



Artificial Intelligence in Compliance

Key Success Factors and Getting Started

Matt Scarborough, CEO of Bridgeforce Data Solutions

Presented on April 17, 2024, at America's Credit Unions
Regulatory and Compliance school

Session Description

This session will focus upon how credit unions can use AI tools to enhance the efficiency and effectiveness of compliance and risk management.

It will start with the critical compliance and risk success factors for beginning to use AI generally, then shift to specific examples and considerations for using AI to enhance compliance and risk management.

Learning Objectives

- Understand key compliance and risk success factors for using AI tools
- Understand actions to take now to ensure that your organization is prepared to embrace these evolving technologies, now or in the future
- Understand categories of use cases for using AI tools to enhance compliance and risk management
- Be prepared to not only react to requests from colleagues to use AI, but also to lead the way by using it for compliance

Background

Personal

- Math, tech, & business
- Banking
- Software
- Consulting

Bridgeforce Data Solutions

- Software and analytics on data quality related to credit bureau data and disputes
- Built products without AI, but now embracing & exploring

Framing the Discussion

Three useful comparisons:

- AI to the Internet
- ChatGPT to the early web browsers (e.g., Mosaic, Netscape)
- Generative AI to self-driving cars

Key lessons across my career that are relevant here:

- Basic principles are critically important during periods of highest uncertainty
- Timing matters: yesterday's hype can often become tomorrow's reality (and those changes are happening fast)
- Have the courage to ask questions that expose what you don't know (but you feel like you should)

AI Adoption: Not “If”, but “When & How”

In use by most of us right now

- Neural nets in fraud prevention
- Capabilities like text prediction on our phones and keyboards

Emerging trends

- Using AI in underwriting
- AI in AML/BSA efforts
- Generative AI for internal assistance

For any AI application, data is the key

- Large volume of high-quality input data (that has been appropriately tagged)
- Clear data-driven definition of the good vs. bad outcome

Regulatory Guidance on AI

The tools may be new, but the same old principles apply

Key potential sources of discrimination highlighted in recent Joint Statement:

- Data: unrepresentative or imbalanced
- Model Opacity: explainability
- Design & Use: flawed assumptions

Note on “discriminatory outcomes”

- Many regulatory statements call out the risks of discriminatory outcomes
- I have not seen anything, however, that changes the rules

Getting Started: AI for Predictive Models

(Relevant for fraud, underwriting, AML/BSA, and more)

Key Questions

How do we precisely define the desired outcome?

- Clear and precise definitions are critical for training the model
- Evaluating performance combines those definitions with data on our results

How good is our data?

- Do we have enough to train a model? To evaluate an existing one?
- How clean, consistent, and timely is our data?

How will we ensure appropriate governance?

- Ensure clarity on data usage, and that you are ok with it
- The core principles of model governance have not changed, but they must now be applied to address adaptive models

Special Focus: Explainability

Explainability always matters, but more in some instances than others:

- Proper model governance always matters
- Use cases requiring adverse-action notices require more effort

For use case requiring adverse-action reasons, they must be specific:

- The CFPB is clear. Specific denial reasons must be provided, even if a “Black Box” model is used

Special Focus: Data Quality

High quality data is foundational.
AI models are only as good as the data that feed them

Illogical Data

Dates that don't make sense

Amounts that don't make sense

Incomplete Data

Missing values

Missing history

Misclassified Data

Numbers that are actually attributes

Inconsistent Data

Field used for multiple purposes

Field that changes definitions over time

Insufficient Data

Too small a sample

Not enough history

Forward Looking Data

Data that would have been unknown to the model at the time of the prediction

Example Use Cases Relevant to Compliance

Use Case	Without AI	AI Focus Areas
Lending	<ul style="list-style-type: none">• Lending criteria are developed through manual analysis & policy decisions• Adapting to changes and adding new data sources is a lengthy process	<ul style="list-style-type: none">• Leveraging additional data sources• Developing more focused targeting• Faster adaptation
AML/BSA	<ul style="list-style-type: none">• Huge costs and lots of false alarms• Significant money laundering remains undetected	<ul style="list-style-type: none">• Reduce false positives• Enhance detection
Writing P&P's, job descriptions, letters, summaries	<ul style="list-style-type: none">• Start from scratch, prior version, or internet search• Complete draft and submit for approval• Human search & skimming	<ul style="list-style-type: none">• Rough drafts• Use Gen AI for customized response• AI-powered summaries
Credit Bureau: Furnishing & Response	<ul style="list-style-type: none">• Manual sampling & review• Best case: comprehensive review with rules, then humans sorting out next action	<ul style="list-style-type: none">• Extracting data from unstructured images• Identifying likely root causes of data issues• Predicting the right dispute responses, including any required corrections

First Example: Compliance Lens to AI in Underwriting

Scenario

- Your colleagues want to use AI to increase approval rates
- You want to enable your organization to learn while protecting from multiple risks

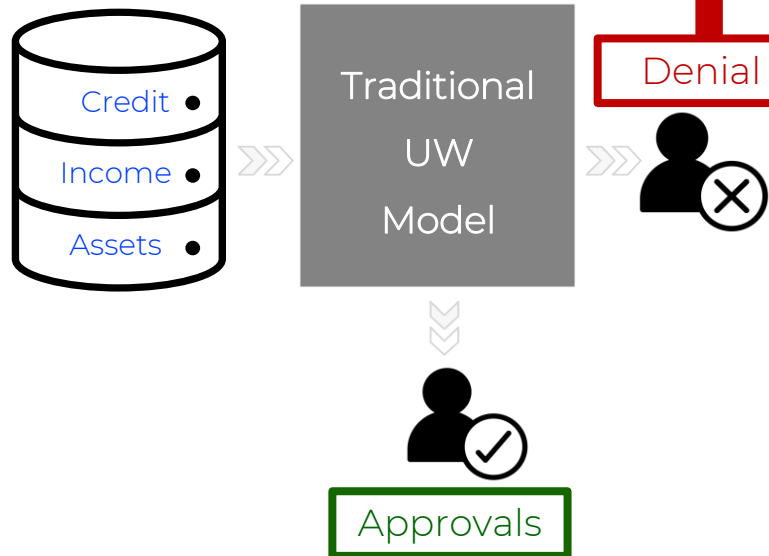
Key Concerns

- Explainability: Declining members (or potential members) without appropriate decline reasons
- Accuracy: Poor decision quality has many implications
- Monitoring: If something is happening that we do not want, how quickly can we catch it and stop it?

Potential Approach in Underwriting

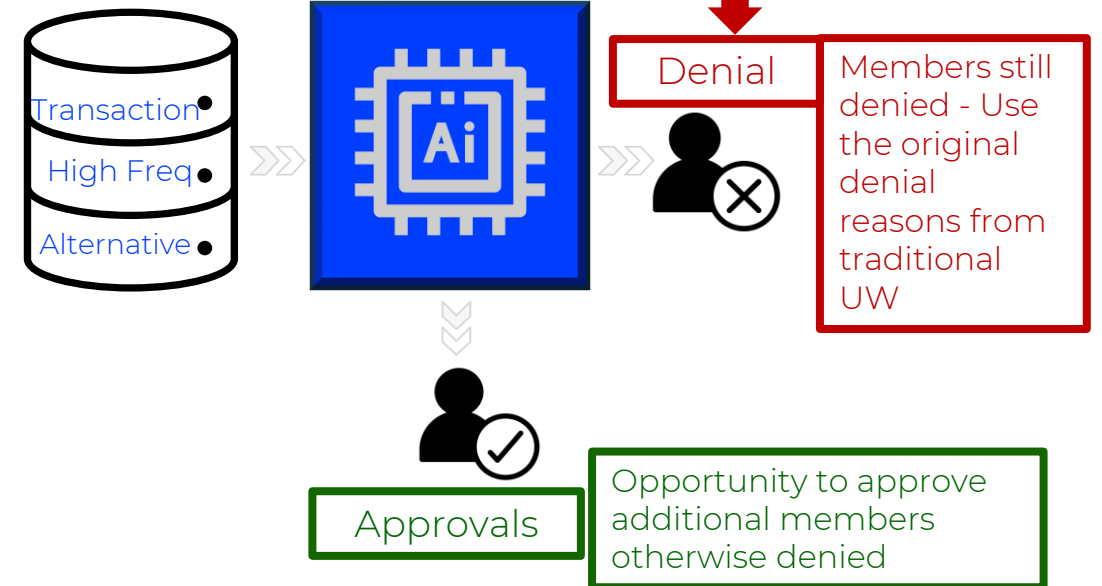
- Run traditional & AI models in parallel
- In first phase, only use the AI model to reverse denials

Use Current Models & Decline Reasons



Run AI Models in Parallel
Pick up supplemental approvals and gather data

Denial Reasons (ECOA)



First Example: How Does This Address Concerns?

- Explainability:
 - Decline reasons sent to members from first phase are fully from existing system
 - Decline reasons provided by new AI tool can be compared to traditional ones for acceptance, rejection, or refinement
- Accuracy:
 - By limiting to the financial risk of potential “bad” approvals, the risk can be easily quantified
 - Set volume limits based upon the test budget and risk appetite
 - Observe the performance of AI-declines that you continue to approve
- Monitoring:
 - Running in parallel with existing process facilitates robust and timely monitoring

Second Example: AML/BSA

- I am not an AML expert, but this strikes me as a compliance area that is ripe for change with AI
 - Some parallels to fraud-prevention efforts that have been using AI tools for decades
 - A current state with very high false-positive AND false-negative rates
 - Lots of attention from vendors and large financial institutions
- Key considerations here will echo considerations for AI in Lending (explainability, accuracy, and monitoring are still important), but with key differences that include:
 - Very different regulatory priorities
 - Different data sources

Third Example: Gen AI for writing

Scenario

- You and your colleagues want to use Gen AI to save time in the drafting of documents, including but not limited to letters, job descriptions, policies, and summaries of content.

Key Concerns

- **Data & Intellectual Property:** Would the prompts that you would like to write include any confidential information? If so, does the tool you are using protect that information?
- **Accuracy & Hallucinations:** Would there be checks in place to ensure that outputs are only used as a rough draft, not a final product?
- **Consistency:** For the use-case you have in mind, are you ok with the output being different for the same prompt asked at different times?
- **Appropriately targeted summaries:** Are the results in the format you want, with the right context?

ChatGPT example 1

Prompt:

“Write a paragraph on key considerations for the use of AI in consumer lending from a compliance and risk perspective”

Response:

“The integration of AI in consumer lending brings both promise and challenges, necessitating careful consideration from a compliance and risk standpoint. First and foremost, data privacy and security regulations must be rigorously adhered to, with a focus on safeguarding sensitive customer information. Fair lending practices are paramount, requiring AI algorithms to be free from biases that may discriminate against certain demographic groups, as well as transparent enough for regulatory scrutiny. Regular model validation and monitoring are essential to ensure the AI system's accuracy and fairness over time, as consumer lending regulations evolve. Additionally, robust documentation and audit trails should be maintained to demonstrate compliance efforts. Moreover, contingency plans must be in place to address potential algorithmic errors or systemic risks, protecting both borrowers and lenders from undue harm. Lastly, fostering open collaboration between regulatory bodies, industry stakeholders, and technology developers can facilitate the establishment of comprehensive guidelines that balance innovation and consumer protection in the rapidly evolving landscape of AI-driven consumer lending.”

ChatGPT example 2

(same prompt, different response)

Prompt:

“Write a paragraph on key considerations for the use of AI in consumer lending from a compliance and risk perspective”

Response:

“When integrating AI into consumer lending practices, compliance and risk management are paramount considerations. Firstly, ensuring adherence to regulatory frameworks like the Fair Credit Reporting Act (FCRA) and Equal Credit Opportunity Act (ECOA) is essential. AI algorithms must be transparent and auditable, guarding against discrimination and bias in lending decisions. Secondly, robust data privacy measures, aligned with regulations such as the General Data Protection Regulation (GDPR), must be implemented to safeguard consumer information. Regular monitoring and validation of AI models are vital to detect and mitigate potential biases or errors. Additionally, clear communication with consumers about how AI influences lending decisions fosters trust and transparency. Lastly, building a culture of compliance within the organization, coupled with continuous training for staff, reinforces ethical AI usage in consumer lending, promoting fairness and accountability.”

Fourth Example: Credit Bureau Data & Disputes

Scenario

- You want to leverage AI to help identify root causes of data issues (including causes outside of your control) and respond to disputes with greater accuracy and efficiency

Key Concerns

- **Data & Intellectual Property:** The data includes PII, and therefore must be protected in ways consistent with PCI DSS and SOC 2
- **Accuracy & Hallucinations:** How do you ensure that the data extraction and recommendations is more accurate and consistent than XX% of human agents? (what % seems “reasonable”?)
- **Monitoring & evidence:** How do you make sure that the process continued to adapt and improve in ways you can evidence to regulators and auditors?
- **Cost Efficiency:** How do you improve quality and speed while lowering costs?

Key Points about the Gen AI Landscape

- ChatGPT (and the OpenAI foundation models behind it) started the buzz and still gets the most consumer press, but there are dozens of powerful Large Language Models (LLMs) on the market
- For use with any confidential information, it is critical to pick a tool that enables you to ensure your data stays confidential (Amazon Bedrock, for example)
- That, in turn, enables you leverage the combination of your data and the capabilities of the LLM most effectively

Using GenAI effectively: Two Key Concepts

- Prompt Engineering involves optimizing the prompts going into a LLM to get the most useful outputs that you can. Examples of it include:
 - Specific instructions regarding desired content and format, like “write a paragraph...” vs. “Return a bulleted list of the names of the people mentioned in the above email, with a brief description of the sentiment expressed about each person”
 - Providing statements that give the model context
 - Using a series of questions, instructions, or statements
- Retrieval-Augmented Generation (RAG) takes this to another level by leveraging separate data sources (which may contain, for example, your proprietary information)
 - This can greatly reduce the risk of hallucinations and provide much more usable results
 - This requires more technical effort than manual prompt engineering, but much less than training a model

What Should We Do With Generative AI?

State a usage policy for generative AI by your colleagues

- Use only approved AI tools, and only within use-case guidelines
- Guidelines and approval process must carefully consider data & IP risk
- Closely monitor outputs for “hallucinations”

Build familiarity

- Generative AI is not ready for member-facing communications, but it is evolving quickly
- If you have never used it, start now to familiarize yourself
- Identify potential internal applications for limited pilots

Monitor for the tipping point for your organization

Next Steps

Ensure that your organization's conversations about the use of AI are framed around “how” and “when”, not “if”

- Identify potential use cases and estimate their potential impact
 - *Feasible now*
 - *Feasible soon (and what is missing)*
- Understand the availability and quality of existing data

Proactively address potential uses

- Have a policy on AI usage by your colleagues
- Ensure that vendor evaluation and oversight includes understanding how your data would be used (and that you are ok with it)

Stay in touch: mscarborough@bridgeforcedatasolutions.com

Appendix

Special Focus: Data Quality

High quality data is foundational.
AI models are only as good as the data that feed them.

Illogical Data

Dates that don't make sense
(ex. account closed date precedes account open date)

Amounts that don't make sense
(ex. A past due amount reflected on an account that has never been delinquent)

Incomplete Data

Missing values
(ex. No original loan amount available)

Missing history
(ex. Only having 12 months of history on an account that was booked 36 months ago)

Misclassified Data

Numbers that are actually attributes
(ex. Zip codes, Routing numbers)

Inconsistent Data

Field used for multiple purposes
(ex. Customer income where for some it is annualized and others it is monthly)

Field that changes definitions over time
(ex. An account type field where old values

Insufficient Data

Too small a sample
(ex. Less than 100 records are available to train a model)

Not enough history
(ex. Only having 3 months of history, so the model cannot account for seasonality)

Forward Looking Data

Data that would have been unknown to the model at the time of the prediction
(ex. Credit score from June 2023 when the model is predicting an event that happened in May 2023)

Actions To Take Now: Preparing



USE CASES

Assess potential applications of AI.

Assign actions to these use cases – what do you want to do?

Prioritize Select those use cases likely to add the most value to start.



DATA

Inventory – know what data you have, the sources of the data, the volume and history of data, and frequency of update

Evaluate Quality of the available data. Identify areas where data may be incomplete or inaccurate

Clean - Select those use cases likely to add the most value to start.



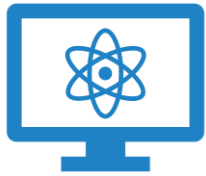
COLLABORATION

Compliance – Prepare to communicate they keep compliance needs for AI use cases

Operations – know the areas where AI can add the most value, and the specific actions that can be taken

Data Governance – know how data can be used, and any limitations to be aware of

Actions To Take Now: Embracing



TOOLS

Evaluate potential suppliers of AI engines. Consider their ability to support your efforts to document methodologies and outcomes

Decide on the tool that best fits your needs

Implement the tool, ensuring connectivity to the relevant data and ability to trigger actions



PEOPLE

Data Scientists to properly construct a model, with expert consideration for model factors and potential biases

SMEs to understand the real-world consequences of AI model predictions

IT – To execute the actions resulting from AI models



MONITORING

Compliance – ensure that the model is being continually monitored for regulatory compliance issues as it learns and evolves over time

Risk – ensure that models are performing their intended functions within the guardrails set

Data Governance – ensure that changes to model inputs or data are well understood

A Few Useful Links

- [CFPB Issues Guidance on Credit Denials by Lenders Using Artificial Intelligence | Consumer Financial Protection Bureau \(consumerfinance.gov\)](#)
- [Generative AI exists because of the transformer](#)
- [Transformer: A Novel Neural Network Architecture for Language Understanding – Google Research Blog](#)
- [Google Cloud Unleashes AI On Money Laundering Activities After Successful HSBC Trial \(forbes.com\)](#)
- [The Flaw That Could Ruin Generative AI - The Atlantic](#)
- OCC Guidance on Model Governance
 - [OCC 2011-12: Sound Practices for Model Risk Management: Supervisory Guidance on Model Risk Management](#) (reg guidance: 21 pages)
 - [Comptroller's Handbook: Model Risk Management | OCC](#) (Examiner's handbook: 109 pages)