AI Stock Price Prediction

Why predict the market with AI?

- Financial markets = vast, complex, fast-moving data.
- Accurate forecasts can give investor an edge.
- Prices reflect both **hard data** (earnings, rates) and **human psychology**.

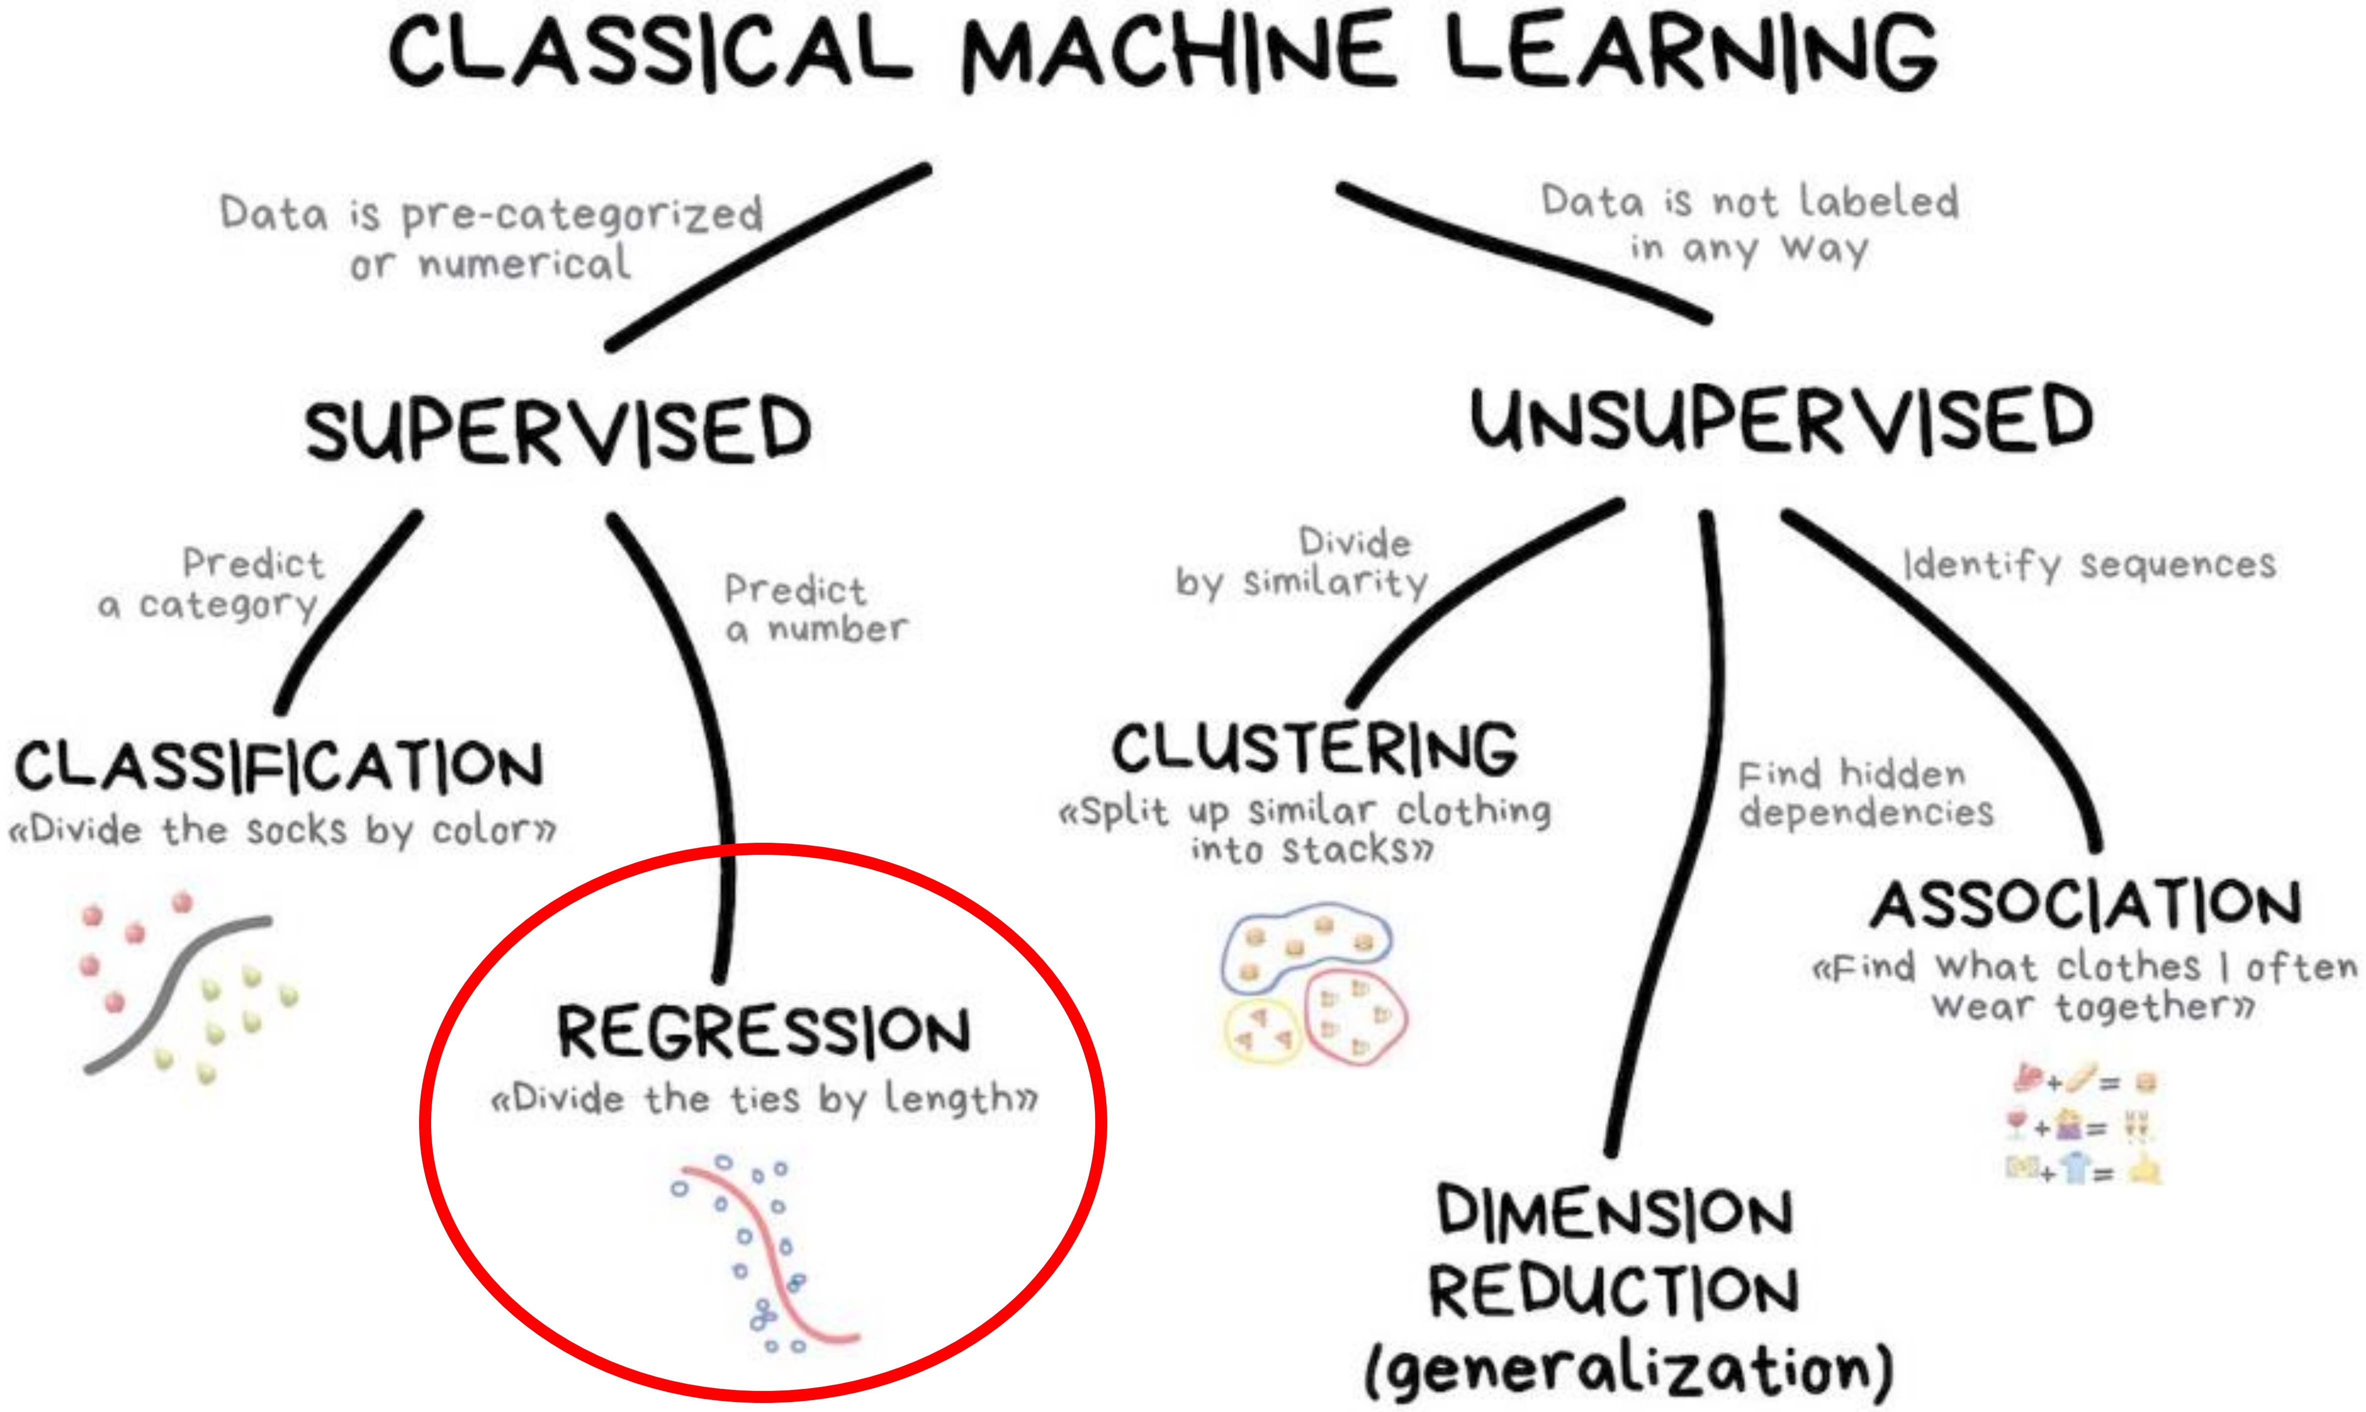
CLASSICAL MACHINE LEARNING
Data is pre-categorized or numerical
Data is not labeled in any way
SUPERVISED
UNSUPERVISED
Predict a category
Predict a number
Divide by similarity
Identify sequences
CLASSIFICATION
«Divide the socks by color»
REGRESSION
«Divide the ties by length»
CLUSTERING
«Split up similar clothing into stacks»
Find hidden dependencies
ASSOCIATION
«Find what clothes I often wear together»
DIMENSION REDUCTION (generalization)

Where AI is Taking Stock Prediction?

Classical Machine Learning → Deep Learning → Sentiment Analysis

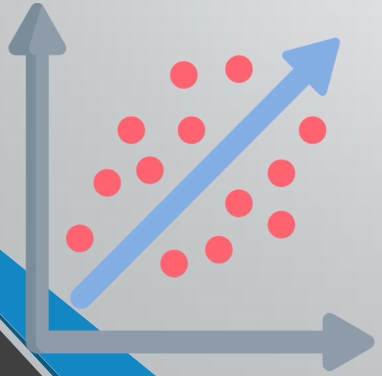

How Do Studies Evaluate AI Models

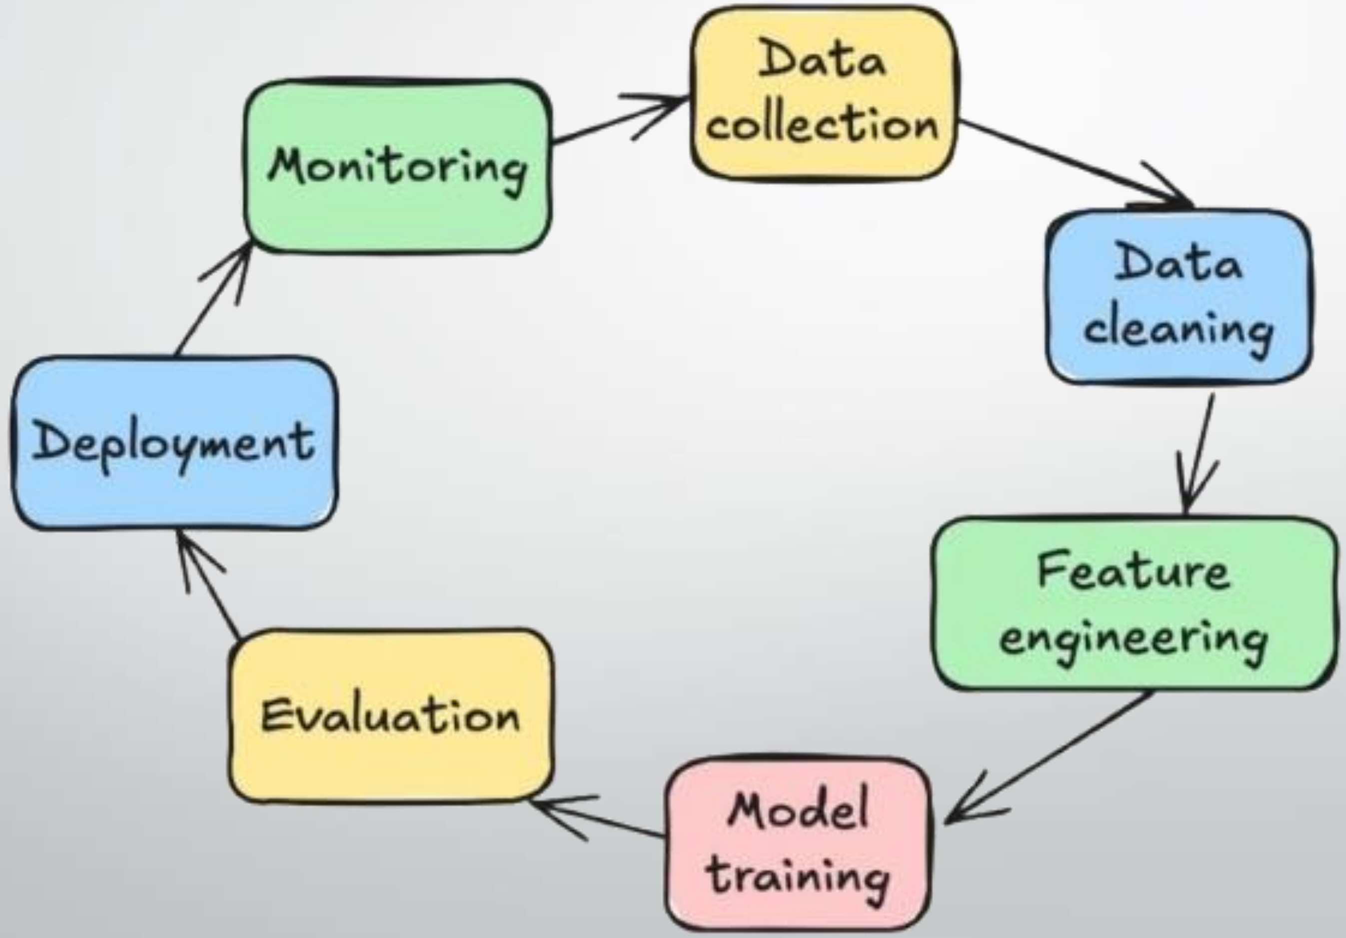

Traditional Approaches

Context

- **Objective:** To compare seven traditional machine learning algorithms to find the most accurate one for predicting the next day's rise or fall of major stocks around the world.

- **Can these traditional models, using only price history, reliably predict market direction?**

Methodology

- **Data Used:** 10 years of daily historical data (2012-2021) for 7 major G-7 indices.
- **Seven different algorithms:** ANN, SVM, Random Forest, Decision Trees, KNN, Naïve Bayes, and Logistic Regression.
- **Data was split:** 80% for training the models, 20% for testing their accuracy.
- **Evaluation Metric:** Prediction Accuracy - the percentage of days the model's forecast was correct.

Key Findings

- **Top 3 Algorithms:** ANNs (83.4%), Logistic Regression (82.6%), and SVMs (79.4%).
- **No Single "Best" Model:** The top algorithm varied by market, suggesting country-specific dynamics are important.
- **Random Forest:** Surprisingly performed poorly.

STOCK INDEXES	DT	RF	KNN	NB	LR	SVMs	ANNs
NYSE100	0.5337	0.6409	0.5298	0.6409	0.8214	0.7222	[a]**0.8373**
NIKKEI225	0.5414	0.5717	0.4848	0.5899	[a]0.8162	0.8000	**0.8101**
FTSE100	0.5889	0.5929	0.5119	0.6067	0.8498	0.8162	[a]**0.9348**
CAC40	0.5957	0.6270	0.5098	0.6406	[a]0.8359	0.8105	**0.8301**
DAX30	0.5474	0.5059	0.4960	0.6225	0.8083	0.8162	[a]**0.8182**
FTSEMIB	0.5656	0.6282	0.4834	0.6380	0.8219	0.8004	[a]**0.8513**
TSX	0.5916	0.5637	0.5438	0.6434	[a]0.8645	0.8466	**0.8586**
AVARAGE	0.5663	0.5900	0.5026	0.6260	0.8256	0.7943	[a]**0.8343**

Implications

- **"Black Box" Problem:** ANNs is the hardest to interpret which can be a big risk for investors
- **Accuracy ≠ Profitability**
- **Financial Decision-Making:** These models can inform investment, corporate, and economic strategies.

Deep Learning Approaches

Comparison with Classical Models

- Addition of in-built memory into the models
  - Backtracking
  - Reference to previous patterns
- Hypothesis is that it will lead to better pattern recognition
  - More accurate stock price predictions

Deep Learning Approaches

Transformers

Gated Recurrent Units (GRUs)

Long Short-Term Memory (LSTM)

The Study

- Use a time-series dataset of Tesla stock prices from 2015-2024
- Input product information, user reviews and sales data
- Train each of the three models
- Evaluate the models
  - Given a timestamp, predict the stock price for the next 30 days
  - Compare predictions with real historical stock prices

Transformers

- Attention-based neural network
- 0.80 R^2 value
- Lowest root mean square error (RMSE)

Gated Recurrent Units (GRUs)

- Simpler version of LSTM
- Controls the flow of information by resetting and updating doors
- 0.85 R^2 value

Long Short-Term Memory (LSTM) model

- Most accurate model of the three
- Specifically designed to process and predict time-series data
- 0.98 R^2 value (measurement of regression fit)

Limitations

Analysis based solely on Tesla stock

Not transferrable to other stocks

Only one industry explored

No ability to track multiple stocks at once (like an indexed fund)

Models predict closing stock price at end of day

Always predicting the price at the same time of day

Intra-day fluctuations missed

Intra-day Stock Price Predictions

How ChatGPT Can Decode Financial News to Predict Short-Term Market Moves

ChatGPT Sentiment Score Applications

4,600+ NEWS ITEMS

1,200+ DIVIDEND ANNOUNCEMENTS

394 S&P 500 COMPANIES

Core Question

Can ChatGPT's analysis of financial news predict short-term stock movements after a major event like a dividend announcement?

Machine Learning Models

- Ridge Regression
- Logistic Regression
- Random Forest
- XGBoost
- LSTM

Key Result

Sentiment is a statistically significant predictor of intraday returns.

The Findings

- More positive news led to higher abnormal returns.
- LSTM and Ridge Regression models were the most effective.
- **Peak Performance:**
  - Predictive power was strongest in the **first 2 hours** after news broke.
  - Directional accuracy hit **82.2%** in the initial trading window.

The Bottom Line

A trading strategy based on these sentiment signals outperformed benchmarks.

Turning Sentiment into Profit

- The LSTM model's strategy captured a **cumulative abnormal return of 0.756%** in the first 30 minutes post-announcement.
- It demonstrates a concrete, profitable application.

Main Contribution

- This study provides a practical blueprint for quantifying market psychology.
- **Before:** Sentiment was a vague, "fuzzy" idea.
- **Now:** It's a measurable, actionable input for a trading algorithm.
- Rigorous out-of-sample and placebo tests confirmed the results weren't just luck.

The Future

Trading is moving towards AI that can interpret market psychology as it happens.

Ethical Considerations in AI Trading Systems

Ethical Areas

- Transparency & Explainability → black-box models (ANN, LSTM, Transformers)
- Fairness → institutional advantage vs. retail access
- Privacy → use of alternative data (social media, news, geolocation)
- Accountability → who is responsible if AI causes market disruption
- Regulation → EU AI Act, China's Fintech Plan classify financial AI as high-risk

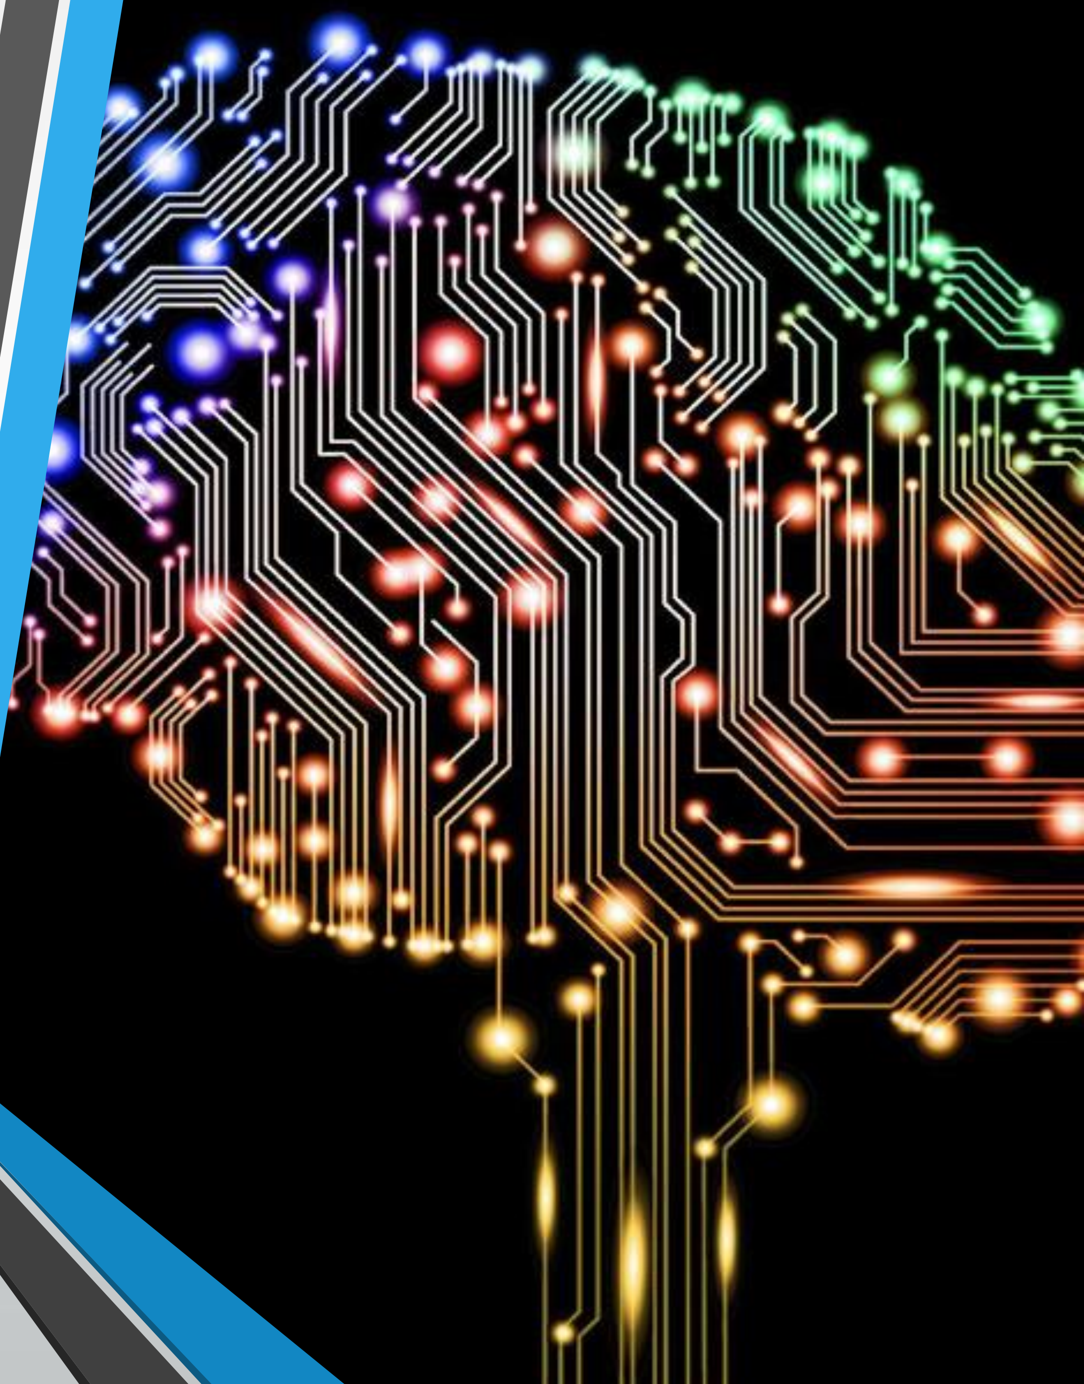

Systematic Exploration Framework

- **Normative ethical theory application** to algorithmic decision-making in financial markets
- **Multi-dimensional moral analysis** examining fairness, transparency, and accountability
- **Stakeholder impact assessment** evaluating effects on different market participants

Methodological Approach

- Applies established moral philosophy frameworks (utilitarian, deontological, virtue ethics) to real-world AI trading scenarios
- Uses existing ethical theories to provide practical guidance rather than creating new frameworks
- Tests how ethical principles can be integrated into current trading system design

Systematic Ethical Paradigm Integration

UTILITARIAN/CONSEQUENTIALIST FRAMEWORK

DEONTOLOGICAL FRAMEWORK IMPLEMENTATION

VIRTUE ETHICS INTEGRATION

Utilitarian/Consequentialist Framework

- **Aggregate Welfare Maximization**: AI trading creates collective benefits through enhanced liquidity and price discovery
- **Distributional Justice**: Tension between market efficiency gains and equitable benefit distribution
- **Long-term Assessment**: Improved capital allocation efficiency provides sustained societal benefits
- **Cost-Benefit Analysis**: Balance immediate trading benefits against potential systemic risks

Deontological Framework

- **Categorical Imperative**: Universal moral principles mandate algorithmic transparency regardless of competitive advantages
- **Rights-Based Access**: Fundamental equality principles require fair market participation for all stakeholders
- **Information Equity**: Moral duties prevent exploitation of informational asymmetries
- **Inherent Rules**: Prohibition of market manipulation regardless of profitability

Virtue Ethics

- **Transparency**: Algorithmic decision-making must embody openness and accountability as core characteristics
- **Integrity**: Trading algorithms should reflect consistent moral principles within institutional culture
- **Responsibility**: Clear moral accountability structures linking algorithmic outcomes to human agents
- **Prudence**: Ethical systems demonstrate careful risk assessment and stakeholder consideration

Ethical Framework Applications in AI Trading

Fairness & Justice

- **Market Access**: AI should enhance equitable participation, not create technological divides
- **Bias Prevention**: Identify and correct discriminatory algorithmic behaviours

Transparency & Accountability

- **Explainability**: AI systems must provide clear justifications for trading decisions
- **Stakeholder Rights**: Market participants deserve to understand decisions affecting them

Market Integrity

- **Systemic Impact**: Individual AI systems affect overall market stability
- **Manipulation Prevention**: Prevent algorithmic behaviours that undermine market integrity

- Flash Crash (2010, US markets): Algorithmic trading amplified volatility → trillions lost briefly
- COVID-19 Crash (2020): AI models failed to adapt to unprecedented events → large sell-offs
- AI-Driven Funds: Outperform benchmarks in some periods, but also increase systemic risk
- Lesson: AI improves speed & efficiency, but unchecked use can destabilise markets

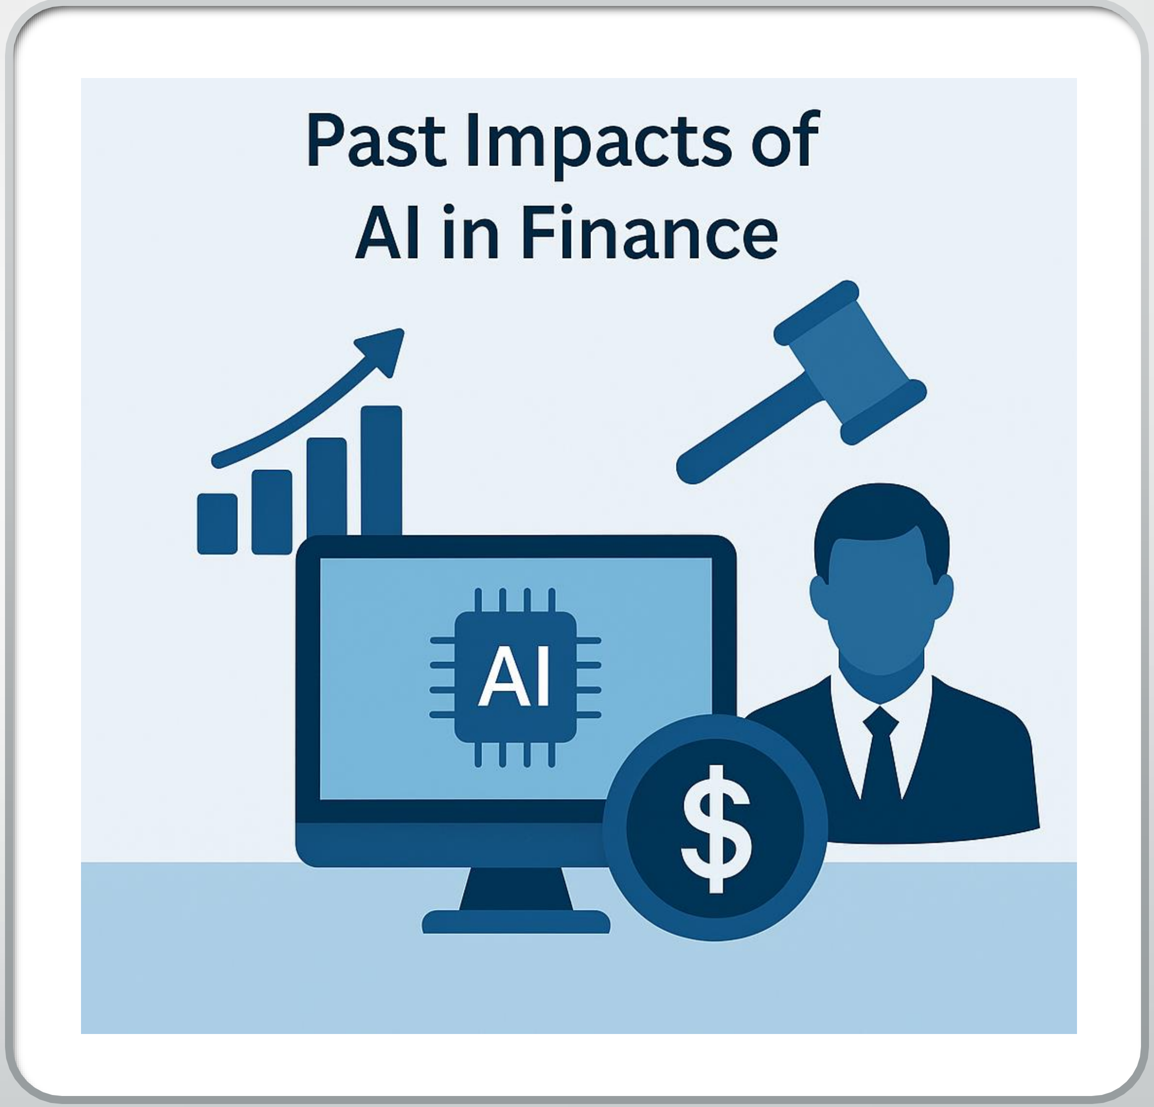

Practical Application for Stakeholders

- Regulators: Require explainability & auditing for AI trading systems
- Financial Institutions: Implement ethical frameworks (fairness, privacy, accountability) in model design
- Retail Investors: Demand transparency in robo-advisors & AI-driven investment products
- Technology Developers: Balance accuracy with interpretability & resource sustainability

- Explainable AI (XAI): Making black-box models interpretable for regulators & investors
- Fair AI Access: Reduce institutional vs retail investor divide by promoting affordable tools
- Privacy-First AI: Strict limits on alternative data (social media, geolocation) usage
- Sustainable AI: Energy-efficient models to reduce environmental impact of large-scale training

CONCLUSION

- Classical ML → useful & interpretable, but limited in complex markets
- Deep Learning → higher accuracy, costly & less transparent
- Sentiment & LLMs → capture market psychology, intraday value
- Ethics → accuracy alone is not enough; fairness & regulation are vital
- Future → Hybrid, responsible AI systems for fair & stable markets