

Federal Housing Finance Agency
Office of Inspector General



Enterprise Use of Artificial Intelligence and Machine Learning

White Paper • WPR-2022-002 • September 19, 2022



WPR-2022-002

September 19,
2022

Executive Summary

The field of **artificial intelligence** (AI) traces back to the 1950s. More recently, the increased availability of both data and computing power has bolstered rapid adoption of AI and **machine learning** (ML) solutions (collectively, AI/ML).

In February 2022, the Federal Housing Finance Agency (FHFA) issued Advisory Bulletin 2022-02, *Artificial Intelligence/Machine Learning Risk Management* (the Advisory Bulletin) to Fannie Mae and Freddie Mac (the Enterprises). In conjunction with the Advisory Bulletin, FHFA's Office of Minority and Women Inclusion issued a supervisory letter providing additional guidance and establishing Agency expectations regarding diversity and inclusion in Enterprise use of AI/ML (the Supervisory Letter). The Advisory Bulletin discusses key risks in Enterprise use of AI/ML and considerations for effective risk management. According to FHFA, AI/ML allows the Enterprises to process large amounts of data and identify complex relationships, adding the opportunity for improved "efficiencies and operations with reduced error and cost." However, AI/ML poses heightened model, data, regulatory and compliance, and other operational risks. AI/ML can also inject or reinforce bias. Although complex, AI/ML bias generally refers to inaccurate or inappropriate methods leading to unfair outcomes. It may manifest as social, economic, political, or cultural biases and can be explicit or implicit. The Agency expects each Enterprise to be guided by a set of "AI/ML Core Ethical Principles" that address these various types of biases with regard to regulatory and compliance risks.

Both Enterprises report a cautious approach to adopting AI/ML. Fannie Mae told OIG that it is proceeding slowly in its adoption. The Enterprise noted the need for additional capabilities to properly manage associated risks and establish confidence in the tradeoffs between risks and rewards. Freddie Mac has identified numerous opportunities for AI/ML use in the mortgage process but, as an Enterprise official told OIG in May 2022, it remains in the early stages of adoption. Freddie Mac observed increased performance and some risk reduction after introducing AI/ML into certain models. These benefits help offset risks heightened by AI/ML, according to an Enterprise official.

The complex AI/ML field continues to rapidly evolve. Rapid technological innovation and use of new modeling approaches can create uncertainty. FHFA provides detailed guidance to the Enterprises on managing the associated risks, and this topic merits continued monitoring and agility considering the interconnected and evolving risks posed by AI/ML.

TABLE OF CONTENTS

EXECUTIVE SUMMARY2

ABBREVIATIONS4

BACKGROUND5

ENTERPRISE USE OF AI/ML.....6

 Fannie Mae7

 Freddie Mac.....7

RISK CONSIDERATIONS.....8

 Model Risks.....9

 Data Risks.....9

 Regulatory and Compliance Risks.....10

 Other Operational Risks11

CONCLUSION.....12

OBJECTIVE, SCOPE, AND METHODOLOGY13

ADDITIONAL INFORMATION AND COPIES14

ABBREVIATIONS

Advisory Bulletin	Advisory Bulletin 2022-02, <i>Artificial Intelligence/Machine Learning Risk Management</i>
AI	Artificial intelligence
Enterprises	Fannie Mae and Freddie Mac
FHFA or Agency	Federal Housing Finance Agency
ML	Machine learning
NIST	National Institute of Standards and Technology
OIG	Federal Housing Finance Agency Office of Inspector General
Supervisory Letter	Office of Minority and Women Inclusion Supervisory Letter on Artificial Intelligence/Machine Learning

BACKGROUND

The field of **artificial intelligence** traces back to the 1950s. More recently, the increased availability of both data and computing power has bolstered rapid adoption of AI and **machine learning** solutions.¹ Generally speaking, an AI/ML system consists of a network of **algorithms** “trained” on data. In a November 2017 report, the Financial Stability Board attributed increasing AI/ML use to the combination of advancements in data collection, declining data storage costs, and the rise of cloud computing—particularly in financial services.

According to FHFA, AI/ML allows the Enterprises to process large amounts of data and identify complex relationships, adding the opportunity for operational efficiencies “with reduced error and cost.” A Freddie Mac official explained that the Enterprise’s use of machine learning algorithms improves its ability to make business decisions and identify new offerings and solutions for clients. According to a Freddie Mac Senior Vice President, “[m]achines—more specifically big data-driven analytics—are better than people at making predictions and do it [in] a fraction of the time.”² Fannie Mae told us the broad benefits of AI/ML include features such as “enhanced interaction with mortgage community and consumers” and “cyber security vigilance.” The Consumer Financial Protection Bureau also describes AI/ML cybersecurity opportunities, such as discovering threats and malicious activity.

Artificial intelligence broadly describes the ability of computers and technology systems to perform tasks that typically require human intelligence and decision-making.

Machine learning comprises a class of artificial intelligence in which algorithms automatically “learn” and optimize based on patterns detected in processing large amounts of data and with limited or no human intervention.

An **algorithm** refers to the steps followed or calculations made based on a set of rules.

¹ There are no standard definitions for artificial intelligence and machine learning. Our descriptions stem from information provided by FHFA, Fannie Mae, and Freddie Mac, and definitions from the Financial Stability Board. See Financial Stability Board, [Artificial Intelligence and Machine Learning in Financial Services](#) (Nov. 2017).

² Freddie Mac, [Playing the Long Game with Big Data and Artificial Intelligence](#) (Aug. 21, 2019).

ENTERPRISE USE OF AI/ML

A **model** processes input data into estimates based on a set of assumptions or theories.

Current Enterprise use of AI/ML almost always involves **models**.³ Models are theoretical constructs that explain relationships and serve as a foundation for each of the Enterprises’ businesses. Fannie Mae says it relies “extensively” on models in business decisions,

risk management, and financial reporting. A Freddie Mac official told OIG that the Enterprise depends on models to make critical business decisions, such as which loans to fund and how much it costs Freddie Mac to do so.

Both Enterprises began applying AI/ML with automated property valuation. Freddie Mac told us that it first applied machine learning in 2017 for its automated valuation model, Home Value Explorer, which canvases data to rapidly produce property value estimates.⁴ A Freddie Mac official noted that the Enterprise began with property valuation primarily because of the vast available data; the more data, the more the tool can “learn.” According to Fannie Mae, AI/ML improves accuracy in its property valuation, compared to prior approaches.⁵

FHFA reported to OIG that in late 2018 and early 2019, the Agency’s ongoing monitoring noticed increasing information on Enterprise incorporation of AI/ML into their businesses. The Agency identified AI/ML as an emerging risk around this time, and, as reported to OIG, it increased its monitoring and assessment of Enterprise AI/ML use as a result.⁶

³ FHFA maintains that AI/ML use is not confined to models. The agency directs each Enterprise to “determine the degree to which it needs to identify and document AI/ML techniques in addition to use cases, understanding that AI/ML can be embedded in models, applications, systems, platforms, tools, and services—either developed in-house or procured from third-party vendors.” For more information, see FHFA, Advisory Bulletin 2022-02, [Artificial Intelligence/Machine Learning Risk Management](#) (Feb. 10, 2022) [hereinafter the Advisory Bulletin].

⁴ An automated valuation model estimates property value through computerized data modeling, as compared to an appraisal report where an appraiser estimates value based on observation and market analysis. For more information, see OIG, [An Overview of Enterprise Appraisal Waivers](#) (Sept. 14, 2018) (WPR-2018-006).

⁵ Fannie Mae did not provide OIG with the specific year it began using AI/ML. The Enterprise reported that the “use of AI/ML techniques in Fannie Mae can be traced back to the early 2000’s but supporting documentation from the era is sparse. This is an approximate estimation based on employee knowledge.”

⁶ In July 2022, FHFA established the Office of Financial Technology to, among other things, support the Agency in addressing emerging risks related to the adoption and deployment of financial technology, including AI/ML.

Fannie Mae

A July 2018 Fannie Mae internal report indicates that the Enterprise explored AI/ML in its efforts to improve operational efficiency, reduce risk, and “simplify housing finance.” By 2019, Fannie Mae reportedly conducted over 10 related pilots and moved into what it called an “active” stage in its AI maturity. In June 2021, Fannie Mae described its use of AI/ML at the time as “limited,”

and it leveraged AI/ML in applications such as property valuation, fraud detection, post-purchase loan reviews, and language processing. In May 2022, Fannie Mae reported to OIG that the Enterprise also uses **chatbots** to answer customer questions regarding the mortgage process and has since retired AI/ML use in post-purchase loan reviews.

Chatbots are computer programs that understand and respond to questions and commands. They are used to provide human-like interactions.

Fannie Mae told OIG that it is proceeding slowly in its adoption of AI/ML. The Enterprise noted the need for additional capabilities to properly manage associated risks and establish confidence in the tradeoffs between risks and rewards. AI/ML presents Fannie Mae with the opportunity to analyze prior results and predict future results, which could reduce risk and improve earnings, according to the Enterprise.

Freddie Mac

In March 2018, Freddie Mac internally reported on AI/ML opportunities, including capabilities to understand customer sentiment, predict borrower behavior, and assist in customer service. By August 2019, the Enterprise “increasingly” adopted machine learning tools. At that time, Freddie Mac also reported developing governance protocols and revising standards to manage the emerging risk presented by its growing use of AI/ML.

Freddie Mac maintains a central repository that tracks AI/ML tools and techniques across the Enterprise. The Enterprise’s internal AI Standard charges a working group with oversight of appropriate risk governance process application to each instance. In early May 2022 the repository comprised 74 items and included uses for language processing, predictive analytics, chatbots, and anomaly detection.

Freddie Mac has identified numerous opportunities for AI/ML use in the mortgage process, but, as an Enterprise official told OIG in May 2022, it remains in the early stages of adoption. Freddie Mac observed increased performance and some risk reduction after introducing AI/ML into certain models. These benefits help offset the risks heightened by AI/ML, according to an Enterprise official.

RISK CONSIDERATIONS.....

In February 2022, FHFA issued Advisory Bulletin 2022-02, *Artificial Intelligence/Machine Learning Risk Management* to the Enterprises.⁷ The Advisory Bulletin discusses key risks in Enterprise use of AI/ML and considerations for effective risk management. FHFA describes the Advisory Bulletin as less prescriptive than the Agency’s other advisory bulletins. FHFA explained to OIG that it chose a more “principles-based” approach partly because AI/ML comprises a broad topic that can cover many areas across an Enterprise.

The Advisory Bulletin highlights model, data, regulatory and compliance, and other operational risks (e.g., information security and third-party providers) in Enterprise use of AI/ML. It discusses bias as a component of both model and data risks. Although complex, AI/ML bias generally refers to inaccurate or inappropriate methods leading to unfair outcomes. It may manifest as social, economic, political, or cultural biases and can be explicit or implicit.⁸ The Agency expects each Enterprise to be guided by a set of “AI/ML Core Ethical Principles” that address these various types of biases with regard to regulatory and compliance risks. Further, FHFA’s Office of Minority and Women Inclusion issued a supervisory letter in conjunction with the Advisory Bulletin directing Enterprise use of AI/ML “be designed to promote [diversity and inclusion] and address explicit and implicit biases to ensure equity in AI/ML recommendations.”⁹

In addition to discussing AI/ML key risks, the Advisory Bulletin intersects with various other FHFA guidance that discusses other risks implicated by Enterprise use of AI/ML. These include, but are not limited to, guidance on model risk, information security, cloud computing, third-party providers, fair lending and fair housing compliance, and business resiliency. FHFA cites the Enterprises leveraging existing risk management and control frameworks “to the extent practicable” as a key principle of the Advisory Bulletin.

⁷ See FHFA, the Advisory Bulletin, *supra* note 3, at 6.

⁸ The National Institute of Standards and Technology (NIST) has done extensive work on identifying and managing bias in AI. FHFA references this work in its guidance to the Enterprises. NIST identified three categories of AI/ML bias: systemic, statistical and computational, and human biases. Systemic bias describes when an institution operates in a way that favors certain social groups and disadvantages others. Statistical and computational biases arise when AI/ML systems use data not representative of the population or the algorithm trains on one type of data but cannot extrapolate beyond that data. NIST reports that human biases are often implicit and reflect the individual, group, or institutional decisions in introducing and implementing AI/ML. For more information, see NIST Special Publication 1270, [Towards a Standard for Identifying and Managing Bias in Artificial Intelligence](#) (Mar. 16, 2022).

⁹ See FHFA, [Office of Minority and Women Inclusion Supervisory Letter on AI/ML](#) (Feb. 10, 2022).

Below, we describe AI/ML risk considerations, including FHFA’s guidance on and Enterprise views of these interconnected risks.¹⁰

Model Risks

AI/ML primarily serves as a modeling tool for the Enterprises, as mentioned earlier. The Advisory Bulletin lists various types of **model risks** heightened with the use of AI/ML,¹¹ with “black box risk” topping the list. Black box risk describes the lack of **explainability**, interpretability, and transparency. AI/ML increases a model’s complexity; the black box essentially means it is very difficult to understand how the model arrived at its results.

FHFA defines **model risk** as the risk of loss resulting from model errors or the incorrect use or application of model output.

In AI/ML, **explainability** refers to clarity in how a model uses inputs to produce outputs.

Black box risk can increase the uncertainty of whether the AI/ML approach is conceptually sound or suitable. Broadly, lack of transparency limits AI/ML auditability and creates a barrier in ensuring equitable treatment of those affected by these systems, according to the Government Accountability Office. Freddie Mac officials highlighted black box risk as the Enterprise’s primary risk concern in AI/ML. According to a Freddie Mac official, black box risk slowed the company’s adoption of AI/ML, but the Enterprise has been working since 2016 to make its use of machine learning algorithms more transparent and explainable. Fannie Mae has also pinpointed explainability as an obstacle in using machine learning in certain areas, such as for credit decisions.

FHFA also expects the Enterprises to manage potential bias in AI/ML models. The Advisory Bulletin explains that bias in AI/ML models “contributes to poor predictability and can lead to discriminatory or unfair outcomes that benefit or harm some individuals, groups, or communities disproportionately.” It directs that the Enterprises consider tools and techniques to assist in AI/ML model bias detection.

Data Risks

Data drives AI/ML. FHFA expects the Enterprises to use appropriate and accurate data to ensure their AI/ML solutions produce reliable results. Both Enterprises agree the quality of an AI/ML solution ties directly to the quality of data used. Freddie Mac’s internal Model Risk Standard details multiple levels of verification required to assure data integrity. The

¹⁰ We also note some ways the Enterprises may mitigate risks, but we did not assess the adequacy or implementation of those mitigants.

¹¹ For more information on Agency guidance on model risk in general, see FHFA Advisory Bulletin 2013-07, [Model Risk Management Guidance](#) (Nov. 20, 2013).

Standard notes that even accurate data may not be appropriate to use in certain models. Designated model owners are charged with periodically performing “data appropriateness” checks.

As in model risks, bias can heighten data risk with AI/ML use. The Supervisory Letter notes that AI/ML concerns may arise due in part to “data embedded with historical, social, economic, and cultural biases.” FHFA reported to OIG that the Enterprises’ first line of defense in addressing AI/ML bias is effective data risk management. NIST explains, AI systems “sometimes do not operate as intended because they are making inferences from patterns observed in data rather than a true understanding of what causes those patterns.” The Agency expects the Enterprises to develop and execute strategies to address potential bias in data. For example, the Enterprises could strategize for diversity in teams and ensure representative datasets, according to FHFA.

Regulatory and Compliance Risks

In the Advisory Bulletin, FHFA alerts the Enterprises to monitor their use of AI/ML for compliance with laws and regulations, including those related to **fair lending**,¹² consumer protection, privacy, and employment discrimination.

Fair lending prohibits discrimination based on race, color, national origin, religion, sex, familial status, disability, or other protected classes during the mortgage process.

AI/ML use in credit underwriting can generate fair lending concerns. Because AI/ML approaches may be viewed as less explainable than traditional statistical methods, issues may arise with justifying and validating AI/ML outcomes. A Freddie Mac compliance official explained to OIG that the Enterprise needs to identify what happens when introducing AI/ML into credit underwriting to understand the potential impact on fair lending. The Advisory Bulletin states the Enterprises may need to adapt “existing regulatory and compliance risk management practices and controls to accommodate AI/ML associated risks,” including “integrating fair lending reviews and testing, as appropriate, through all lifecycle stages.” As reported to OIG in May 2022, neither Enterprise currently uses AI/ML directly in underwriting; however, Freddie Mac hopes to incorporate an AI/ML credit scoring solution currently under development in its automated underwriting system.

¹² Federal fair lending laws that apply to the Enterprises include the Fair Housing Act, 42 U.S.C. 3601 et seq., the Equal Credit Opportunity Act, 15 U.S.C. 1691 et seq., and section 4545 of the Federal Housing Enterprises Financial Safety and Soundness Act, 12 U.S.C. 4501 et seq. For more information, see FHFA Advisory Bulletin 2021-04, [Enterprise Fair Lending and Fair Housing Compliance](#) (Dec. 20, 2021).

Other Operational Risks

The Advisory Bulletin lists additional operational risks implicated in Enterprise use of AI/ML, including: information technology infrastructure, information security, business continuity, and use of AI/ML through third-party providers. We detail below the interconnection of these risks. And as previously discussed, the Advisory Bulletin references separate but related guidance for these risk areas. In addition to leveraging existing risk management frameworks, FHFA expects the Enterprises to enhance policies and standards to address the heightened risk posed by AI/ML.

- **Information technology infrastructure:** FHFA warns that AI/ML computing networks can degrade due to insufficient resources. The Agency advises the Enterprises to consider scalable infrastructure to meet AI/ML’s data storage and computing needs. Both Enterprises reported to OIG that cloud computing has helped provide the infrastructure to meet these needs.
- **Information security:** The Advisory Bulletin states AI/ML systems “may pose risks to existing processes that can compromise the confidentiality, integrity, and availability of information.” The Enterprises use cloud-based AI/ML tools. While not unique to AI/ML, FHFA considers cloud information security a “critical” issue given the shared responsibility model that cloud service providers use.¹³ The Agency noted that the cloud may support more efficient use of AI/ML but also may “create a dependency on those tools.”
- **Business continuity:** The Advisory Bulletin also notes potential business continuity risks when functions supported by AI/ML feed into other business processes and systems. According to FHFA, “significant” disruptions across an Enterprise can arise “if AI/ML performance is degraded or compromised.”
- **Use of AI/ML through third-party providers:** The Enterprises rely on third-party cloud providers to meet AI/ML’s data storage and computing needs. Additionally, the Enterprises use third-party machine learning platforms for data processing and model development. These third-party uses present additional risk concerns, including information security, business resiliency, and concentration.¹⁴

¹³ When using a third-party cloud provider, the Enterprises must share responsibility for safeguarding information and systems in the cloud environment. For more information, see OIG, [An Overview of Enterprise Use of Cloud Computing](#) (Mar. 11, 2020) (WPR-2020-002).

¹⁴ The Advisory Bulletin states the potential for business resiliency and concentration risks if only a few vendors provide AI/ML services.

CONCLUSION.....

Enterprise use of AI/ML is rich with reported benefits and opportunities while also possessing a myriad of attendant risks; AI/ML covers many interconnected areas across the Enterprises and poses heightened model, data, regulatory and compliance, and other operational risks. AI/ML can also inject or reinforce various types of biases. Competitive pressures also shadow these risks. Fannie Mae identified a “growing danger” that it’s proceeding too slowly in building up AI/ML capabilities and this could hurt its “competitive position (profitability, cost efficiency of operations, etc.).”

Freddie Mac previously observed there are “very few accepted AI standards for the mortgage industry.” Further, the complex AI/ML field continues to rapidly evolve. Rapid technological innovation and use of new modeling approaches can create uncertainty. FHFA provides detailed guidance to the Enterprises on managing the associated risks, and this topic merits continued monitoring and agility considering the interconnected and evolving risks posed by AI/ML.

OBJECTIVE, SCOPE, AND METHODOLOGY

The objective of this white paper was to provide information on risks associated with Enterprise use of artificial intelligence and machine learning. To achieve this objective, we reviewed internal FHFA and Enterprise documents, as well as publicly available documents. We interviewed FHFA and Enterprise officials and requested and reviewed information from the Agency and Enterprises.

We provided FHFA with the opportunity to respond to a draft of this white paper. We appreciate the cooperation of FHFA staff, as well as the assistance of all those who contributed to the preparation of this white paper.

ADDITIONAL INFORMATION AND COPIES.....

For additional copies of this report:

- Call: 202-730-0880
- Fax: 202-318-0239
- Visit: www.fhfaoig.gov

To report potential fraud, waste, abuse, mismanagement, or any other kind of criminal or noncriminal misconduct relative to FHFA's programs or operations:

- Call: 1-800-793-7724
- Fax: 202-318-0358
- Visit: www.fhfaoig.gov/ReportFraud
- Write:

FHFA Office of Inspector General
Attn: Office of Investigations – Hotline
400 Seventh Street SW
Washington, DC 20219