SERC

SMITH ENTERPRISE RISK CONSORTIUM

AI in Banking: Balancing Innovation, Risk, and Responsibility

Dr. Clifford Rossi, Director, SERC
Professor of the Practice & Executive in Residence
Robert H. Smith School of Business

AGENDA

1. AI Critical Path Enablers & Risks

2. AI Tools and Tech in Financial Services

4. AI Governance & Risk Management

3. AI Applications & Use Cases

Pre-AI World

Post-AI World

What Could Possibly Go Wrong?

5. Takeaways and What it Means for Your Business

AI Opportunities in Banking

Customer Growth
& Earnings

Risk Reduction

Process
Automation

Increased Risk-adjusted Returns

Identification of AI Critical Path Enablers and Risks

- **Leverage and Scale Enablers**
  - Existing technology systems and infrastructure capabilities
  - Employee domain expertise (partially drives build vs buy decisions)
  - Enterprise data architecture and management
    - Structured
    - Unstructured
- **Assess and establish necessary guardrails prior to deployment** – (*absolutely essential*)
  - Policies and procedures (Enterprise Risk AI Policy and framework and enhancement to Model Risk Management Policy)
  - Information security
  - Monitoring and feedback loops
- **Board, Executive Committee and regulatory engagement and acceptance**

AI Use Cases and Potential Sequencing*

- Focus on broad reusable applications than on narrow use cases for maximum AI impact
- Project sequencing based on business value and technical level of effort

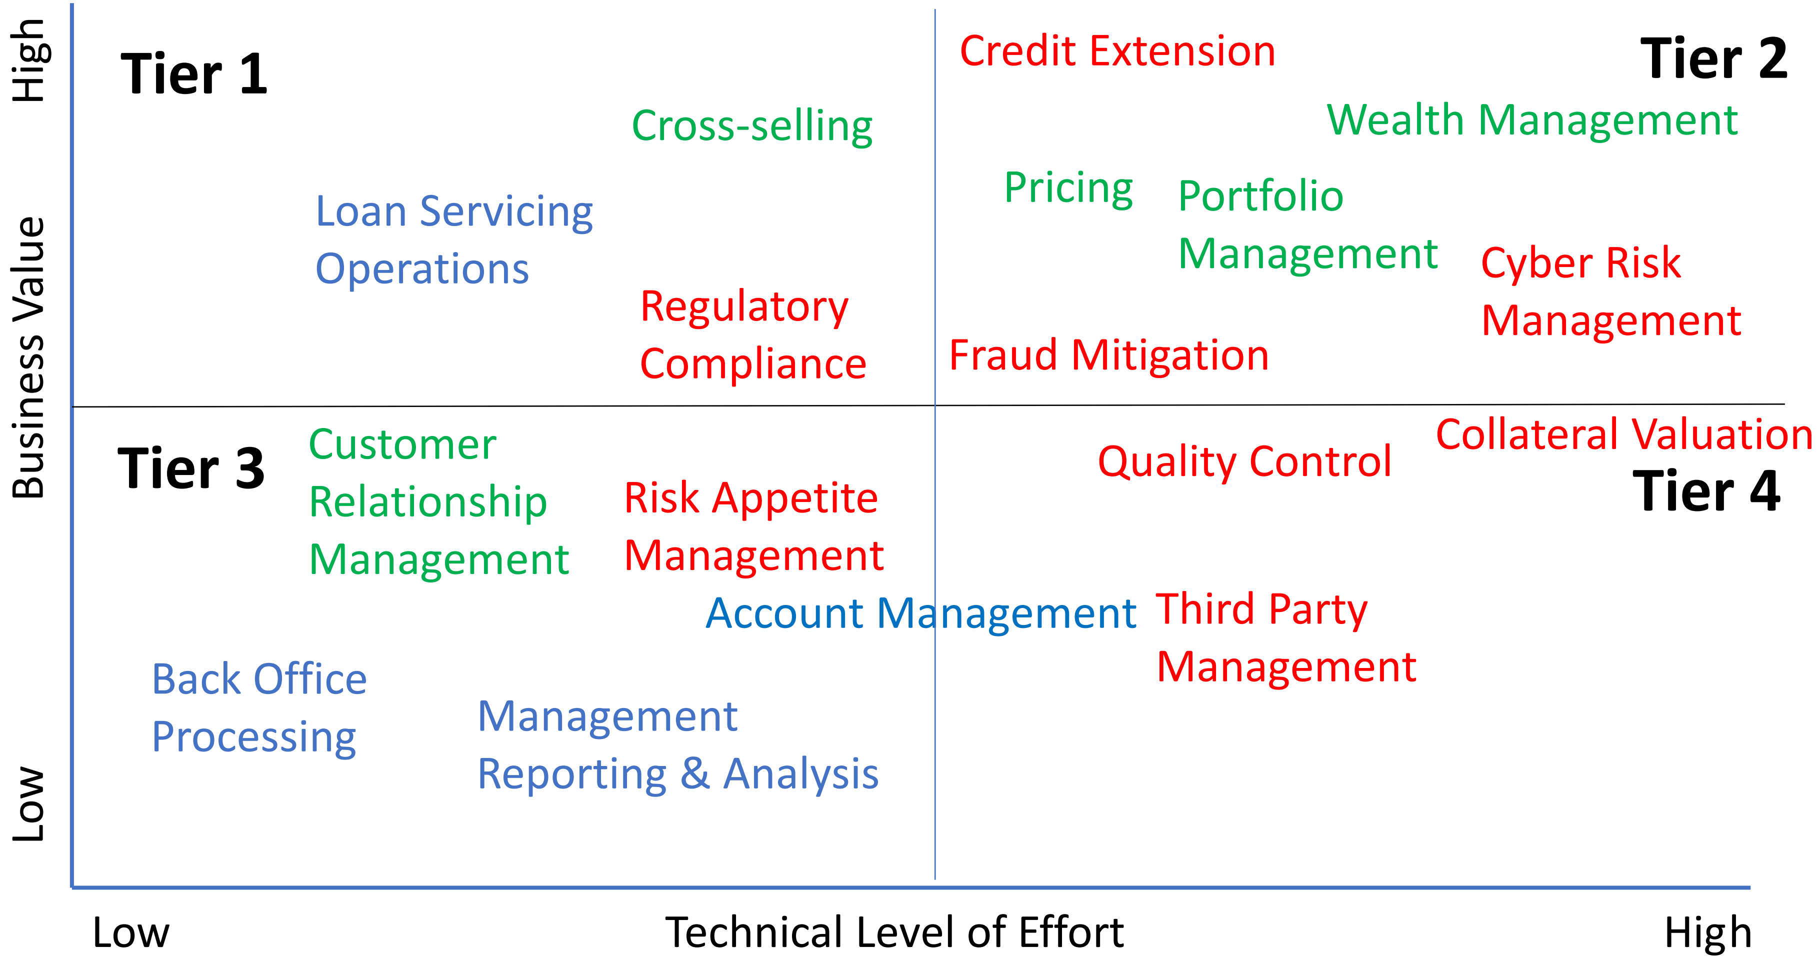

- Sequencing of AI projects might be as follows:
  - Tier 1 – 1st
  - Tier 2 – 2nd
  - Tier 3 – 3rd
  - Tier 4 – 4th

Business
Risk
Process

* Not exhaustive list and illustrative placement of use cases

Building an AI Centric Organization

Acitvity	Subactivity	Home Loans	Business Line Treasury Mgt	Auto	Wealth Mgt
Customer Engagement					
	Acquisition	X		X	X
	CRM	X		X	X
	Cross selling	X		X	X
	Retention	X		X	X
	Default Management & Servicing	X		X	
Evaluation/Assessment					
	Pricing/Valuation	X	X	X	X
	Underwriting	X		X	
	Portfolio Allocation	X	X	X	
	Collateral Valuation	X		X	
Processing					
	Loan onboarding	X		X	
	Payments	X		X	X
	Transactions		X		X
	Account Management		X		X
Reporting					
	Regulatory & Compliance	X	X	X	X
	Management & Board	X	X	X	X
Enabler					
Risk Management					
	Policies & Procedures	X	X	X	X
	Third Party Risk Management	X	X	X	X
	Model Risk Management	X	X	X	X
	Cyber Risk Management	X	X	X	X
	Fraud Management	X	X	X	X
	ALM		X		
	Regulatory & Compliance	X	X	X	X
	Credit	X		X	X
	Legal	X	X	X	X
	Operational Risk Management	X	X	X	X
Data Management					
	Structured	X	X	X	X
	Unstructured	X	X	X	X
Technology/Systems					
	AI Hub	X	X	X	X
	Information Security	X	X	X	X
	Cloud	X	X	X	X
	API architecture	X	X	X	X

Note: Figure illustrative

- Rewiring the organization for AI requires detailed mapping of core activities and subactivities for each business line and its suitability for an AI solution
- Risks must be identified for each business line and associated requirements determined in advance of any AI rollout
- Data management needs for AI tools in each business line must be determined and developed (e.g., structured and unstructured data)
- Technology, systems and platforms must be evaluated across the enterprise to ensure that AI applications can be hosted with minimal disruption
- Centralized AI hub ensures tools are consistent, and managed at an enterprise level for quality, third party and model risk purposes
- AI tools for each subactivity can then be determined for suitability; e.g., LLMs for regulatory and compliance reporting vs machine learning tools for predictive analytics

Considerations in Executing the AI Strategy

- Test drive the process on a limited basis before moving to widespread rollout
- Follow an interdisciplinary strategy; involve business, risk, technology and other key players in formulating the strategy
  - Consider creating an AI Competitive Advantage Team (AICAT) charged with executing senior management AI strategy across the company
- Executive accountability; identify who is responsible for developing and leading the strategy to realization
  - Business should drive strategy with Risk and Technology as key partners
- Create detailed AI execution plan
  - Timelines
  - Targets/Milestones
  - Assignments and accountabilities

AI Risks in Banking

Regulatory & Compliance Risk

Model Risk

Reputation Risk

Operational Risk

Market Risk

Legal Risk

Information Security Risk

Credit Risk

3rd Party Risk

What Exactly is AI? The AI Ecosystem

Supervised Learning

- Uses labeled data to make predictions or identify relationships
- Logit credit underwriting model

Unsupervised Learning

- Leveraging unlabeled data to identify patterns without direct guidance
- Cluster analysis of borrowers for risk segmentation
- Deep Learning is a subset

Reinforcement Learning

- Autonomous problem-solving
- Agentic AI to make a credit decision

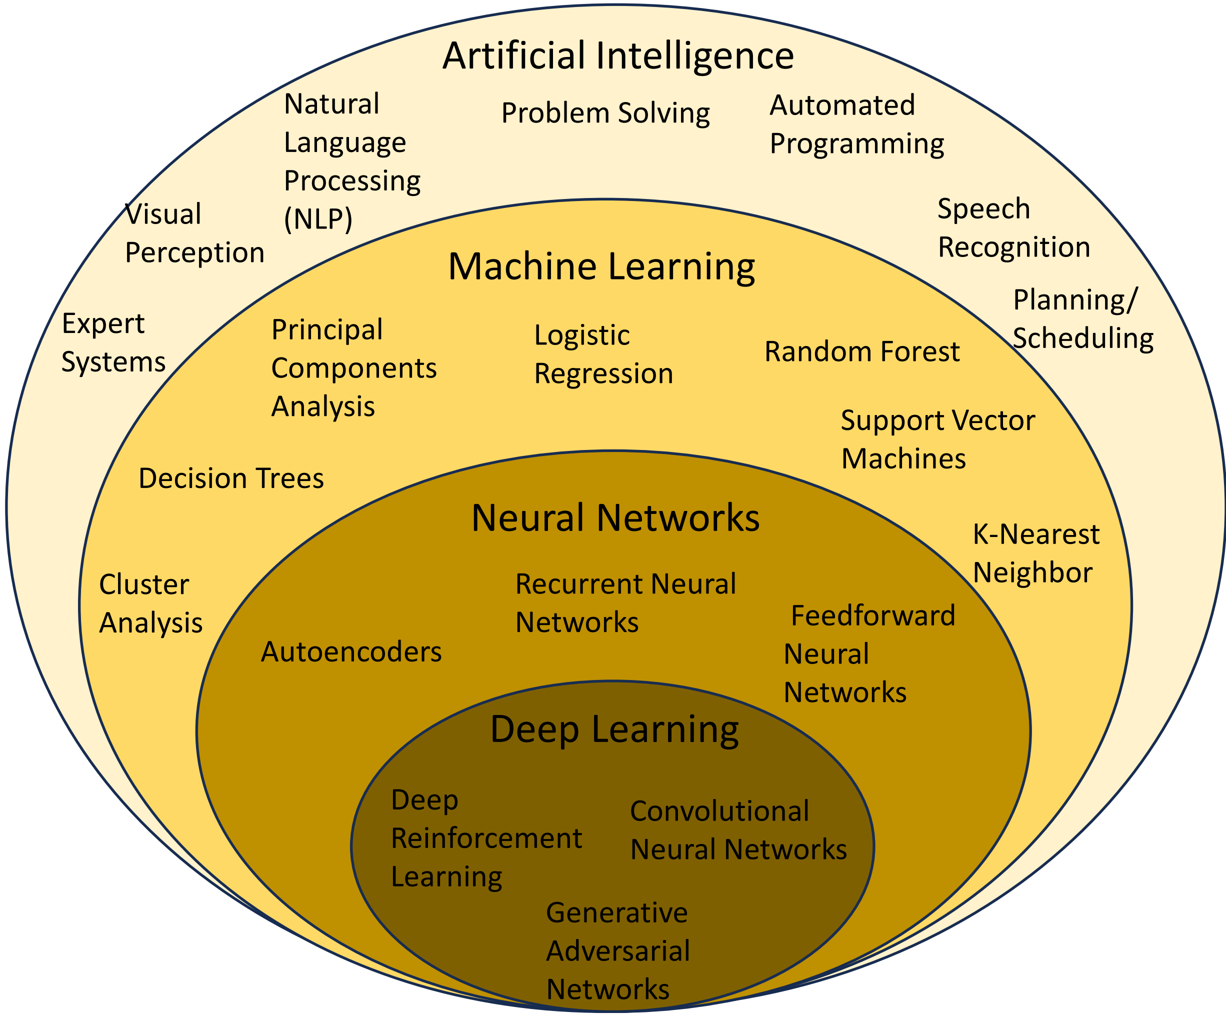

AI Use Cases for Discussion

Use Cases for Discussion

1. Auto Loan Credit Underwriting
2. Collateral Valuation
3. Fraud Detection - Credit Cards
4. Asset-Liability Management (ALM)
5. Regulatory & Compliance
   - Home Mortgage Disclosure Act (HMDA)
6. Robo-Advising

AI Techniques

Decision Trees

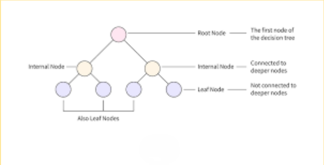

Convolutional Neural Network
Natural Language Processing

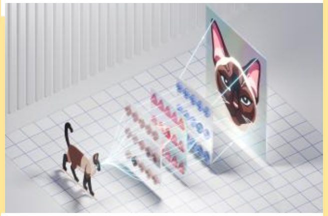

Generative Adversarial Networks (GANs)

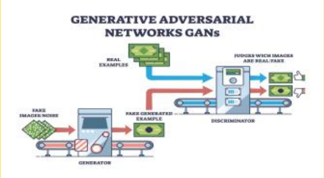

Agentic AI
Deep Reinforcement Learning

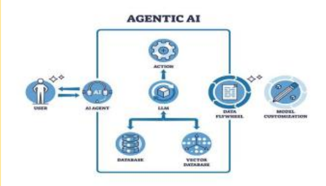

Large Language Model (LLM)

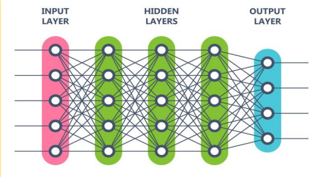

Recurrent Neural Network
LLM

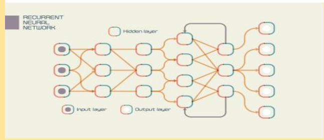

Use Case #1: Auto Loan Credit Underwriting

Pre-AI World – Logit model at a Bank

Borrower Attributes
- Credit score
- Debt-to-income ratio
- Loan-to-value ratio

Collateral Attributes
- Model
- Make
- Year

Loan Attributes
- Loan size
- Terms

Generate
Score
Pr(Default)

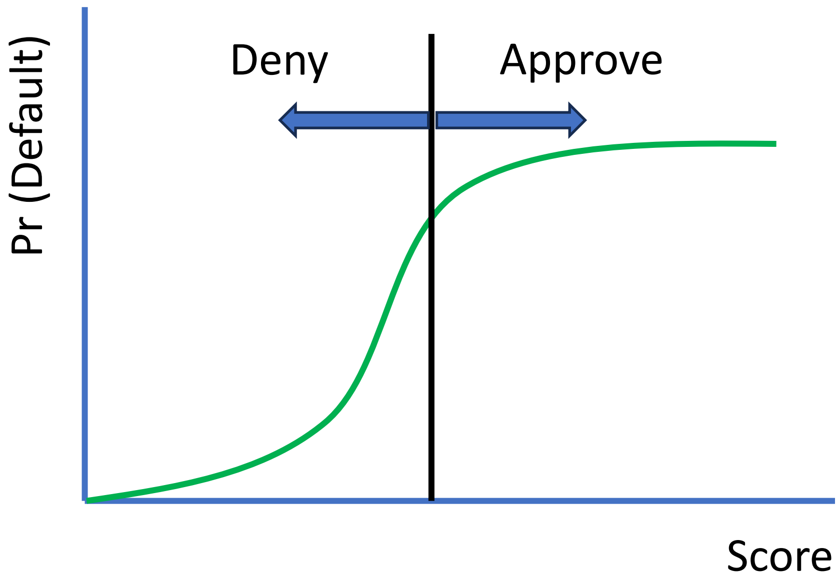

Use Case #1: Auto Loan Credit Underwriting

- Post-AI World – Decision Tree Model (e.g., Random Forest) at a Fintech

Same data used as before, but AI model identifies patterns not inherently understood or easily explained (lack of transparency)

Example: DTI >40% **and** LTV > 90%

Increases predictive power over Logit…

But at what cost?

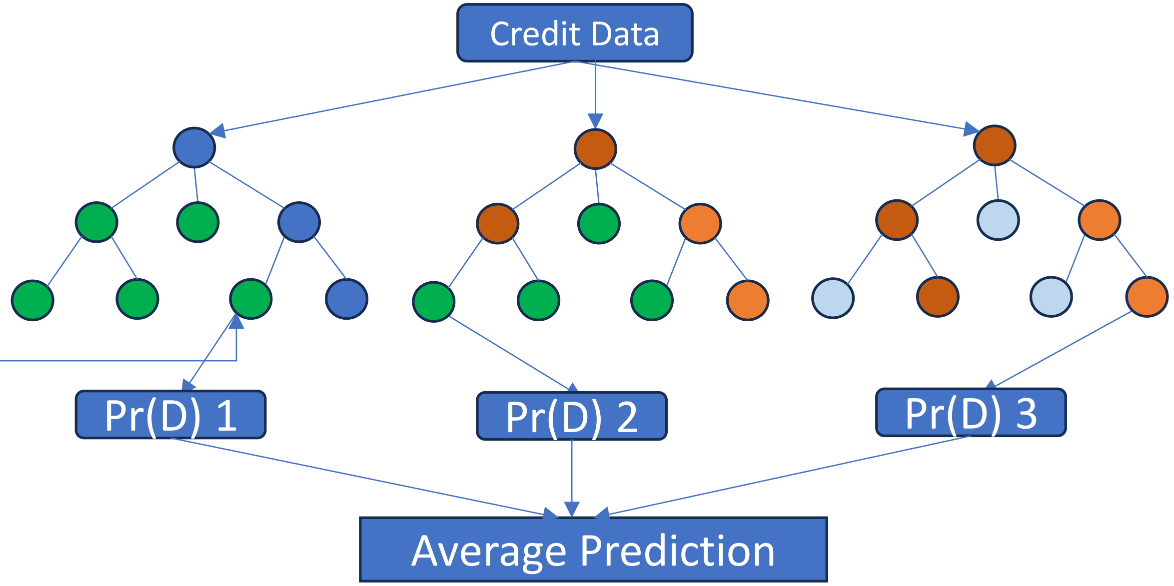

Use Case #1: Auto Loan Credit Underwriting

- AI Identifies Hidden Value

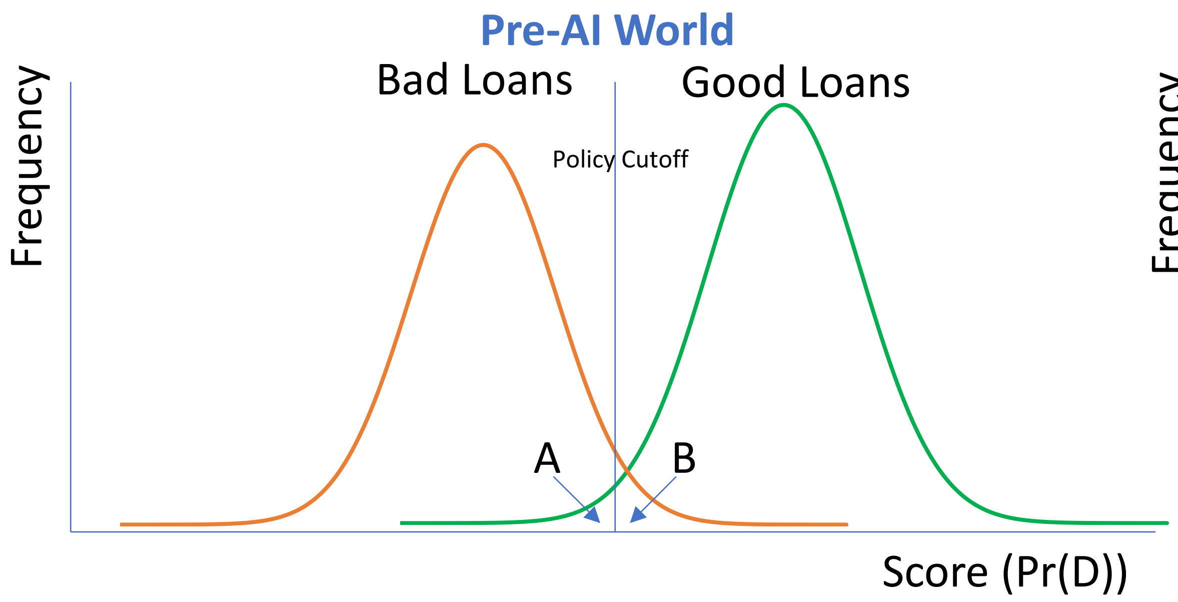

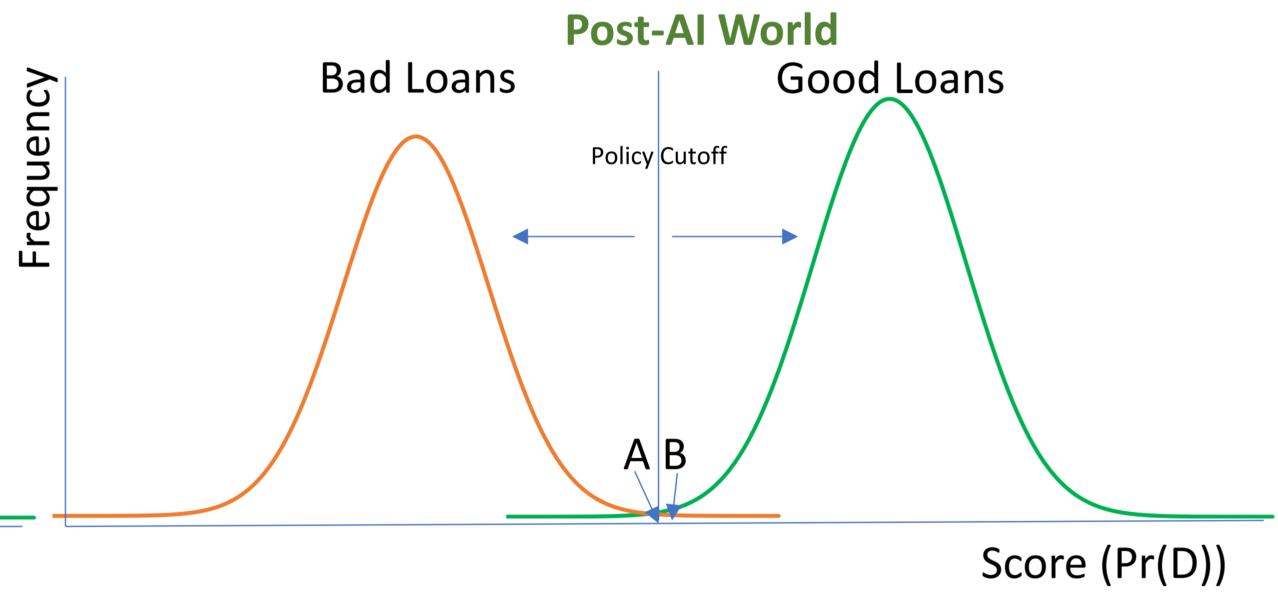

Area A represents good loans that were rejected – lost revenue
Area B represents bad loans that were accepted – credit loss

AI model identification of other hidden relationships allow better separation between good and bad loans (i.e., bad and good loan distributions push further apart), thus reducing lost revenue and credit losses

Use Case #1: Auto Loan Credit Underwriting

- What Could Possibly Go Wrong?
  - Additional hidden relationships such as DTIs >40% and LTVs >90% could lead to potential ECOA Reg B violations for disparate treatment
    - Lack of explainability and transparency compared to regression-based analysis
    - Overrides imposed from AI model rejecting more protected classes
    - Scores determining credit may also wind up rejecting more protected class borrowers due to risk factor interactions correlated with a protected class feature
    - Data used in training the model could have some inherent bias
    - Loan pricing could be higher for protected classes based on additional insights gained from the AI model
  - The model may still be acceptable based on business judgment, but additional testing would be required

Use Case #2: Collateral Valuation

- Pre-AI World

MLS and County Property-level Data Inputs
- # bedrooms
- #bathrooms
- Square footage
- Lot size
- Age
- Location
- Home type
- Neighborhood features

→

Regression Analysis (Hedonic and Repeat Sales) Generates Automated Value w/ Forecast Deviations

→

Applications
- Credit extension
- Portfolio valuation
- Investment

Problem: no recognition of differences in condition or marketability

Use Case #2: Collateral Valuation

- Post-AI World

Property Attribute Data
+
Interior/exterior photos
Realtor (MLS) listing description

AI-based Analysis – e.g., Convolutional Neural Network (CNN)

- Image Segmentation/Classification

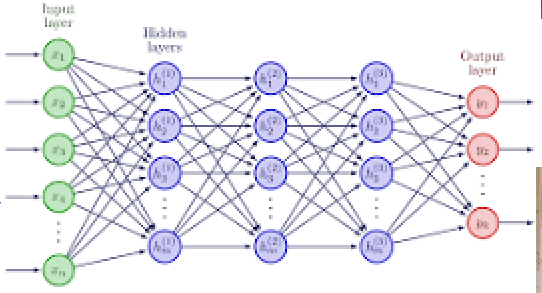

or

Generates Automated Value w/ Forecast Deviations Inclusive of Condition & Marketability

Natural Language Processing (NLP) model

- Optical Character Recognition
- Text mining

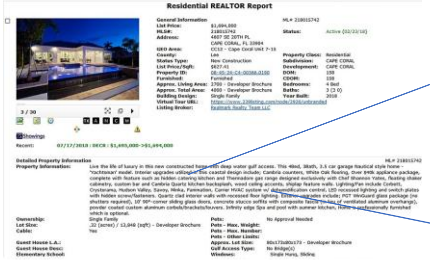

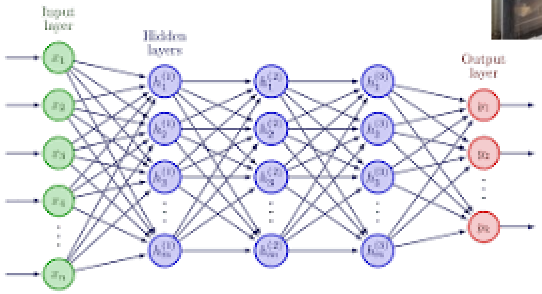

New granite countertops

Use Case #2: Collateral Valuation

- What Could Possibly Go Wrong?

AI model introduces bias driven by misinterpretation of features
Or bias from historical undervaluation in certain neighborhoods

Underestimated Value

Model misinterprets condition and upgrades from realtor description and photos

Credit not extended (AI generates higher LTV)

Business Implication: Lost revenue

Accurate Value

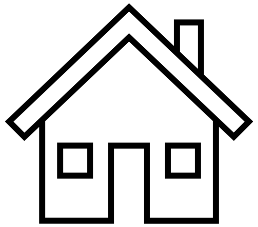

Realtor's property listing accurately depicts condition and upgrades

Credit extended

Use Case # 3: Fraud Detection - Credit Cards

- Pre-AI World

Transaction, Behavior and Other Data

$ Amount
Frequency
Time & Date
Payment Type
Purchase History
Location
Merchant Code
Recency
Account Activity
IP Address
ID verification

Techniques

- Rules-based
- Statistical model (e.g., logit)

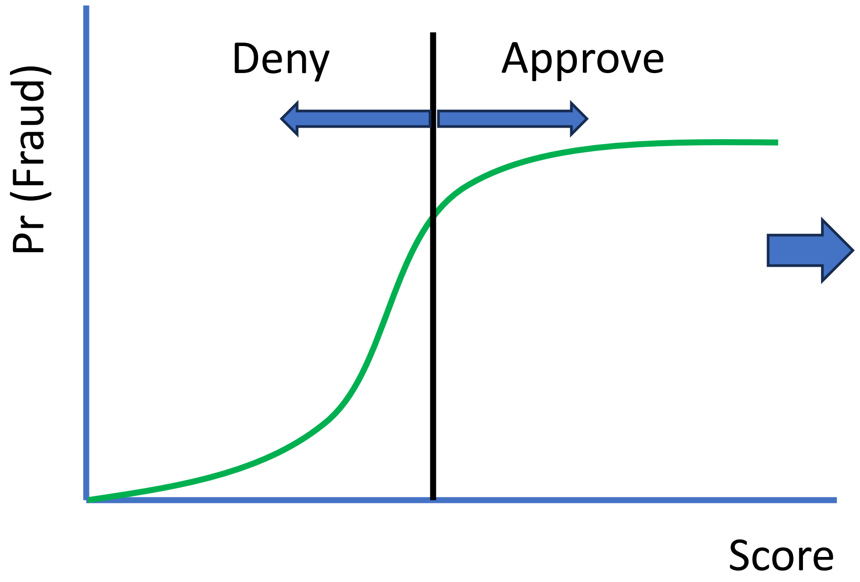

Customer X

Credit Card Transaction 1
- $35 Gas purchase
- Mobil Station
- 1 mile from home
- 10 similar transactions last 3 months
- Decision - **Approve**

Credit Card Transaction 2
- $2,000 e-bike purchase
- Bike store
- Yuma, AZ (3,000 miles away)
- No other transactions in area (e.g., airline tix)
- Decision – **Verify or Deny**

★ Imbalanced data; i.e., very few fraudulent transactions compared with the sample make it difficult for models such as logit to accurately identify fraud

Use Case #3: Fraud Detection –Credit Cards

- Post AI World – Generative Adversarial Networks (GANs)

z represents all the latent features of a fraudulent transaction

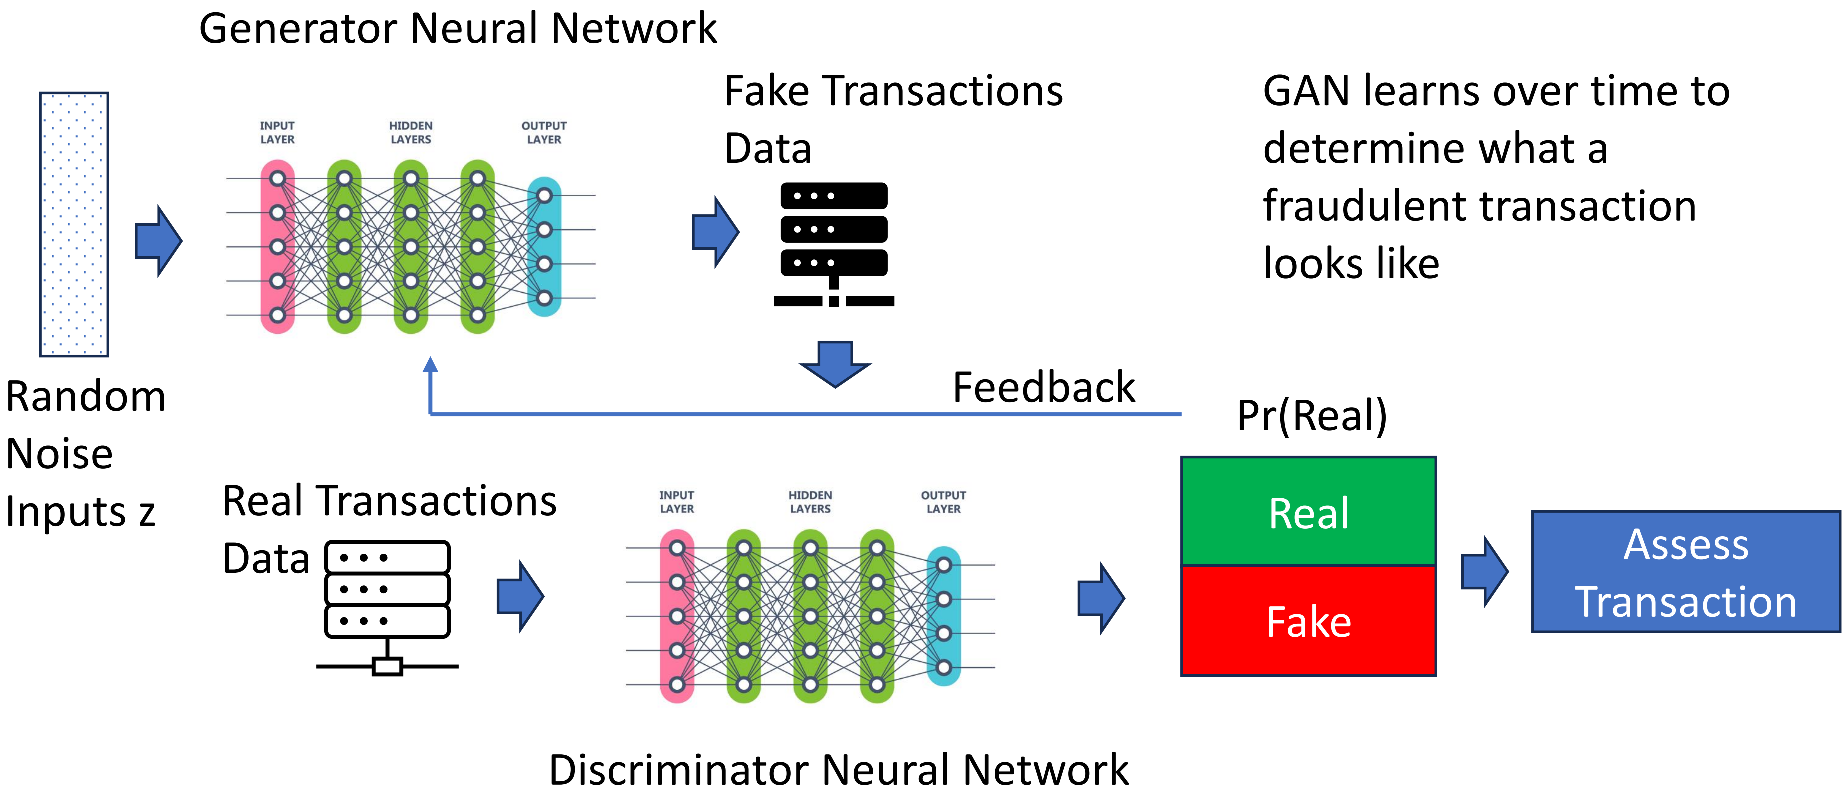

GAN learns over time to determine what a fraudulent transaction looks like

Use Case 3: Fraud Detection – Credit Cards

- What Could Possibly Go Wrong?

Training Dataset

Bias?

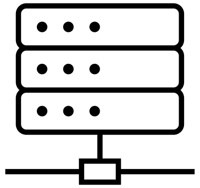

Instability?

Discrimination?

Regulatory/Legal/Reputation Impacts

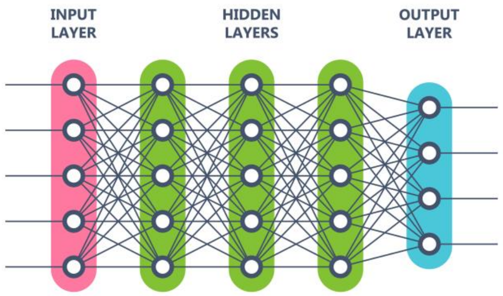

Overfitting?

Explainability?

Fraud Detection Errors?

- Flagging transaction fraud when none exists
- Approving fraudulent transaction

Financial/Business Impacts

Use Case 4: Asset-Liability Management

- Pre-AI World

Banks are naturally exposed to interest rate risk and must balance that risk between their assets and liabilities

Bank assets tend to be more sensitive to changes in interest rates than liabilities

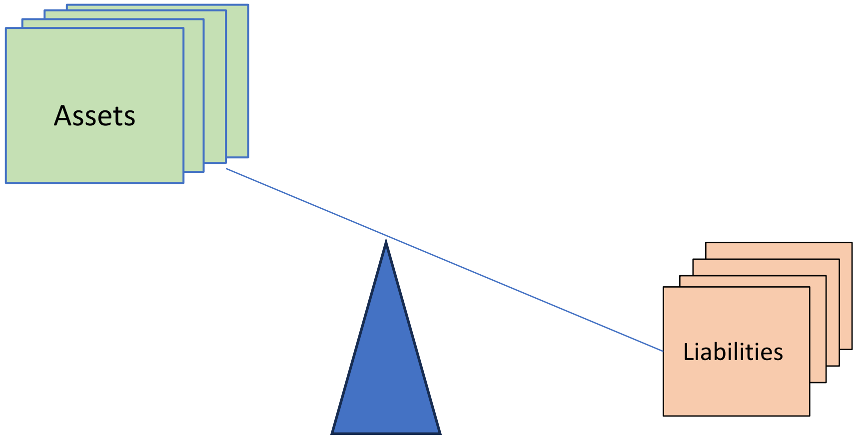

To address potential imbalances, banks will attempt to manage the duration of assets and liabilities to some target level via a process called portfolio immunization

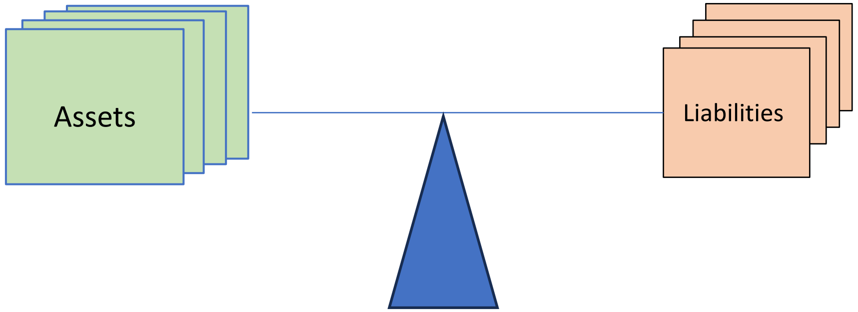

- Portfolio immunization combines elements of technical analysis with judgment
- Balancing multiple objectives; interest rate and liquidity risk management along with risk-adjusted return maximization becomes difficult to accomplish and this is where AI can lend a hand

Use Case 4: Asset-Liability Management

- Post-AI World

Recurrent Neural Network allows for time-dependent relationships between states (duration of assets and liabilities) and actions (weights of assets and liabilities)

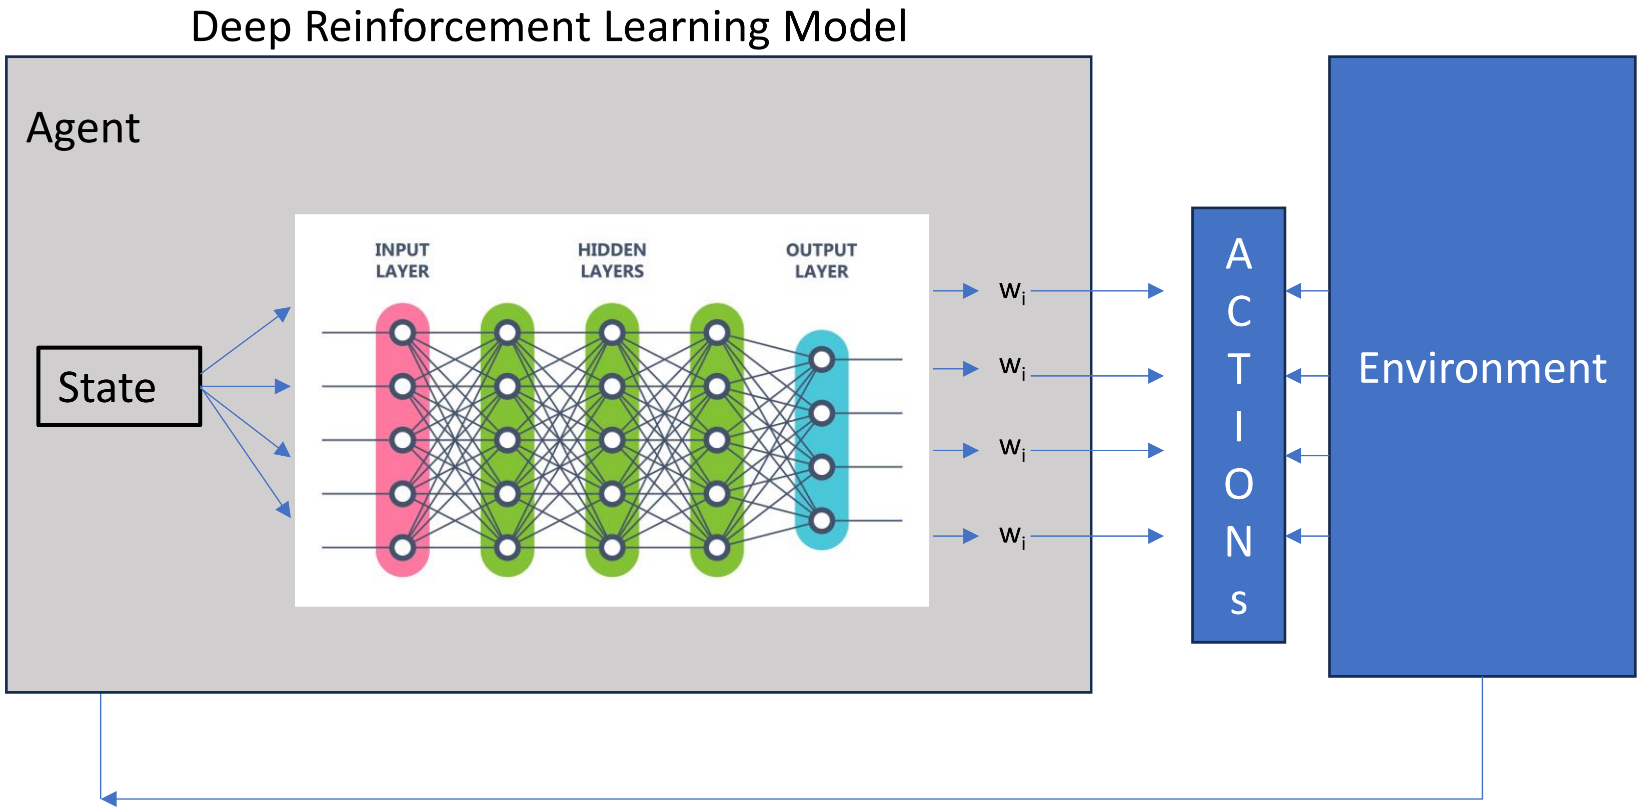

Interest rates and other factors in the ALM environment determine states (durations) for which autonomous policies (portfolio composition) take place to meet the objective

Objective is to identify best policies that minimize penalty function (e.g., $PV_{Equity} = PV_{Assets} - PV_{Liabilities}$)

Use Case 4: Asset-Liability Management

- What Could Possibly Go Wrong?

1. Training data might not be reflective of prospective environment - Think pre-March 2022 period of interest rates

   Result: Portfolio misallocation and lower PV_{Equity}

2. Model complexity risk

   Result: Black box nature of a neural network model makes it extremely difficult to understand how key factors such as interest rates determine asset and liability weightings

3. Penalty function might not be sufficiently specified to capture other important aspects of the objective – e.g., preserving a certain level of liquid assets

   Result: Potential portfolio misallocation and greater liquidity risk

4. Potential for overfitting

   Result: Model is unable to generalize to different conditions or environments

Use Case 5: Regulatory & Compliance

- Pre-AI World
  - Since 2008, number of new laws and regulations in banking skyrocketed and along with that the cost of compliance.
  - [Bank Policy Institute reports](#) that the number of hours of bank employees devoted to compliance-related activities rose over 60% between 2016 and 2023.

For mortgage companies, this heightened compliance environment has become especially burdensome

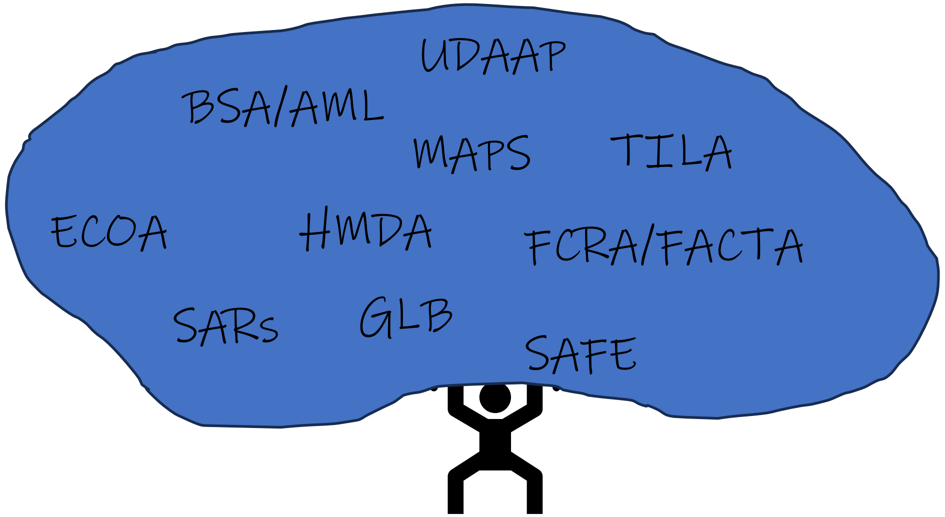

Across all steps in the mortgage origination process, compliance personnel ensure that companies adhere to all federal and state laws and regulations

While this has become increasingly automated, AI is becoming a game changer in significantly reducing compliance costs and enhancing the process

Use Case 5: Regulatory & Compliance

- Post-AI World

  Large Language Model (LLM) used to determine compliance with specific regulations (e.g., Home Mortgage Disclosure Act (HMDA))

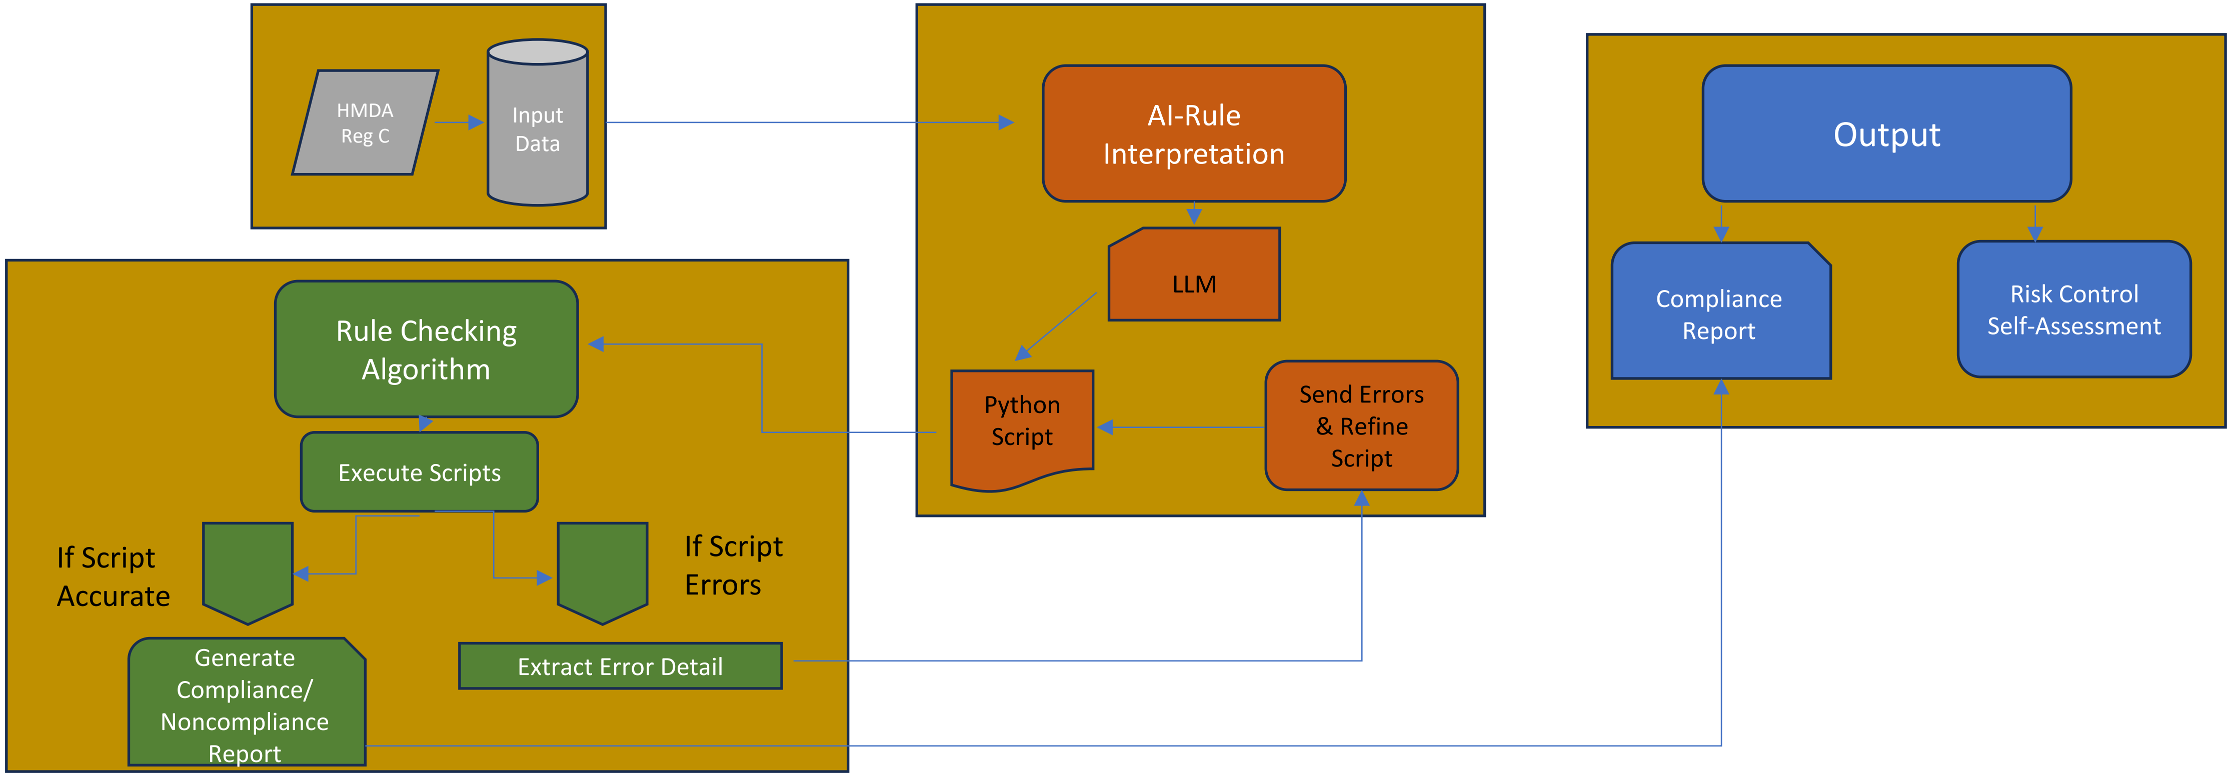

Use Case 5: Regulatory & Compliance

- Post-AI World – Rules, Prompts and Scripts

Example HDMA Rule Checks

- Confirm that the date reported on the LAR matches the actual date the institution received the complete application
- Ensure the amount reported on the LAR matches the loan amount listed in the loan closing documents.
- Verify the action taken (e.g., approved, denied, withdrawn) and the date of the action taken are correctly reported based on the official decision or applicant's action

Example Prompts

- Prompt 1 – Write a Python script that verifies the Uniform Residential Loan Application (URLA) for the loan application date against the LAR reported application date
- Prompt 2 – Write a Python script verifying the loan amount reported on the loan closing document matches the loan amount in dollars on the LAR
- Prompt 3 – Write a Python script verifying the action taken on the loan matches that on the LAR

Example Python Script

```
import pandas as pd
def verify_loan_amount(closing_doc_csv, hmda_lar_csv, id_column, closing_amount_column, hmda_amount_column):
  """
  Verifies that the mortgage loan amount on the closing document matches the HMDA LAR reported amount (in dollars).
    Parameters:
    closing_doc_csv (str): Path to CSV file containing closing document data.
    hmda_lar_csv (str): Path to CSV file containing HMDA LAR data.
    id_column (str): Name of column containing the unique loan identifier (must be in both files).
    closing_amount_column (str): Column name for loan amount in closing doc (in dollars).
    hmda_amount_column (str): Column name for loan amount in HMDA LAR (reported in thousands).
     Returns:
    DataFrame: Merged DataFrame showing match/mismatch status.
  """
    # Load the datasets
  closing_df = pd.read_csv(closing_doc_csv)
  hmda_df = pd.read_csv(hmda_lar_csv)
    # Merge on loan ID
  merged_df = closing_df.merge(hmda_df, on=id_column, suffixes=('_closing', '_hmda'))
    # Convert HMDA loan amount from thousands to dollars
  merged_df['hmda_amount_dollars'] = merged_df[hmda_amount_column] * 1000
    # Check for match
  merged_df['amount_match'] = merged_df[closing_amount_column] == merged_df['hmda_amount_dollars']
    return merged_df[[id_column, closing_amount_column, 'hmda_amount_dollars', 'amount_match']]
# Example usage:
if __name__ == "__main__":
  results = verify_loan_amount(
    closing_doc_csv="closing_disclosure.csv",
    hmda_lar_csv="hmda_lar.csv",
    id_column="LoanID",
    closing_amount_column="LoanAmount",
    hmda_amount_column="LoanAmount_HMDA"
  )
    print(results)
  # Optionally save to file
  results.to_csv("loan_amount_verification_results.csv", index=False)
```

Use Case 5: Regulatory & Compliance

- What Could Possibly Go Wrong?

1/ Prompt engineering is not an exact science!

Result: Ineffective prompts can create errors; false negatives and positives

2/ LLM model hallucinations

Result: Data not available in the training set causes LLM to provide answers completely made up and unreliable

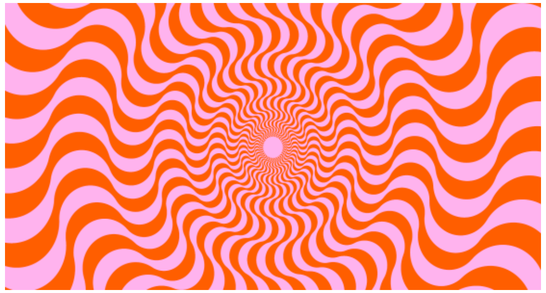

3/ Limited knowledge

Result: Data or changes to regs since the model was trained could be overlooked in the compliance check process

4/ Potential data privacy breaches

Given the personally identifiable information contained in loan applications and related mortgage documents, use in training an LLM could expose that data to misuse and unauthorized access

Use Case 6: Robo-Advising

- Pre-AI World

Professional Investment Advisor

Or,

Develop Client Financial Profile

Establish Risk Tolerance (Questionnaire)

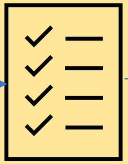

Determine Investments & Allocate

Monitor Portfolio & Update

Identify Client Goals

Tax Planning

Retirement Planning

Estate Planning

Philanthropy

Limitations
- Fees (Cost)
- Minimum investment requirements

Self-Directed Investment

Limitations
- Knowledge
- Time

Use Case 6: Robo-Advising

- Post-AI World

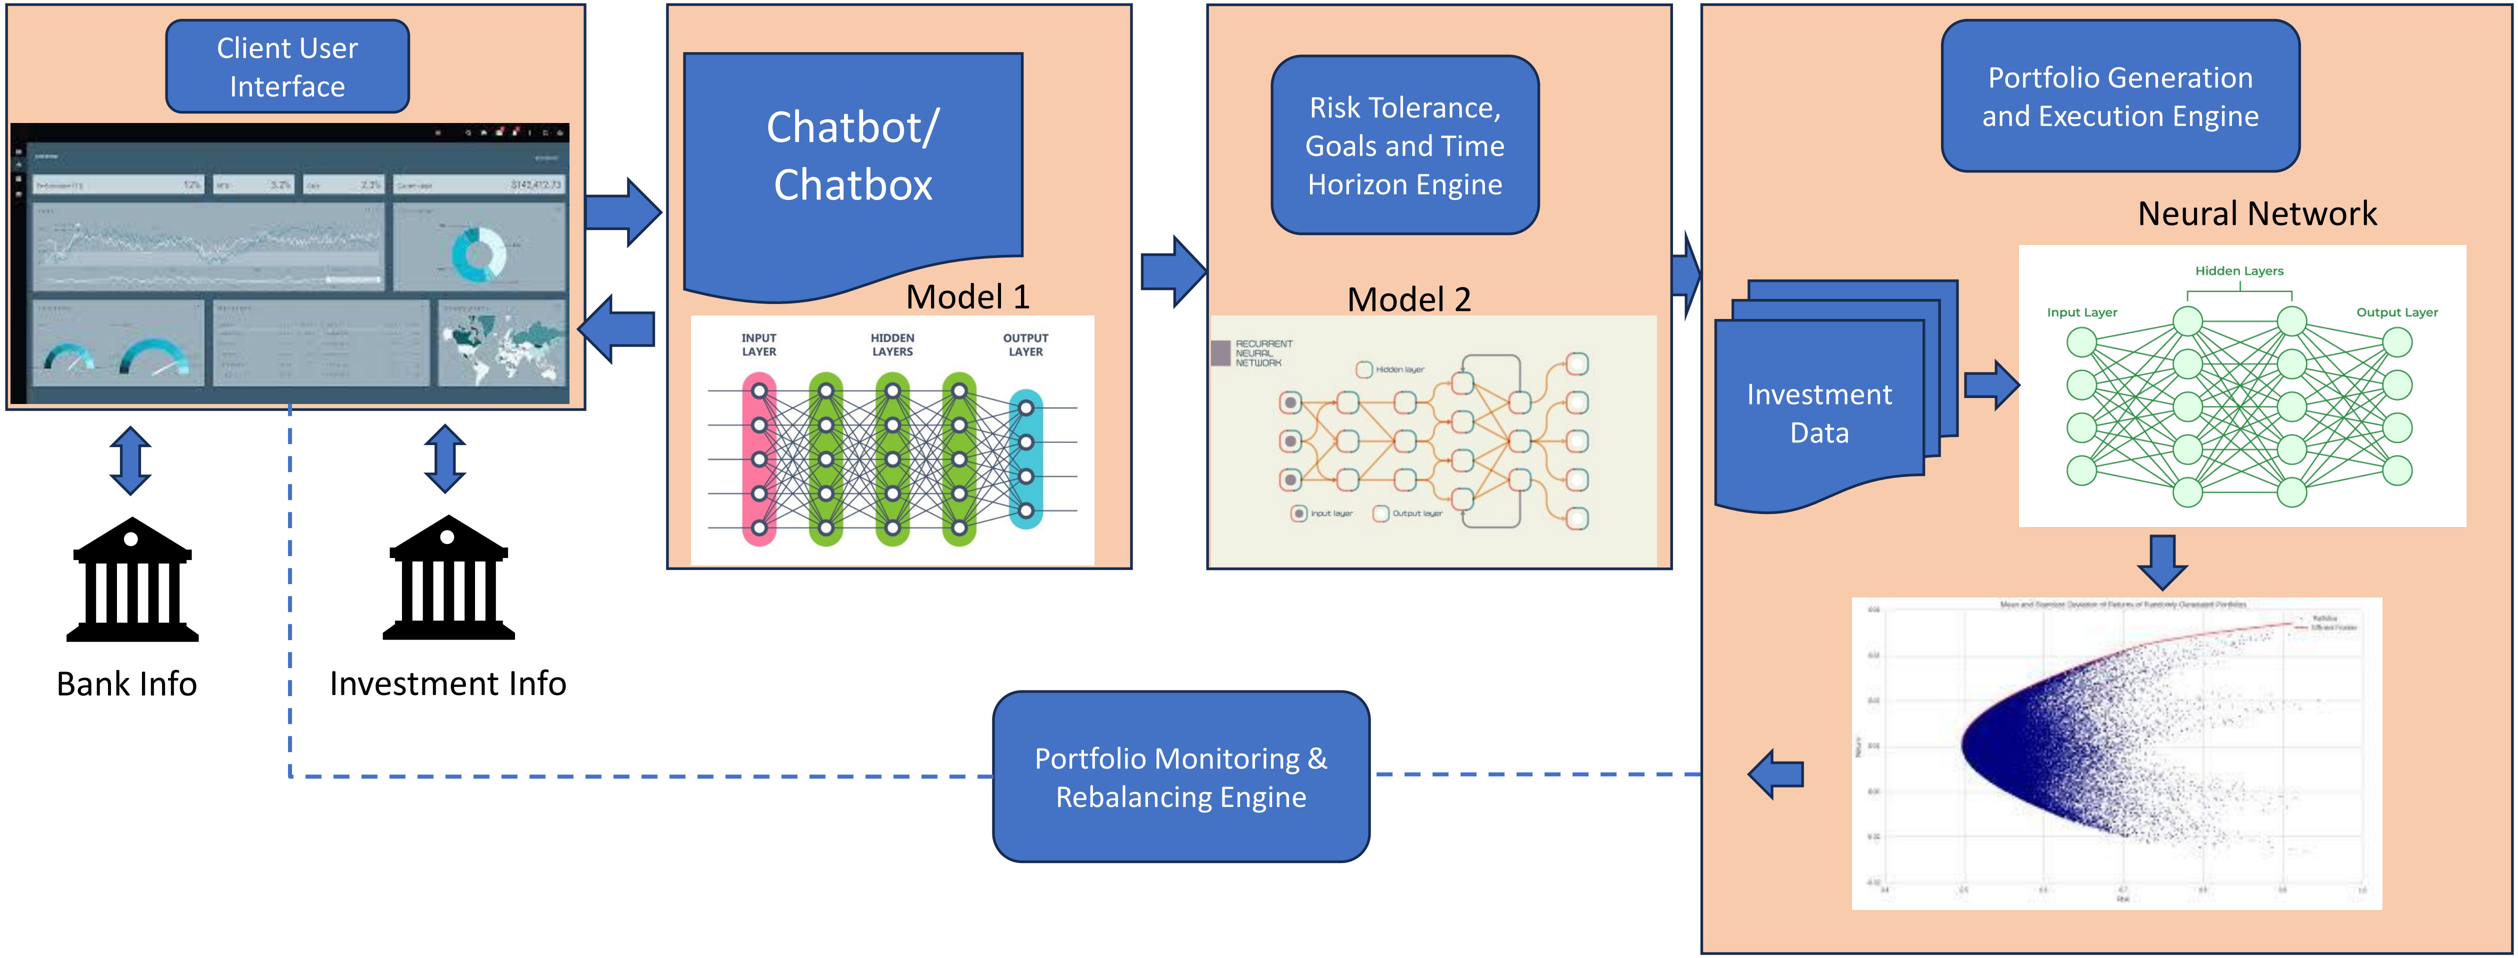

Use Case 6: Robo-Advising

- What Could Possibly Go Wrong?

1/ Algorithmic bias - Historical data used in determining portfolio allocation does not align with current and future market conditions

Result: Portfolio misallocation, underperformance

2/ Data privacy and security

Result: Potential misuse and breach of personal data and records

3/ Fiduciary obligations

Result: Potential litigation risk if not fulfilling fiduciary responsibilities or failing to meet standards of care for clients

4/ Nonhuman frailties

Result: Potential for repetitive interactions, nonresponses, lack of empathy, hallucinations

AI Guidance & Oversight

- What is Best Practice?

Industry-wide Resources – National Institute of Standards & Technology (NIST)

- NIST AI Risk Management Framework (AI RMF)
- NIST AI RMF Playbook

The 4 pillars of the AI RMF and described in detail in the AI RMF Playbook distill the oversight and management of AI risk into themes common to banking but applicable to less regulated entities such as fintech companies and independent mortgage banks, for example

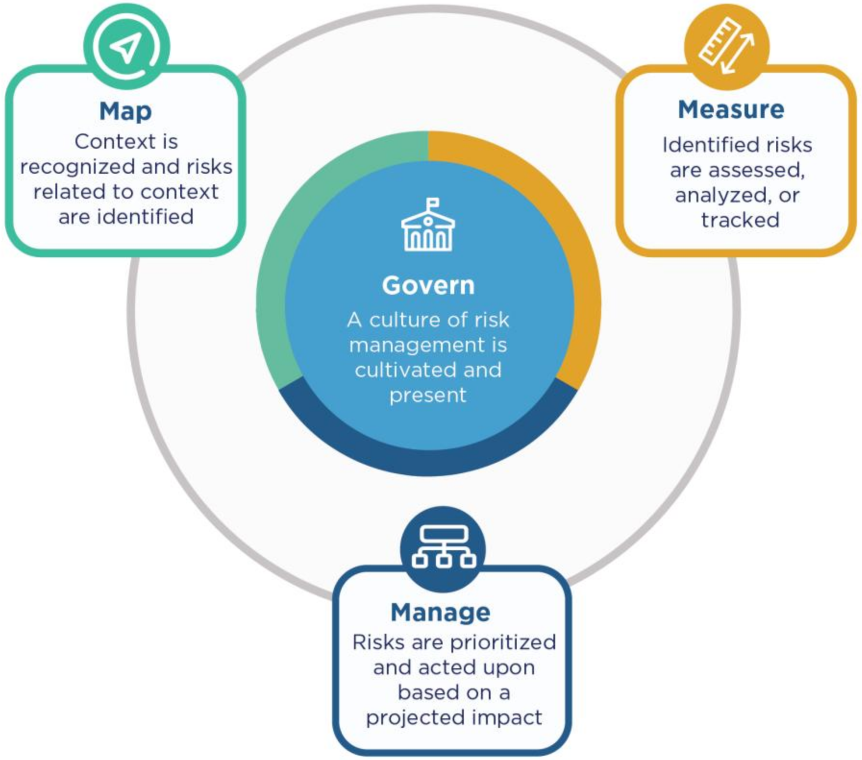

Source: NIST AI RMF

Summary & Key Takeaways

- AI models are proliferating rapidly across the financial services sector
- AI is ushering in the next industrial revolution that will have game changing impacts to financial services
  - Efficiency
  - Loss
  - Revenue and profit generation
- However, the diversity of AI techniques, level of sophistication, data and systems intensity and applications greatly amplifies inherent model and third-party risks across the enterprise
- Boards, management and staff must ultimately be self-aware of "shiny object bias" and become skeptical champions of AI in the organization
- Strong oversight and risk management practices are critical to avoiding a costly headline, legal or regulatory event down the road

AI Governance & Oversight

- Checklist
  - ✓ Catalog and define AI terms and techniques in your organization
  - ✓ Inventory applications across the enterprise
  - ✓ Establish the company's risk appetite for AI risk
  - ✓ Identify all inherent risks from each AI application, determine appropriate controls and assess residual risks
  - ✓ Create an Enterprise AI Policy or integrate into existing model risk management policy
  - ✓ Ensure Boards are educated on AI uses, applications and risks at the organization
  - ✓ Create a path of review, escalation and monitoring of AI risks and activities via management, senior management and board committee structures
  - ✓ Ensure AI tools and models are rigorously validated by independent experts, reviewed and approved by a model risk management committee prior to deployment
  - ✓ Develop a process for independent review and assessment of third-party developed AI tools
  - ✓ Ensure effective data governance and infrastructure are established to support AI data needs
  - ✓ Create an interdisciplinary AI Competitive Advantage Team (AICAT) comprised of seasoned business, risk, compliance, legal and IT professionals to provide ongoing direction and management over enterprise-wide AI applications

AI Governance & Oversight

- Bank-specific AI Risk Guidance
  - Governance and Risk Management
    - OCC Handbook on Corporate and Risk Governance
    - OCC Guidelines Establishing Heightened Standards for Certain Large Insured National Banks, Insured Federal Savings Associations, and Insured Federal Branches; Integration of Regulations, 79 Fed. Reg. 54,518 (Sept. 11, 2014)
    - US Treasury, Artificial Intelligence in Financial Services, (December 2024)
    - FHFA, ADVISORY BULLETIN AB 2022-02: Artificial Intelligence/Machine Learning Risk Management
  - Model Risk Management
    - Federal Reserve, SR 11-7, Supervisory Guidance on Model Risk Management (Apr. 4, 2011)
    - OCC, Comptroller's Handbook on Safety and Soundness: Model Risk Management Version 1.0 (Aug. 2021)
  - Third Party Risk Management
    - Federal Reserve, FDIC, OCC, Interagency Guidance on Third-Party Relationships: Risk Management (June 7, 2023)

About SERC

- Initiative of the Robert H. Smith School of Business at the University of Maryland
- Strive to cultivate best practices in Enterprise Risk Management across industries
- Sample of 2026 events & training for industry practitioners
  - AI & Risk Workshop
    - March 26-27, Washington, DC
  - Risk Academy: *Art of Risk-Informed Decisionmaking*
    - June 1-5, College Park, MD
  - CRO Risk Summit: *Risk Management in an Age of Change*
    - Sept. 17th, Washington, DC