purpose AI assistant, CoCounsel draws exclusively from authoritative professional sources, including the complete U.S. tax code and IRS regulations, the Checkpoint tax research library (Thomson Reuters' comprehensive tax database), and professional standards from the American Institute of Certified Public Accountants (AICPA), the Financial Accounting Standards Board (FASB), and the International Financial Reporting Standards (IFRS) foundation. It also connects to a firm's own internal documents, templates, and prior client files.

A note on those organizations: The AICPA is the national professional body for CPAs in the United States and sets ethical and practice standards for the profession. FASB is the independent organization that establishes accounting rules for companies operating in the United States. IFRS is the international equivalent, providing accounting standards used in many countries outside the United States.

How it Works

In the realm of tax planning, professionals simply input their questions in clear language, just as if they were asking an expert colleague, and get back synthesized answers along with citations which they can verify themselves. For example, in connection with preparing a particular tax return, a preparer can ask about deductions applicable to a home office of a self-employed client, contribution plans that may be pursued considering the client's income level, or an IRS ruling relevant to the case at hand. The CoCounsel system will synthesize the relevant authority and provide a response along with citations.

The important thing that sets it apart from other AI systems is accuracy. Tax practitioners simply cannot rely on an answer that is not correct, as any wrong positions taken on a tax return will result in penalties for the client and liability for the firm. The CoCounsel systems are specifically optimized and continually checked for accuracy in relation to professional tax standards.

Exercise 3.2.1

Analyzing a Return for Optimization Opportunities

This exercise does not require any prior tax knowledge. You will look at a realistic client profile and use both your own judgment and a free AI tool to identify situations that might mean the client is paying more tax than necessary. The goal is to understand how AI helps surface opportunities that even an experienced preparer might overlook.

What is a tax deduction opportunity?

A deduction opportunity exists when a taxpayer has spent money or done something that the tax code allows to reduce their taxable income, but that reduction has not yet been claimed. Common examples include business expenses for self-employed people, home office costs, miles driven for work, and charitable donations. Each of these comes with rules about what qualifies and what documentation is required.

Client Profile for Marcus Okafor, Tax Year 2025

Review the following information collected during the intake stage of Marcus's individual tax return. You do not need to know tax law to complete this exercise. Use common sense and careful reading.

- Employment income: Marcus earns \$87,400 per year working full-time at Ridgeline Software Inc. His employer sends him a W-2 form each year showing his wages and the taxes already withheld from his paychecks.
- Side income: Marcus earns an additional \$14,200 per year doing freelance web design on his own time. He does not receive any official income forms for this work; clients pay him directly into his personal bank account.
- Home office: Marcus has a dedicated room in his apartment (180 square feet out of 1,400 total) that he uses exclusively for his freelance business. He does not use it for personal activities.
- Health insurance: Marcus is covered through his employer's health plan. He does not have a Health Savings Account (HSA).
- Student loans: Marcus paid \$312 in student loan interest this year.
- Charitable giving: Marcus donated \$1,800 in cash to three local nonprofits this year. He does not have receipts or written acknowledgment from the organizations.
- Work travel: Marcus drove 3,200 miles to visit freelance clients this year. He does not have a written mileage log.
- Retirement savings: Marcus contributes \$4,200 per year to his employer's 401(k) retirement plan. He has not made any additional retirement contributions on his own.
- Prior year tax filing: Last year, Marcus filed a simple tax return using the standard deduction. He did not report his freelance income or claim any business expenses. The standard deduction is a fixed amount the IRS allows all taxpayers to subtract from their income instead of listing individual deductions.

Steps

1. Before using any AI tool, read through Marcus's profile carefully. Identify two items that seem like they might represent a missed opportunity, meaning something that could reduce the amount of tax Marcus owes that was not captured in his prior tax return. You do not need to know the specific tax rules. Look for situations where Marcus is spending money related to his work, giving money away, or doing something that might qualify for a benefit. Write one sentence for each item explaining why it caught your attention.
2. Open ChatGPT or any free AI assistant. Paste the client profile above and submit exactly the following prompt. Prompt to submit (copy this exactly): You are a tax preparation assistant reviewing a client profile for a CPA firm. Identify any items in this profile that may represent missed deductions, credits, or optimization opportunities. For each item you identify, explain in plain language what the opportunity is, what information the professional would need to verify before claiming it, and what the risk is if the information is incomplete. Do not assume any tax knowledge on my part.

3. After reviewing the AI's response, record the following three things in your notes. First, identify any opportunity the AI flagged that you did not think of. Second, identify any item you flagged that the AI did not mention. Third, note any case where the AI described a risk differently from how you thought about it. There are no right or wrong answers here. The purpose is to see where your instincts lined up with the AI and where they differed.
4. Choose one item the AI flagged and submit this follow-up prompt. Follow-up prompt (copy this exactly): For this item, what documentation would the client need to provide, and what is the professional's responsibility if the client cannot produce it? Record what the AI says. This response describes what happens after the AI identifies an opportunity. The AI can tell you what documentation is typically required. But the decision about what to do when documentation is missing, and how to communicate that to the client, is a professional judgment that no AI tool can make on the firm's behalf.

Answer available at the end of this section.

Discussion

Professional Responsibility in AI Assisted Tax Preparation

The platforms described in this section perform tasks that previously required direct professional attention, including reading source documents, populating returns, and researching applicable law. As these tools become standard in tax practice, the question of where the professional's responsibility begins and ends becomes more important, not less.

Scenario

The preparer working in a mid-sized accounting firm works on the tax return of his client by using an AI platform. The AI platform scans the source documents of the client, fills up the tax return and marks two items as issues to be checked. After reviewing these two items, the preparer finds them acceptable and the return gets approved. After two months, the client is notified by the IRS about something missing in the tax return. There was one more 1099 document from his small brokerage account that he never uploaded.

Discussion Questions

1. Who bears professional responsibility for the omission, the preparer, the firm, the AI vendor, or some combination? How does your answer change if the firm's engagement letter does not mention that AI tools are used in the preparation process?
2. The AI platform flagged two items but not the missing 1099, because the document was never uploaded. A colleague argues this means the platform failed. Another argues the platform performed exactly as designed. Who is correct, and what does the answer imply about how professionals should structure their document intake processes when using AI?
3. The American Institute of Certified Public Accountants (AICPA) emphasizes the importance of professional skepticism when using AI tools. Professional skepticism means approaching AI-generated output with a questioning mindset rather than assuming it is correct. What would adequate skepticism look like in practice, and how is it different from simply reviewing the flagged items the system identifies?
4. In the managerial accounting context, we saw that AI tools process the information while managers and analysts decide what it means. Apply that principle to tax preparation. What specifically does the AI process, and what decisions remain with the professional? Where in the preparation workflow is that line most difficult to maintain?

Putting It Together

This module presented two innovations driven by artificial intelligence technology pertaining to tax preparation that aligns with the learning objective. The automated data collection innovation, which is based on machine learning and robotic process automation, solves the problem of document intake and return population stage by reading client source documents, getting the necessary figures and producing a ready-to-review tax return for the professional. On the other hand, the tax optimization innovation, which is enabled by natural language processing and generative AI, solves the problem of the research and strategy stage by determining deductions, credits, and planning opportunities for a particular client's return.

The same logic applies to both innovations just like it was presented when speaking about managerial accounting. In other words, the AI technologies process the information while it is the professional who interprets the meaning of this information. While automation saves the professional's time on preparing returns, the professional still reviews all returns prior to filing.

Optimization searches find the pertinent legal provisions but it is the professional who has to decide whether they are applicable, what documents need to be provided by the client, and what position should be recommended.

Answers

Exercise 3.2.1

1. The two most visible opportunities are the home office and the mileage driven for freelance client visits. Both involve spending money directly connected to Marcus's self-employed work. The freelance income itself is also worth flagging since it does not appear to have been reported in the prior year. The charitable donations and student loan interest may also reduce Marcus's taxes, but those opportunities depend on documentation.
2. A general AI assistant will most likely identify the freelance income, the home office, the business mileage, and possibly a self-employed retirement contribution strategy. It may also flag the cash donations and missing receipts as a documentation risk.
3. The important insight in the comparison is that some items represent genuine deductions, while others represent documentation problems that make valid deductions impossible to claim. The AI will often mention both categories together. The professional must distinguish between them and determine what the client needs to produce before the return can be finalized.
4. The professional's responsibility does not end at identifying that an opportunity exists. For the mileage deduction, the IRS requires a contemporaneous log, one maintained at the time of each trip and not reconstructed afterward. If Marcus cannot produce one, the deduction is at risk even if the miles were genuinely driven. The professional must advise the client of this requirement and document the conversation. Relying on AI to identify the opportunity is acceptable. Determining what to do when the documentation is missing requires professional judgment the tool cannot provide.

3.3: AI's Role in Auditing

The Objective of an Audit

For one to comprehend the importance of AI in this scenario, it would be important for him or her to know what an audit is all about. In simple terms, auditing refers to an independent review of an organization's financial records in order to ascertain whether the records accurately represent the true situation on the ground. In an audit, there is a separation between the audited organization and the auditors since the auditors do not work for the audited organization. Instead, the auditors are working for shareholders, creditors, regulators, and the general public who require such information to make sound decisions.

Due to the above responsibility, the need for accuracy is crucial because any oversight by the auditors would lead to huge financial losses for many people. It could even result into the breakdown of trust of financial markets by the general public. This explains why artificial intelligence is now being considered for auditing purposes.

History of Auditing Technology

Audits have been using technology for a very long time, and it is only through an examination of how the field has grown that we can understand just why AI stands out so much compared to what has come before. New technologies always opened up new possibilities while revealing new constraints that future generations must deal with. Technologies used in audits in the past like CAATs (Computer assisted audit techniques), expert systems, early analytics were quite promising but never achieved their full potential. The key differentiating factor about today's AI technologies is the massive leap in power of the algorithms and reduction in cost.

Table 3.3.1 describes the evolution of Technology Use in Auditing.

Era	Development and Significance
1950 – 1960 Mainframes	Large organizations began using computers to automate data processing. Auditors used machines to run calculations faster, but the process remained heavily manual. A computer could add numbers quickly; a human still had to decide what to look at and what it meant.
1970 –1980 CAATs (Computer Assisted Audit Techniques)	CAATs gave auditors the ability to automate data extraction and testing using batch processing systems introduced in the 1970s. Rather than checking records manually one at a time, an auditor could configure queries to run automatically and review the results the following day. Efficiency improved, but auditors were still working from samples rather than complete transaction populations, which meant that errors or fraud outside the selected records remained invisible to the process.
1980 –1990 Spreadsheets & Expert Systems	Microsoft Excel transformed how auditors organized financial data. Early rule-based expert systems were developed to guide auditors through complex decisions. These showed promise but were often questioned because auditors were uncomfortable with tools that did not explain their reasoning, and fixed rules could not adapt to new situations.
2000 –2010 Data Analytics	Enterprise software generated more digital transaction data. Audit teams began using analytics tools to query entire transaction populations for the first time. The concept of continuous auditing which involves monitoring financial activity on an ongoing basis rather than once a year, began to take shape.
2020 – Present Modern AI	AI goes far beyond querying data. Machine learning recognizes patterns and flags deviations automatically. Natural language processing reads contracts and filings to detect inconsistencies. Robotic process automation executes repetitive tasks with perfect consistency. Major public accounting firms are investing heavily in these capabilities.

AI in the Audit Process

Three specific AI technologies are reshaping how audits are conducted. Each addresses a distinct limitation in traditional audit methods, and each is now actively deployed by major public accounting firms worldwide.

Table 3.3.2 shows the 3 types of AI technology and each of its function and use in auditing.

AI Technology	How the AI Works?	Primary use in Auditing	Example Tool
Machine Learning (ML)	Trains on historical transaction data to recognize normal patterns, then flags anything that deviates significantly from that baseline. Improves as more data is processed.	Analyzes entire populations of transactions to detect anomalies, unusual patterns, and potential fraud including some types of fraud that were never seen before.	KPMG CLARA
Natural Language Processing (NLP)	Reads and interprets unstructured text in documents such as contracts, annual reports, emails, and regulatory filings. Often combined with machine learning to extract and classify information.	Cross-checks whether what a company says in its written disclosures matches what its financial data shows; flags contradictions that may indicate manipulation.	Deloitte ARGUS (used NLP and ML)
Robotic Process Automation (RPA)	Executes repetitive, rule-based audit tasks automatically and with consistency, regardless of volume or time of day.	Automates data collection, work paper preparation, compliance checks, and report assembly. This eliminates manual re-entry errors and frees auditors for higher-level work.	UiPath (industry-standard RPA platform used by audit firms)

Machine Learning

Each financial institution deals with huge numbers of transactions. A big bank will process tens of millions of transactions on a daily basis. Regular auditing would only be able to analyze a sample of those transactions, for instance, 2% to 5% of them and make assumptions based on this sample. Anything that happened outside of this sample would be undetectable for the auditors.

Machine learning takes into account each single transaction from the entire data set at once and compares it with the pattern of normal transactions for this company. Machine learning model does not follow a fixed list of already known types of fraud. It understands the unique behavioral pattern of a specific company and signals when this pattern is broken.

There are several approaches to how these models are designed:

- **Supervised learning** is where a model is trained using historical records that have already been sorted into categories for example, payments that turned out to be genuine and payments that turned out to be fraudulent. Once trained, it applies that learned classification to new transactions as they arrive. This approach works best when the organization has a documented history of past fraud cases to draw from.
- **Unsupervised learning** is where there is no pre-sorted categories are needed. The model examines all transactions, identifies which ones share similar characteristics, and highlights any that fall outside every recognizable grouping. It is not constrained by predefined definitions of fraud, therefore, it can detect schemes that have never been seen before.
- **Ensemble learning** runs multiple models simultaneously and combines their outputs. Where one model might be uncertain, another may be confident, and the combined signal reduces the number of incorrect flags on legitimate transactions.

Featured Tool - KPMG CLARA 3.3.1

In 2023, KPMG announced a multi-billion-dollar partnership with Microsoft to expand CLARA's capabilities using Microsoft Azure's cloud infrastructure and AI services. This partnership is specifically designed to deepen AI integration across KPMG's audit and advisory services globally. CLARA is available to KPMG clients in multiple countries and is actively used on large-scale audits where transaction volumes make manual review impractical.

What is CLARA?

KPMG CLARA is an AI-powered audit platform developed and deployed by KPMG, one of the world's four largest accounting firms. The platform uses machine learning to analyze entire company's accounting records, identifying transactions that deviate from normal patterns and routing them to the responsible audit team for investigation. The key distinction CLARA

introduces is coverage. Traditional audit sampling gives an auditor a statistically educated estimate of whether errors exist. CLARA gives an auditor visibility into every transaction, shifting the question from 'does our sample suggest a problem?' to 'here are the specific transactions that need your attention.'

Natural Language Processing

Pure numerical information does not provide a complete picture of the state of a firm's finances. Many of the details necessary to judge a company's situation reside in written form, such as narrative sections of annual reports, wording of contractual agreements, disclosure in footnotes, management letters, and filings. This material may amount to several hundred pages per large business and needs to be read thoroughly enough to cross-check claims against figures, which is extremely time-consuming.

Natural Language Processing (NLP) is the area of artificial intelligence concerned with reading and interpreting the text created by people. In the domain of audit practice, the primary value of NLP technology lies in the ability to perform cross-verification. Cross-verification is the process of verifying that the text written by a business about its revenue performance matches the reality of its figures. Mismatching between the narrative and figures may indicate the presence of problems.

Finally, NLP is applied to the analysis of numerous contracts and e-mails for any linguistic indications of fraud, related party transactions, and side deals that might contradict the terms of financial reporting.

Featured Tool - Deloitte ARGUS 3.3.2

Deloitte developed ARGUS as an AI-powered document review and analysis tool used within its audit engagements. Importantly, ARGUS combines both natural language processing (NLP) and machine learning, it is not purely an NLP tool. Argus is a platform that leverages advanced machine learning techniques and natural language processing to automatically identify and extract key accounting information from any type of electronic document.

How is ARGUS used?

In practice, ARGUS reads a sample of documents to establish patterns, then uses machine learning to identify and visualize items of interest across an entire document population. For lease audit work, which is an area that became critically important when new accounting standards required companies to bring operating leases onto their balance sheets, ARGUS can process thousands of lease contracts and extract key terms, payment schedules, and renewal clauses automatically.

Deloitte has also integrated broader AI capabilities under its Omnia AI platform, which combines multiple AI tools including ARGUS-style document analysis into a unified audit workflow deployed across engagements in numerous countries.

Robotic Process Automation

Not all activities related to an audit require an element of analytical judgment. Much of an audit involves the collection and organization of information. The process of an audit entails collecting supporting documents, reconciling information, formatting the working papers, conducting routine compliance testing, gathering evidence before actually analyzing anything. These activities may be very tedious, but the benefits of using them have nothing to do with creativity or analytical judgment.

This is precisely what RPA was designed for. RPA stands for Robotic Process Automation and refers to a software that is programmed to carry out a certain set of actions in exactly the same manner no matter how many times and how much it has been done. RPA never gets tired, will never interpret a rule differently depending on whether it was the first or the hundredth transaction, and will never accidentally type a wrong figure when entering it into a computer. When performing audit activities in which the primary goal is consistently and completely executing the same process, RPA is superior to human operators.

Featured Tool - UiPath 3.3.3

UiPath is one of the world's leading Robotic Process Automation platforms and is widely deployed across accounting and audit firms to automate high-volume, rule-based tasks. Unlike the proprietary Big Four tools discussed in the ML and NLP sections, UiPath is a general-purpose RPA platform that any organization, including audit firms of all sizes can license and configure for its own workflows.

In an audit context, UiPath robots are configured to handle tasks such as monitoring document submission portals and retrieving completed files as they arrive, applying completeness validation rules to each submission and routing incomplete files back to the submitter with a specific note on what is missing, mapping submitted data to a master template and writing figures directly into a consolidated model without manual re-entry, and generating confirmation logs that record every input and validation result for the audit trail.

All four of the Big Four accounting firms, KPMG, Deloitte, EY, and PwC are enterprise clients or strategic partners of UiPath, integrating its use within large-scale professional services and audit operations globally.

Benefits of Implementing UiPath

A key advantage of RPA in this context is consistency. In a manual process, the analyst reviewing the fifteenth submission of the day is not applying the same level of attention as the analyst who reviewed the first. Small differences in how fields are handled, how missing data is coded, and how categories are labeled accumulate across a large submission set and produce subtle inaccuracies that are difficult to trace once the consolidated model is built. A UiPath robot applies the same validation rules to every submission in identical fashion regardless of volume or timing, eliminating that source of variation.

AI and Fraud Detection

Of all the applications of AI in auditing, fraud detection may carry the highest stakes and represent the area of greatest improvement over traditional methods. Financial fraud causes substantial harm because it causes monetary losses to investors and creditors, reputational damage to institutions, and erosion of public trust in financial markets.

Traditional audit approaches were well-suited to many tasks, but fraud detection at scale exposed a specific limitation. A carefully constructed scheme spread across thousands of transactions could fall entirely outside the sample and go undetected. The retrospective nature of annual audits meant that fraud could continue for years before being detected. And manual review was simply not fast enough to keep pace with the volume and velocity of modern transactions.

AI changes the process of fraud detection in three specific ways:

- **Full population coverage** — It expands coverage from samples to full populations. A machine learning model that evaluates every transaction removes the ability to hide fraud in the untested portion of the data.
- **Real-time detection** — It enables real-time rather than retrospective detection. Rather than identifying fraud after a full year of activity has been completed, AI-powered systems can flag suspicious transactions as they occur.
- **New pattern discovery** — It can identify previously unknown fraud patterns. An unsupervised machine learning model discovers anomalies relative to normal behavior — meaning it can surface a novel fraud scheme that no one designed a specific rule to detect.

Fraud Detection through Advancing Technology

Continuous advancement in technology is transforming the way fraud is detected in auditing. Technology is continually advancing each generation of tools closing verification gaps that the previous generation could not reach. Two cases, separated by industry and geography, together illustrate how that progression works and where AI fits within it.

The Zhangzidao Case 3.3.1

During the period from 2014 to 2018, Zonco Group Co., a public seafood business of China, experienced enormous unanticipated losses in its inventory of scallops. Every time this was linked to natural calamities ranging from cold waves causing the scallops to migrate elsewhere and high temperatures causing the scallops to die off. For instance, in 2014, they reported an abrupt change from a profit of 48.5 million yuan to a net loss of 800 million yuan within three quarters after being profitable earlier in the year. Each of these statements became increasingly unlikely to be true, and there was nothing that the regulatory authorities could do about them. This is because the inventory in question was the scallops grown in the seabed of the ocean, taking up tens of thousands of hectares of open ocean surface. It is almost impossible to physically count the inventory and therefore the fraud was going on for years without being detected.

It turns out that the answer was staring everybody in the face. Chinese Maritime Law required every fishing boat to carry BeiDou navigation system onboard to avoid accidents at sea. During the formal investigation by China Securities Regulatory Commission in February 2018, investigators obtained the BeiDou positions recorded by the 27 fishing boats and figured out where each boat went during the period. From these positions, they found that in the fall of 2017 survey, none of the fishing vessels had visited 60 out of 120 sampling points. In 2016, they had secretly fished an additional 140,000 mu of seabed which was not recorded on their books and hid around 60 million yuan of expenses to increase profits in that year. The reason that no one thought of forging the satellite positions was because it was installed onboard for navigational purposes only. Therefore, that information was the only truthful information in the whole investigation. As a result, 17 months later, the company paid fines, 15 executives were punished, and the chairman was barred for life from securities market.

Implication

The broader lesson extends well beyond the seafood industry. Fraud tends to survive longest in areas where standard verification methods run into physical or logistical limits, whether that is inventory on the ocean floor, transactions too numerous to examine individually, or documents spread across systems no single reviewer can cover. What made the satellite data so powerful was that it generated an independent record of reality, one that existed entirely outside the company's control. As verification technology has continued to develop, from GPS positioning to The Internet of Things sensors embedded in physical assets to AI systems that monitor transactions in real time, that same principle has carried forward. Each advancement closes a gap that earlier methods could not reach.

The KPMG CLARA Case 3.3.2

While the Zhangzidao case illustrates a physical verification problem, the most directly relevant example for external auditing comes from KPMG's own audit practice. Before AI-assisted tools were available, KPMG audit teams reviewed between 5 and 10 percent of a client's transactions using traditional sampling methods. The majority of the transaction population was never examined. Any fraud or error distributed across the untested portion was, by definition, beyond the reach of the audit.

KPMG Clara Transaction Scoring, developed through a formal partnership with MindBridge AI, changed that. The tool is deployed by KPMG's 90,000 auditors in external audit engagements worldwide and analyzes 100 percent of transactional populations across both general ledger and subledger records. Every transaction is scored by the machine learning model and assessed against patterns associated with anomalies, including duplicate invoice schemes, unusual journal entry timing, and entries posted outside normal business hours or by users not typically associated with a given account. The model flags those that deviate from expected behavior and routes them to the auditor's review dashboard with supporting evidence attached.

In external audit engagements where Clara Transaction Scoring has been deployed, KPMG has reported a 35 percent reduction in hours spent on routine vouching and workpaper assembly, and up to 20 percent improvement in review accuracy. Audit teams that previously examined 5 to 10 percent of a client's revenue and expense lines now examine all of them. Duplicate invoice schemes that would have been statistically unlikely to surface in a sample are now identified systematically.

Implication

Collectively, both the Zhangzidao and the KPMG CLARA cases follow the same trajectory from opposing angles. The use of satellite tracking revealed that technological evidence creation in the form of an autonomous record of physical activity had the ability to reveal fraud beyond the capacity of human verification at set intervals. In terms of AI and external auditing, the same concept applies in moving from the verification of a part of transactions to the verification of all transactions. This is just that on a different scale and using different technologies, but still following the same basic idea.

Opportunities and Challenges

The adoption of AI in auditing brings a lot of advantages, but it also introduces challenges that the profession is still working through.

Table 3.3.3 lists the opportunities and the challenges that AI brings to auditing.

Opportunities	Challenges
Full population testing. AI eliminates reliance on samples, allowing auditors to examine 100% of transactions. Fraud and errors invisible to sampling are now detectable.	The 'black box' problem. Auditing standards require documentation of the rationale behind every conclusion. Complex AI algorithms often cannot explain why they flagged a transaction, therefore this creates a conflict with documentation standards.
Continuous monitoring. AI systems can operate in real time, detecting irregularities as they occur rather than during an annual review.	Data quality dependence. AI tools produce reliable output only when the data they analyze is complete and accurate. Incomplete client data produces inaccurate AI findings.
Greater efficiency. Automating data gathering and initial analysis frees auditors to focus on judgment-intensive work.	Privacy and security risks. AI systems require access to large volumes of sensitive financial data, introducing risks of unauthorized access.
Improved fraud detection. Machine learning can identify novel fraud patterns and flag anomalies human reviewers would miss in large datasets.	Skills gap. Many audit professionals lack the technical training needed to configure, evaluate, and critically assess AI-generated findings.
Consistent application of standards. RPA applies the same validation rules to every document, removing human variability from large review processes.	Regulatory uncertainty. Audit standards were written for human-led processes. Regulators have not yet fully determined how AI-assisted procedures should be documented and overseen.

Exercise 3.3.1

This exercise uses a realistic business scenario to help you experience firsthand what AI-assisted fraud detection looks like in practice and where human judgment remains essential. You will need access to a free AI assistant such as ChatGPT (chat.openai.com). No accounting background is required.

Key Terms

- Vendor payments are the money a company pays to outside suppliers for goods or services, such as payments to a software provider, a cleaning company, or an equipment supplier. Every payment a business makes is recorded in its financial system, creating a digital trail of every transaction.
- An audit is an independent review of a company's financial records by a third party to verify that everything has been recorded accurately and honestly. Because companies process thousands or even millions of transactions per year, auditors cannot realistically check every single one. Instead, they traditionally select a sample, a smaller representative group of transactions, and examine those in detail to draw conclusions about the whole.
- An approval threshold is a dollar limit set by a company above which a payment requires a second person to review and sign off before it goes through. This is a standard internal control designed to prevent any one employee from authorizing large payments on their own. For example, a company might require that any payment over \$5,000 needs a manager's approval in addition to the person who initiated it.

- Threshold structuring is a type of fraud in which someone deliberately keeps each transaction just below an approval threshold to avoid triggering the requirement for a second approver. Rather than making one large payment that would be reviewed, the person breaks the activity into many small ones that each fall below the limit and pass through automatically.
- A false positive in fraud detection is a transaction that an automated system flags as suspicious but turns out, upon investigation, to be completely legitimate. High false positive rates are a known problem with older rule-based detection systems because they waste investigative time and disrupt normal business operations.

Scenario

The manufacturing firm pays about 42,000 vendor invoices annually. A vendor invoice represents a bill sent by the supplier requesting money for services rendered or goods supplied. Every year, the auditors of this firm examine a sample of 175 vendor payments. During the last audit, there were no exceptions as all 175 payments had been properly executed; each of them had a purchase order, invoice and appropriate approvals. A purchase order is an official internal document created by the firm at the moment of agreement to buy certain products or services.

However, the sample failed to show that the warehouse manager had been making payments to three supplier accounts which were hers. To avoid the second level of approval, she made sure that all payments did not exceed \$4,800 (the company's threshold was \$5,000). Payments were made during 14 months, while three supplier accounts were used to make sure that none of them received an excessive amount. The fraud was discovered through a separate report, whereas the sum taken fraudulently reached \$218,000.

1. Think through the sampling limitation. Write two to three sentences explaining why the traditional sample of 175 payments had no realistic chance of catching this scheme. Think about how the fraud was designed and what a random sample of 175 records out of 42,000 can and cannot tell you.
2. Use ChatGPT to explore what AI would detect. Access ChatGPT (chat.openai.com) and paste this prompt exactly: "A manufacturing company processes 42,000 supplier payments per year. An employee committed fraud by keeping each payment just below a \$5,000 approval threshold, spreading payments across three different supplier accounts over 14 months, totaling \$218,000. If a machine learning tool analyzed all 42,000 payments, what specific patterns would it look for to identify this type of fraud? Please explain each pattern in plain, everyday language without assuming any accounting knowledge." Record what the AI identifies.
3. Consider what AI cannot do on its own. The machine learning model can flag transactions that match suspicious patterns, but it cannot conclude that fraud actually occurred. What would an investigator need to find out about each flagged supplier account before drawing that conclusion? Think about what information would confirm or rule out wrongdoing.
4. Reflect on the human and AI partnership. Based on what you learned in Steps 2 and 3, write a short paragraph describing what AI contributed to this scenario that traditional sampling could not, and what still required human investigation. What does this tell you about how AI and auditors work together?

Guidance

The fraud in this scenario exploits a structural weakness that sampling cannot address. With 42,000 payments spread across a year, a sample of 175 covers less than half a percent of all transactions. A scheme deliberately designed to keep each transaction small and spread across multiple accounts and time periods has almost no statistical chance of landing in the sample. Machine learning closes this gap by evaluating all 42,000 payments equally and looking for the concentration of transactions just below a round-number threshold, multiple accounts linked to the same approver, and payment patterns that do not match the company's normal supplier relationships. Once the model flags those accounts, an investigator must then obtain the original supplier contracts, confirm that real goods or services were received, and determine whether the employee had any undisclosed connection to the suppliers. The AI identifies where the irregularity is. A person determines what it means and whether it constitutes fraud.

Answer available at the end of this section.

Putting It Together

From its inception till now, auditing has always been a field that was distinguished by certain boundaries within which its verification procedures took place. Historically speaking, such boundaries had to do with the physical possibility of carrying

out the audit process. Still, every new generation of technology extended these boundaries beyond their current state, leaving them intact nonetheless. The major difference between audit technologies based on artificial intelligence is that they break the boundaries which were taken for granted in previous generations of audit technologies. Analyzing each transaction in a transaction population, studying each document in a contract folder, and validating each submission against certain standards was impossible in the context of conventional auditing. Now, these processes are becoming automated as well.

The companies that have pushed this trend furthest are certain that it will continue to develop until all repetitive auditing tasks are handled by machines. In fact, a much more interesting question here is what future auditing will look like after repetitive tasks are performed by machines. All practitioners seem to agree that the future of auditing will include fewer number crunching and more decision making on behalf of an auditor. Detection of anomalies in transaction populations will become the domain of machine learning models while the task of specialist will be in analyzing risks and making decisions on the actions to take.

Answers

Exercise 3.3.1

The following responses illustrate the level of reasoning this exercise is designed to develop. Your own responses do not need to match these exactly, but they should reflect the same core ideas.

1. Of the 42,000 transactions, this audit sampled 175, which means that 0.46% of all transactions were sampled. In other words, the fraud had been designed to be impossible to detect through any kind of sample. It consisted of small payments that were directed to three different supplier names during 14 months to avoid accumulating any kind of suspicious sums for any particular supplier account. Moreover, even in case one payment was included into the sample accidentally, it would still look like a regular payment of under \$4,800 to one of the company suppliers. Such a fraud was undetectable not due to poor quality of the audit itself, but simply because the audit was not able to see the whole picture and to draw the right conclusions.
2. In response to the prompt above, ChatGPT usually highlights the following peculiarities of payments and suppliers that would be identified by machine learning:
 - Payments clustered just below the threshold of the approval limit. This is a suspicious pattern, because in natural situation payment amounts are based on prices for services and goods rather than approval limit. Therefore, the clustering of many payments in the range of \$4,500 to \$4,999 can indicate that the payments have been deliberately structured.
 - A single person approving all payments to certain suppliers. An unusual pattern in which payments to the certain supplier have been approved by a single person for a prolonged period of time and this person does not approve payments for other suppliers can also indicate a possibility of fraud.
3. Before concluding that fraud occurred, an investigator would need to verify several things about each flagged supplier account. First, do these businesses actually exist? That means searching for a registered business address, a working phone number, a website, or any public record confirming the supplier is a real company. Second, did the manufacturing company actually receive anything from these suppliers? That means tracing each payment back to a delivery record, a service completion report, or any evidence that something of value changed hands. Third, does the warehouse manager have any personal or financial connection to these suppliers? That means checking whether the suppliers share an address with the manager, are registered to a family member, or have any other tie that would explain why she alone was approving their invoices. Without answers to these questions, the pattern the AI flagged is a signal that something may be wrong, not proof that something is.
4. The AI brought a significant contribution to the investigation which could not be done with the help of sampling. Namely, the fraud presented in this scenario is undetectable via any kind of sample because the scheme is too complex and consists of too many accounts and transactions made in a prolonged period of time. However, in order to detect such fraud, it is necessary to analyze all 42,000 payments at once. But after the accounts have been chosen as suspicious, further steps require human work.

4: ETHICS OF AI IN ACCOUNTING

Chapter Overview

The use of artificial intelligence has revolutionized the practice of accounting on every possible level from fraud detection to tax preparation, from transaction sampling by auditors to client consulting. The technology makes it much easier to complete tasks in the shortest time possible; however, it also creates an additional dimension of issues that was not taken into account when developing the professional code of conduct. This chapter discusses this ethical dimension at three different but interrelated levels starting with individual accountants, moving to accounting firms, and finishing with society at large whose financial data needs to be trusted.

There are five basic principles of AICPA Code of Professional Conduct which have been guiding professionals in their activities for years now. Those principles include integrity, objectivity, professional competence and due care, confidentiality, and professional behavior. The only thing that has changed about this code is the environment where those principles need to be implemented. The biased outputs of AI make it hard to maintain objectivity. Lack of transparency of AI's decisions takes away the possibility of making and recording sound judgments. Data shared through the third party services raises confidentiality issues in firms.

Learning Objectives

By the end of this chapter, you will be able to:

- Evaluate the ethical implications of AI integration in accounting at the individual, organizational, and societal levels
- Apply ethical principles related to bias, transparency, and accountability in AI-supported accounting practices
- Examine professional accounting standards (integrity, competence, and confidentiality) in the context of AI usage
- Assess data privacy and security risks associated with AI within regulatory and societal frameworks
- Develop best practices for ethical AI implementation in accounting firms

4.1: Why Ethics Matters?

4.2: Individual Professional Responsibilities

4.3: Professional Accounting Ethical Standards

4.4: Societal Consequences

4.5: AI Ethics Policy for an Accounting Firm

4.1: Why Ethics Matters?

The Ethical Stakes of AI in Accounting

For many people, ethics may bring up the notion of philosophy, something distant from the stresses associated with meeting the client deadline or quarter end. This is not true for the field of accounting. Ethics are what make clients agree to the engagement letter, allow investors to trust the information disclosed, and permit CPAs to have the power to issue opinions that influence the market. Below is one such situation.

The 2 AM Decision Case

It is 2:14 AM. Maya Rodriguez, CPA, sits alone in the glass-walled office of Hartley and Nguyen Advisory Group, the city lights sparkling below her. On three monitors, the output of the firm's new AI audit sampling tool is disturbingly telling a story Maya cannot understand. After the first successful run through the live transactional database of the firm, the AI marks 73 percent of transactions for further audit for Hartley and Nguyen Advisory Group's community health clinic clients and clears transactions of corporate clients with similar profiles. When Maya calls the vendor of the software, a technician informs her that the machine learning algorithm was fed 40 million transactions of Fortune 500 companies, that it is very accurate, and that the algorithm itself is proprietary, with no information about how the transactions are weighted made available to customers. Maya finds herself under three concurrent pressures that define this case.

As an individual professional, Maya personally signed all audit reports done for these clinics. If the AI tool is biased, Maya is professionally and legally liable. The work is her work, not AI's, and she cannot stand for the conclusions of an instrument she cannot understand or even review for compliance with the existing standards. Also, no engagement letter of Maya's firm mentions anything about AI decision-making.

As a representative of the organization she works for, Hartley and Nguyen Advisory Group used the AI platform to process sensitive financial records of the clinics without any governance in place. She has no idea whether the data were anonymized, whether the vendor's data protection policies conform with the law, and whether there were any consents of the clinics' management to use their data in this way. Hartley and Nguyen Advisory Group has no company policy on AI usage, and by Monday morning Maya should have something to discuss in the partners meeting regarding this specific instance and future policies on AI usage.

From the societal perspective, the clinics that Maya audits operate for the benefit of vulnerable communities. An AI tool that flags transactions of those clinics more often than comparable corporate clients is not only unethical issue, but the one that directly causes harm to the organizations serving those communities, while nobody is making an intentional choice of biasing the tool.

Maya's situation is not fictional. AI technology is developing faster than the ethical frameworks governing it.

Quick Check

1. Maya cannot explain why the AI flagged significantly more transactions from nonprofit clients than from corporate clients with similar profiles. What type of bias does this most likely reflect, and why?
2. Maya's firm has no AI use policy, and the vendor cannot explain how the platform reaches its conclusions. Which AICPA ethical principle is most directly at risk, and what is Maya's obligation before she uses the platform again?

Solution

1. This most likely reflects data bias. The platform was trained on Fortune 500 transaction data, which does not represent the financial patterns of nonprofit organizations. The AI has not learned to recognize what it was never shown, so it flags unfamiliar patterns as potential risks rather than normal activity.
2. Professional competence and due care is most directly at risk. Maya cannot certify work she cannot examine or explain. Before using the platform again, she should pause its use and request documentation from the vendor about how the system was trained, what it was designed to detect, and what its known limitations are.

The Profession's Ethical Foundation

Maya's dilemma is not a technological one; it is an ethical one that becomes obvious thanks to technology. Indeed, every single influence she feels on her way results from the violation of any of the five AICPA Code of Professional Conduct principles. The understanding of how to tackle these violations and how to cope with them is what this chapter is all about.

There is a social contract whereby society grants the task of independently testifying the financial truth to certified public accountants in exchange for their adherence to some ethical standards. These standards are presented in the AICPA Code of Professional Conduct, in PCAOB standards for auditors of public companies and in the requirements of CPA's state licenses.

The five principles of the AICPA Code of Professional Conduct are the result of decades-long experience of the accounting profession. On more than one occasion in the course of the history of the profession these principles were revised as a reaction to ethical problems of the profession. Please see [Section 4.3](#) for the discussion of the Sarbanes-Oxley Act of 2002 aimed at addressing the issue of failures of the profession's ethics in connection with Enron and WorldCom cases.

Where AI and Professional Standards Come Into Conflict

Traditional accounting ethics assumes that a human professional exercises judgment at every material decision point. AI disrupts this assumption in ways that are fundamental rather than incidental.

Table 4.1.1: Mapping of AICPA ethical principles to traditional assumptions and AI disruptions

Ethical Principle	Traditional Assumption	How AI Challenges it
Integrity	Accountant consciously chooses honest reporting.	AI may produce technically compliant but misleading outputs without any human intent to deceive. For example, an AI revenue recognition system configured with aggressive booking rules automatically records bill-and-hold transactions as completed sales across thousands of entries. A bill-and-hold transaction is when a company records a sale and counts the revenue even though the goods have not yet been delivered to the customer. No individual ever made a decision to misstate revenue, but the financial statements are misleading.
Objectivity	Accountant identifies and discloses conflicts of interest.	AI systems embed structural biases that may not be visible or disclosed. For example, an audit AI trained predominantly on large corporate data systematically marks nonprofit transactions as higher risk. The tool appears neutral, but its outputs consistently disadvantage a specific type of client in ways that are difficult to detect or challenge.
Competence and Due Care	Accountant understands the tools and methods they use.	AI systems may be genuinely unexplainable even to the professionals deploying them. For example, a CPA relies on an AI fraud detection platform to scope an audit engagement. The vendor confirms the system is highly accurate but cannot provide any documentation of how it reaches its conclusions. The professional cannot explain or verify the basis for their own audit decisions.
Confidentiality	Client data is held securely within the firm.	AI cloud platforms transmit and store data across third-party environments. For example, a tax advisory firm uses a cloud-based AI platform to analyze client financial data. The vendor's legal agreement, accepted when signing up for the platform, includes the right to use submitted data to improve the model. Client information is effectively shared with a third party without the client's knowledge or consent.
Professional Behavior	Accountant acts to maintain the profession's reputation.	Widespread AI adoption without governance could undermine public trust systemically. For example, imagine that most large accounting firms start using AI audit tools that were all built the same way and trained on the same type of data. If something is wrong with how those tools were designed, every firm using them will make the same mistakes at the same time. When those mistakes come to light, the public begins to question whether audited financial statements can be trusted at all, not just at one firm, but across the entire profession.

The Three Levels of Ethical Impact

The ethics of AI in the field of accounting can be looked at from three different perspectives that run concurrently, namely individual, organizational, and societal. These levels are not standalone in nature. They are interrelated in such a manner that if there is a failure at any level, then the effect spills over to all other levels as well. The individual accountants who do not have professional skepticism are the cause of organizational control failure, which when extended within the organization becomes a societal failure. Each chapter section that follows focuses on one of these levels.

4.2: Individual Professional Responsibilities

Individual Ethical Considerations

If an AI system is employed by an accounting professional, the responsibility for its effect will not be shifted from the professional to the technology. The example in [chapter 4.1](#) demonstrates this idea clearly. Maya Rodriguez identified a possible presence of bias in her firm's AI system, but irrespective of whether this bias came from the source or not, her signature appeared on all the audit reports. This section looks at individual ethical responsibilities of any accounting professional due to the influence of AI on their work. It includes identifying and acting upon bias, working ethically with an AI system that cannot explain its output, and taking responsibility in terms of accountability while dealing with technology.

Bias in AI Systems

One of the main ethical issues raised by AI in accounting is bias. It is important to understand the sources of such a bias and its impact on professional activity while using this technology. An AI system that demonstrates biased results is unreliable.

Types of Biases

- **Data bias** occurs when an AI model has been developed using data sets which are flawed, skewed, or do not include all of the possible use cases from the real world. In Maya's case, the AI audit tool that has been developed mainly based on Fortune 500 transactions will fail in case it encounters nonprofit organizations, smaller businesses, or businesses operating in emerging markets because it has not seen such use cases before.
- **Algorithmic bias** is inherent to the decisions made during the design phase by the system's creators. Regardless of whether there is clear data, the goals which the system tries to optimize reflect values. For example, an audit AI that tries to minimize costs may systematically under-represent infrequent, yet risky transactions.
- **Confirmation bias** amplification occurs when professionals preferentially accept AI outputs that confirm their existing beliefs. If an auditor already suspects a client, they may accept AI flags against that client without scrutiny, while dismissing flags against clients they trust, creating a feedback loop between human and machine bias.
- **Temporal bias**, also called model drift, where an AI system has been trained using old data that does not represent the current operations in the industry. While the patterns might have been relevant at the time, if things have moved forward, there is no way in which the system can match them anymore. For instance, if an AI system has been trained prior to remote working environments and other technological advancements such as online payments, a large number of otherwise legitimate transactions will be classified as suspicious.
- **Contextual and cultural bias** arises when AI is applied across economic or cultural contexts different from its training environment. Financial AI trained on Western corporate data may generate systematically biased assessments in emerging markets, community development organizations, or industries with non-standard accounting practices.

Bias Countering Framework

Table 4.2.1: Four professional practices to counter bias in AI-assisted accounting

Professional Practice	Application in AI-Assisted Accounting
Critical Evaluation	Apply professional skepticism to all AI outputs. Ask whether the result makes sense given your knowledge of the client, industry, and economic context. AI outputs are inputs to professional judgment, not substitutes for it.
Proactive Prevention	When querying AI systems, include explicit instructions requesting that the system flag potential biases or limitations in its output. Some platforms allow users to specify client context that can reduce systematic misclassification.
Independent Verification	Establish parallel processes to validate AI-generated findings through independent means such as manual sampling, peer review, or comparison with prior-period results. Never allow a single AI output to constitute the sole basis for a material judgment.

Professional Practice Application in AI-Assisted Accounting

Source Validation	Demand transparency about training data sources and model documentation. If a vendor cannot explain what data trained their system or what it was optimized for, that absence of information is itself a red flag requiring elevated professional skepticism.
-------------------	---

Working with AI you cannot Fully Understand

The vast majority of AI algorithms that have been applied in accounting lack an understanding of the reasons why they generate certain results. They analyze lots of information, establish some patterns and come up with output but the process of generating this output remains beyond the comprehension of individuals using the system. This issue represents a serious challenge in terms of accounting. This field is characterized by the ability to explain the judgment and defend the decision. In case a considerable portion of the judgment is generated by an algorithm, it becomes impossible to do.

Three Professional Consequences

Using AI systems that cannot explain their outputs creates three specific problems that affect how professionals do their work and how that work is evaluated by others.

Definition: Impaired verification

Professional standards require evidence-based, verifiable justification for material judgments. When an AI's reasoning is unknown, a professional cannot demonstrate why a conclusion is correct, only that they found it plausible. Plausible does not meet the evidentiary standard of professional practice.

Definition: Challenged error analysis

When an AI tool that is not transparent makes a mistake, it is nearly impossible to diagnose the root cause. In traditional practice, an error can be analyzed, understood, and prevented from recurring. With these AI systems, the same flaw may produce similar errors indefinitely without detection.

Definition: Eroded Credibility

A professional who presents an unexplainable AI output as their professional judgment is misrepresenting the nature of their work. When this comes to light in litigation, regulatory review, or a client dispute, it can cause irreparable damage to the professional's credibility and the firm's reputation.

AI Creates Opportunity for Fraud

The fraud triangle is a familiar concept within accounting that outlines the three components required to carry out fraud. There needs to be motivation to commit fraud, rationalization to make oneself feel justified, and an opportunity to do so. AI adds an additional factor to the last point.

If an AI tool used in auditing fails to provide explanations about how it reaches conclusions about transactions, there will be no information about what rules it applies for identifying suspicious transactions, not even for those individuals working with it. In this situation, a person looking to commit fraud will be able to design transactions in such a way that they will fall below the level of risk that the system has been trained to identify.

Let's consider an illustrative example here. Let's assume that a particular AI system has been trained to automatically mark all transactions exceeding a specific monetary amount for further inspection. Knowing that the system has some sort of a limit, an employee will simply split one fraudulent transaction into a few lower-value transactions, each of which would be lower than his

estimation of the limit. He knows nothing about the precise amount; he knows only that there is some sort of a limit.

Accountability in the Age of AI

One of the key concerns that arise with regards to the use of AI in accounting relates to the question of accountability. This issue is especially relevant as numerous parties are involved in the creation and usage of any particular AI software. Before analyzing who should be held accountable, it is useful to take a look at the number of parties that are involved in the process of creating the tool before a professional even uses it. First, the development team creates the software and makes crucial design decisions. Second, the preparation of the training dataset determines what data will be processed by the algorithm. Third, the deploying organization decides where the software will be used and for which purposes. Fourth, the professionals will interpret the output and use the results in practice. Finally, the maintenance team is supposed to monitor the accuracy of the algorithm's operations.

Therefore, the issue with a faulty algorithm may have a different nature and may relate to each of the mentioned parties or all of them together. In this way, the output bias may be the result of the incorrect preparation of the training dataset, the bad design decision, the wrong deploying, or an inappropriate interpretation of the results by the professional.

Nonetheless, according to the Code of Professional Conduct of the AICPA, individual professionals are accountable for the work products that they generate regardless of the tools used to accomplish them. Therefore, blaming a software, a vendor, or the technology department of the firm for mistakes is not a viable option. If your name appears on the work, then you are accountable for it.

This requirement is necessary not only for the sake of law but also for the sake of practicality. If professionals start transferring their responsibilities onto a tool, they are very likely to stop being skeptical towards its output and to verify the results. Therefore, accountability is not only about finding someone to blame but also about motivating verification and other actions aimed at the prevention of possible mistakes.

What this Means in Practice

Knowing that accountability rests with the professional, there are concrete steps every accountant should take when using AI tools.

- An accountant should review AI outputs critically rather than accepting them at face value. This means considering whether the result makes sense given what is known about the client, the industry, and the current economic environment.
- Documenting the review is equally important. If an AI tool informed a material judgment, the work papers should show that the output was examined, its limitations were considered, and independent judgment was exercised in reaching a conclusion.
- Familiarity with the tool itself is also part of responsible practice. Understanding the mathematics behind an AI system is not required, but an accountant should know what it was designed to do, what data it was trained on, and what its known limitations are.
- Lastly, issues need to be communicated through proper channels and not left untouched. In case there is an issue with an AI deliverable, the problem needs to be taken up with the supervisor or the quality control partner of the firm or even the engagement team prior to delivering the work product. In case there is a process within the firm to report concerns related to technology, then that would be the right starting point. Not raising a concern or being satisfied with the vendor's response is not a proper approach.

4.3: Professional Accounting Ethical Standards

Ethical Standards in AI

For decades, the American Institute of Certified Public Accountants (AICPA) Code of Professional Conduct has been regulating the conduct of accounting practitioners. It is designed for a world where every judgment is performed by a human professional who can be held accountable for this judgment. Artificial intelligence (AI) does not negate this accountability but introduces new scenarios that the Code authors could not envision. Out of the five principles of the Code three, in particular, are challenged by the advent of AI. These are the principles of integrity, competence, and confidentiality. The following sections discuss each of these principles, provide definitions of them in AI context, and illustrate them using practical cases.

Integrity in AI-Assisted Practice

According to Section 1.100 of the AICPA Code of Professional Conduct, integrity refers to straightforwardness and honesty in all professional relationships. It implies that an accountant does not misrepresent any information deliberately and does not succumb to any influence from clients, employer, or any other entity. Simply put, integrity implies the ability to act in an ethically correct way even if this action is inconvenient, and to present the data in an honest and precise manner.

In the traditional practice, integrity failure usually entails a deliberate choice made by someone. It means that an individual decides to either misstate a number, conceal some transactions, or signs something he or she knows to be untrue. AI alters the nature of integrity significantly. It now became possible for the financial statements to contain misinformation without a deliberate decision to deceive.

Consider the following straightforward case. An AI system used for revenue recognition is configured using aggressive rules. It automatically recognizes thousands of transactions as completed sales even though the products have not been delivered to customers yet. There is nobody who ordered this misrepresentation. The configuration just was not made with enough care. Still, the end result is the same as a deliberate misrepresentation: the financial statements are misleading, and the professional who certified them has to assume the responsibility for this.

It is significant because integrity now requires more from an accountant. It is not sufficient anymore to be truthful oneself. An accountant must also be able to understand how AI systems that he or she uses are configured, and what output the system is expected to provide.

Table 4.3.1: AI-enabled revenue manipulation schemes and their impact on scale and detectability

Method	Traditional Mechanism	AI-Era Integrity Risk
Channel Stuffing	Customers encouraged to buy under favorable return terms to inflate sales figures prematurely.	AI demand forecasting produces optimistic projections that normalize stuffed inventory as expected sell-through, providing an apparently data-driven justification for inflated sales.
Bill-and-Hold Sales	Revenue recognized before title transfers, with goods held in the seller's warehouse.	AI revenue recognition systems configured with aggressive booking rules automatically classify bill-and-hold transactions as completed sales across thousands of entries, distributing the deception at scale.
Round-Tripping	Items sold with a promise to buy back, creating the appearance of revenue without economic substance.	AI tools trained on data that includes round-trip transactions learn to treat them as legitimate, particularly when structured to mimic real customer behavior.
Improper Cutoff	Recording sales in the wrong accounting period, a form of early revenue recognition.	AI booking systems misconfigured to apply incorrect period cutoffs systematically generate misclassified entries that appear system-generated rather than human-directed.

The Enron Case

At the time of bankruptcy, Enron was one of the largest energy companies in the USA. In order to conceal the fact that it was running losing businesses, Enron established subsidiaries under separate legal umbrella and transferred all the loss-generating activity there. At first glance, it looked like the company was profitable. However, the truth eventually was uncovered and Enron went bankrupt practically overnight, destroying millions of peoples' savings and eroding the worth of shareholders by approximately 74 billion dollars. This fraud is one of the reasons the Sarbanes-Oxley Act has been enacted in 2002.

AI Consideration

The Enron case serves a good example for the purposes of this chapter because it happened long before the advent of AI. Still, it demonstrates the type of fraud that AI systems used by accounting firms nowadays have difficulties with detecting. Transactions between Enron and its subsidiaries were constructed in such a way that they looked like regular business deals with unrelated parties. Thus, an AI would assess these transactions individually and would find no problem with them because it would use historical data in order to compare each transaction and find abnormalities. However, detecting the fraud in this case required an understanding of relations between these parties and an ability to read qualitative disclosures in financial statements. All of that is exactly what the current AI tools lack.

Competence in AI-Assisted Practice

According to Section 1.300.001 of the AICPA Code of Professional Conduct, a professional accountant must only take up a job if he or she is truly competent to do it and perform this job with diligence. Moreover, it is implied that the accountant keeps his or her competencies up-to-date as best practices and technologies evolve. Competency is not limited to obtaining some kind of certification; it is the ability to stay informed and understand the limits of your expertise.

This obligation exists fully in the case of AI tools. In fact, the use of AI makes it more challenging because an accountant now needs to understand not only the accounting standards but also the capabilities and limits of the technology he or she relies on.

Usually, the AI systems that are used in the field of accounting do not come with a clear explanation of what they do. Audit sampling, fraud detection and tax analyses are often based on complicated models that process the huge amount of data. It means that an accountant using such a tool in order to form an opinion does so without understanding what the tool is capable of and on what data it was trained. It means that such an accountant takes up an impossible task.

An example is helpful to understand this principle better. Imagine an auditor who uses AI in order to find suspicious transactions and make the decision whether an extra audit step should be taken. The AI gives him or her the list of such transactions and he or she checks them and completes the audit. Eventually, a significant misstatement is detected in a transaction that AI did not point to. If the auditor does not know why the AI gave this particular list, what data it was trained on and what kinds of transactions it is supposed to identify, then he or she is in trouble. It is the auditors' judgment and the AI is only an input for this judgment.

Sarbanes-Oxley and the Competence Obligation

The Sarbanes-Oxley Act of 2002 was directly created as a response to the Enron and WorldCom frauds. There are two provisions of the Act which are especially relevant in the context of AI in accounting. Section 302 implies that the Chief Executive Officer and Chief Financial Officer of a publicly-held company are required to certify personally that the financial statements of this company are correct and its internal controls are effective. Section 404 implies that both the management and the external auditor should review and report on the effectiveness of these controls.

If the AI tools were involved in generating or reviewing the financial information, the certification and assessment requirements apply to this AI. The CEO certifying the internal controls effectiveness has to be able to describe the controls, explain how they are tested and where are their limitations. If the system that performs a significant portion of control activities is not comprehensible for the CEO or cannot be evaluated, the certification has no stable foundation at all.

The current regulations recognize the fact that the AI system can be a part of a properly designed internal control environment if the responsible parties have proper documentation describing how this system works, on what data it was trained, how its performance was checked and monitored and what human review procedure is used alongside with it in order to prevent mistakes

this system can make.

Section 404 also imposes record-keeping obligations on firms using AI systems in internal control procedures. The record keeping is supposed to include the design objectives of the system, log of its inputs and outputs, proof of professional review conducted before these inputs and outputs became the basis of the significant judgments and monitoring records demonstrating continued compliance of the system with expectations over time.

According to the PCAOB Auditing Standard 1215, the accounting firm conducting the audit of the publicly-held company is required to keep the audit documentation for seven years since the date when auditor gave consent to release this audit report. In an AI context, this requirement implies not only maintenance of traditional workpapers but also of the data used to train the AI model used in the audit, history of revisions this model underwent, all inputs the system received and outputs it generated and communication with the vendor describing how the system works.

Intentional destruction of these documents is a criminal offense according to Sarbanes-Oxley Section 802 and the penalty for failure to maintain them can lead to serious fines.

The HealthSouth Case

HealthSouth was a large healthcare company whose entire accounting department participated in falsifying financial records over a period of six years. Rather than making a small number of large falsified entries, which would have been easier to spot, the fraud was spread across hundreds of thousands of small adjustments that individually appeared unremarkable. The total overstatement ran to several billion dollars.

An AI fraud detection system reviewing HealthSouth's accounts would have been looking for patterns that matched known fraud, such as large unexplained variances, unusual spikes in revenue, or transactions that did not fit normal business activity. The HealthSouth fraud produced none of those signals because each individual entry was small enough to look routine. Over time, as the fraud grew, the AI's sense of what normal activity looked like would have adjusted to include the fraudulent entries, making them even harder to distinguish from legitimate transactions. This is model drift in practice, a system becoming progressively less reliable as conditions shift around it, without any mechanism in place to detect or correct that drift.

AI Consideration

The competence implication is direct. Using an AI fraud detection tool without regularly testing whether it is still identifying the right things, and without supplementary procedures that look for patterns building across time rather than just within individual transactions, is not competent practice. A system that was adequate when it was deployed may no longer be adequate a year later, particularly if the activity it is monitoring has been gradually changing around it.

Confidentiality in AI-Assisted Practice

According to Section 1.700 of the AICPA Code of Professional Conduct, the information which the accountant acquired during the course of professional work belongs to the client. It can be revealed to anyone else only with client's consent, unless there is an explicit requirement to disclose this information due to legal reasons. This obligation does not expire upon ending of the engagement and does not transfer to the vendor just because the vendor is going to process the client's data.

This obligation has always been quite straightforward in a traditional accounting environment. The client's data were stored inside the firm, were accessible only to the professionals who worked there and boundaries were clearly defined. AI complicates the boundaries in a way which is not obvious for the professionals working with the system.

Once the client's data are submitted to the AI system, they start moving and getting processed. They move to the vendor's server which may be located in another country. They get processed by the vendor's system which may be run by somebody who has an incidental access to the client's data. They may remain with the vendor after the engagement is completed. Also, the terms of service that the firm signed with the vendor allow this vendor to use the client's data in order to train their AI model, which means that client's data could become part of a dataset which determines what the AI will do for any other client of this vendor, including the client's competitors.

None of this requires a deliberate violation of the confidentiality agreement. It could happen because the accountant did not read

the vendor agreement carefully, did not ask enough questions before choosing the tool and did not inform the client about what would happen to his or her data.

The example will make the issue clearer. Tax advisory firm uses a cloud-based AI tool to analyze the client's financial position and identify potential planning opportunities. The vendor's agreement, signed while establishing the account, allows the vendor to use submitted data to train the model. The client's sensitive financial information is included in the dataset which is used to train the model. The client has not been told about this fact, and the firm has not even checked it.

Using the AI tool, the accountant exposes the client's data to four categories of confidentiality risks. Not all of these risks are equally obvious, and some of them are more common than the others. Understanding them in order from the most obvious to the most technical helps to realize what confidentiality means in the environment of AI.

1. The most immediately recognizable risk is **retention after an engagement ends**. Most people understand intuitively that when a professional relationship concludes, the information shared during that relationship should not be held indefinitely by someone else. In traditional practice, a firm controls its own records and can manage how long they are kept. When a vendor's platform has been used to process client data, the firm no longer controls the data. Many vendors retain data submitted to their platforms for extended periods, sometimes without a clear end date. Once an engagement concludes and the firm moves on, the vendor may continue holding client financial records in ways the firm can no longer monitor or influence. The client has no way of knowing this is happening, and the firm may no longer be in a position to do anything about it.
2. The second risk, **vendor employee access**, is also relatively familiar because it involves the idea of unauthorized people having access to private information. Technical staff at AI vendors often have legitimate access to customer data for the purposes of maintaining and troubleshooting the system. This access is necessary for the platform to function, but it means that people outside the accounting firm, who have no professional relationship with the client and no obligation under the AICPA Code, can view sensitive financial information. Without contractual protections that specifically restrict and govern that access, the firm has no meaningful control over who at the vendor organization can see client data or under what circumstances.
3. The third risk involves **storage of data in several jurisdictions**, but in order to explain this point, it is necessary to provide some background information. Data is not usually stored on one server. Rather, cloud platforms normally store it on several servers located in different countries in order to allocate the capacity better and maintain the availability of the service. Each country has its own legislation concerning the right of the government to access data stored in it. In other words, the client's bank accounts that were supposed to be stored safely could be stored on the servers located in the jurisdiction where the government bodies had the right to access this data without notifying anyone about it.
4. The final, and perhaps least obvious risk of those discussed here, is the **use of client data for the purposes of training AI models**. It is certainly the one that needs to be explained most thoroughly because there is no way one can run into this problem anywhere else but in AI. Indeed, many vendors of AI software put clauses in their user agreements that enable them to use customer data to improve their algorithms. This is how many platforms evolve. As far as accounting firms are concerned, the implications of this issue are quite serious. Submitting the records of your clients to such a platform will make it possible for you to use the information about these records as a part of the training dataset which determines the output of the platform for all the other firms that use it, including the competitors of the client you serve.

Confidentiality in Practice

The AICPA's guidance on third-party service providers makes one thing clear. When a firm uses an outside vendor to process client data, the firm does not hand over its confidentiality responsibility along with the data. The obligation to protect client information stays with the firm, regardless of whose platform is being used to do the work.

What does this mean in practice? Before submitting any client data to an AI platform, the accountant should have a conversation with the client. The client needs to know which tool is being used, who the vendor is, and what will happen to their information once it enters that system. This is not a formality. It is the client's right to know how their information is being handled, and the accountant's obligation to tell them.

Before any client data is submitted, the vendor agreement should also be read carefully. Many firms accept these agreements quickly without reviewing them in detail, which is how confidentiality breaches occur without anyone making a deliberate mistake. Where the agreement includes provisions that are inconsistent with the firm's obligations, such as allowing the vendor to use client data for model training, those provisions should be negotiated or the tool should not be used for client work.

A straightforward way to reduce confidentiality risk is to submit only the data the AI tool actually needs for the specific task at hand. If the platform requires transaction dates and amounts, there is no reason to submit names, addresses, or other identifying

details alongside them. Providing less information reduces the potential consequences if something goes wrong.

Confidentiality in an AI environment is no longer maintained simply by professional habit and goodwill. Regulators are paying increasing attention to how firms manage technology in their work, and documentation of those practices is becoming a baseline expectation across the profession. An accountant who can demonstrate that they reviewed the vendor agreement, obtained client consent, and kept clear records of how client data was handled is in a far stronger professional position than one who relied on the assumption that the platform was safe to use.

4.4: Societal Consequences

Data Privacy and Security in AI

We examined how accountants should adhere to their professional responsibilities when working with AI solutions. However, it is worth adding that these professional responsibilities do not stop at the internal workflow. In the course of processing client data with the help of AI solutions, an accountant must comply with specific requirements related to the handling of this information.

This problem is particularly relevant today when the accounting professionals use various platforms, which perform the processing operations. In this connection, client information may be sent to the vendor's server, stored in multiple jurisdictions, retained after the completion of the engagement, and used to train models for other users. Moreover, the output of the model may serve to make audit decisions or to decide whether the client has tax liabilities. All these factors create the risk of privacy violations, legal actions, and damage to the professional reputation of the accountant. The accountant is responsible for how the client information is processed, and data privacy and security belong to the list of the main problems of today's practice.

The Data Privacy Regulatory Framework

There are several legislative acts, which regulate data privacy and data security in relation to personal and financial data. At the same time, it is rather problematic to find common requirements for all cases because the rules vary depending on where the client lives, what kind of data should be protected, and what AI platform is used.

The General Data Protection Regulation, or GDPR, is a law, which is active in the European Union and regulates the use of personal data. Although this act is created for European customers, it applies to any organization, which holds data of individuals, living in the EU, regardless of the geographical location of the firm. The GDPR gives individuals the right to obtain the information about how their data is being processed, the right to deletion in some situations, and the right to explanation, in the case of automated decision-making. In such a way, when using an AI tool to process data of European clients, an accountant should comply with these provisions.

As for the privacy obligations in the USA, they are regulated by both federal and state laws. According to the Gramm-Leach-Bliley Act, the firm should protect customer financial information and also analyze the practices of third-party vendors in order to ensure data privacy and security in the case of outsourcing. In the case of using an AI solution, an accountant cannot simply rely on the vendor's statements.

Also, California has enacted two important pieces of legislation, which guarantee the privacy of personal data. The California Consumer Privacy Act and the California Privacy Rights Act provide Californians with the right to know what information is being collected, the right to deletion, and the right to opt out of the sale or sharing of personal information. In many cases, the privacy protection extends to the employee data, which means that when providing HR services to a California resident, an accountant should consider these laws. Virginia, Colorado, and Connecticut have also adopted similar laws, and more states are expected to enact such legislation. The practical outcome is that the accountant cannot rely on a single data privacy standard. Which act is applicable in a particular case depends on where the client and his data are located.

Data Governance Obligations for AI

However, the fulfillment of data privacy requirements is not the end of the discussion. There is a need for proper data governance, which includes the policies and practices, used in the process of management of data throughout its whole lifecycle. In the environment of AI usage, these practices should be not just written. They should be implemented in everyday operation.

An important principle in the sphere of data governance is the data classification. Some information is more sensitive than others. Personal identifiers in the financial records, such as name, tax identification number, banking details, etc., should be handled differently compared to the operational data. Correct data classification allows determining how data should be stored, accessed, and what legal or ethical requirements apply to this data. Another important principle is the purpose limitation. Data should be processed only for the reasons for which it was initially collected. If a client record was obtained in the course of the audit engagement, this record should not be further processed by AI to improve a vendor's model, unless the client explicitly agrees with this purpose. Thus, it is possible to protect the client from unexpected use of the information.

Also, the vendor due diligence is an important aspect. Before selecting an AI platform, and before processing any client data with its help, an accountant should learn how this vendor manages the information, what measures are taken in terms of security, how data is retained and deleted after the expiration of the period of retention. A vendor that cannot clearly answer these questions cannot be considered safe to handle the sensitive data. The data minimization is another important practice. An accountant should give the AI only the data it needs to perform the task. If a tool needs transaction dates and amounts, there is no need to share names, addresses, or any other personal identifiers.

The retention and deletion of the data completes the list. The accountant should know how long data will be kept in the system, where it will be stored, and how it will be permanently deleted. Without the proper policies in the sphere of retention, the firm will store the information for a longer time than is legally or ethically justified.

The Wojeski & Company Case

Wojeski & Company, a New York accounting firm, suffered two cybersecurity incidents in 2023 and 2024 that exposed clients' personal information, including Social Security numbers and financial account numbers. The New York Attorney General said the firm failed to adequately protect the data and delayed notice to affected individuals for more than a year, and the firm agreed to pay \$60,000 and strengthen its security practices.

Connection to AI in Accounting

The point is broader than this one breach. Whether the threat comes from phishing, weak encryption, or an AI vendor relationship, accounting firms still have the same core obligations to safeguard sensitive client data and respond quickly when something goes wrong.

Cybersecurity Risks in AI Accounting Systems

The AI systems for accounting purposes are vulnerable to attacks because they unite sensitive information, financial decision-making, and connected software systems. They may contain huge volumes of confidential information, are connected to the accounting platforms directly and produce the output, which is used in the process of professional judgment. This makes these systems interesting for attackers, who want to steal this information or to alter the output.

First, there is a danger of model poisoning. In this case, the attacker gets unauthorized access to the learning process of an AI and inputs false or altered data to force it to work in the interest of the attacker. For example, a model, which is used for the purpose of fraud detection, may be trained to ignore certain types of transactions, creating a blind spot that may be used for further fraudulent transactions. Another threat is the adversarial input. In this case, the attacker alters the data to be entered into the system in a way, which will affect the output in a favorable way for him. Such alterations are usually small and may pass undetected by a person, but the person, who knows the mechanism of work of the model, will be able to use this method to alter the results.

The theft of data is another significant risk. AI systems are usually used to process a huge volume of sensitive data. However, unlike theft in traditional sense, data theft in this case may be hidden for a long period of time. Later, the stolen information may be used in other organizations' models or may appear in the possession of criminal groups, thus making the source of leakage practically impossible to track. The Application Programming Interface (API) exploitation is yet another threat to AI systems. APIs are technical connections between the AI tool and the financial software, the client database, and other systems. If these connections are not protected, an attacker may extract data from the system, observe how it works, or feed the system with the false data.

The Legacy Professionals Case

Legacy Professionals is an Illinois certified public accounting firm that serves labor unions, nonprofit organizations, and employee benefit plans. In late April 2024, the firm detected suspicious activity on its network and later determined that an unauthorized actor had removed files containing sensitive personal information, including names, Social Security numbers, and medical and health insurance information. Reports say the breach affected about 216,752 individuals and that the firm notified affected people in February 2025, roughly ten months after the incident was first detected, with multiple class action lawsuits following.

Connection to AI in Accounting

Legacy Professionals did not experience this breach through an AI system. But the case demonstrates the scale of harm that can result when an accounting firm holding large volumes of sensitive personal information does not have adequate controls in place and does not respond promptly when those controls fail. When AI platforms are part of a firm's workflow, these risks grow more complex. Client data submitted to a third-party vendor may be held across multiple systems in multiple locations, the firm may have limited visibility into what happened if a breach occurs, and direct control over the vendor's response cannot be assumed. Reviewing vendor agreements carefully, limiting the data submitted to what is strictly necessary, and maintaining documented security standards are the practical steps that allow a firm to meet its obligations to clients when something goes wrong.

Societal Consequences

The data privacy and security considerations mentioned above do not have importance only for the accountant and the client. They relate to larger issues of social justice and the role of accounting in society.

When the AI tools, used in financial reporting or tax administration, produce a biased output, it means that these biases are incorporated in the financial decisions of real clients. Accounting does not have only a task to record the transactions. It has much bigger role – to make decisions about the allocation of capital and public resources and to provide trustworthy financial information. When AI tools are applied in accounting without sufficient knowledge and understanding of the limitations of these tools, the consequences are likely to extend far beyond the individual firm.

The experience of the Internal Revenue Service illustrates this issue perfectly. The IRS uses automated tools to select tax returns for audit. Returns receive risk scores and those returns with high scores are subject to further inspection. In January 2023, a research team from Stanford University, University of Michigan, University of Chicago, and the United States Treasury Department published the study, according to which Black taxpayers are audited roughly three to five times more often than non-Black taxpayers, despite the fact that the automated selection process does not involve any reference to the race of the applicant. Moreover, the bias was highest among those taxpayers who received Earned Income Tax Credit, which is a tax benefit designed to assist low and moderate income workers. While Black taxpayers comprised 21 percent of all taxpayers receiving EITC, 43 percent of audits fell on them. Also, it was stated that the difference could not be explained by the variations in income, family size, household composition, or the frequency of tax underreporting. The Government Accountability Office, which is an independent agency that conducts the audit of the functioning of the federal government, analyzed the IRS automated systems separately and found out that the automated selection program of the IRS had not been comprehensively checked for biases, some of the risk scores in this program depended on the applicant's gender, and some parts of this program had not been updated since 2001.

The bias in this case was apparently not caused by any intentional discrimination. The more probable explanation is that the IRS automated tool was built using historical data about past audit results, which reflected patterns of discrimination in past enforcement practices of financial institutions and government. The algorithm learned from that history and replicated it, selecting taxpayers for audits based on the patterns of past behavior, not because of any errors in the return itself. The consequences of this bias may be serious. Public resources may be allocated incorrectly, not because of any malicious actions of individuals, but because the system was built on the basis of incomplete or biased data.

This case is directly relevant for the accounting professionals because tax preparation, tax advisory and audit support are core services provided by accounting firms. The use of biased systems in the course of rendering of such services means unfair consequences for clients. An individual taxpayer being selected for audit based on biased algorithm will bear unnecessary costs and stress. The accounting professionals, relying on AI tools for assessment of audit risk, detecting potential compliance problems, and advising clients on tax positions are fully responsible for checking the origins of such a tool and the possible impact of its outputs on their clients.

In addition, there is a collective risk for the profession. If many firms use similar tools, based on the same design and the same training data, they will inherit similar biases. The 2008 financial crisis showed exactly what this means. Banks and other financial institutions were issuing mortgages to the borrowers, who did not have reliable incomes. Those mortgages were assembled in financial products and sold to investors worldwide. Despite the evident risks, the products received strong rating and were purchased by many banks and investment firms. When the house prices started to fall, and many borrowers could not pay their loans, the problem became global. Banks lost money, jobs, pensions and houses of millions of Americans were lost because of the collapse of the mortgage market. This crisis shows the collective vulnerability of the market. Many institutions had invested in similar types of products and believed in the same myths and assumptions. The same thing happens when AI is used in

accounting practice. If many firms use similar tools, they will be vulnerable to the same problem. Once it appears, the problem will spread across the profession and will affect the entire financial reporting ecosystem, which investors, regulators, and the public rely on.

Professional Responsibility in a Connected System

Thus, the proper use of AI in accounting requires not only skills to work with the tool itself. The accountant has to check how the tool reaches its conclusion, what data the tool was trained on, and if it produces the fair results for all clients. In practice, before adopting an AI platform, the firm has to carefully analyze the vendor agreement in order to understand how the client data will be stored, used and deleted. While using the tool, during the engagement, only the necessary data should be provided to the tool. After completing the task, the firm should document the used tools and data. Finally, on a regular basis, the firm has to check if the tool is producing reliable results for all kinds of clients. If the tool produces reliable results for one category of clients and unreliable for the other, this tool will require additional measures before its use.

Accounting is the profession, which is based on trust. The trust of the client who provides his financial information, the trust of investors who relies on the information about the company, and the trust of the regulator who analyzes the reports is given to the accountant. The professional is responsible for verification of technological outputs.

4.5: AI Ethics Policy for an Accounting Firm

Why a Policy Framework Matters

An AI ethics policy provides guidance on the manner in which the organization manages artificial intelligence within its entire professional practices. It provides guiding principles on the values of the organization and how the organization intends to adopt technologies in line with these values. The policy on AI ethics is not merely about compliance. Rather, it is an expression of the commitment of the organization to integrity and professionalism. In its absence, the use of artificial intelligence will be without any standards and individual judgment calls.

A well-designed policy serves four key purposes:

- It protects clients by ensuring their data is handled ethically and securely.
- It protects the firm by documenting standards against which conduct can be evaluated.
- It protects professionals by providing guidance for ethical decision-making in complex situations.
- It protects the accounting profession by reinforcing public trust through responsible AI adoption.

The Seven Components of an AI Ethics Policy

Developing an AI ethics policy requires a firm to make deliberate decisions across seven areas, each addressing a different dimension of responsible AI use in professional practice.

Component 1: Scope and Applicability

The policy must clearly define its scope, covering all AI-assisted tools used in professional work. This includes language models for client communications, audit sampling platforms, tax analysis software, and predictive practice management tools. It should identify which tools and personnel are covered (including contractors) and note special compliance requirements for engagements involving regulated entities or sensitive client data.

Component 2: Approved AI Use and Pre-Adoption Assessment

The firm should not permit ad hoc adoption of AI by individual practitioners without organizational review. The policy should establish a pre-adoption assessment process that evaluates any new AI tool on all five criteria.

- Vendor data practices: Does the vendor provide clear, contractually binding commitments about how client data is handled, stored, and protected?
- Bias assessment: Has the tool been tested for bias relevant to the firm's client base? Can the vendor provide validation documentation?
- Explainability: Can the tool provide explanations for its outputs at a level sufficient to support professional judgment and documentation requirements?
- Regulatory compliance: Does the tool's audit trail functionality meet applicable record-retention and internal control documentation requirements?
- Security certification: Has the tool been independently assessed for cybersecurity vulnerabilities relevant to financial data?

Component 3: Human Oversight Requirements

All tools used in the professional accounting engagements need to be evaluated and approved by a qualified professional before any decision is made based on the output. The policy needs to describe the points in the process where the evaluation of the output becomes mandatory. Every judgment that would be included in the audit opinion, financial statement, or tax return needs to be evaluated by the professional regardless of how the output has been obtained. The same applies to the risk assessments that define the scope of the engagement, to every output that looks unusual in the context of what the professional knows about the client, and to every communication from the client that has been produced or significantly influenced by an AI tool. In all these cases the professional is not merely signing off on the output. He or she is applying the judgment that the client, regulator,

and the profession expects from him or her.

Component 4: Data Privacy and Security Standards

The policy should set up clear minimal standards for handling client data whenever AI tools are used. Only the platforms that have been evaluated and approved by the firm prior to the adoption can be used for processing client data. Prior to uploading the data on an external platform, any personally identifiable details must be stripped from it where such action would still allow completing the task. If the usage of AI tools entails processing client data for other purposes than just the current engagement, the consent of the client must be obtained in advance and properly documented and stored in the file. Prior to engaging the vendor, the firm should get the confirmation in writing for how long it would retain the client data and how it would permanently delete it after the engagement is finished.

Component 5: Documentation and Audit Trail Requirements

The policy should describe the type of documentation that is needed if AI tools are used for the engagement. At a minimum, it should include the name and version of the AI tool used, data inputs to the tool and its outputs, the professional's evaluation of the AI output including concerns (if any) and how they have been addressed, and occasions when the professional did not follow the recommendation of the tool and why.

Component 6: Training and Competency Requirements

The AICPA Code of Professional Conduct, Section 1.300.001, requires that the accounting professionals possess sufficient skills and knowledge to perform the work and continuously maintain their competency. Where AI tools are becoming a part of the work, competence includes the knowledge of how these tools operate and which of their outputs can be trusted and which need to be further evaluated.

The policy should mandate that every professional using an AI tool in client work would complete the training on that tool before starting to use it. This training should be refreshed whenever the tool changes or a new tool is introduced. Besides the training specific to particular tools, every professional should get trained on how to detect bias in AI outputs and fulfill his or her responsibilities regarding data privacy when using external platforms. The policy should provide a mechanism for raising concerns and reporting any ethical issues related to the AI tool use with the guarantee that no professional reprisals would ensue.

Component 7: Governance, Review, and Enforcement

A written policy only functions if someone is responsible for maintaining and enforcing it. The firm should designate a governance body with representation from its technical, legal, and practice leadership. That group should own the policy, monitor how it is being followed, and have the authority to act when it is not.

The policy should be reviewed on a regular cycle, with a full review taking place at least once a year. If a significant incident occurs or a major regulatory change takes effect, a review should happen immediately rather than waiting for the scheduled cycle. Where violations occur, the consequences should be clearly defined and applied in proportion to the seriousness of the breach. The firm should also maintain a straightforward process for staff to raise ethical concerns about AI use, one that is genuinely accessible and that professionals feel safe using.

A case from one of the world's largest accounting firms illustrates what is at stake when firms do not establish clear boundaries around AI use in professional development.

The KPMG Australia Case

In February 2026, a partner at KPMG Australia was fined more than A\$10,000 after using an AI tool to complete an internal training course on the use of AI. The partner uploaded firm reference materials into an AI platform to help answer exam questions, a direct breach of the firm's internal AI use policy. The misconduct was not self-reported. It was detected by monitoring software the firm had introduced specifically to identify unauthorized AI use in testing.

Following an internal investigation, the partner was required to redo the assessment and self-report the misconduct to Chartered Accountants ANZ, which opened its own examination of the case. The Australian Securities and Investments Commission was also notified. The case was not isolated. KPMG Australia identified 28 separate instances of staff using AI tools to complete internal assessments in the same period, involving employees at manager level and below as well as the senior partner.

KPMG has since committed to reporting AI-related misconduct cases separately in its annual results, an acknowledgment that transparency around AI governance has become a professional expectation.

Connection to AI Ethics Policy

AI misuse does not occur only in client-facing work. It can also appear within the professional development processes that firms rely on to build competence and maintain standards. Seniority does not lessen individual accountability. In one instance, a partner faced formal disciplinary action and was required to self-report to a professional body. This demonstrates that the responsibility to use AI appropriately applies to everyone within a firm.

The KPMG Australia case shows that an effective AI training and competency policy must clearly define boundaries, communicate those boundaries transparently, and have monitoring systems in place to enforce them. A firm that claims to be a leader in AI advisory services but fails to manage internal AI use effectively faces a credibility issue that no amount of external marketing can fix.

Implementation of an Ethics Policy

Writing a policy creates the standard. Putting it into practice changes how the firm actually operates.

Table 4.5.1: Implementation stages and indicators of successful completion

Stage	What it Involves	How to know it is Accomplished
Foundation	Create a complete inventory of every AI tool currently in use across the firm, including tools adopted informally by individual staff. Assess each tool against the pre-adoption criteria in Component 2. Prioritize addressing gaps based on risk level, starting with tools that handle the most sensitive client data.	A documented inventory exists. Each tool has been assessed. A remediation plan is in place for any tool that did not meet the criteria.
Training	Train partners and managers first so they can reinforce the standards with their teams. All other staff should complete training before using any approved AI tool in client work. Keep training records and establish a schedule for annual refresher training from the outset.	All personnel have completed training before using AI in client work. Training records are maintained. An annual refresher schedule is in place.
Process Integration	Update engagement templates to document which AI tools were used, what data was submitted, and how outputs were reviewed. Revise client engagement letters to address AI use. Add checkpoints to quality review checklists that confirm the policy was followed before work is finalized.	Engagement templates and client letters have been updated. Quality review checklists include AI oversight checkpoints.
Monitoring	Establish a regular schedule for reviewing AI tool performance and checking whether the policy is being followed consistently. Set up a clear communication channel for staff to report concerns, and test it before it is needed. Schedule the first annual policy review on the day the policy is adopted.	A monitoring schedule is operational. A concern reporting channel has been tested. The first annual review is scheduled.

5: EMERGING TRENDS IN AI AND ACCOUNTING SOFTWARE

Chapter Overview

The accounting field is at a turning point. Artificial intelligence is the third big technological change for accountants, after electronic spreadsheets and cloud computing. It's expected to be the most significant change yet. Unlike previous innovations that mainly made things more efficient, AI completely changes how accounting work is done, who does it, and what skills are most important. This new technology is not just about making things easier, it's about changing the way accountants work and what they need to know to do their jobs well.

This chapter examines the cutting edge of AI in accounting. We explore emerging technologies including robotic process automation, blockchain integration, and decentralized AI systems. More importantly, we address how accounting professionals can harness these tools responsibly while maintaining the trust, accuracy, and ethical standards that define the profession.

Learning Objectives

Every business relies on accounting to track financial activity, report performance, and maintain compliance. As a business professional, you will interact with accounting systems, use financial information for decision-making, and operate in environments where artificial intelligence (AI) is reshaping accounting processes and roles.

This chapter introduces key AI technologies used in accounting, examines their impact on accounting work and professional roles, and explores how organizations can implement AI in a trustworthy and responsible way. By the end of this chapter, you will be able to:

- Examine emerging AI technologies in accounting, including robotic process automation, blockchain, and agentic AI
- Evaluate how AI is transforming accounting processes and its impact on the accounting profession
- Explain how organizations can establish trust, accuracy, and verification in AI-driven accounting systems
- Apply strategies for implementing ethical and reliable AI practices in accounting environments

5.1: Emerging AI Trends in Accounting

5.2: Evaluation of the Impact of AI in the Accounting Profession

5.3: Creating and Implementing Guardrails

5.1: Emerging AI Trends in Accounting

Robotic Process Automation (RPA)

RPA is software that performs repetitive, rules-based digital tasks by mimicking human interactions with computer systems. Unlike general artificial intelligence, which can learn and adapt to new situations, RPA follows predetermined rules and excels at executing structured processes with consistency and speed.

Robotic Process Automation, or RPA, is a technology that's really making a difference in accounting right now. The best part is, it doesn't need a whole new system to work - it just uses the existing software and does tasks that people would normally do by hand. Accounting is a great fit for RPA because it involves a lot of repetitive tasks that follow rules, like matching invoices with orders, reconciling accounts, filing reports, and moving data between systems. These are all things that bots can do faster and more accurately than people. The good news is, this means two big benefits: first, the routine work gets done right, and second, people have more time to focus on the important stuff, like analyzing data, coming up with strategies, and making tough decisions. This way, accountants can use their skills and judgment to make a real difference, instead of just doing the same tasks over and over.

Note: RPA vs. AI

RPA automates structured, rule-based tasks. It does not "think." When combined with machine learning or natural language processing, the resulting system is often called *intelligent process automation (IPA)* and that combination is where the most significant accounting transformation is occurring.

Accounting Applications of RPA

RPA is already deployed across accounting functions.

Table 5.1.1: Robotic process automation across the four core accounting functions

Accounting Functions	What RPA Does	Outcome
Transaction Processing	Extracts data from invoices, validates against purchase orders, records transactions in the accounting system, flags exceptions.	Faster processing and fewer manual data-entry errors.
Account Reconciliation (verifying that records from two different sources match)	Matches transactions across systems, identifies discrepancies, prepares reconciliation reports.	Faster month-end close (the process of finalizing financial records at the end of each period), consistent application of matching rules.
Compliance and Reporting	Handles routine regulatory filings, generates standard financial reports, and maintains audit trails.	Reduced compliance risk and as a result, professional time is freed for more complex judgments.
Data Migration	Transfers historical data during system implementations, maps accounts across structures.	Reduced conversion errors and provide documented process trail.

The transformative power of RPA lies in scalability. A single software robot processes thousands of transactions with the same care it applies to one, without fatigue. Modern RPA platforms also allow accountants to configure processes using visual, drag-and-drop interfaces, reducing dependence on IT departments for routine automation.

However, RPA has limits. It struggles with unstructured data and cannot exercise professional judgment. A system that processes structured invoices smoothly will fail when it encounters a handwritten note, an unusual contract term, or a situation that falls outside its programmed rules. These limitations are precisely why RPA's future lies in combination with AI's analytical capabilities.

Blockchain Technology

Blockchain is a type of distributed ledger technology, meaning it is a database stored and maintained simultaneously across many computers rather than on one central server. Its most well-known application is as the foundation of crypto currencies like Bitcoin, but its implications for accounting extend far beyond digital currency.

Note: How Blockchain Works

Transactions are grouped into blocks that are cryptographically linked to the previous block, forming a chain. Each participating computer, called a node, holds a complete copy of this chain. To alter a historical record, an attacker would need to simultaneously alter every copy on the network, a task that is computationally impractical in large, well-established networks.

Triple-Entry Accounting

The concept of triple-entry bookkeeping is a game-changer for accounting, and it's all thanks to blockchain technology. This idea was first introduced by Yuji Ijiri back in 1986, and later built upon by Ian Grigg in 2005. So, what's the big deal about triple-entry bookkeeping? Well, traditional double-entry accounting only records transactions twice, once as a debit and once as a credit, within a single company's books. But triple-entry accounting takes it to the next level by adding a third entry on a shared, immutable blockchain ledger. This ledger is accessible to both parties involved in the transaction, as well as their auditors, who can all verify the information independently. The third entry serves as a cryptographically sealed receipt that can't be altered after the fact, providing an extra layer of security and transparency. This means that both parties can trust that the transaction has been recorded accurately, without the need for intermediaries or centralized authorities. It's a pretty significant development, and one that could potentially revolutionize the way we do accounting.

Smart Contracts

Smart contracts are like digital agreements that automatically do what they're supposed to when certain conditions are met. They're written in code and stored on a blockchain, which is a special kind of database. So, for instance, let's say you have a smart contract that says payment will be made as soon as a delivery is confirmed. As soon as that delivery is verified, the payment gets transferred, the transaction gets recorded in both parties' books, and any tax liabilities get updated - all without anyone having to lift a finger. This makes it a lot less likely for disputes to happen, and it also means that businesses or individuals on either side of the transaction don't have to waste time double-checking everything with each other.

Implications for Audit

Blockchain creates enhanced, tamper-resistant audit trails. Every transaction includes metadata about its origin, timing, and content, and that metadata cannot be retroactively modified. This holds significant implications for auditors, who currently rely heavily on reconstructing transaction histories from records that could, in principle, be altered.

Current Limitations

Despite its theoretical promise, widespread blockchain adoption in mainstream accounting remains limited as of 2025. Significant barriers exist, including scalability challenges, high energy consumption (particularly for proof-of-work blockchains), complex integration with existing accounting software, and unresolved questions about regulatory treatment and privacy. Solutions like private blockchains, where access is restricted and controlled rather than fully open, are emerging to address some privacy concerns, but they introduce tradeoffs around decentralization. Students should understand blockchain's potential without overstating its current adoption.

Decentralized AI and Federated Learning

Most accounting systems that use artificial intelligence today are controlled by one main server. This means that all the data is sent to a central location, often a cloud service provided by a vendor, where it's processed and then sent back. This way of doing things can cause problems because it goes against the basic rule of accounting, which is to keep client information private.

Federated learning is a way to deal with the problem of sharing data. It lets an AI model learn from data that is spread out across many places, without actually moving the data from where it is. Each group that is taking part in the system trains the model using their own data, and then they share the changes they made to the model with a main person in charge. This main person puts all the changes together to make the model better, and then sends the new model back to the groups. This way, the data stays where it is, and only the changes to the model are shared.

In an accounting context, consider a tax AI that becomes smarter by learning from patterns across thousands of firms' returns without any single firm ever transmitting confidential client records to a third party.

Why Federated Learning Matters for Accountants

Companies can use artificial intelligence without sharing private client information. This is because they can work together to train AI systems without revealing sensitive details. This approach also helps businesses follow strict data protection rules. For example, the General Data Protection Regulation (GDPR) in Europe and the California Consumer Privacy Act (CCPA) in the US have strict guidelines on how personal data is handled. By using federated learning, companies can improve their AI systems while still complying with these regulations. Smaller companies can also benefit from this approach, as they can access AI advancements made by others in the industry without having to compromise their own client relationships. This creates a level playing field, where all companies can compete fairly.

Decentralized Autonomous Organizations, commonly referred to as DAOs, represent a related frontier. A DAO is an organization that runs through rules encoded in smart contracts on a blockchain, replacing traditional management with collective stakeholder voting. No single person or executive team is in charge; instead, decisions are made automatically based on agreed-upon code and community input. While DAOs in accounting are largely experimental, they illustrate a direction in which AI governance and accounting service delivery could evolve, particularly for decentralized financial platforms and crypto-native businesses, meaning companies whose operations are built entirely on blockchain and digital asset infrastructure.

Agentic AI

Agentic AI refers to artificial intelligence systems that can accomplish a specific goal with limited human supervision. According to IBM, these systems consist of AI agents, which are machine learning models that mimic human decision-making to solve problems in real time. In a multi-agent system, each agent handles a specific subtask, and their efforts are coordinated through a process called AI orchestration. The word "agentic" captures what makes these systems distinct from earlier AI tools. They have agency, meaning the capacity to act independently and purposefully toward a goal.

Traditional artificial intelligence systems work in a pretty straightforward way, they follow a set of rules and wait for a person to tell them what to do next. But Agentic AI is different. It can take a big project, break it down into smaller tasks, find the information it needs, take action, and adjust its plan if something doesn't work out. And it can do all of this without a person needing to manage every single step. Instead, the person can just review the results and make sure everything is going as planned. This way, Agentic AI can work more independently and get things done more efficiently. It's like having a helper that can take care of a lot of the details, so you can focus on the bigger picture.

Note: Agentic AI vs Generative AI

Both types of AI use large language models, but they are built for different jobs. Generative AI is reactive: it waits for a prompt, processes it, and produces an output such as text, an image, or code. Agentic AI is proactive: it is built on top of generative AI models but adds layers of reasoning, planning, and autonomous action. Where a generative AI responds to what you ask, an agentic AI pursues a goal. A generative AI can draft a vendor payment summary if asked. An agentic AI can monitor vendor invoices, flag discrepancies, initiate approvals, and post entries without being asked at each step.

Table 5.1.2: Five characteristics of agentic AI and their applications in accounting

Characteristic	What It Means	Accounting Application
Autonomous	Performs tasks without constant human oversight and tracks progress toward long-term goals over time.	Monitors a vendor portfolio continuously rather than waiting to be asked.

Characteristic	What It Means	Accounting Application
Proactive	Searches the web, interacts with external software interfaces, and queries databases to gather information before acting.	Pulls new tax law changes and flags which business transactions they might affect.
Specialized	Individual agents focus on specific tasks; complex systems use a conductor agent to coordinate the work of simpler ones.	One agent extracts invoice data, another validates it, and a third posts entries and flags exceptions.
Adaptable	Learns from experience, takes in feedback, and adjusts behavior over time with appropriate guardrails in place.	Improves matching accuracy over time by learning from corrected exceptions.
Intuitive	Users interact through natural language rather than specialized interfaces, making complex tools more accessible.	A manager asks a plain-English question and receives an immediate, sourced answer.

How Agentic AI Work

Most Agentic AI frameworks follow a similar operational sequence. The system begins with perception, collecting data from external sources such as software interfaces, databases, or user input. It then moves to reasoning, interpreting that data to detect patterns and understand context. From there it sets goals, developing a strategy to accomplish the task at hand. In the decision-making step it evaluates possible actions and selects the best one. Execution follows, where the system acts on external tools or delivers output to a user. Finally, in the learning step, it evaluates what happened and refines its approach for future tasks.

Running across all of these steps is AI orchestration, the coordination and management of multiple agents working together. Orchestration platforms automate AI workflows, track progress, manage resources, monitor data flow, and respond to failures. With the right architecture, dozens, hundreds, or theoretically thousands of agents can work in concert toward a shared goal.

Video Title

Video Length: 12 minutes 58 seconds. Watch on YouTube for accessible close captions.

Context: This clip from The Accounting Podcast puts agentic AI to a real-world test: can ChatGPT-5's Agent Mode log into Xero and reconcile a checking account on its own, without a human typing a single entry? Hosts Blake Oliver and David Leary run the experiment live, narrating what the agent does well, where it stumbles, and what is the realistic boundary between automation and professional judgment today. The video offers a candid, unscripted look at where agentic AI in accounting currently stands and what it still cannot do without human oversight.

This interactive element is not available in this format. Visit the link or scan the QR code to access it online.

[View this YouTube Video online](#)



Reflection Questions

1. Blake and David watched the agent the entire time it worked. If an accounting firm deployed this agent for clients, who should be responsible if the agent makes an error, will it be the firm, the software vendor, or the AI? How does this compare to errors made by a junior staff accountant?
2. The episode mentions that 58% of employees are using AI despite company bans. If you were starting at an accounting firm today, how would you navigate a situation where AI tools could make you more productive but the firm has not officially approved them?

EY Agentic Platform 5.1.1

In March 2025, EY launched its EY.ai Agentic Platform in collaboration with NVIDIA. According to EY's published announcement, the platform deploys 150 specialized AI agents supporting 80,000 EY professionals in handling compliance workflows, autonomously managing data collection, document analysis, and filings across income and indirect tax domains. EY reports that the platform targets surpassing 3 million tax compliance outcomes and redefining more than 30 million tax processes annually across their tax practice. Note: These figures are vendor-reported by EY and have not been independently audited by a third party. They illustrate the scale of enterprise agentic deployment but should be read as company disclosures.

Impact

The shift from tools that assist to agents that execute represents a fundamental change in how accounting work will be organized, and in what human professionals are needed to oversee.

Challenges and Risks of Agentic AI

When we give AI systems the power to make decisions on their own, it can be really powerful, but it also comes with some big risks. A lot of these systems use something called reinforcement learning, which is like a reward system that teaches the AI to do what we want. But if we don't design that reward system carefully, the AI might find ways to cheat and get the reward without actually doing what we want. And because these systems are often made up of many different parts working together, one small mistake can cause big problems that spread quickly. For example, an AI that's supposed to make money from trading stocks might take risks that are bad for the market. An AI that's supposed to get people engaged might promote fake or misleading news. And an AI that's supposed to keep people safe from harmful content might end up censoring things that are actually okay. When an AI is only focused on one goal and doesn't have any checks in place, it can start to do things that aren't what we really want. That's why we need to make sure we build good governance into these systems from the very beginning, not just try to fix things after something goes wrong.

Discussion

The following questions are designed to move your thinking beyond definitions and toward application. Unlike review questions that ask you to recall what you read, reflection questions ask you to reason through a problem using the concepts you have learned. There are no single correct answers. What matters is that your response demonstrates understanding of the concept, engages seriously with the complexity of the situation, and supports your position with logical reasoning.

1. RPA automates structured tasks but cannot exercise judgment. Accounts payable is the department responsible for managing the money a business owes to its suppliers and vendors. When a company receives an invoice for goods or services it has purchased, the accounts payable team verifies the invoice, matches it to the original purchase order, approves it for payment, and processes the transaction. It is one of the most document-heavy and repetitive functions in any accounting department. With that in mind, identify three tasks in a typical accounts payable department that RPA could handle and two tasks that would still require a human accountant. Explain what distinguishes the tasks a machine can perform from those that require human involvement.
2. Triple-entry accounting using blockchain would theoretically allow auditors to verify transactions in real time rather than after the fact. If this became standard practice, how might the role of the external auditor change? Consider both what auditors might no longer need to do and what new responsibilities might emerge.
3. Federated learning allows AI to improve without centralizing data. What professional ethics obligations would an accounting firm need to consider before participating in a federated learning network that uses client data, even in anonymized form? Think about consent, transparency, and the limits of anonymization as a protection.

5.2: Evaluation of the Impact of AI in the Accounting Profession

Process-Level Transformations Across Accounting Functions

The transaction processing workflow has undergone the most visible transformation. AI systems that combine optical character recognition (OCR), machine learning, and Robotic Processing Automation (RPA) can now handle the full process from start to finish. They extract data from invoices in any format, validate information against purchase orders and historical patterns, interpret coding instructions through natural language processing, record transactions in the accounting system, initiate payments, and flag anomalies for human review.

The practical result is not simply speed but a shift in skill requirements. The function no longer primarily demands data entry proficiency. It now demands analytical review of exceptions and judgment about ambiguous cases.

Financial Reporting and Close Processes

AI accelerates the financial close process. Automated systems consolidate financial records across multiple business units, apply intercompany eliminations (the removal of transactions recorded between the company's own subsidiaries before combining financial records), identify potential errors through pattern analysis, and generate standard footnote disclosures (explanatory notes attached to financial statements) and management commentary. Processes that historically required weeks can now occur in days.

Artificial intelligence is not just about doing things faster, it's also about changing what accounting teams can actually do. For instance, natural language generation can take financial data and turn it into stories that make sense. This means accountants can give management and boards a better understanding of what's going on. Variance analysis is another area where AI helps, it compares what actually happened to what was planned, and flags any differences that need looking into. Then there are predictive models that use data to forecast how things will go in the future, based on what's driven performance in the past. All this means accountants can offer a lot more to management and boards, and really help them make informed decisions.

Audit and Assurance

Traditional audit methodology relies on sampling, which means testing a subset of transactions to draw conclusions about a population. AI makes continuous auditing and population-level testing possible. Rather than analyzing 5 to 10 percent of transactions, AI systems can analyze 100 percent, identifying anomalies that sampling would miss.

AI is changing the way auditors work. They don't have to spend as much time checking every little transaction. Instead, they can focus on looking into things that don't seem right and checking the assumptions that managers make when they put together financial statements. These assumptions are like educated guesses that managers use to make sure the financial statements are accurate. By using AI, auditors can make better judgments about whether the financial statements are reliable or not. This means they can give a clearer picture of a company's financial health.

Agentic AI Application

Major audit firms are deploying AI agents for vendor compliance monitoring. Rather than manually reviewing vendor documentation on a periodic cycle, AI agents continuously monitor risk indicators across entire vendor portfolios. Auditors then focus their expertise on investigating elevated-risk vendors and exercising professional judgment on complex situations. The role shifts from routine review to risk-focused oversight.

Tax Preparation and Planning

Tax is a great fit for AI because it's based on rules and is actually pretty complex. AI systems can take in financial data, apply

the right tax rules, find the best deductions and credits, fill out forms, and even explain things. When they come across a new tax situation, AI can quickly look up the relevant laws and draft a preliminary report for a professional to review. This makes the whole tax process more efficient and accurate. AI can handle a lot of the routine work, like preparing forms and researching laws, which frees up professionals to focus on more complicated issues. Overall, AI is really changing the way we do taxes, making it faster and more reliable.

Tax experts can now use special tools to figure out how much tax someone might owe in different situations. This helps them plan ahead and avoid problems, rather than just reacting to things after they happen. It's a big change in what tax professionals can do to help their clients.

Advisory Services

As AI does more routine compliance work, accounting professionals are repositioning toward advisory services. This is not simply a shift in service mix. It is a fundamental change in professional identity, from backward-looking scorekeepers to forward-looking strategic partners. AI-enabled advisors can build scenario models that integrate financial, operational, and market data to inform decisions about acquisitions, pricing strategy, and capital allocation (how to distribute the company's financial resources). These capabilities were previously impractical for most firms.

The Evolving Accounting Role

The skill profile for accounting professionals is shifting in identifiable ways.

Table 5.2.1: Changing importance of key accounting skill areas with the rise of AI

Skill Area	Direction of Change	Reason
Technical accounting knowledge	Essential but somewhat less differentiating	AI can apply rules, but professionals must evaluate whether the right rules were applied.
Analytical thinking and problem framing	Increasing importance	AI provides answers and humans must ask the right questions.
Technology literacy	Increasing importance	Must understand capabilities and limitations to use and oversee AI effectively.
Communication and client relationships	Differentiator	As AI handles technical execution, interpersonal skills become more valuable.
Ethical judgment and professional skepticism	Increasing importance	AI can calculate but it cannot determine what is right or appropriate.
Continuous learning capacity	Increasing importance	AI capabilities evolve rapidly and static knowledge becomes obsolete.

Strategic Implications

AI is fundamentally changing the economics of accounting firms. Organizations that built their business around efficiently executing routine compliance work now face pressure as AI performs that same work faster and at lower cost. The firms best positioned to adapt are those repositioning toward higher-value advisory work, deep industry specialization, or a quality of client experience that technology alone cannot replicate.

These shifts are also changing how firms are staffed. The traditional model relies on large numbers of junior professionals handling routine work, supervised by a smaller group of senior staff. As AI absorbs that junior-level work, the shape of the organization changes and billing models built around hours of human labor become harder to sustain. These changes are already underway and will continue to unfold over the coming years.

When jobs change or go away, companies have a duty to be upfront with their employees about what's happening. They should invest in helping their workers learn new skills and avoid making it seem like every change is just about cutting costs. Accounting firms, professional groups, and regulators all have a stake in making sure this transition happens in a fair way that gives workers real opportunities. They need to work together to make sure workers are treated fairly and have chances to succeed.

Discussion

1. A large audit firm announces it has implemented AI to perform 100% testing of all client transactions rather than sampling. What are the benefits of this approach? What new risks does it introduce that auditors must address?
2. If AI handles routine tax compliance, and accountants increasingly focus on advisory services, what does this mean for entry-level positions in accounting firms? How should accounting education respond?
3. The move from the billable-hour model to value-based or subscription pricing is accelerating as AI changes cost structures. What are the advantages and risks of this transition for (a) accounting firms, (b) clients, and (c) the profession's public trust role?

5.3: Creating and Implementing Guardrails

The Importance of Public Trust

Accounting plays a vital role in our society. The information in financial statements helps people make smart investment decisions that can impact a lot of lives. Tax returns are also crucial as they provide the funds needed for the government to operate. Moreover, audit opinions give people the confidence they need to participate in capital markets, where stocks and bonds are traded. Because accounting has such a big impact on the public, it's essential for the profession to adopt artificial intelligence in a responsible way. This is not just a good idea, but a professional obligation. By doing so, accounting can continue to serve the public's interest and maintain the trust that is so essential to its function.

A Framework for Responsible AI Implementation

A practical guardrails framework operates across three layers: input controls (what goes into the AI), process controls (how the AI operates and is monitored), and output controls (how results are verified before use).

Layer 1 — Input Controls

Table 5.3.1: Controls for ensuring accurate, secure, and trustworthy data inputs in AI systems

Control	Purpose	Implementation
Data Validation	Ensure accuracy and completeness before feeding AI.	Automated error checks; test AI systems with good data to establish accuracy baselines.
Privacy Protection	Prevent unauthorized exposure of client information.	Remove personally identifiable information before uploading to cloud-based AI; verify vendor data use contracts.
Access Controls	Limit AI access to necessary data only.	Role-based permissions; principle of least privilege (only granting access to the minimum data and systems a user actually needs) applied to AI system access.
Source Authentication	Verify data origin reliability.	AI cannot distinguish genuine invoices from fraudulent ones; human verification remains critical.

Layer 2 — Process Controls

Table 5.3.2: Controls for AI monitoring, transparency, and human review

Control	Purpose	Implementation
Algorithm Transparency	Understand why AI produced a result.	Prefer AI tools that explain their reasoning; be cautious with "black box" systems where outputs cannot be explained.
Audit Trail Maintenance	Reconstruct what AI did and how humans reviewed it.	Document prompts used, data provided, AI outputs generated, and human modifications or validations.

Control	Purpose	Implementation
Version Control	Ensure reproducibility.	Track which AI model version produced specific results; AI model updates can change outputs.
Human-in-the-Loop Requirements	Define when human review is mandatory.	High-dollar transactions, unusual circumstances, and first-time scenarios should trigger mandatory human review before action.

Layer 3 — Output Controls

Table 5.3.3: Controls for AI output verification, bias detection, and source validation

Control	Purpose	Implementation
Cross-Verification	Validate AI results against alternative methods.	Compare AI cash flow projections against historical patterns and industry benchmarks; reconcile AI-generated balances to source documents.
Reasonableness Review	Catch conclusions that don't make business sense.	Professional judgment about whether outputs are plausible remains the accountant's responsibility because AI cannot assess its own reasonableness.
Bias Detection	Ensure consistent, fair AI decisions.	Regularly audit AI outputs across different client types, transaction sizes, and demographic categories.
Citation Verification	Catch AI hallucinations in cited authorities.	Always independently verify any accounting standard, tax code provision, or case law cited by AI because AI systems can generate plausible but nonexistent citations.

Did you know?

We need to be careful when using artificial intelligence in accounting because it can sometimes provide false information that sounds true. This is called hallucination, where the system makes up facts and presents them with confidence, including fake references to accounting rules, tax laws, or court decisions. If we rely on this fake information, it can lead to serious mistakes, legal problems, and trouble with regulators. This is not just a theoretical risk, it's something that has actually happened with current AI systems. So, it's crucial that we never use any accounting rule or standard given by an AI without checking it against the original source to make sure it's accurate.

Governance and Policy

Having a solid organizational structure in place is crucial for controls to actually work. Effective governance is all about bringing together a team of people from different areas like accounting, technology, legal, and compliance to work together. This cross-functional team plays a vital role in evaluating AI tools before they're adopted, giving the green light for implementations, keeping an eye on performance over time, and making updates to policies as the technology continues to evolve. If you don't have this kind of ongoing oversight, individual controls can start to drift and become applied inconsistently, which can lead to all sorts of problems. It's really important to have a team that can provide continuous monitoring and guidance to ensure that controls are working as they should be. Without it, things can quickly fall apart.

Companies should have clear guidelines on when to use artificial intelligence and when not to. This includes situations where AI results need to be checked by a professional before being sent to a client. All employees working with AI must undergo training to understand what these tools can and can't do, as well as the company's rules. Firms need to be transparent about their use

of AI with clients and regulators. This transparency should be part of the formal agreements that outline the work being done for each client. It's essential to have these policies in writing to avoid any confusion or misunderstandings. By having these policies in place, companies can ensure that AI is used responsibly and effectively, and that clients are well-informed about how their data is being handled. This is crucial for building trust and maintaining a good reputation. Moreover, these guidelines can help prevent potential issues, such as AI results being misinterpreted or misused. By requiring professional review in certain situations, companies can minimize the risk of errors or inaccuracies. Overall, clear policies and procedures are essential for the responsible use of AI in business. They help protect clients, employees, and the company itself, while also promoting transparency and accountability.

Before adopting any third-party AI tool, the firm should evaluate the vendor's data security practices, how they use training data, how they handle updates, and what contractual protections they offer. Once a tool is in use, ongoing tracking of its accuracy, speed, and cost-effectiveness should be routine, with documented steps for responding when performance falls short. EY's published Responsible AI Framework illustrates this approach, describing governance that integrates guardrails at the individual agent level and emphasizes continuous monitoring rather than one-time implementation across cloud, on-premises, and edge environments. The underlying principle is that AI systems are not static. They are updated, retrained, and expanded, and oversight structures must keep pace.

Managing Specific AI Risks

Data privacy is the most immediate risk for any organization uploading confidential information to a cloud-based AI tool. Client data may be used for model training or exposed to unauthorized parties unless the firm has contractual guarantees preventing this. The practical response is to anonymize data wherever possible, use enterprise tools with enforceable data-use restrictions, and audit access controls on a regular basis.

AI systems trained on historical data can also perpetuate existing biases, producing outputs that discriminate certain types of clients or demographic groups in tax or financial analysis contexts. Mitigating this requires consistently testing AI outputs across varied scenarios and demographic groups, incorporating diverse perspectives in the evaluation process, and documenting the results.

Over-reliance on AI is a more subtle risk, but it's a real one. When professionals hand too much over to automated systems, their own judgment and skills start to fade. And the oversight that's supposed to catch AI's mistakes only works if people actually use it, the less that judgement is exercised, the less it catches. Firms can guard against this by keeping technical training active, rotating staff between AI-assisted and manual work, and treating independent judgment as something that has to be maintained, not something you can just assume will still be there.

Transparency with clients matters here too. If firms present AI-generated work as if a person did it, clients can end up with the wrong idea about what's actually behind a professional opinion. Firms need clear policies on when and how AI involvement gets disclosed, and those policies should be spelled out in the engagement letter and signed off on by both parties, not just live in an internal memo.

Finally, there's the growing threat of deep fakes. AI can now fake audio, video, and documents convincingly enough to enable real impersonation fraud, aimed at both firms and their clients. Staff need training to spot the warning signs, financial authorizations should always require multi-factor verification, and anything unusual should be double-checked through a separate channel before anyone acts on it.

Responsible AI Implementation Roadmap

Responsible AI adoption is a phased process, not a single deployment decision. The roadmap reflects a realistic implementation sequence for an accounting firm or department.

Table 5.3.4: Responsible AI implementation roadmap across four phases

Step	Description
Team Formation & Tool Selection (Months 1–2)	Establish a cross-functional AI project team. Define objectives, timelines, and measurable success criteria. Research and select AI tools appropriate to firm size, practice areas, and technical capacity. Provide initial training on tool operation and responsible-use standards.

Step	Description
Process Assessment & Use Case Development (Months 2–3)	Inventory existing processes to identify AI opportunities. Interview process owners about pain points. Map use cases to AI capabilities and prioritize based on impact, implementation complexity, and risk profile. Develop detailed specifications for pilot projects.
Pilot Implementation (Months 3–6)	Launch pilots in controlled, lower-risk settings. Implement the three-layer control framework (input, process, output). Document all procedures for AI use, output verification, and exception handling. Gather user feedback. Measure against established metrics. Refine based on lessons learned before expanding.
Scaling & Continuous Improvement (Months 6–12+)	Roll out successful pilots more broadly. Continue monitoring performance and gathering feedback. Adjust policies as AI capabilities and regulatory guidance evolve. Expand into additional use cases incrementally. Establish routine review cycles for AI performance, policy effectiveness, and emerging risks.

Building an AI-Ready Culture

Technology adoption succeeds or fails on culture as much as on the technical rollout itself. No guardrail framework works if staff are afraid to report AI errors, if leadership sees AI mainly as a way to cut costs rather than build capability, or if deadline pressure keeps squeezing out the time proper review actually takes.

People need to genuinely feel safe raising concerns about AI outputs, without worrying it'll be held against them. When staff stay quiet because they're afraid of the fallout, problems don't go away, they just pile up quietly until they turn into real ones. That kind of openness can't just be written into a policy. Leadership has to actually model it.

Ongoing learning also needs to be built into the job, not something people are left to squeeze in on their own time. AI capabilities move fast, and a team that stops learning falls behind its own tools. Firms that make training a regular part of how work gets done will end up in a much better position than the ones that treat it as optional.

And ultimately, the people closest to the work are usually the first to notice when something's off. Everyone, at every level, should feel able to speak up if an AI tool seems to be misused, producing skewed results, or drifting outside what it was actually built for. A culture that welcomes those conversations is one that can actually be trusted.

Discussion

1. Say your firm's AI tool is 99.7% accurate. Sounds great, right? But if you're processing 200,000 transactions a year, that 0.3% still adds up to 600 errors. You obviously can't have a human check all 200,000 transactions by hand. But you also can't just shrug off 600 mistakes. How would you design a system to catch and fix those errors efficiently, and who should be responsible for reviewing them?
2. A client tells you flat out: they don't want AI touching any part of the work your firm does for them. Here's the catch: your firm uses AI behind the scenes to pull and organize data and run a first pass of analysis, but a licensed professional reviews and signs off on every final document before it goes out. How do you respond to the client, and what would you put in the engagement letter so both sides are clear on how the work actually gets done?
3. Most professions have ethics rules that require practitioners to stay competent, protect client information, and question results rather than simply accept them. How do each of those three obligations apply when AI tools are doing part of the work? If an AI produces a result that looks correct but you have no way to verify how it reached that conclusion, what is your responsibility as the professional signing off on it?

INDEX

No index terms found.

GLOSSARY

S

Sample Word 1 | Sample Definition 1

DETAILED LICENSING

Overview

Title: Artificial Intelligence in Accounting

Webpages: 40

All licenses found:

- [CC BY 4.0](#): 75% (30 pages)
- [Undeclared](#): 15% (6 pages)
- [Public Domain](#): 10% (4 pages)

By Page

- Artificial Intelligence in Accounting - [CC BY 4.0](#)
 - Front Matter - *Undeclared*
 - Title Page - *Undeclared*
 - InfoPage - *Public Domain*
 - Program Page - *Public Domain*
 - Table of Contents - *Public Domain*
 - Licensing - *Public Domain*
 - Accessibility Statement - [CC BY 4.0](#)
 - About the Book - [CC BY 4.0](#)
 - About the Author - [CC BY 4.0](#)
 - 1: The roles of intelligence and artificial intelligence (AI) in accounting - [CC BY 4.0](#)
 - 1.1: Introduction and Technology in Accounting Today - [CC BY 4.0](#)
 - 1.2: The Concept of Intelligence - [CC BY 4.0](#)
 - 1.3: Different types of Artificial Intelligence - [CC BY 4.0](#)
 - 1.4: The Historical Development of Intelligence in Accounting - [CC BY 4.0](#)
 - 1.5: Analyzing Artificial Intelligence Strengths and Weaknesses in Accounting - [CC BY 4.0](#)
 - 1.6: Detailed Text Descriptions for Chapter 1 - [CC BY 4.0](#)
 - 2: AI in Financial and Managerial Accounting - [CC BY 4.0](#)
 - 2.1: Distinguish the roles of financial and managerial accounting purposes - [CC BY 4.0](#)
 - 2.2: AI-Driven Tools in Financial Accounting - [CC BY 4.0](#)
 - 2.3: AI's Role in Managerial Accounting - [CC BY 4.0](#)
 - 2.4: Opportunities and Challenges in AI Applications - [CC BY 4.0](#)
 - 3: Applications of AI in Audit and Tax - [CC BY 4.0](#)
 - 3.1: AI's Evolution in Audit and Tax - [CC BY 4.0](#)
 - 3.2: AI as a Tax Preparation Assistant - [CC BY 4.0](#)
 - 3.3: AI's Role in Auditing - [CC BY 4.0](#)
 - 4: Ethics of AI in Accounting - [CC BY 4.0](#)
 - 4.1: Why Ethics Matters? - [CC BY 4.0](#)
 - 4.2: Individual Professional Responsibilities - [CC BY 4.0](#)
 - 4.3: Professional Accounting Ethical Standards - [CC BY 4.0](#)
 - 4.4: Societal Consequences - [CC BY 4.0](#)
 - 4.5: AI Ethics Policy for an Accounting Firm - [CC BY 4.0](#)
 - 5: Emerging Trends in AI and Accounting Software - [CC BY 4.0](#)
 - 5.1: Emerging AI Trends in Accounting - [CC BY 4.0](#)
 - 5.2: Evaluation of the Impact of AI in the Accounting Profession - [CC BY 4.0](#)
 - 5.3: Creating and Implementing Guardrails - [CC BY 4.0](#)
 - Back Matter - *Undeclared*
 - Index - *Undeclared*
 - Glossary - *Undeclared*
 - Detailed Licensing - *Undeclared*

This page titled [Artificial Intelligence in Accounting](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by LibreTexts.