

***“Intelligence is the computational part of the ability to achieve goals in the world.***

*Varying kinds and degrees of intelligence occur in people, many animals and some machines.”*

**John  
McCarthy**

**1956**

**Stanford  
University**

*“When I say AI — what’s the first word that comes to your mind?”*



*Uncertainty*



*Navigation*



*Possibility*

## Five eras · One Continuous Argument

**1956 – 1980s**

**Artificial  
Intelligence**



**The Perceptron**

Rosenblatt, 1957

First rule-based  
execution systems

**1970s – 1980s**

**Expert Systems  
Knowledge Eng.**



**Inference Engine**

IF / THEN logic

Black Monday 1987  
Program trading

**1990s – 2000s**

**Machine Learning  
Data Mining**



**Support Vector Machine**

Renaissance Medallion

Statistical Arbitrage

**2010s**

**AI Revival  
Deep Learning**



**LSTM / CNN**

Flash Crash 2010  
JPM LOXM, 2017

**2022 – Present**

**Generative AI  
Agentic AI**



**Transformer / Large  
Language Model**

LLM signal  
generation  
Autonomous agents

**Same Underlying Methods · 1956 to Present**

**New executive generation — no memory of prior cycles**

*AI winter · 1970s  
Label abandoned*

*Second winter · 1980s  
Systems brittle*

*2008 crisis  
Regime blindness*

*85% failure rate  
documented*

*Same mathematics.  
New label.*

*If the math has been continuous for 70 years — why does it keep failing?*

---

85%

of enterprise AI initiatives  
fail to scale beyond pilot

#### NINE YEARS OF EVIDENCE • 2016–2025

##### RAND Corporation

Root causes of failure for AI projects — anti-patterns of AI

##### Gartner Research

30% of generative AI proof-of-concepts abandoned by 2025

##### McKinsey & Company

State of AI: how organizations are rewiring to capture value


##### MIT Sloan Mgmt Review

The data problem stalling AI

# The Relabeling Trap


*The vocabulary changes.  
The mathematics does NOT.  
The failure mode remains.*

*Agentic AI 2025 -> Beyond?*



# The Three Frameworks of AI

*What AI actually does — beneath every label*



## DESCRIPTIVE ANALYTICS

### Unsupervised Learning

*What patterns exist in the data?*

Market regime detection · security clustering ·  
anomaly detection in price behavior

*Est. in finance: 1990s*




## PREDICTIVE ANALYTICS

### Supervised Learning

*What will happen next?*

Return prediction · factor signals · earnings  
surprise modeling · volatility forecasting

*Est. in finance: 2000s*



## RECOMMENDATION ANALYTICS

### Reinforcement Learning

*What action maximizes the objective?*

Execution optimization · dynamic portfolio  
rebalancing · adaptive order routing

*Est. in finance: 2010s+*

## Generative AI

*Cross-Paradigm Discipline*

**Unsupervised** scale + **Supervised** precision + **RL** alignment = **Generative AI**

*Enabled by NLP and Deep Learning at scale*




# WHERE AI FAILS ON THE TRADING FLOOR

*Six failure modes your allocators won't tell you about*

01

## Overfitting Theater



Backtests shine. Live trading collapses. The model memorized the exam — it never learned the subject.

02

## Label Leakage



Future data buried in training sets. The model appears to predict — it is cheating. Until capital is committed.

03

## Regime Blindness



Low-vol models meet March 2020. No algorithm survives a regime shift without human judgment layered in.

04

## The Committee Problem



40 AI vendors. Zero coordination. Conflicting signals across desks. Technology isn't the bottleneck — coordination is.

05

## Benchmark Theater



"Our AI outperforms." Outperforms what? On whose data? In which time period? Against which drawdown?

06

## Algorithm Age Argument



Linear regression (1800s) still beats neural nets in most live production systems. Complexity  $\neq$  capability.

**The constraint isn't the model. It's the organization around it.** → Governance next.



# GOVERNANCE & REGULATION

*The regulator is not behind the technology*

REGULATORY BODY	SPECIFIC RULE / REQUIREMENT	STATUS	IMPLICATION FOR THIS ROOM
 <b>SEC</b>	Predictive Data Analytics Rule — conflict-of-interest mitigation when AI optimizes for firm benefit over client benefit	<b>In Force 2023</b>	<i>Objective misalignment is now a compliance violation, not just a design flaw</i>
 <b>FINRA</b>	Rule 3110 + 2026 GenAI Report — formal review, approval, ongoing monitoring; agentic AI guardrails; mandatory kill-switches	<b>Active + 2026 Update</b>	<i>If your agentic system has no kill condition, you are non-compliant — not just poorly governed</i>
 <b>OCC + Fed Reserve</b>	Model Risk Management — validation, documentation, and stress-testing now explicitly extended to ML models previously exempt	<b>Expanding Ongoing</b>	<i>Models your quants built without MRM documentation are now in scope</i>
 <b>SEC / FINRA</b>	Explainability Mandates — algorithmic decisions affecting clients or markets must be explainable to regulators on demand	<b>Emerging Now</b>	<i>Models built without explainability architecture face costly retrofit or replacement</i>

**KEY INSIGHT** FINRA reviews hundreds of billions of market events daily using AI.

**Governance is not compliance overhead — it is the architecture of the 15% that succeed.** → A.I.R. Framework next.



# THE A.I.R. FRAMEWORK™

Assess · Implement · Revolve

## A

### ASSESS

*What problem are you actually solving?*

- Define success before the model runs
- Name your kill condition in advance
- Align objective with stakeholder intent

## I

### IMPLEMENT

*What does success look like before deployment?*

- Pre-defined metrics — not post-hoc ones
- Live-friction paper trading period
- Explainability architecture built in

## R

### REVOLVE

*What triggers your kill condition?*

- Model decay thresholds set before go-live
- Regime-shift monitoring loop
- Continuous ROI validation — not one-time

*“Better models won’t save a bad problem definition.”*

# What will you ask differently on Monday morning?

## John Thomas Foxworthy

Founder & Chief AI Officer, Global Institute of Data Science

AI / ML Instructor at Caltech CTME and UCSD Extended Studies

Global Institute of Data Science

Officially sponsored by Carnegie Mellon University's Heinz College

### GET IN TOUCH

● [johnthomas@gids.co](mailto:johnthomas@gids.co)

● [LinkedIn](#)

<https://www.linkedin.com/in/john-thomas-foxworthy-m-s-data-science-1718073/>