AI-Powered Fraud Detection: Beyond Rules-Based Systems to Machine Learning

Presented by: Paul E. Zikmund

IOFM

Do you need NASBA CPE credits?

- Navigate to website: **iofm.cnf.io**

  or scan the QR code →

- Check-in and check-out of your sessions to track your attendance for NASBA CPEs

- Certified with IOFM? No need to check-in and out of sessions. Self-report CEUs on IOFM.com instead after the event!

IOFM Spring
CONFERENCE & EXPO

AP in the crosshairs

AP controls the final mile of cash movement

AP handles high volumes under time pressure

AP interacts with external parties daily

AP often processes exceptions and urgent requests

Vendor data changes can look routine

Staff are expected to be responsive and service-oriented

IOFM

Weaponization of AI

76% of Organizations: According to the 2026 AFP Payments Fraud and Control Survey, over three-quarters of U.S. firms experienced attempted or actual payments fraud in 2025.

The Median Loss: Recent data from H1 2026 shows that the median loss for a single digital fraud incident has climbed to $2,307, though Business Email Compromise (BEC) events in the enterprise space often average $120,000+ per successful heist.

Volume Spike: Phishing attacks reached a record high in early 2025 driven largely by automated AI tools.

184-Day Dwell Time: Fraudsters now remain undetected in standard AP systems for over 6 months, using that time to map out executive approval patterns before striking.

61% Mid-Market Targeting: Smaller and mid-sized firms are increasingly the primary targets due to a perceived "Security Maturity Gap" compared to large enterprises.

90% of Attacks: Cyber-defense data confirms that approximately 90% of successful breaches begin with AI-enhanced phishing or spoofing.

The Sophistication Gap: Fraudsters are now using Generative AI to create "Synthetic Invoices" and Deepfake audio overrides. 1 in 6 consumers and businesses reported losing money to these increasingly precise, AI-driven scams in the past year.

Speed of Attack: Criminals can now generate thousands of unique, personalized fraudulent emails in seconds—a task that previously took weeks of manual "social engineering."

IOFM

Business Landscape

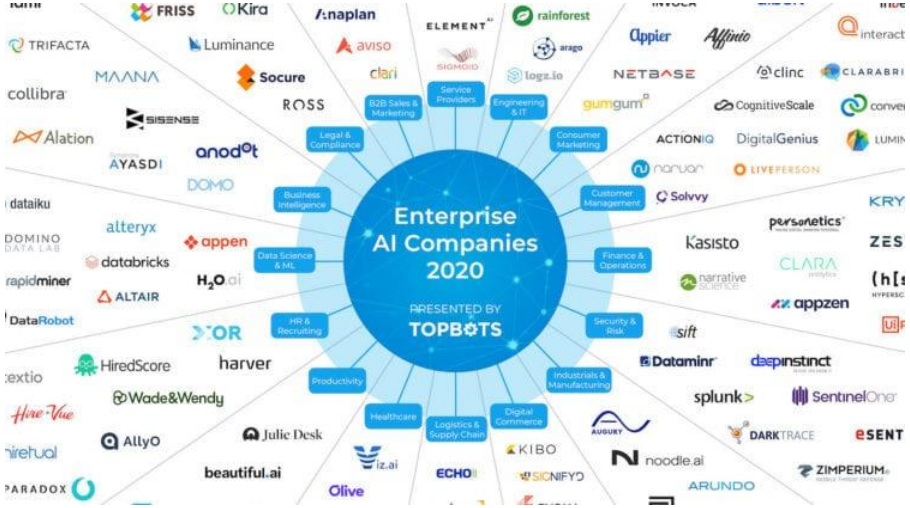

The Adoption Lag: While fraud is rampant, a significant gap remains. As of early 2026, research shows that only 7% of AP processes currently utilize true AI.

The 2026 Turning Point: There is a massive shift underway; 40% of businesses have stated they plan to implement AI-driven AP automation and fraud detection before the end of 2026.

Market Growth: The AP automation market is projected to reach nearly $8 billion by the end of 2026, reflecting a 14% annual increase as companies scramble to modernize their defenses.

Quantifiable Gains: 1 in 4 S&P 500 companies—and 40% of financial services firms—now report quantifiable financial returns from their AI deployments as of Q1 2026.

Efficiency Boost: Organizations using AI for AP have seen a 50% reduction in invoice processing time and a 70-90% decrease in manual audit hours.

The "Profit Center" Shift: By catching duplicate payments and fraud before they occur, automated AP systems can reduce overall processing costs by up to 80%, directly impacting the bottom line.

IOFM

High Risk Areas

- Vendor onboarding
- Vendor master updates
- Invoice submission channels
- Payment approval workflows
- Manual exception handling
- Urgent/off-cycle payment requests
- Email-based approvals
- Supplier self-service portals

IOFM

Rules Based vs. Machine Learning

Feature	Rules-Based Systems (Traditional)	AI & Machine Learning (Modern)
Logic	Static, pre-defined by humans.	Dynamic, learns from data patterns.
Detection	Catches known errors (e.g., exact duplicates).	Catches "fuzzy" anomalies (e.g., slight variations).
Adaptability	Requires manual updates for new threats.	Self-updates as it sees new fraud tactics.
False Positives	High; flags many legitimate exceptions.	Low; understands context and vendor history.

Key Takeaway: *While a rule-based system flags a $5,000 invoice because it's "large," ML flags it because the vendor bank account was changed 2 hours before the invoice was submitted from an unfamiliar IP address.*

IOFM

Rule-Based vs. LLM-Based AI Agents: A Side-by-Side Comparison

Feature	Rule-Based AI Agents	LLM-Based AI Agents
Operation	Executes predefined rules and logic structures.	Generates responses based on learned patterns from training data.
Decision Process	Deterministic—same input always produces the same output.	Probabilistic—responses depend on context and training data.
Flexibility	Limited to predefined cases, cannot handle unknown inputs.	Can adapt dynamically to various types of input.
Complexity Handling	Struggles with ambiguity and unstructured data.	Excels in processing complex and nuanced information.
Scalability	Becomes difficult to scale as the number of rules grows.	Easily scales to handle large datasets and diverse queries.
Transparency	Highly transparent and easy to debug.	Opaque decision-making process, often seen as a black box.
Learning Ability	No learning—static rules must be manually updated.	Can be trained on additional data to improve performance.
Computational Requirements	Low, does not require intensive processing power.	High, requires advanced hardware and infrastructure.
Use Case Examples	Form validation, compliance checking, rule-based chatbots.	Conversational AI, content generation, AI-powered virtual assistants.

Training Models

Data Ingestion: The AI looks at 3–5 years of historical payment data, vendor master files, and employee records.

Baseline Creation: It establishes "Normal Behavior" for every vendor (e.g., Vendor A usually bills on the 15th, amounts vary by 5%, and they use a specific Chase bank account).

Feature Engineering: The system identifies variables—such as invoice metadata, submission time, and font consistency—to create a "risk score."

Feedback Loops: When an AP clerk marks an alert as "Safe" or "Fraud," the AI updates its logic. This is Reinforcement Learning.

The Problem: AI cannot read "messy" data.

The Guidance: Use AI-based OCR **(like Rossum or Kofax)** to convert PDFs and paper into structured data. Ensure your Vendor Master File (VMF) is the Golden Record

Supervised vs. Unsupervised Learning:

- **Supervised:** You feed the AI 500 examples of *known* past fraud and 5,000 examples of *clean* invoices. It learns to "see" the difference.
- **Unsupervised:** The AI looks at your data with no labels and says, "I don't know what fraud is yet, but these three payments to a vendor in Estonia are 4 standard deviations away from anything you've done in 5 years."

The "Human-in-the-Loop" (HITL) Feedback:

- When an auditor confirms a flag is a "False Positive," the AI adjusts its **Confidence Score**. Over 6–12 months, the system becomes highly tailored to your specific industry risks.

IOFM

Agentic Alerting

Tiered Alerting:

- Low Risk: Auto-flagged for a "soft review" by an AP Clerk.High Risk: The AI autonomously freezes the payment in the ERP (SAP, Oracle, or NetSuite) and triggers a Multi-Factor Authentication (MFA) request to the Vendor's registered controller.

Contextual Explanation:

- Modern tools (like AppZen or Overland) no longer just say "Flagged." They provide a Natural Language Summary: "Flagged because: Vendor 'X' changed their IBAN 48 hours ago; this bank is located in a high-risk jurisdiction; and the invoice total is 40% higher than the 3-year average."

Violations versus Intent

- The "Fuzzy" Match Logic: Traditional systems fail if a vendor name is "Global Tech Inc" but an invoice comes in as "Global-Tech Incorporated." AI uses Natural Language Processing (NLP) to recognize these are the same entity—or, more importantly, to flag if a new bank account is suddenly associated with a slightly misspelled name.
- Temporal Analysis: AI analyzes the timing of AP actions. It flags "Batching Anomalies"—if 50 invoices are approved in 10 minutes at 11:00 PM on a Sunday by a user who usually works 9-5, the system identifies a compromised credential or an internal "smash and grab" attempt.

Basic Automation vs. Agentic AI

Basic Automation (often called Robotic Process Automation or RPA) is strictly rule-based. It follows a "If This, Then That" logic. If the input deviates by even a fraction from the programmed path, the process breaks.

- **How it works: You tell the system exactly what to do. If the invoice amount matches the purchase order within 1%, then approve it. If it does not match, then stop and flag it for a human.**
- **Limitations: It cannot handle ambiguity. If a vendor sends an invoice in a slightly different format, or if the price is different because of an unstated shipping change, the automation halts.**
- **The "Human-in-the-loop" reality: A human must intervene to "fix" the error, interpret the discrepancy, and decide how to proceed.**

Agentic AI uses Large Language Models (LLMs) to reason. It doesn't just follow a script; it understands context, sets its own sub-goals, and uses tools to solve problems. It can handle "fuzzy" logic and unexpected variables without human intervention.

- **How it works: You give the agent a goal. Goal: "Ensure all vendor invoices match our POs and resolve discrepancies."**
- **Contextual Understanding: It reads the invoice, realizes the shipping charge is missing, finds the updated shipping policy email from the vendor, checks if it's within the approved budget, and decides to approve the invoice anyway.**
- **Autonomous Resolution: If there's an error, it doesn't just stop. It drafts a professional, contextual email to the vendor asking for clarification or a corrected invoice.**
- **Continuous Learning: It remembers that this specific vendor often forgets to list their tax ID and updates its internal profile for that vendor to watch for it next time.**

IOFM

Basic vs. Agentic AI

Invoice Ingestion and Data Extraction

- Basic Automation: Can pick up a PDF from an email and move it to a folder. It uses OCR (Optical Character Recognition) to find a date, but if the date is in a different spot than usual, it fails.
- Agentic AI: Reads the invoice like a human. It understands that "Net 30" implies a due date even if one isn't explicitly listed. It can extract data from a blurry photo of a handwritten receipt just as easily as a digital PDF.

Fraud Detection

- Agentic AI excels here by spotting patterns that don't fit "standard" behavior. While basic automation only checks if a vendor is on an "approved list," an AI agent might notice a vendor suddenly changed their bank account details to a high-risk jurisdiction and proactively pause the payment for a security review.Summary of Impact

The Three-Way Match

- Basic Automation: Compares the Invoice, Purchase Order (PO), and Receiving Report. If the numbers don't match exactly (e.g., a $0.01 rounding difference), it flags it for a human.
- Agentic AI: Recognizes the $0.01 is a tax rounding error. It checks the historical relationship with that vendor, sees this has happened before, and decides to approve the payment anyway within its "discretionary limit," or writes a polite email to the vendor asking for clarification.

Communication and Dispute Resolution

- Basic Automation: Can't "talk." It can only send a generic "Status: Pending" notification.
- Agentic AI: Acts as a first-line support agent. It can read a vendor’s angry email about a late payment, check the ERP system to find the bottleneck, realize a manager hasn't signed off yet, and ping that manager with a summary of why this needs urgent attention.

IOFM

Case Examples

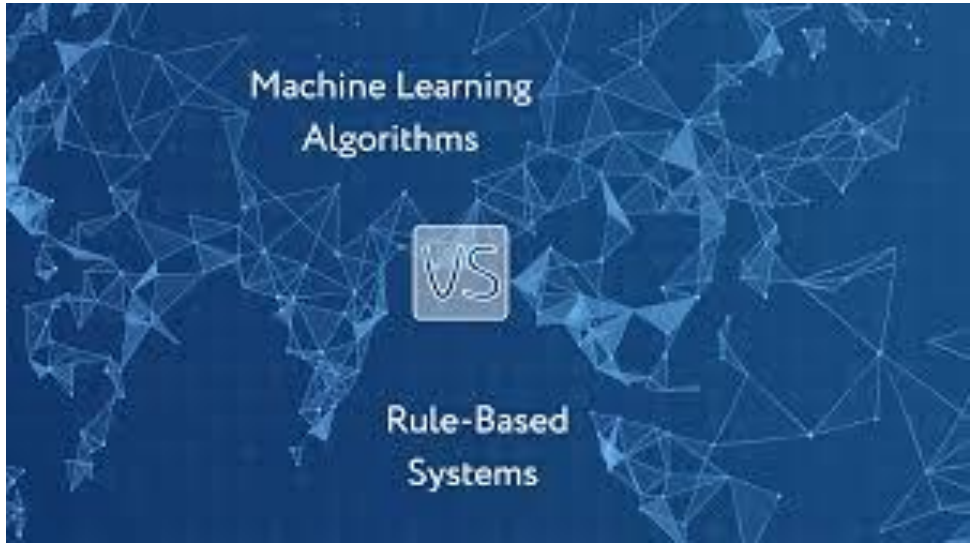

The "Slow Leak" (Internal Fraud)

- The Scenario: An employee at a mid-sized manufacturing firm created a shell company. Over 18 months, they submitted 40 invoices for "consulting services," all under the $2,500 threshold that required executive approval.
- The Detection: A rules-based system missed this because every invoice was "within policy." The ML Model flagged it because the "Vendor" had the same ZIP code as the employee and the invoice numbers were perfectly sequential (001, 002, 003), suggesting no other customers.
- Result: Saved $92,000 in future losses.

The Sophisticated Business Email Compromise (BEC)

- The Scenario: A high-volume vendor’s email was hacked. The hacker sent a "Notice of Bank Change" that looked identical to previous correspondence.
- The Detection: The AI flagged the request because the digital footprint (email headers) originated from a different geographic region than usual and the new bank account had never been seen in the "industry cluster" the vendor belongs to.
- Result: Stopped a $1.2M wire transfer.

IOFM

Case Examples

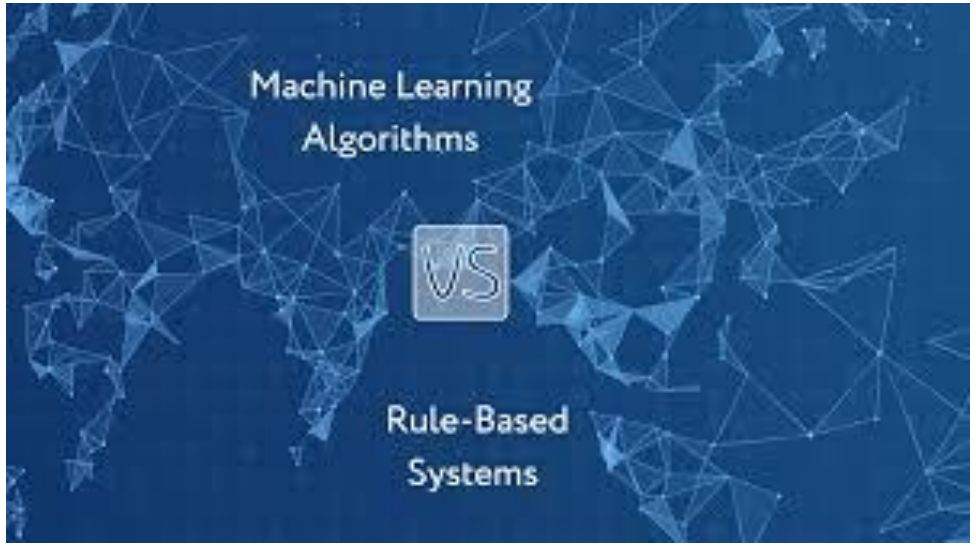

The "Ghost Vendor" & Sequential Invoicing

- The Trap: A trusted long-term employee created a vendor record that mirrored a real supplier’s name but used their own bank details. They submitted small, monthly "maintenance fees.
- The AI Catch: The AI noticed the Invoice Numbering Pattern. While real vendors have gaps in invoice numbers (e.g., #102, #145, #210), this vendor's invoices were #001, #002, #003. The AI flagged the "lack of other customers" as a sign of a shell company.
- ROI Impact: Stopped a 3-year leakage totaling $240,000.

The "Deepfake" Executive Override

- The Trap: An AP staffer received a "Deepfake" audio message from the CFO requesting an urgent, out-of-cycle payment to a "new M&A partner.
- The AI Catch: Because the AI was integrated with the Treasury Management System, it flagged the payment as "Out of Policy" because the destination account had a low Trust Score and the "Urgency" factor didn't match the historical M&A payment workflow.
- ROI Impact: Prevented a one-time loss of $4.5M.

Tools

Spend Auditor

- Tool: AppZen
- The AI Edge: It uses Computer Vision and NLP to "read" an invoice like a human would, but at scale. It doesn't just look at the numbers; it checks the vendor's website, regulatory filings, and even news reports to see if the company is legitimate.
- ROI Factor: Drastic reduction in "Audit Cycle Time." Instead of humans auditing 10% of invoices, the AI audits 100% in seconds, allowing the team to only touch the "high-risk" 1%.

Fraud Prevention

- Tool: Tipalti
- The AI Edge: It specializes in Payee Intelligence. Before a payment is even processed, its AI checks the recipient against 26,000+ global watchlists (OFAC, SDN, etc.) and uses "Frictionless Onboarding" to validate bank details in real-time.
- ROI Factor: Avoidance of massive regulatory fines and "payment error" losses that occur when dealing with complex international currency and tax regulations.

IOFM

Tools

Behavioral Analyst

- Tool: Oversight Systems
- The AI Edge: This tool looks for Collusion Patterns. It analyzes the relationship between the person who created the vendor, the person who approved the invoice, and the person who authorized the payment. It flags "Split Purchases" (breaking a $10k invoice into four $2,500 chunks to avoid approval workflows).
- ROI Factor: Detecting "Theft of Opportunity" by employees and high-risk vendors that traditional systems ignore.

Duplicate Detective

- Tool: FISPAN or Directly Integrated ERP AI (e.g., SAP Cash Application)
- The AI Edge: It catches "Fuzzy Duplicates." A rule-based system catches Invoice #123 and Invoice #123. AI catches Invoice #123-A, Invoice #I23 (using an 'I' instead of '1'), and the same invoice submitted twice with slightly different dates.
- ROI Factor: Direct "Hard Dollar" recovery. For a company with $500M in spend, even a 0.1% duplicate rate is $500,000 in found cash.

IOFM

Challenges

High volume: "we have too many invoices to look at."

Global complexity: "we pay people in 20 countries and it's a mess."

Accuracy/duplicate issues: "we keep paying the same thing twice."

Compliance/audit: "our internal auditors are always finding errors after the fact."

Data management: “our data is everywhere and we cannot get IT to assist us.

Knowledge gap: “we do not have the required SMEs to implement AI solutions”

ROI: “we have not been able to demonstrate the ROI for these investments in AI”

Data overload: “we are unable to effectively manage the overwhelming number of alerts”

IOFM

Applications

Fraud Prevention and Anomaly Detection

- Traditional systems look for exact duplicates. AI looks for behavioral anomalies.
- Example: An AI system flags a sudden change in a long-term vendor's bank account details. It cross-references this with global databases to determine if the change matches known "business email compromise" patterns before the payment is processed.

Dynamic Approval Routing

- Instead of rigid "if-then" rules, AI learns the organizational structure and behavior.
- Example: If an invoice is coded to "Marketing" but exceeds a specific budget threshold for that quarter, the AI automatically reroutes the approval to the CFO instead of the Department Manager, proactively enforcing compliance.

Global Compliance (E-Invoicing)

- With 2026 seeing a surge in global e-invoicing mandates (like ViDA in Europe), AI is used as a "compliance gatekeeper."
- Example: A company operating in France or Poland uses AI to automatically format and validate outgoing and incoming invoices to ensure they meet specific government XML standards (e.g., Peppol), avoiding heavy non-compliance fines.

Supplier Sentiment and Inquiry Management

- Example: Using Generative AI, companies like Adyen or Billerud have implemented AI "agents" that monitor AP email inboxes. These agents answer 80% of supplier questions regarding "When will I be paid?" by looking up the real-time status in the ERP and responding in natural language.

IOFM

Red Team/Blue Team Exercise

- Your group is a sophisticated criminal syndicate. You have three pieces of information: the name of a real vendor your company uses, the name of your Controller, and the fact that your company is currently undergoing a busy software migration. How do you steal $50,000 without triggering a basic 'rules-based' alert?
- Scenarios to Consider:
  - The Velocity Attack: Submitting 20 small invoices just under the approval threshold.
  - The Identity Pivot: Sending a bank change request from a domain that is one letter off (e.g., @vendor-inc.com vs. @vend0r-inc.com).
  - The Social Engineer: Using a Deepfake audio clip or "Urgent" email from an executive to bypass standard PO matching.
  - Other examples

IOFM

Detection Strategy Matrix

Fraud Method	The "Red Team" Strategy	Why Traditional Rules Fail	AI "Blue Team" Defense
The "Ghost" Incrementalist	Submitting monthly "service fees" just below the $500 threshold that triggers manual review.	Rules are set to ignore anything under a specific dollar amount.	**Anomaly Detection:** AI identifies that a "Marketing" vendor has no physical address or LinkedIn presence and is the only vendor with 100% "round number" billing.
The Domain "Cousin" Attack	Registering acme-corp-global.com instead of the real acme-corp.com to send bank change requests.	Email filters see a "valid" domain; humans miss the subtle character additions.	**Metadata & NLP Analysis:** The AI flags that the sender's IP address is in a new country and the "Age of Domain" is only 48 hours old.
The "Synthetic" Invoice	Using Generative AI to create a pixel-perfect invoice with a real logo but a fraudulent IBAN.	OCR reads the text correctly; rules see a valid-looking invoice that matches a PO number.	**Computer Vision:** The AI detects "Digital Artifacts" (subtle layout inconsistencies) or notices that the invoice font/template doesn't match the last 5 years of history.
The "Split-Purchase" Shuffle	Breaking a $20,000 capital expense into eight $2,500 payments to avoid "Executive Approval" workflows.	Each individual transaction is "compliant" with the $3,000 limit rule.	**Relational Clustering:** The AI links these transactions by date, vendor ID, and GL code, flagging the "Cluster" as a single attempt to bypass internal controls.
The "Dormant Vendor" Hijack	Hacking the email of a vendor you haven't paid in two years and submitting an "overdue" notice.	The vendor is already in your VMF (Vendor Master File), so it bypasses "New Vendor" red flags.	**Temporal Behavioral Analysis:** The AI flags a sudden "High-Velocity" interaction with a "Cold" vendor, requiring a secondary biometric or MFA verification.
The "Price Gouging" Leak	Gradually increasing the unit price of a common item (e.g., from $10 to $12) over six months.	Rules check if the price matches the current PO, not the historical market trend.	**Price Benchmarking:** AI compares the price against historical data and "Peer Community Data" to flag that you are paying 20% above the market average.

IOFM

Fraud Maturity Curve

Level 1 (Manual): We check things by hand and hope for the best.

Level 2 (Rules-Based): Our ERP flags duplicates and over-limit invoices.

Level 3 (AI-Assisted): We use ML to score risk, and our team only reviews the "Red" flags.

Level 4 (Agentic): The AI autonomously blocks suspicious payments and requests MFA from vendors.

IOFM

Shift from Pilot Projects to Foundational Standards

Moving beyond simple data extraction toward Agentic AI—systems that don't just flag problems but can autonomously communicate with suppliers and resolve discrepancies

Capability	AI Technology Used	Operational Impact
Intelligent Capture	NLP & Computer Vision	Extracts data from "unstructured" formats like handwritten receipts, blurry PDFs, or complex email bodies with ~99% accuracy.
Autonomous Coding	Machine Learning (ML)	Predicts General Ledger (GL) codes and cost centers based on historical patterns, eliminating manual data entry.
Three-Way Matching	Pattern Recognition	Automatically reconciles invoices against Purchase Orders (POs) and Goods Received Notes (GRNs).
Agentic Exception Handling	Generative AI	If a price mismatch occurs, the AI can draft and send an email to the vendor to request a credit memo or clarification.
Predictive Treasury	Predictive Analytics	Forecasts future cash requirements by analyzing historical payment cycles and seasonal trends.

IOFM

Challenges

- Data Quality: 57% of organizations admit their data is not yet "AI-ready.
- The Skills Gap: Finding talent that understands both finance and AI remains a top priority for 52% of CFOs.
- Legacy Systems: 41% of early-stage adopters cite outdated technology stacks as the primary barrier to implementing modern AI tools
- Poorly designed goals and objectives: AI projects are not well defined or the solutions do not produced the intended ROI.
- Scope creep: initial projects are enhanced or changed midway through, which creates additional workload and confusion for implementation teams.
- New shiny object: AI is developing so rapidly that managers want the latest and greatest tools.
- Lack of security: AI solutions are not considering the CIA triad and involving the appropriate SME's to manage security controls.

Shadow AI

- Use of artificial intelligence tools, models, or browser extensions by employees without the formal approval, oversight, or even knowledge of the IT and Finance leadership.
- Occurs when a team member uses a personal LLM (like a public ChatGPT account) to analyze company spreadsheets or draft vendor emails.
  - An analyst pasting a messy vendor spreadsheet into ChatGPT to "clean up the formatting.
- Represents the "invisible" side of the transition to agentic systems. While you are busy building a formal AI framework, your team may already be using AI in ways you can't see.

Shadow AI Risks

- Data Leakage & Model Training: Most free AI tools use your inputs to train their next model. If an employee pastes an invoice with a vendor’s private bank details or a proprietary pricing contract into a public LLM, that data is now "out in the wild" and could theoretically resurface in a competitor’s query.
- The "Black Box" Audit Failure: If an unapproved AI makes a recommendation (e.g., "This invoice looks okay to pay") and a human follows it, there is no audit trail. If a payment is later found to be fraudulent, you cannot prove why the decision was made.
- Regulatory Non-Compliance: Regulations like GDPR, SOC 2, and the 2026 AI standards require strict data handling. Shadow AI bypasses these entirely, exposing the company to massive fines and "unplanned liabilities" during M&A or annual audits.
- Inaccurate Logic (Hallucinations): Public AI tools are prone to "hallucinating" facts. An unvetted tool might "invent" a tax ID or miscalculate a currency conversion that looks correct but is mathematically wrong, leading to overpayments.

AI Lifecyle

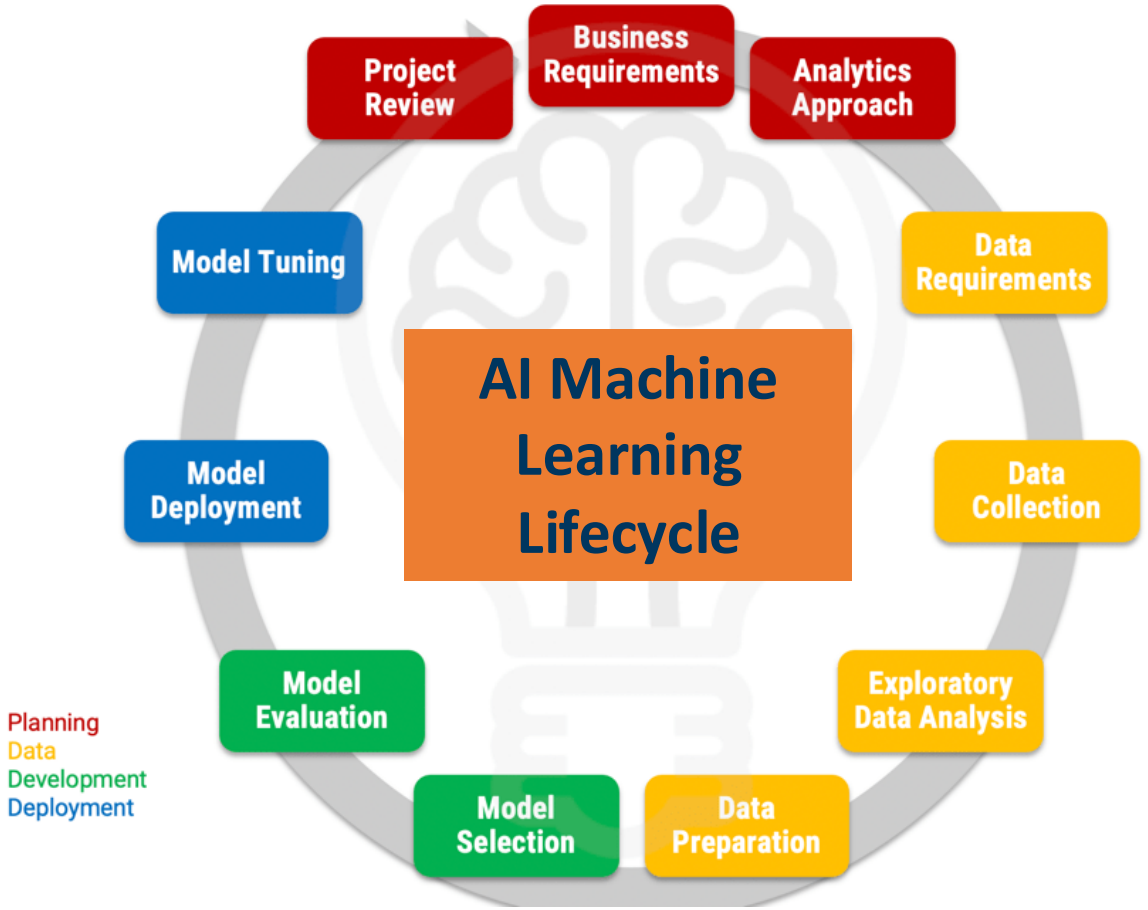
Business Requirements
Analytics Approach
Data Requirements
Data Collection
Exploratory Data Analysis
Data Preparation
Model Selection
Model Evaluation
Model Deployment
Model Tuning
Project Review
AI Machine Learning Lifecycle
Planning
Data
Development
Deployment

AI Governance & Risk Management

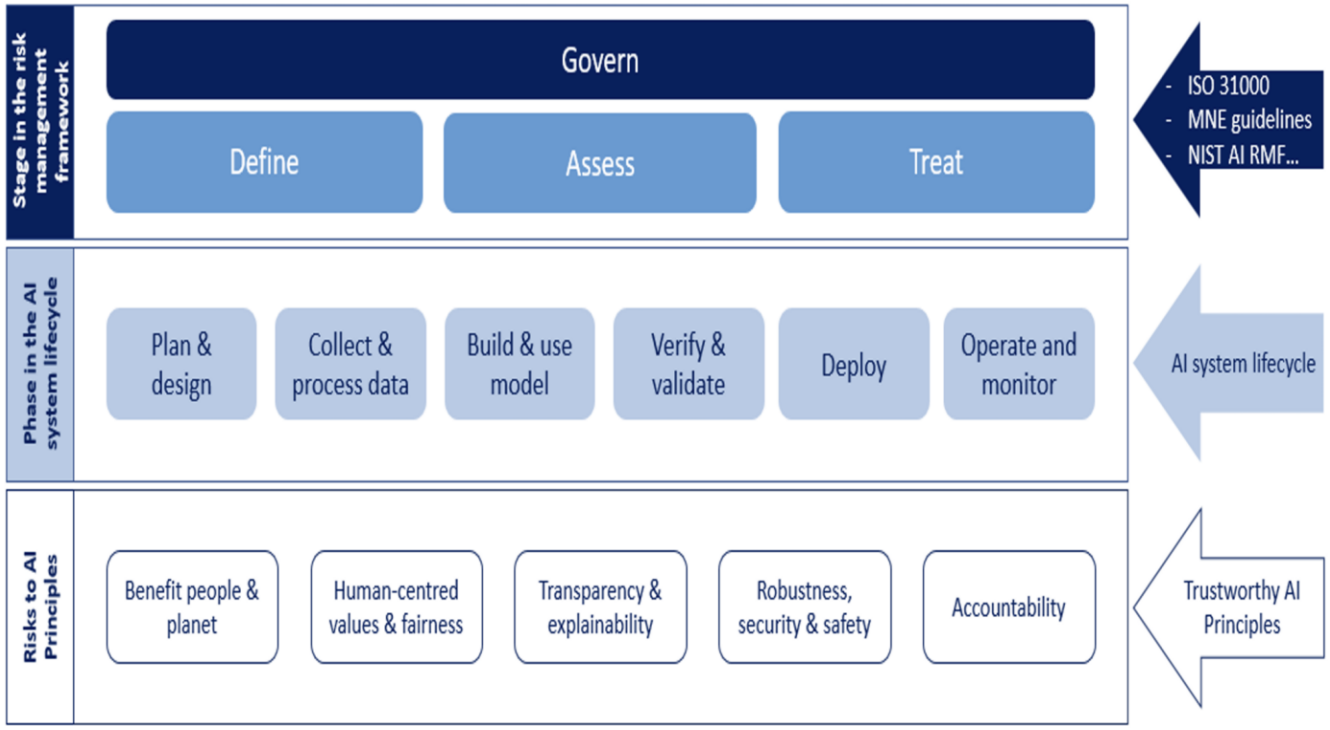

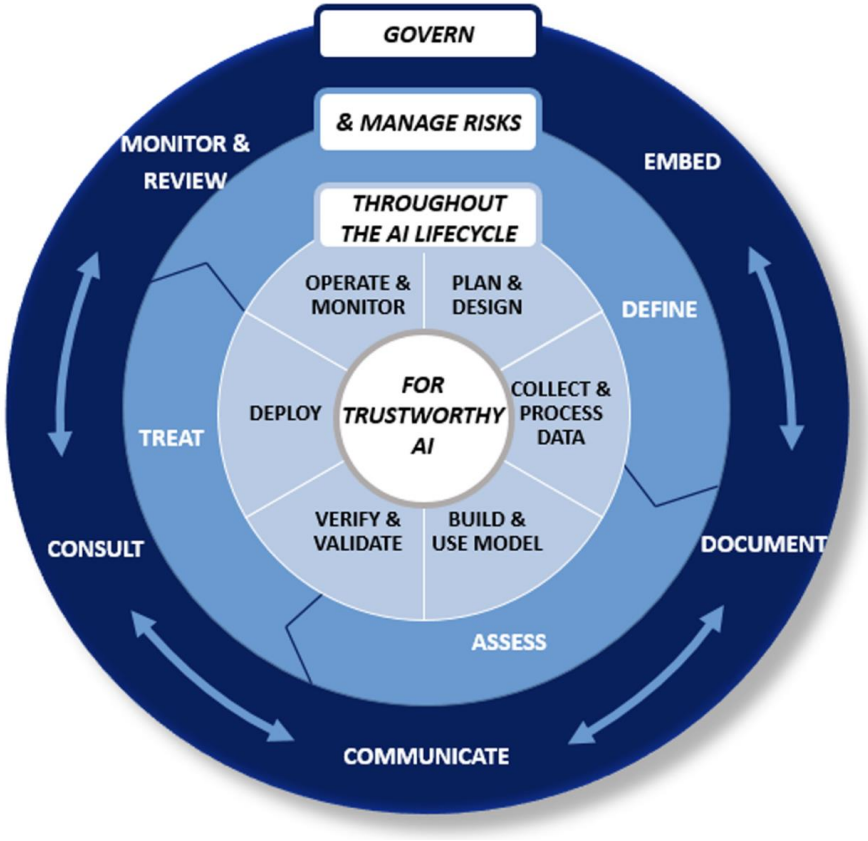

Conclusion

- Document and "Clean" Before You Automate
- Adopt a Phased "30-60-90 Day" Roadmap
- Re-Skill, Don't Just Replace
- Optimize the "Procure-to-Pay" (P2P) Loop
- Prioritize "Integrations" Over "Features"
- Establish a "Vendor Partnership" Program
- Monitor "Decision Drift"

IOFM

Please tell us what you think!

- Please scan this QR code using your mobile to access a short feedback survey →
- Also accessible via the mobile app