

# AI Value vs. Hype: An Objective Analysis of How AI Is (and Isn't) Creating Institutional Impact

## Executive Summary

Artificial intelligence (AI) has moved from a theoretical frontier to a core component of institutional strategy. In 2024, corporate AI investment reached a record \$252.3 billion, and enterprise adoption has surged, with over 78% of organizations now using AI in at least one business function.<sup>1</sup> This frenzy of investment, largely focused on massive data center infrastructure, has led some economists to warn of a speculative "AI bubble" where valuations have detached from the technology's current ability to deliver profit.<sup>33</sup> A significant "value gap" persists. Recent studies reveal that a staggering 95% of generative AI pilots fail to yield tangible business impact, and the average enterprise-wide AI initiative achieves a meager 5.9% ROI against a 10% capital investment.<sup>3</sup>

This report provides an objective, evidence-based analysis of where AI is creating measurable value and where it is failing to meet its hyped potential. It is guided by the critical lens of *AI Snake Oil*, a framework developed by Princeton computer scientists Arvind Narayanan and Sayash Kapoor, which distinguishes between trustworthy AI and products that are ineffective or harmful.

### Key Findings:

1. **Value is Concentrated in Specific Domains:** AI is delivering substantial ROI in well-defined, data-rich environments. High-value use cases are clustered in **process automation** (e.g., logistics, administrative tasks), **pattern recognition** (e.g., fraud detection, medical imaging), and **human augmentation** (e.g., AI copilots for coding and writing).<sup>5</sup> Sectors like finance, healthcare, and supply chain management are realizing the most significant gains.<sup>8</sup>
2. **"AI Snake Oil" is Prevalent and Costly:** AI applications that claim to predict complex human social outcomes—such as job performance from video interviews or criminal recidivism—are often scientifically baseless and perform little better than random chance. High-profile failures at Amazon (hiring), Zillow (home-buying), and IBM Watson (oncology) underscore the immense financial and reputational risks of deploying AI in

domains it is not suited for.<sup>9</sup>

3. **The "Pilot Purgatory" is Real:** The primary barrier to ROI is not the technology itself but a failure to scale. An estimated 46% of AI proofs-of-concept are scrapped before production.<sup>12</sup> Root causes include poor data quality, a lack of integration into core workflows, and a failure to connect AI initiatives to clear business metrics.<sup>3</sup>
4. **A New Governance Imperative:** Successful AI adoption requires a strategic shift. "AI Leaders"—the top 15% of organizations—are distinguished not by superior algorithms but by superior governance.<sup>14</sup> This includes C-suite oversight, a focus on a few high-impact use cases, and a disciplined process for measuring ROI.<sup>13</sup>

This report concludes with a decision-making framework for executives, designed to instill a culture of critical evaluation and ensure that future AI investments are driven by tangible value rather than technological hype.

---

## 1. The Critical Lens: Identifying "AI Snake Oil"

The term "AI Snake Oil" was coined by Arvind Narayanan and Sayash Kapoor to describe AI products that are sold on the back of overblown claims but are ultimately ineffective or harmful. Just as fraudulent elixirs were once marketed as cure-alls, many modern AI tools are positioned as solutions to complex problems they are fundamentally incapable of solving. For leaders, adopting this critical lens is the first line of defense against wasted investment.

### The Core Arguments of the *AI Snake Oil* Thesis:

- **Not All AI is Equal:** Lumping all AI together is a critical error. The authors make a crucial distinction between different modes of AI, primarily **predictive AI** (forecasting outcomes) and **generative AI** (creating content). These technologies have vastly different capabilities, limitations, and risk profiles.
- **Predictive AI for Social Outcomes is Untrustworthy:** The book's strongest critique is aimed at AI systems that claim to predict complex human futures. Applications in hiring (e.g., analyzing video to assess personality), criminal justice (e.g., predicting recidivism), and insurance (e.g., predicting medical needs) are often no better than "elaborate random-number generators". These systems fail because human behavior is inherently unpredictable and the historical data they are trained on is often riddled with societal biases.
- **Generative AI is a Powerful Tool, Not a Mind:** Tools like ChatGPT and Gemini are genuinely useful for augmenting knowledge workers in tasks like drafting, summarizing, and coding. However, they are prone to "hallucinations" (producing confident but false information) and raise complex issues of copyright and data privacy. Worryingly, this issue is not improving with scale; internal tests of OpenAI's much-anticipated GPT-5

model show it still 'hallucinates' in approximately one out of every ten responses on certain factual tasks.<sup>33</sup> They are powerful pattern-matching engines, not reasoning minds.

- **The Real Danger is Misuse, Not "Rogue AI":** The fear of a sentient, world-ending AI is a distraction from the immediate, tangible harms already occurring: biased algorithms making life-altering decisions, the misuse of AI for surveillance, and the concentration of power within a few unaccountable tech companies.

This framework provides executives with the vocabulary to challenge vendors and internal teams, shifting the conversation from marketing claims to a demand for empirical evidence.

---

## 2. Value Creation: Where AI is Delivering Measurable ROI

Despite the hype and high failure rates, AI is creating significant institutional value when applied to the right problems in the right sectors. Success is concentrated in areas where AI can automate well-defined processes, identify patterns in massive datasets, or augment human expertise.

### Top Proven, High-Value Use Cases of AI:

| Use Case                                     | AI Type(s) Used        | Description & Value Proposition                                                                                                                                                                            | Key Sectors               |
|----------------------------------------------|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------|
| <b>Fraud Detection &amp; Risk Management</b> | Predictive AI, ML/DL   | Analyzes millions of transactions in real-time to identify anomalies and patterns indicative of fraud, reducing losses by 15-25% in the first year and cutting false positives by up to 70%. <sup>16</sup> | <b>Finance, Insurance</b> |
| <b>Medical Image Analysis</b>                | ML/DL, Computer Vision | Detects signs of diseases like cancer or diabetic                                                                                                                                                          | <b>Healthcare</b>         |

|                                                        |                     |                                                                                                                                                                                                             |                                                     |
|--------------------------------------------------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------|
|                                                        |                     | retinopathy in MRIs, CT scans, and X-rays with accuracy rates often exceeding 90%, augmenting the capabilities of human radiologists. <sup>7</sup>                                                          |                                                     |
| <b>Demand Forecasting &amp; Inventory Optimization</b> | Predictive AI, ML   | Analyzes historical sales, market trends, and external factors to reduce forecasting errors by up to 50%, leading to a 65% decrease in lost sales and a 20-50% reduction in inventory levels. <sup>18</sup> | <b>Supply Chain, Retail</b>                         |
| <b>Predictive Maintenance</b>                          | Predictive AI, ML   | Uses sensor data from machinery to predict equipment failures before they occur, reducing downtime by up to 50% and maintenance costs by 25%. <sup>19</sup>                                                 | <b>Manufacturing, Logistics, Energy</b>             |
| <b>AI Copilots / Digital Assistants</b>                | Generative AI, LLMs | Augments knowledge workers by automating tasks like writing code, drafting documents, and summarizing information. Can reduce time spent on administrative                                                  | <b>All (especially Tech, Professional Services)</b> |

|                                           |                      |                                                                                                                                                                                     |                                                  |
|-------------------------------------------|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------|
|                                           |                      | tasks by up to 60%.                                                                                                                                                                 |                                                  |
| <b>Logistics &amp; Route Optimization</b> | Predictive AI, ML    | Analyzes traffic, weather, and delivery schedules to optimize routes, reducing fuel costs by 15-30% and improving on-time delivery rates by 15-25%. <sup>20</sup>                   | <b>Transportation, Logistics</b>                 |
| <b>Automated Document Processing</b>      | NLP, Computer Vision | Extracts and processes information from invoices, contracts, and claims, reducing manual review time and saving hundreds of thousands of work hours annually. <sup>6</sup>          | <b>Finance, Insurance, Legal</b>                 |
| <b>Customer Service Automation</b>        | NLP, LLMs            | Deploys chatbots and virtual assistants to handle routine customer inquiries, resolving up to 70% of queries autonomously and reducing call center loads by over 15%. <sup>21</sup> | <b>All (especially Retail, Telecom, Finance)</b> |
| <b>Drug Discovery &amp; Development</b>   | Generative AI, ML/DL | Analyzes biological data to identify promising drug candidates, reducing drug discovery timelines by over 50% in some cases. <sup>22</sup>                                          | <b>Pharmaceuticals, Healthcare</b>               |

|                                           |                   |                                                                                                                                       |                                             |
|-------------------------------------------|-------------------|---------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------|
| <b>Cybersecurity<br/>Threat Detection</b> | Predictive AI, ML | Monitors network traffic and user behavior to identify and neutralize threats in real-time, enhancing security posture. <sup>14</sup> | <b>All (especially Finance, Government)</b> |
|-------------------------------------------|-------------------|---------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------|

## Sector Spotlight: Measurable ROI

- **Financial Services:** This sector is a leader in AI adoption, driven by clear ROI in fraud detection, algorithmic trading, and compliance. Over 88% of financial institutions report that AI has increased their revenues.<sup>8</sup>
- **Healthcare:** The healthcare AI market is projected to nearly double from \$20.65B in 2023 to \$38.66B in 2025.<sup>8</sup> Value is concentrated in diagnostics (FDA-approved AI medical devices grew from 6 in 2015 to 223 in 2023) and administrative automation, which can save large organizations millions annually.<sup>1</sup>
- **Supply Chain & Logistics:** AI-driven forecasting and optimization are delivering significant cost savings. Companies report up to 50% reductions in operational costs, 18% fewer forecasting errors, and 15% more on-time deliveries.<sup>5</sup>

## 3. Failure Cases and Overhyped Applications

The gap between AI's promise and its reality is most evident in its high-profile failures. These cases are not merely technical errors; they are strategic miscalculations, often stemming from a misunderstanding of AI's fundamental limitations.

### High-Profile AI Failures: A Reality Check

|            |         |                             |                       |
|------------|---------|-----------------------------|-----------------------|
| Case Study | AI Type | The Promise vs. The Reality | Root Causes & Lessons |
|------------|---------|-----------------------------|-----------------------|

|                                    |                   |                                                                                                                                                                                                                                                                                                                                                         | Learned                                                                                                                                                                                                                                                     |                                                                                                                            |
|------------------------------------|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| <b>Amazon's AI Recruiting Tool</b> | Predictive AI, ML | <p>Promise: An automated tool to screen resumes and identify top talent, removing human bias.</p> <p>Reality: The tool learned from 10 years of Amazon's hiring data and systematically penalized resumes containing the word "women's" and those from all-women's colleges. It was scrapped after engineers could not remove the bias.<sup>9</sup></p> | <p><b>Root Cause: Data Bias.</b> The AI simply automated and amplified the existing biases present in the historical training data.</p> <p><b>Lesson:</b> AI is not a solution for deeply ingrained societal biases; it is a mirror that reflects them.</p> |                                                                                                                            |
| <b>Zillow Offers (iBuying)</b>     | Predictive AI, ML | <p>Promise: Use the "Zestimate" algorithm to accurately predict home values, enabling Zillow to buy properties and resell them for a profit.</p>                                                                                                                                                                                                        | <p><b>Root Cause: Model Brittleness.</b> The model was trained on past correlations and broke when market conditions shifted. It was a failure to account for</p>                                                                                           | <p><b>Lesson:</b> Correlation-based models are unreliable for high-stakes financial decisions in dynamic environments.</p> |

|                                |                    |                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                            |                                                                                                                                     |
|--------------------------------|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
|                                |                    | Reality: The algorithm failed to adapt to market volatility, leading Zillow to overpay for thousands of homes. The business unit was shut down after a write-down of over \$500 million. <sup>10</sup>                                                                                                                                   | causality and real-world context. <sup>26</sup>                                                                                                                                                            |                                                                                                                                     |
| <b>IBM Watson for Oncology</b> | NLP, Expert System | <p>Promise: A revolutionary AI that would read medical literature and patient data to recommend optimal, personalized cancer treatments.</p> <p>Reality: The system often gave erroneous and unsafe recommendations. It struggled with the nuance of medical language and was trained on a small set of hypothetical cases, not real</p> | <b>Root Cause: Over-Scoping &amp; Lack of Context.</b> The problem ("solve cancer") was too broad and complex. The AI could not grasp the nuanced, contextual judgment required in oncology. <sup>11</sup> | <b>Lesson:</b> AI struggles with tasks that require deep, contextual human expertise and cannot be reduced to simple data patterns. |



|                                           |                          |                                                                                                                                                                                                                                                                                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                       |  |
|-------------------------------------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--|
|                                           |                          | <p>patient data. The project was a commercial failure after billions in investment.<sup>11</sup></p>                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                                                                       |  |
| <p><b>COMPAS Recidivism Algorithm</b></p> | <p>Predictive AI, ML</p> | <p>Promise: An algorithm to predict the likelihood of a defendant re-offending, used to inform bail and sentencing decisions. Reality: Independent audits found the tool was only slightly more accurate than a coin flip and was racially biased, incorrectly flagging Black defendants as high-risk at twice the rate of white defendants.<sup>17</sup></p> | <p><b>Root Cause: Flawed Data &amp; Unsolvable Problem.</b> The AI was trained on arrest data, which reflects historical policing biases, not inherent criminality. Predicting human behavior is fundamentally unreliable.</p> <p><b>Lesson:</b> Applying AI to complex social problems without addressing the underlying data flaws and systemic issues will automate injustice.</p> |  |

## The Macro Risk: A Speculative Bubble and Hidden Costs

Beyond individual project failures, there is a growing concern that the entire AI sector is experiencing a speculative investment bubble. The current AI boom is fueled by hundreds of billions of dollars in spending on data centers and other infrastructure, based on the belief that this scaling will unlock unprecedented productivity gains and corporate profits.<sup>33</sup>

However, evidence suggests that we may have reached "peak GenAI" in terms of what current Large Language Models (LLMs) can achieve. Simply scaling up—building more data centers and using more chips—appears to be yielding rapidly diminishing returns and is not leading to the promised goal of Artificial General Intelligence (AGI).<sup>33</sup> This disconnect between massive capital outlay and stagnating technological progress is a classic indicator of a bubble poised to burst.<sup>33</sup>

Furthermore, this infrastructure race comes with significant and often overlooked environmental costs. AI-focused data centers consume enormous amounts of electricity and water.<sup>34</sup> The U.S. Department of Energy estimates that data centers, which consumed 4% of the nation's electricity in 2022, could see that share rise to 12% by 2028.<sup>34</sup> In some water-scarce regions, the strain on local water supplies is already causing environmental concerns. One study estimates that generating a single 100-word email on ChatGPT uses around 1.5 liters of water.<sup>34</sup> These hidden environmental externalities are a real cost of the AI boom that rarely appears on a project's balance sheet.

## Executive Red Flags: How to Spot "AI Snake Oil"

Leaders should be highly skeptical when a vendor or internal team proposes an AI solution that exhibits these warning signs:

- **Claims to Predict Complex Human Social Outcomes:** Any tool claiming to predict job success, criminality, personality, or other nuanced human behaviors is a major red flag.
- **Lacks Independent, Peer-Reviewed Validation:** Marketing materials and vendor-supplied case studies are not sufficient evidence. Demand rigorous, independent proof that the tool works as advertised.
- **Operates as a "Black Box":** If the vendor cannot explain *how* the algorithm makes its decisions in a clear and understandable way, it is too risky for high-stakes applications.<sup>17</sup>
- **Promises to Solve a Vague, Overly Broad Problem:** AI excels at narrow, well-defined tasks. Be wary of solutions that promise to "revolutionize sales" or "solve healthcare".<sup>11</sup>
- **Relies on Biased or Incomplete Historical Data:** The adage "garbage in, garbage out" is paramount. An AI system's output is only as good as the data it is trained on.

---

## 4. AI's Impact on Institutional Strategy & Governance

The rise of AI is forcing a fundamental rewiring of organizations. Simply purchasing AI tools is not enough; institutions must adapt their strategy, infrastructure, and governance to capture value and mitigate risk.

### The State of AI Adoption and the ROI Challenge

Recent industry reports paint a picture of rapid adoption colliding with significant execution challenges:

- **Adoption is Accelerating:** In 2024, 72% of organizations reported using AI, up from around 50% in previous years. Generative AI use nearly doubled in less than a year, with 65% of organizations now using it regularly.<sup>28</sup>
- **Investment is Soaring:** Global private AI investment hit a record \$109.1 billion in the US alone in 2024.<sup>1</sup> 92% of companies plan to increase their AI investments over the next three years.<sup>29</sup>
- **But Value Remains Elusive:** Despite the spending, most organizations are struggling to see a return. A 2025 MIT report found 95% of GenAI pilots yield no tangible P&L impact.<sup>3</sup> A BCG study found that only 26% of companies have the capabilities to move beyond pilots and generate value.<sup>13</sup> This disconnect fuels arguments that the AI industry is in a speculative bubble, where massive capital expenditure on infrastructure is predicated on future productivity gains that have yet to materialize.<sup>33</sup> Some analyses suggest that current Large Language Models have reached "peak GenAI," where simply scaling up with more data centers and chips yields rapidly diminishing returns, making the path to profitability even more uncertain.<sup>33</sup>
- **Public Sector Challenges:** Governments face similar hurdles, with most AI efforts stuck in pilot phases. Common challenges include skills gaps, poor data quality, risk aversion, and difficulty demonstrating ROI.<sup>30</sup>

### The Governance Imperative: What "AI Leaders" Do Differently

The difference between success and failure often comes down to governance and strategy.

Research from firms like McKinsey and BCG identifies several characteristics of high-performing organizations <sup>13</sup>:

1. **C-Suite Ownership:** CEO oversight of AI governance is one of the factors most correlated with achieving bottom-line impact from AI. <sup>15</sup> 72% of "AI Leaders" report full alignment between the C-suite and IT leadership. <sup>14</sup>
2. **Strategic Focus:** Leaders pursue half as many AI opportunities as their peers. They focus investment on a few high-priority use cases tied directly to core business processes, where they generate 62% of the value. <sup>13</sup>
3. **Emphasis on People and Process:** Leaders follow a "10-20-70" rule: 10% of resources go to algorithms, 20% to technology and data, and 70% to people and process redesign. The value of AI comes from rewiring how work gets done. <sup>13</sup>
4. **A Robust Data Foundation:** AI is only as good as the data it runs on. Leaders prioritize data governance, quality, and accessibility as prerequisites for any AI initiative. <sup>14</sup>
5. **Disciplined ROI Measurement:** Leaders relentlessly measure the return on their AI investments, using pilot programs to validate assumptions before scaling. <sup>12</sup>

## 5. Conclusion: A Framework for Evaluating AI Proposals

To navigate the landscape of AI value and hype, executives must move from being passive consumers of technology to active, critical evaluators, especially in an economic climate characterized by a potential AI investment bubble. <sup>33</sup> The "Normal Technology" perspective argues that AI, like electricity or the internet, will be transformative, but its adoption will be slow, gated by economic and social realities. This calls for a patient, disciplined, and value-focused approach.

The following framework, inspired by the *AI Snake Oil* critique and the *Pillar Framework* for technology investment, provides a set of critical questions for leaders to ask of any AI proposal.

### An Executive Framework for Evaluating AI

| Category                   | Key Question                             | What to Look For (Green Flags)                | Red Flags to Watch For                                    |
|----------------------------|------------------------------------------|-----------------------------------------------|-----------------------------------------------------------|
| 1. <b>Problem-Solution</b> | Is this a problem AI is good at solving? | <b>Narrow, well-defined tasks:</b> Automation | <b>Broad, ambiguous social problems:</b> Predicting human |

|                                        |                                                                                    |                                                                                                                                                                                     |                                                                                                                                                      |
|----------------------------------------|------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Fit</b>                             |                                                                                    | of repetitive processes, pattern recognition in large datasets, content generation for augmentation.                                                                                | behavior, assessing personality, solving "crime" or "poverty."                                                                                       |
| <b>2. Strategic Alignment</b>          | How does this create measurable business value?                                    | <b>Direct link to a core KPI:</b><br>Increasing revenue, reducing specific costs, improving efficiency of revenue-generating teams, or mitigating quantifiable risks. <sup>17</sup> | <b>Vague benefits:</b><br>"Improving synergy," "driving innovation," "being more data-driven," without a clear financial metric attached.            |
| <b>3. Evidence &amp; Validation</b>    | What is the independent proof that this works?                                     | <b>Peer-reviewed studies,</b> successful and instrumented internal pilots, and transparent data on accuracy and error rates.                                                        | <b>Vendor-supplied case studies only,</b> "proprietary" algorithms that cannot be explained, claims of near-perfect accuracy (e.g., "95% accurate"). |
| <b>4. Data &amp; Bias</b>              | Is the training data robust, relevant, and unbiased?                               | <b>Transparent data provenance,</b> clear documentation of data sources, and explicit steps taken to identify and mitigate biases.                                                  | <b>"Black box" data sources,</b> reliance on historical data known to contain societal biases (e.g., historical hiring or arrest records).           |
| <b>5. Implementation &amp; Scaling</b> | How will this integrate into our existing workflows, and who will own the results? | <b>A clear plan for workflow redesign,</b> a designated business owner                                                                                                              | <b>A purely technical implementation plan</b> that treats the AI as a plug-and-play tool                                                             |

|  |  |                                                                                                  |                                                           |
|--|--|--------------------------------------------------------------------------------------------------|-----------------------------------------------------------|
|  |  | accountable for ROI, and a focus on people and change management (the "70%" rule). <sup>13</sup> | without considering how people and processes must change. |
|--|--|--------------------------------------------------------------------------------------------------|-----------------------------------------------------------|

By systematically applying this critical framework, leaders can cut through the hype, avoid costly "snake oil" investments, and focus their resources on AI initiatives that deliver real, sustainable institutional value.

## Works cited

1. The 2025 AI Index Report | Stanford HAI, accessed September 22, 2025, <https://hai.stanford.edu/ai-index/2025-ai-index-report>
2. Economy | The 2025 AI Index Report | Stanford HAI, accessed September 22, 2025, <https://hai.stanford.edu/ai-index/2025-ai-index-report/economy>
3. Why most AI products fail: Key findings from MIT's 2025 AI Report, accessed September 22, 2025, <https://www.mindtheproduct.com/why-most-ai-products-fail-key-findings-from-mits-2025-ai-report/>
4. How to maximize ROI on AI in 2025 - IBM, accessed September 22, 2025, <https://www.ibm.com/think/insights/ai-roi>
5. AI in Supply Chain Management: Real Use Cases & ROI - CE Interim, accessed September 22, 2025, <https://ceinterim.com/ai-in-supply-chain-management/>
6. AI in Finance Use Cases: 10 Ways Artificial Intelligence Delivers Better ROI - Medium, accessed September 22, 2025, <https://medium.com/@kanerika/ai-in-finance-use-cases-10-ways-artificial-intelligence-delivers-better-roi-9016bd0e2f64>
7. 5 Game-Changing AI Healthcare Use Cases with Real ROI - SR analytics, accessed September 22, 2025, <https://sranalytics.io/blog/ai-healthcare-use-cases/>
8. AI Industry Outlook 2025: Global Trends in the US, Europe, and Asia, accessed September 22, 2025, <https://www.britopian.com/wp-content/uploads/2025/03/AI-Industry-Outlook-2025-Global-Trends.pdf>
9. Amazon's sexist hiring algorithm could still be better than a human - IMD Business School, accessed September 22, 2025, <https://www.imd.org/research-knowledge/digital/articles/amazons-sexist-hiring-algorithm-could-still-be-better-than-a-human/>
10. Trust Incident Zillow - iceberg.digital, accessed September 22, 2025, <https://iceberg.digital/trust-incident-zillow/>
11. 5.4 | IBM Watson in Oncology | LEAP for PIT - Sites at Penn State, accessed September 22, 2025, <https://sites.psu.edu/leapforpit/artificial-intelligence/ibm-watson-in-oncology/>
12. AI project failure rates are on the rise: report | CIO Dive, accessed September 22,

- 2025, <https://www.ciodive.com/news/AI-project-fail-data-SPGlobal/742590/>
13. AI Adoption in 2024: 74% of Companies Struggle to Achieve and Scale Value | BCG, accessed September 22, 2025, <https://www.bcg.com/press/24october2024-ai-adoption-in-2024-74-of-companies-struggle-to-achieve-and-scale-value>
  14. AI in Action 2024 Report - IBM, accessed September 22, 2025, <https://www.ibm.com/think/reports/ai-in-action>
  15. The State of AI: Global survey | McKinsey, accessed September 22, 2025, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
  16. AI Fraud Detection in Banking: Real-Time Use Cases and Applications Explained, accessed September 22, 2025, <https://appwrk.com/insights/banking-use-cases-of-ai-in-fraud-detection/>
  17. Barna Modulo - CHRIS CORCINO.pdf
  18. Machine Learning in Logistics and Supply Chain [7 Use Cases Included] - Acropolium, accessed September 22, 2025, <https://acropolium.com/blog/adopting-machine-learning-in-supply-chain-and-logistics-for-successful-automation/>
  19. AI in Logistics 2025: Real Use Cases & Industry Results - Noloco, accessed September 22, 2025, <https://noloco.io/blog/ai-in-logistics>
  20. The True ROI of AI in Logistics: Unlocking Intelligent, High-Impact Supply Chain Operations, accessed September 22, 2025, <https://finmile.co/resources/the-true-roi-of-ai-in-logistics-unlocking-intelligent-high-impact-supply-chain-operations>
  21. 10 ROI of AI case studies show results - BarnRaisers, LLC, accessed September 22, 2025, <https://barnraisersllc.com/2025/06/20/10-roi-of-ai-case-studies-show-results/>
  22. 2025 AI Business Predictions: PwC, accessed September 22, 2025, <https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-predictions.html>
  23. 23 Healthcare AI Use Cases with Examples - Research AIMultiple, accessed September 22, 2025, <https://research.aimultiple.com/healthcare-ai-use-cases/>
  24. Why Amazon's Automated Hiring Tool Discriminated Against Women | ACLU, accessed September 22, 2025, <https://www.aclu.org/news/womens-rights/why-amazons-automated-hiring-tool-discriminated-against>
  25. How Zillow could have avoided its \$500M AI mistake - SphereOI, accessed September 22, 2025, <https://sphereoi.ai/how-zillow-could-have-avoided-its-500m-ai-mistake/>
  26. Zillow — A Cautionary Tale of Machine Learning - causalens, accessed September 22, 2025, <https://causalai.causalens.com/resources/blog/zillow-a-cautionary-tale-of-machine-learning/>
  27. M. D. Anderson Breaks With IBM Watson, Raising Questions About Artificial Intelligence in Oncology | JNCI: Journal of the National Cancer Institute | Oxford Academic, accessed September 22, 2025,

- <https://academic.oup.com/jnci/article/109/5/djx113/3847623>
28. The state of AI in early 2024: Gen AI adoption spikes and starts to generate value - McKinsey, accessed September 22, 2025,  
<https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-2024>
  29. AI in the workplace: A report for 2025 | McKinsey, accessed September 22, 2025,  
<https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/superagency-in-the-workplace-empowering-people-to-unlock-ais-full-potential-at-work>
  30. Implementation challenges that hinder the strategic use of AI in ..., accessed September 22, 2025,  
[https://www.oecd.org/en/publications/governing-with-artificial-intelligence\\_795de142-en/full-report/implementation-challenges-that-hinder-the-strategic-use-of-ai-in-government\\_05cfe2bb](https://www.oecd.org/en/publications/governing-with-artificial-intelligence_795de142-en/full-report/implementation-challenges-that-hinder-the-strategic-use-of-ai-in-government_05cfe2bb)
  31. AI: public sector must address data challenge - THINK Digital Partners, accessed September 22, 2025,  
<https://www.thinkdigitalpartners.com/news/2025/05/29/ai-public-sector-must-address-data-challenge/>
  32. Generative AI in organizations 2024 - Capgemini, accessed September 22, 2025,  
<https://www.capgemini.com/insights/research-library/generative-ai-in-organizations-2024/>
  33. The AI Bubble and the U.S. Economy: How Long Do “Hallucinations” Last?, accessed October 4, 2025,  
<https://www.ineteconomics.org/perspectives/blog/the-ai-bubble-and-the-u-s-economy-how-long-do-hallucinations-last>
  34. The AI Bubble and the U.S. Economy: How Long Do 'Hallucinations' Last?, accessed October 4, 2025,  
[https://www.ineteconomics.org/uploads/papers/WP\\_240-Storm-AI.pdf](https://www.ineteconomics.org/uploads/papers/WP_240-Storm-AI.pdf)