WEBINAR

Governance of AI & ML in Financial Services

January 17, 2024 | 11:00 AM

Aaron Johnson, CFA, FRM, FDP
Senior Examiner, Federal Housing Finance Agency

Kathryn Wilkens, PhD, CAIA
Founder of Pearl Quest LLC,
FDP Curriculum Advisor

Register now

FDP INSTITUTE by CAIA

Governance of AI & ML in Financial Services

Welcome
We will start promptly at 11 AM ET

Please add your questions to the chat box, we will address them during the Q&A.
If you are unable to hear speakers just let us know in the chat box.

Following the webinar, you will receive a link to the recording.
You can find a copy of all our webinar recordings at

www.caia.org/caia-infoseries
www.fdpinstitute.org/webinars

1

Financial Data Professional Institute

FDP Institute provides world-class training and education to financial professionals to meet the accelerating needs of digital transformation in the industry.
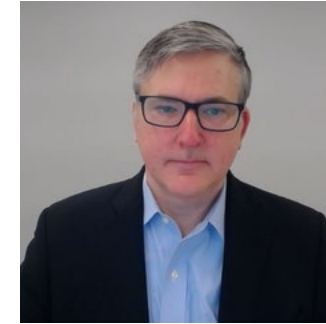
www.FDPInstitute.org

Introductions

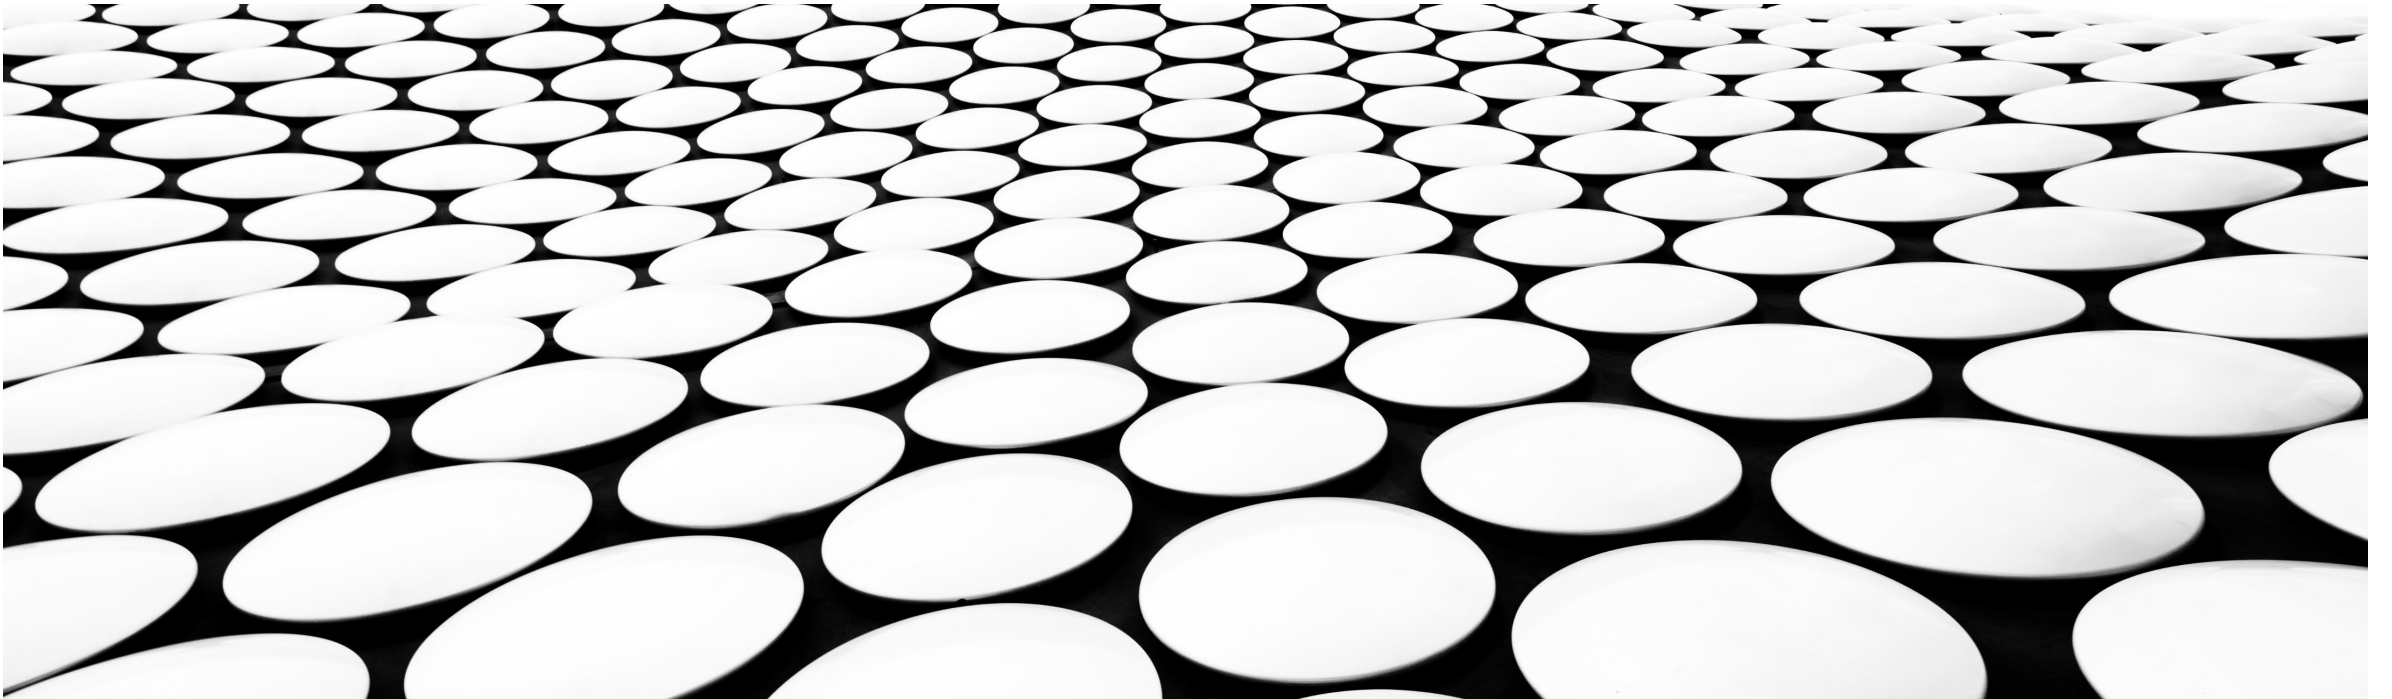Kathryn Wilkens, Ph.D., CAIA
Founder,
Pearl Quest LLC

Aaron Johnson, CFA, FRM, FDP
Senior Examiner,
Federal Housing Finance
Agency

Today's Topic:
Governance of Artificial Intelligence & Machine Learning In Financial Services

GOVERNANCE OF ARTIFICIAL INTELLIGENCE IN FINANCIAL SERVICES

AARON JOHNSON ⋮⋮⋮ JANUARY 17, 2023

CONTENTS

- Some preliminary definitions

- Risks from AI

- Trustworthy AI

- Some authoritative statements on AI governance

- Practical considerations in AI governance

- Parting thoughts

- Appendix

DISCLAIMER

- The views expressed in this presentation should not be construed to be those of my current employer, the Federal Housing Finance Agency, nor of any prior employer. Nothing in this presentation should be construed to be legal or financial advice.
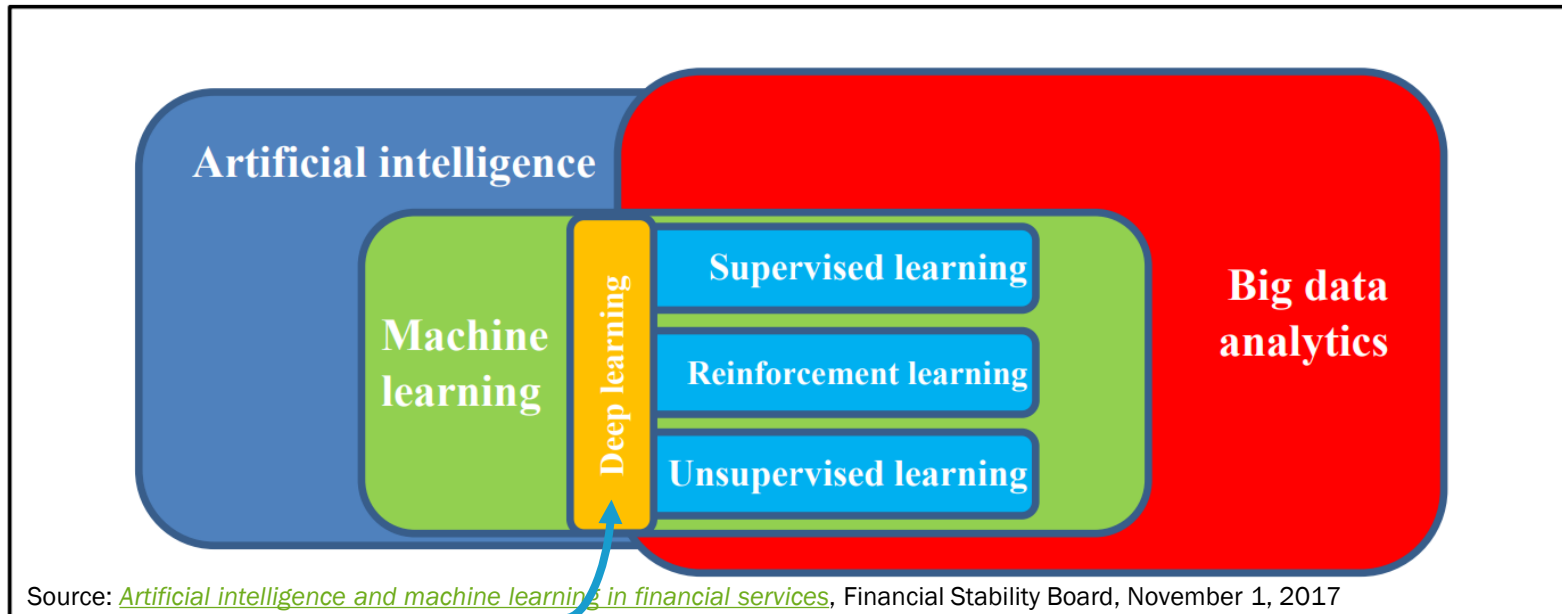
SOME PRELIMINARY DEFINITIONS

MODELS, ARTIFICIAL INTELLIGENCE, MACHINE LEARNING, ETC.

A SIMPLE CATEGORIZATION OF AI

Source: *Artificial intelligence and machine learning in financial services*, Financial Stability Board, November 1, 2017

(Diagram showing nested categories: Artificial intelligence containing Machine learning, Deep learning, Supervised learning, Reinforcement learning, Unsupervised learning, and Big data analytics)

- **Artificial Intelligence:** "the theory and development of computer systems able to perform tasks that traditionally have required human intelligence" (FSB, 2017).

- **Machine Learning:** "a method of designing a sequence of actions to solve a problem, known as algorithms, which optimise automatically through experience and with limited or no human intervention" (FSB, 2017).

- **Deep Learning:** "a form of machine learning that uses algorithms that work in 'layers' inspired by the structure and function of the brain" (FSB, 2017).

- **Large Language Models:** "a machine learning model that aims to predict and generate plausible language" (Google, 2023).
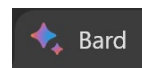
AI/ML

THE FINANCIAL INDUSTRY AND REGULATORS HAVE CONVERGED ON THE DEFINITION OF A MODEL. WHAT ABOUT AI/ML?

- Regulators have largely converged on the definition of model from first promulgated by the Fed and OCC in 2011 as SR 11-7 and Bulletin 2011-12, respectively. The FDIC adopted identical guidance in 2017 (FIL-22-2017). (See FHFA's AB 2013-07, Bank of England PRA SS1/23, Canda's Office of Superintendent of Financial Institutions draft E-23)

- Model (SR 11-7): "a quantitative method, system, or approach that applies statistical, economic, financial, or mathematical theories, techniques, and assumptions to process input data into quantitative estimates."

- Components of a model: Information input + processing + reporting

- SR 11-7 covers qualitative and quantitative approaches, and often covers AI/ML (see p. 7 of the Model Risk Management chapter in the OCC's *Comptroller's Handbook*).

- A good first guess is that all AI/ML techniques are models for the purposes of governance. But because calling something a model means it must be validated, people sometimes try to have their model not called a model.

- Whether called a model, tool, service, or something else, use of AI/ML must be governed.

RISKS FROM AI

CAPABILITIES AND PREVALENCE OF AI CREATES RISKS FOR FINANCIAL INSTITUTIONS

RISKS FROM AI NOW APPROACHING RISKS OF EMPLOYING PEOPLE

	Non-AI/ML Model	AI/ML	Generative AI	People
Data	• Data typically not large	• 3 Vs: velocity, volume, variety increased governance required	• Significant data needed (e.g., the Pile's 825 GB) for training • How to govern?	• Can exfiltrate, poison, and delete data
Opacity, explainability	• Often easy to understand (e.g., regression weights, decision trees)	• Often opaque (e.g., random forest, neural networks) • Explainability techniques (e.g., SHAP, LIME) often required	• Significant efforts to finetune are required (e.g., reinforcement learning from human feedback)	• Preference falsification, lying, forgetting, lack of understanding of own or business's true motivations
Discrimination, Bias, and Regulatory Compliance	• Often easy to assess discrimination risks	• Increased opacity increases fair lending and other risk of harm, such as privacy, employment discrimination	• Even more opaque • Large text and image datasets of unknown or uneven quality can perpetuate bias • May be cost-prohibitive to clean data or train model to limit bias	• Subtle discrimination, bias, unconscious bias, group think, forgetfulness, anchoring, availability bias, base rate neglect, confirmation bias, loss-aversion bias
Privacy and Security	• Risks from using/storing sensitive data • Relatively low demands on IT infrastructure • Even without AI movement to cloud may make sense	• Automated and intelligent detection of vulnerabilities are balanced by increased capacity of threats • Increased complexity and expense of IT	• Generative AI can be tricked into revealing private information • Deep fakes enhance phishing attacks	• Leaking to newspapers or activist organizations, shoulder surfing, document security at home
Performance	• Some risk of overfitting • p-hacking • Replication crisis	• Increased risk of overfitting; bias/variance trade-off	• Unknown risk of overfitting	• 80/20 rule in organizations • Excellent management is hard

RISK MATRIX FOR AI/ML USE CASES

Lower overall risk

- Google search
- Autocomplete
- Voice-to-text transcription
- Sentiment analysis
- OCR for remote deposit
- Biometric authentication (passkeys)
- Microsoft 365 Copilot
- Virus detection

Medium overall risk

- Automated home valuation
- Resume scanning
- Internal chatbots
- Commercial loan underwriting
- Financial crimes detection
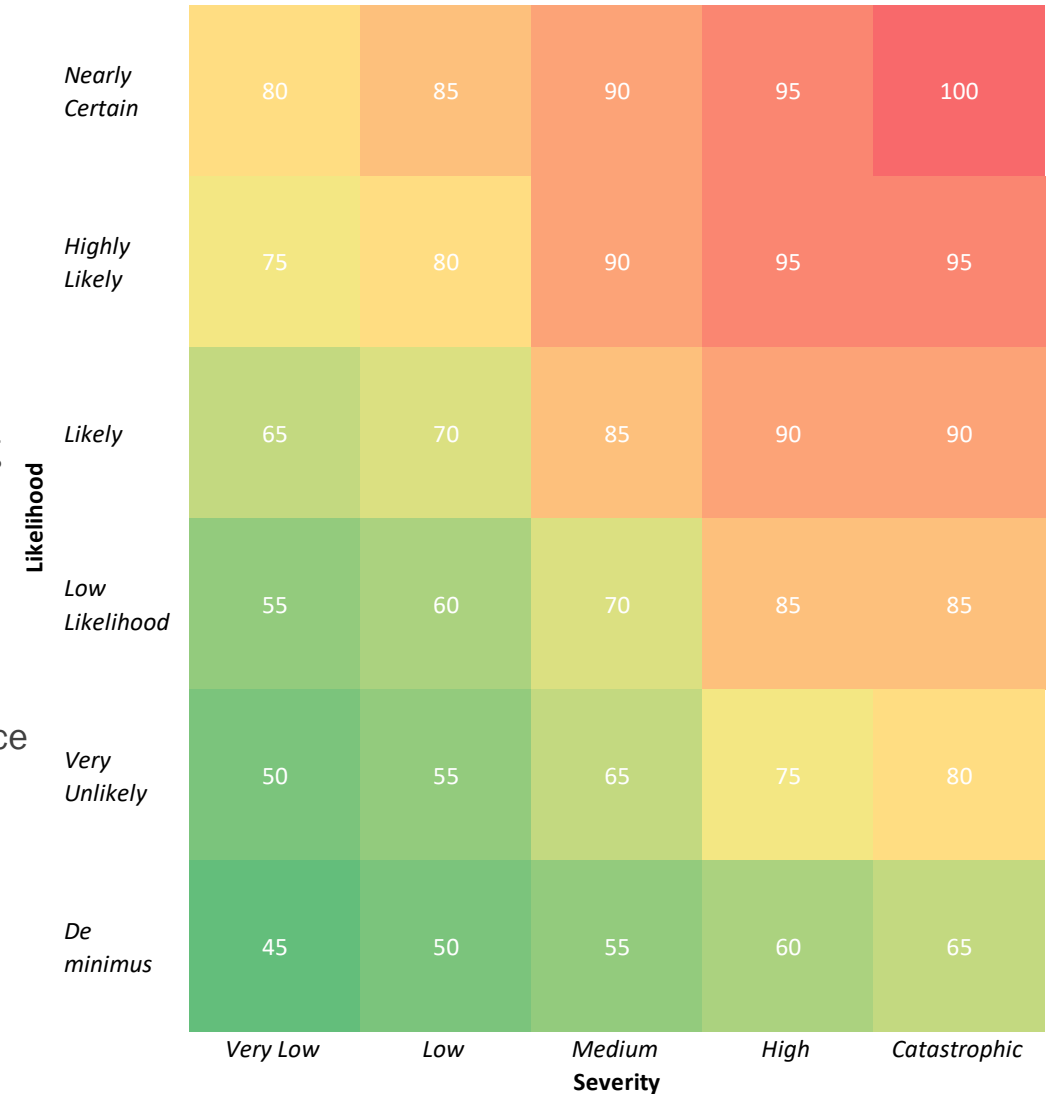- Automated pair coder (e.g., GitHub Copilot)

High overall risk

- Consumer-facing chatbots
- Robo-advisor
- Consumer credit underwriting (e.g., credit cards, mortgages) with ML

Unacceptable overall risk

- Consumer credit underwriting using GenAI

To be determined but technology exists

- Automatic summarization of laws, regulation, and guidance for rapid assessment of regulatory compliance
- Automatic model documentation
- Automatic model validation reports
- LLM-driven consumer marketing
- Internal document summarization
- Writing suspicious activity reports

Likelihood	Very Low	Low	Medium	High	Catastrophic
Nearly Certain	80	85	90	95	100
Highly Likely	75	80	90	95	95
Likely	65	70	85	90	90
Low Likelihood	55	60	70	85	85
Very Unlikely	50	55	65	75	80
De minimus	45	50	55	60	65

Severity

TRUSTWORTHY AI

SOCIETAL AND INDUSTRY CONVERGENCE ON THE GOAL OF AI GOVERNANCE

PRINCIPLED ARTIFICIAL INTELLIGENCE

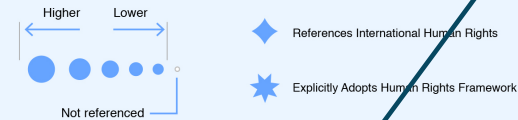A Map of Ethical and Rights-Based Approaches to Principles for AI

Authors: Jessica Fjeld, Nele Achten, Hannah Hilligoss, Adam Nagy, Madhulika Srikumar

Designers: Arushi Singh (arushisingh.net) and Melissa Axelrod (melissaaxelrod.com)

HOW TO READ:

Date, Location
Document Title
Actor

COVERAGE OF THEMES:

Higher — Lower

Not referenced

◆ References International Human Rights

✦ Explicitly Adopts Human Rights Framework

The size of each dot represents the percentage of principles in that theme contained in the document. Since the number of principles per theme varies, it's informative to compare dot sizes within a theme but not between themes.

The principles within each theme are:

Privacy:
Privacy
Control over Use of Data
Consent
Privacy by Design
Recommendation for Data Protection Laws
Ability to Restrict Processing
Right to Rectification
Right to Erasure

Accountability:
Accountability
Recommendation for New Regulations
Impact Assessment
Evaluation and Auditing Requirement
Verifiability and Replicability
Liability and Legal Responsibility
Ability to Appeal
Environmental Responsibility
Creation of a Monitoring Body
Remedy for Automated Decision

Safety and Security:
Security
Safety and Reliability
Predictability
Security by Design

Transparency and Explainability:
Explainability
Transparency
Open Source Data and Algorithms
Notification when Interacting with an AI
Notification when AI Makes a Decision about an Individual
Regular Reporting Requirement
Right to Information
Open Procurement (for Government)

Fairness and Non-discrimination:
Non-discrimination and the Prevention of Bias
Fairness
Inclusiveness in Design
Inclusiveness in Impact
Representative and High Quality Data
Equality

Human Control of Technology:
Human Control of Technology
Human Review of Automated Decision
Ability to Opt out of Automated Decision

Professional Responsibility:
Multistakeholder Collaboration
Responsible Design
Consideration of Long Term Effects
Accuracy
Scientific Integrity

Promotion of Human Values:
Leveraged to Benefit Society
Human Values and Human Flourishing
Access to Technology

Further information on findings and methodology is available in Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-Based Approaches (Berkman Klein, 2020) available at **cyber.harvard.edu**.

BERKMAN KLEIN CENTER
FOR INTERNET & SOCIETY
AT HARVARD UNIVERSITY

The principles within each theme are:

Privacy:
Privacy
Control over Use of Data
Consent
Privacy by Design
Recommendation for Data Protection Laws
Ability to Restrict Processing
Right to Rectification
Right to Erasure

Accountability:
Accountability
Recommendation for New Regulations
Impact Assessment
Evaluation and Auditing Requirement
Verifiability and Replicability
Liability and Legal Responsibility
Ability to Appeal
Environmental Responsibility
Creation of a Monitoring Body
Remedy for Automated Decision

Safety and Security:
Security
Safety and Reliability
Predictability
Security by Design

Transparency and Explainability:
Explainability
Transparency
Open Source Data and Algorithms
Notification when Interacting with an AI
Notification when AI Makes a Decision about an Individual
Regular Reporting Requirement
Right to Information
Open Procurement (for Government)

Fairness and Non-discrimination:
Non-discrimination and the Prevention of Bias
Fairness
Inclusiveness in Design
Inclusiveness in Impact
Representative and High Quality Data
Equality

Human Control of Technology:
Human Control of Technology
Human Review of Automated Decision
Ability to Opt out of Automated Decision

Professional Responsibility:
Multistakeholder Collaboration
Responsible Design
Consideration of Long Term Effects
Accuracy
Scientific Integrity

Promotion of Human Values:
Leveraged to Benefit Society
Human Values and Human Flourishing
Access to Technology

Source: Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-Based Approaches to Principles for AI, Jan 15, 2020

NIST'S AI RMF LIKELY TO BECOME STANDARD ACROSS INDUSTRIES

- The National Institute of Standards and Technology (NIST) released its AI Risk Management Framework (RMF) in 2023. Other work at NIST includes the US AI Safety Institute, managing bias, explainability, and security.

- Leading technology companies have similar principles, typically under the term "responsible AI": Microsoft, Google, Meta, OpenAI, IBM, and Amazon.

Source: National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, p. 12

15

SOME AUTHORITATIVE STATEMENTS ON AI GOVERNANCE

LAWS, REGULATIONS, GUIDANCE, STATEMENTS, ETC.

AN INCOMPLETE AND OPINIONATED AI GOVERNANCE TIMELINE

Pre-2020 (Models and "normal" AI/ML)
- Apr 2011   SR 11-7/ Bulletin 2011-12 Guidance on MRM
- Jun 2017   FDIC adopts FRB, OCC MRM Guidance
- Nov 2017   FSB AI/ML in Financial Services
- Feb 2019   Executive Order 13859 on American Leadership in AI
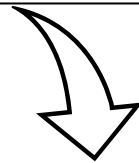
2020 – 2021 (Increasing focus on risks from AI/ML)
- Feb 2020   European Commission Whitepaper on AI
- Dec 2020   Executive Order 13960 on Trustworthy AI
- Jun 2020   Artificial Intelligence in the Securities Industry (FINRA)
- Mar 2021   AI/ML RFI (FRB, OCC, FDIC, NCUA, CFPB)
- Apr 2021   European Commission proposes AI Act
- Jun 2021   GAO Accountability Framework
- Aug 2021   Model Risk Management, Controller's Handbook (OCC)

2022 – 2023 (Cambrian explosion after ChatGPT)
- Feb 2022   FHFA AB 2022-02 and OMWI SL on AI/ML
- Mar 2022   Identifying and Managing Bias in AI (NIST)
- Oct 2022   White House Blueprint for an AI Bill of Rights
- Jan 2023   NIST Risk Management Framework
- Apr 2023   Measures for Management Generative AI (China)
- Nov 2023   Bletchley Declaration
- May 2023   SS1/23 – MRM Principles for Banks (BoE)
- Jun 2023   European Parliament Approves Position on AI Act
- Jun 2023   Chatbots in Consumer Finance (CFPB)
- Jul 2023   Predictive Analytics Proposed Rule (SEC)
- Sep 2023   Guidance on Credit Denials using AI (CFPB)
- Oct 2023   Draft TC260 on generative AI (China)
- Oct 2023   Executive Order 14110 on Trustworthy AI
- Dec 2023   Questionnaire on AI/ML, Quantum Computing (OSFI)
- Dec 2023   Provisional Agreement on EU's AI Act

FHFA'S 2022 GUIDANCE ILLUSTRATES CURRENT REGULATORY THINKING

- The FHFA's guidance, Advisory Bulletin 2022-02, contains four parts. This should be read in conjunction with a Supervisory Letter from the Office of Minority and Women Inclusion.

- This AB is the sole guidance on AI/ML from a US Federal financial regulatory agency, although regulators have made statements, issued requests for information, and updated handbooks.

I. Governance
- AI/ML Core Ethical Principles
  - Transparency
  - Accountability
  - Fairness and equity
  - Diversity and inclusion
  - Reliability
  - Privacy and security
- AI/ML Definitions and Taxonomy
- AI/ML Inventory
- Roles and Responsibilities
- Policies, Standards, and Procedures

II. Risk Identification and Assessment
- Model risks
- Data risks
- Other operational risks
- Regulatory and compliance risks

III. Control Framework
- Model controls
- Data controls
- Other operational controls
- Regulatory and compliance controls

IV. Risk Monitoring, Reporting, and Communication

PRACTICAL CONSIDERATIONS IN AI GOVERNANCE

WHAT STEPS CAN ORGANIZATIONS TAKE NOW AND IN THE NEAR TERM?

DEVELOP AI STRATEGY AND ETHICAL PRINCIPLES

- **Develop an AI Strategy** to ensure alignment across the organization in terms of how AI will be acquired, used, and governed. "The essence of strategy is choosing what not to do." Michael Porter.

  - Risk tolerance/appetite and risk taxonomy

  - Buy or build? Lead or follow? Which products?

  - Ranking use cases and risks

  - Reference to complementary frameworks (e.g., model risk)

  - Clarify roles, responsibilities, and authorities

- **Develop AI Ethical Principles**

  - Are existing company statements adequate, perhaps with some tweaks?

  - Don't just copy/paste. Customize based on strategy, risks, and benefits.

  - Good starting point are transparency, trustworthiness, safety, human in the loop, and non-discrimination

ENSURING EFFECTIVE CHALLENGE OF AI/ML

- The guiding principle is ensuring effective challenge for AI/ML, including GenAI. This may require:

  - Updating the RCSA and doing a risk assessment.

  - Updating the model inventory and coordinating with or rationalizing this and other inventories, like IT. How to handle ubiquitous AI/generative AI?

  - Reconsidering the definition of model. Include all AI/ML?

  - Revising roles and responsibilities (e.g., committees, working groups, business unit leadership, model owners, developers).

  - Communicating updates and conducting role-based training.

  - Updating other frameworks and controls across the organization (e.g., credit risk, vendor, software development lifecycle, data, vendor).

  - Revising management and board reporting and risk limits to address AI and GenAI risks specifically.

ENSURING EFFECTIVE CHALLENGE OF AI/ML, ADDITIONAL STEPS

- Additional steps to ensure effective challenge may include:

  - Considering incremental risks from AI across the model lifecycle.

  - Easing compliance burden and increasing consistency of governance with enablers (e.g., templates, testing scripts).

  - Carefully considering governance for anything deemed not a model under company policy but containing AI/ML, including GenAI.

  - Revising model risk tiering with AI in mind and scaling governance to risk.

  - Managing vendor relationships with a view to validation, ethical principles, etc.

ILLUSTRATIVE PLAN FOR ADOPTION OF GENERATIVE AI

Initiate
- Set project plan and goals
- Get Board and senior management buy-in
- Identify executives to own outcomes
- Sketch roles and authorities
- Determine whether consultants are needed

Communicate
- Set up working groups and committees
- Establish communication plan for Board updates, working groups, committees, and employees

Engage Resources
- Engage and consultants as needed
- Develop hiring plan for strategic areas in AI/ML
- Engage technology vendors to understand capabilities

Revise Governance
- Develop AI strategy
- Identify top use cases
- Identify unacceptable risks, use cases, and technologies
- Update policies, procedures
- Write AI ethical principles

Run Pilot Program
- Start with small set of most promising use cases
- Configure sandbox & identify staff
- Run tests
- Socialize results

Assess Pilot & Iterate
- Track lessons learned
- Optimize pilot process (people, process, technology) for higher throughput
- Move to next set of use cases

Approve and Roll Out
- Finalize review and testing from pilot (engage MRM & others per governance)
- Put into production

PARTING THOUGHTS

PRINCIPLES AND SPECULATIONS

HIGH-LEVEL PRINCIPLES FOR NAVIGATING AI

- If it was illegal before AI, it is illegal with AI.*

    - A corollary: if it was required before AI, it is required with AI.

    - A good example is CFPB's 2023 statement that adverse action notices in consumer credit are required regardless of technology used to make credit decisions.

- While general principles from NIST AI RMF and other organizations are helpful, financial services is already highly regulated. Separate frameworks for AI may not be needed provided existing frameworks are updated intelligently.

- It is impossible to plan for every eventuality. With AGI potentially happening within our lifetimes, people and organizations need to be resilient – meaning sound procedures, good communication, good infrastructure, etc.

- Beware the precautionary principle, panic-driven policies, and box-checking bureaucratization.

- Consider trade-offs in every decision. There are risks from using AI but also risks from not.

- Governance of models and AI is likely to change very rapidly in the near term. (See next slide.)

* I am not a lawyer. This is not legal advice. Seek advice of competent counsel.

SPECULATION ABOUT THE NEAR FUTURE

- FIs are burdened with regulatory compliance and expensive employees. With AI now or soon able to perform as well as the best available human in many tasks, expect FIs to use AI to build, document, implement, and validate models.

  - Pair coding (Amazon CodeWhisperer, GitHub Copilot, Replit) + LLM trained on banking regulations, guidance, model documentation, and model validation reports + AutoML tools (Amazon SageMaker, DataBricks, Dataiku) → BankerGPT that builds and validates models?

- What are we humans to do when AI is at or near human-level performance?

  - Find a government job, preferably one with a union (e.g., IRS, OCC, FHFA, FDIC).

  - Become a super-star model developer or validator who understands the business, can code, knows the math, and can communicate clearly (especially writing) so that your output is the input for training AI.

  - Search the "jagged technological frontier" to identify ways that AI fails (hallucinates convincingly). Become the human in the loop, i.e., leverage your domain expertise to make sure the AI is operating as designed.

APPENDIX

GLOSSARY

- AB   Advisory Bulletin (from FHFA)
- AGI   Artificial general intelligence
- AI/ML   Artificial intelligence/machine learning
- AVM   Automated valuation model (in residential real estate)
- BoE   Bank of England
- CFPB   Consumer Financial Protection Bureau
- ECOA   Equal Credit Opportunity Act
- EU   European Union
- FDIC   Federal Deposit Insurance Corporation
- Fed   The Board of Governors of the Federal Reserve System
- FHFA   Federal Housing Finance Agency
- FI   Financial institution, e.g., bank, broker-dealer, insurance company
- FIL   Financial Institution Letter (from the FDIC)
- FINRA   Financial Industry Regulatory Authority
- FRB   See Fed
- FSB   Financial Stability Board
- GAO   Government Accountability Office

- GenAI   Generative artificial intelligence, e.g., LLMs like ChatGPT
- GPT   Generative pre-trained transformer
- IT   Information technology
- LLM   Large language model
- MRM   Model risk management
- NCUA   National Credit Union Administration
- NIST   National Institute of Standards and Technology
- OCC   Office of the Comptroller of the Currency
- OMWI   Office of Minority and Women Inclusion (at FHFA)
- OSFI   Office of the Superintendent of Financial Institutions
- PRA   Prudential Regulatory Authority
- RCSA   Risk and control self-assessment
- RFI   Request for information
- RMF   Risk Management Framework (from NIST)
- SEC   Securities and Exchange Commission
- SL   Supervisory Letter (from FHFA)
- SR   Supervision and Regulation Letter (from the Fed)

FURTHER READING – WHERE TO START

Model Risk

- Guidance on Model Risk Management (FRB, OCC, FDIC)

- Model Risk Management Guidance (FHFA)

- S1/23 – MRM Principles for Banks (BoE)

AI/ML Risk Management

- AI/ML Risk Management (FHFA)

- FSB AI/ML in Financial Services

- Request for Comment on AVMs (Multiple US Agencies)

- Circular on ECOA Requirements with Black-Box AI (CFPB)

- Identifying and Managing Bias in AI (NIST)

- Risk Management Framework (NIST)

- Chatbots in Consumer Finance (CFPB)

- Predictive Analytics Proposed Rule (SEC)

- Guidance on Credit Denials using AI (CFPB)

NGOs and Other

- Institute for Human-Centered AI

- FinRegLab

- Ethan Mollick, One Useful Thing (Substack)

FDP Candidate Webinars

WEBINAR

FDP Charter Info Session

Join FDP Experts to learn about the FDP Charter, achieving exam success, and more.

January 9 at 11 AM ET

Road map of learning journey & expectations.

Watch the video recording here!

Testing options
January 25 @ 11 AM ET

WEBINAR

Q2-2024 FDP Review of Prometric Test Center and ProProctor (Remote Testing)

Learn about study tools, exam prep and testing options with the FDP team

January 25, 2024 at 11 AM ET

WEBINAR

FDP Q2-2024 Candidate Orientation

March 14 at 11 AM ET

Get ready for YOUR exam day!
March 14 @ 11 AM ET

Other webinars www.fdpinstitute.org/webinars

Q & A

You are invited to join us for our upcoming webinars!
https://fdpinstitute.org/Webinars/

Feb 13 @ 10 AM ET
Register Here

Feb 22 @ 10 AM ET
Register Here

Thank You

Contact Us:

🌐 fdpinstitute.org

✉ info@fdpinstitute.org

🐦 @FDPbyCAIA

in www.linkedin.com/company/fdpinstitute/