Google Cloud

July 1, 2021

Chief Counsel's Office
Attention: Comment Processing
Office of the Comptroller of the Currency
400 7th Street SW, Suite 3E-218
Washington, DC 20219

Comment Intake
Bureau of Consumer Financial Protection
1700 G Street NW
Washington, DC 20552

The Honorable Ann E. Misback
Secretary, Board of Governors of the Federal Reserve System
20th Street and Constitution Avenue NW
Washington, DC 20551

Melane Conyers-Ausbrooks
Secretary of the Board
National Credit Union Administration
1775 Duke Street
Alexandria, VA 22314

The Honorable James P. Sheesly
Assistant Executive Secretary
Attn: Comments
Federal Deposit Insurance Corporation
550 17th Street NW
Washington, DC 20429

Re: ***Request for Information and Comment on Financial Institutions' Use of Artificial Intelligence, Including Machine Learning*** (*OCC*: Docket ID OCC-2020-0049; *Federal Reserve System*: Docket No. OP-1743; *FDIC*: RIN 3064-ZA24; *CFPB*: Docket No. CFPB-2021-0004; *NCUA*: Docket No. NCUA -2021-0023)

Google Cloud welcomes the opportunity to provide comments in response to the "*Request for Information and Comment on Financial Institutions' Use of Artificial Intelligence, Including MachineLearning*" (86 Fed Reg16837) (March 31, 2021) (the "RFI") issued jointly by the Office of the Comptroller of the Currency, Federal Deposit Insurance Corporation, Board of Governors of the Federal Reserve

System, Bureau of Consumer Financial Protection, and the National Credit Union Administration (collectively referred to herein as the "Financial Regulators").

I. Introduction

As the Financial Regulators recognize, the use of artificial intelligence and machine learning (AI/ML) in financial services holds substantial promise in driving efficiency, reducing costs, enhancing compliance, and improving customer experience and outcomes.

Cloud technology underpins these advancements by providing a key platform for innovation in financial services. Cloud provides the compute power, infrastructure, and access to AI/ML tools that may have otherwise been out of reach for all but the biggest and most tech-forward institutions. Financial services firms of all sizes are now capable of leveraging real-time transaction information streamed into large-scale, real-time databases and applying AI/ML models to assess, quantify, and calculate risk.

AI will have a significant impact on society for many years to come. At Google Cloud, we have developed AI Principles (including for identification of applications we will not pursue) to guide our teams on the responsible development and use of AI. These are backed by the operational processes and structures necessary to ensure they are not just words but concrete standards that actively impact our research, products and business decisions to ensure trustworthy and effective AI application.

Similarly, the responsible development and adoption of AI will need to be a collective goal of market participants and regulators in the financial services industry as well. For this purpose, and given the rich frontier of advances that AI/ML can bring to the financial services industry, Google Cloud appreciates the opportunity to provide the following comments in response to the specific questions presented in the RFI.

II. Fostering Adoption of AI/ML to Enhance Key Regulatory Compliance Functions

Use of AI/ML for key regulatory compliance functions is a clear area of promise. One of the most compelling areas of development is in the context of AI/ML for internal risk management and compliance systems, including with respect to anti-money laundering/countering the financing of terrorism ("AML/CFT") compliance and reporting, know-your-customer ("KYC") processes, and fraud detection and prevention.

To take the AML/CFT use-case, current heavily manual and labor intensive approaches to identifying and combating financial crimes have not been effective. Despite the billions of dollars spent by banks and financial institutions to detect and report suspicious activity, a reported 95 percent of alerts are ultimately deemed "false positives" and very few alerts lead to an actual

Suspicious Activity Report (SAR) filing. See e.g., Anti-money Laundering Controls Failing to Detect Terrorists, Cartels and Sanctioned States (March 2018). Money laundering fuels such societal ills as drug trafficking, human trafficking, and terrorist activities.

The methodological challenge with current approaches is that they are largely "rules-based," which makes them inherently brittle and easily circumvented by bad actors who can game these rules and make their transactions look "normal." At the same time, the applicable rules can be overly broad and generate substantial false-positives.

AI-enabled AML/CFT approaches, on the other hand, can develop a much more sophisticated analytic lens capable of ingesting massive volumes of data, in a more timely way, in order to detect new patterns and anomalies that might bypass simple, rules-based logic. These engines can be trained to improve accuracy, reduce false-positives, and help banks perform internal risk assessments and better determine when, amongst millions of legitimate transactions being processed, bad actors are trying to move criminal money. AI can further incorporate more contextual signals and generate more targeted flags for investigators, reducing toil and allowing them to focus on the most serious issues that are identified.

Beyond the AML context, AI/ML tools are further advancing other key compliance and risk mitigation functions. With respect to KYC requirements, for example, AI/ML technology can help automate workflows and improve the customer onboarding experience by decreasing completion times. More specifically, machine learning tools can help automate the processing of KYC data and records, including by powering the search and deduping of such information.

Similarly, for fraud detection, AI/ML technology can help provide real-time, low-latency fraud monitoring that constantly adapts to new fraud patterns, starting with transaction fraud but expanding to other attack surfaces. This same approach can enhance trade and transaction monitoring in order to reduce instances of fraud.

The U.S. government has taken important steps to signal an intent to help enable adoption of innovation in the regulatory compliance space. For example, in a joint statement issued in late 2018 by the Fed, FDIC, OCC, NCUA, and FinCEN, the agencies stated that, "innovations and technologies can strengthen BSA/AML compliance approaches, as well as enhance transaction monitoring systems. The Agencies welcome these types of innovative approaches to further efforts to protect the financial system against illicit financial activity."

Following this example of agency coordination, the OCC in September 2019 issued an interpretive letter in response to a bank request for regulatory clarity on whether it could use software to automatically identify a category of suspicious financial activity and submit an auto-populated SAR. Subject to certain parameters, the OCC concluded that under this particular fact pattern the use of automation to detect suspicious activity "would be permissible under the OCC's SAR regulation" and that the use of automation to populate the SAR form "is legally consistent" with the OCC's requirements. Google Cloud views these kinds of clear signals of regulatory intent to be critically important inasmuch as they provide increased guidance and clarity regarding regulatory expectations.

Google Cloud believes there are also other tools in the regulator's toolkit that can reduce ambiguity that would otherwise chill adoption of technologies capable of increasing the safety and soundness of the financial system and enhancing compliance.

First, given the pace of AI/ML innovation, it is essential for regulators to stay well-informed of technological developments and to incorporate the expertise of technology providers into regulatory processes. Frequent reviews and requests for feedback, like this RFI, are useful. Similar approaches could be applied to other aspects of regulation that significantly impact technology adoption.

One example of this is the guidance that the Federal Financial Institutions Examination Council (FFIEC) issues in relation to cloud computing. Google Cloud strongly supports the FFIEC's efforts to update the Information Technology Handbook to more clearly address cloud and adapt traditional rules of outsourcing and risk management to cloud. We believe, however, that this would be most effectively done through a process that seeks formal engagement around contemplated guidance and provides opportunity for review and comment. Cross-agency coordination and working groups can help facilitate this objective.

Second, industry standards and best practices can serve as a helpful baseline for regulatory expectations and guidance. Even though many in the financial services industry are in early stages of their AI/ML adoption, facilitating the development of standards as has been done in other areas, including around cybersecurity, would be good policy. Formal nods by regulators regarding industry standards, and certifications, or key best practices, could help provide clarity and positive incentives to market participants in their technology adoption journeys. These efforts should also retain flexibility to incorporate new standards as they are developed or as existing standards evolve.

Third, pilots and controlled testing environments for certain regtech applications, including around AML/CFT, can help advance regulators understand and market adoption of AI/ML technologies. Given the increasing role technology plays in our financial system, it is important that regulation recognizes the benefits of iteration, testing and measuring empirical outcomes. It is only through such testing that regulated entities and the regulator can determine whether new tools can drive better compliance, customer, and market outcomes.

Finally, Google Cloud remains supportive of continued efforts by regulators to capacity-build, including through hackathons, internal pilots and innovation competitions. Google Cloud commends agencies involved in this RFI for their leading roles in developing innovation programs and offices, and believes that further build-out of these efforts can enhance regulatory understanding and help the industry to confidently innovate and adopt technologies that benefit consumers, enhance competitiveness, and advance overall safety and soundness.

III. Explainability

Google Cloud recognizes that a key to adoption of more advanced AI/ML technologies will be demonstrating sufficient "explainability" around the models, their decision making, and their operations. In developing regulatory guidance and expectations around "explainability" a number of considerations may be helpful.

Taxonomy: As an initial matter, it may be helpful to develop a taxonomy of "explainability." The term evokes -- and sometimes is used interchangeably with -- a number of different, but related concepts, including "interpretability," "auditability," "traceability," "contestability," "accountability," and "transparency." Having a clear understanding of what is meant by explainability in any given financial services regulatory context, drawing from international standards where consensus develops, will be critical.

Purpose/Audience: The RFI defines "explainability" as "how an AI approach uses inputs to produce outputs." This refers to a core aspect of explainability -- namely, understanding what the model is doing (it's logic and reasoning). A second core aspect of explainability, however, which is not reflected in the RFI's definition, relates to how an explanation is communicated -- i.e., whether insights are effectively communicated to people in a way that they understand and meets their needs. Importantly, different audiences may have different expectations and needs from an explanation. The National Institute of Standards and Technology (NIST) refers to this as "meaningfulness" in its *Four Principles of Explainable AI*(hereinafter "NIST Principles"). Specifically, the NIST Principles recognize that:

> Multiple groups of users for a system may require different explanations. . . . Groups may be defined broadly as the developers of a system vs. end-users of a system; lawyers/judges vs. juries; etc. The goals and desiderata for these groups may vary. For example, what is meaningful to a forensic practitioner may be different than what is meaningful to a juror.

Clarity about the purpose and the audience of the explanation in specific regulatory settings can help ensure that this "meaningfulness" principle can also be met.

Against this backdrop, there are a number of technological approaches that are being developed to advance explainability determinations. From a technological perspective, explainability can be broadly analogized to a tree with two major branches.

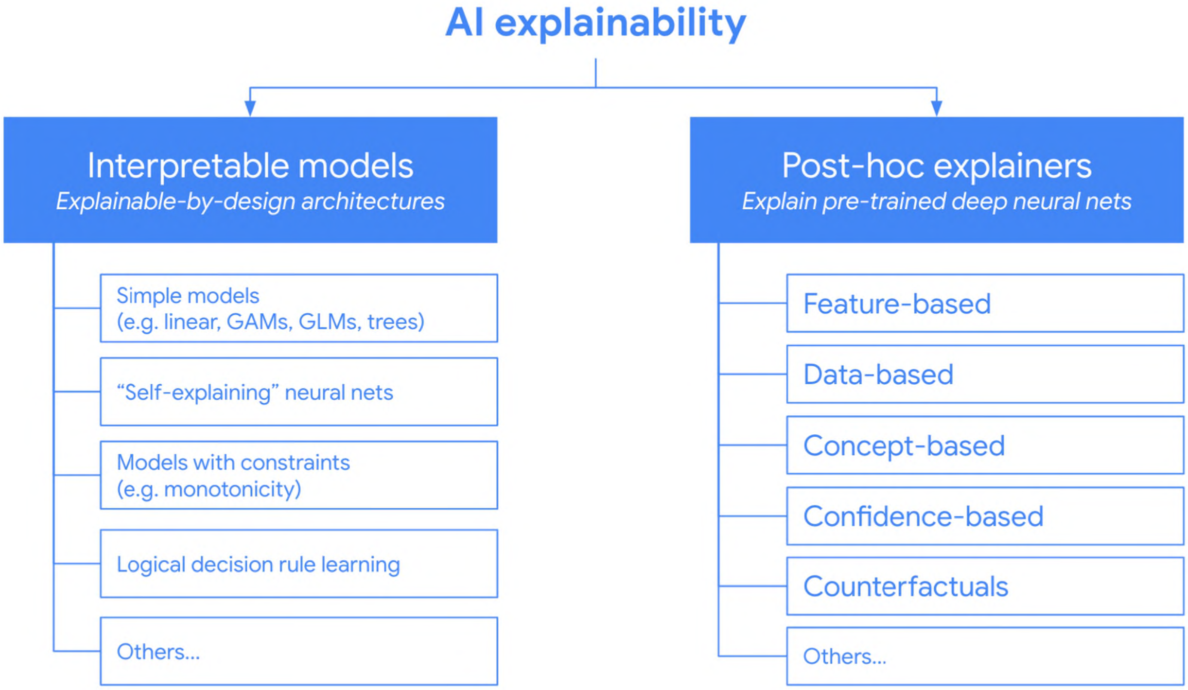
AI explainability
Interpretable models
Explainable-by-design architectures
Simple models (e.g. linear, GAMs, GLMs, trees)
"Self-explaining" neural nets
Models with constraints (e.g. monotonicity)
Logical decision rule learning
Others...
Post-hoc explainers
Explain pre-trained deep neural nets
Feature-based
Data-based
Concept-based
Confidence-based
Counterfactuals
Others...

On one branch, are AI/ML models that are designed and built to be directly interpretable. This includes simple models such as linear regression models, which can be "explained" by looking at the coefficients that exactly describe the relationship between input features and model outputs. Basic decision trees are also able to be explained by looking at the path taken through the tree to arrive at a particular decision - the combination of rules that determine the output. Interpretable models increasingly include more flexible models that have built-in interpretability affordances, including, for example, "self explaining neural nets" that use "attention mechanisms" to track what parts of a datapoint the model is keying in on most to make a particular prediction.

The second branch consists of so-called "post-hoc" explainers, which involves techniques for understanding the behavior of highly complex ML models. The significant predictive power of these models comes from the fact that its internal logic is not easily reduced to simple rules. For these models, many techniques are being developed to understand why a model has made certain predictions.

For example, "feature-based" explanations try to explain why a model has made a particular prediction by quantifying how much each input feature contributed to the model's prediction. For a model trained to predict the likelihood that an airline flight will be delayed, the weather is likely to be a very important input feature, whereas the average age of the passengers is likely not important. In this example, the percentage of a prediction attributable to weather might be a significant percentage contributor to the overall score.

The state of technological development now allows for the identification of methodologies that can help provide insights into the outputs of particular types of AI/ML models. Clarity from regulators regarding the objectives that they are seeking to meet in particular contexts can help with the assessment of which approaches and methodologies are appropriate in any given situation. Providing this kind of clarity could help regulators significantly advance adoption of tools that quantifiably upgrade the effectiveness and usefulness of certain data analytics, including in the financial crime detection context. For example, more complex techniques may need to be embraced appropriately to meet the objectives in complex scenarios (e.g., AML detection and prevention). At the same time, prescriptive or one-size-fits-all requirements are unlikely to be effective -- especially given the fast pace of technological development in this space. A focus on the development of context-based, outcome-based, and principles-based approaches is more likely to be effective.

IV. Cybersecurity Risk

As noted in the RFI, it is critical to assess and mitigate cybersecurity risks to AI/ML models and systems, including the data used by them. For AI/ML tools provided on Google Cloud, customers are able to leverage an infrastructure that is designed, built, and operated with security at the core, protected by more than 900 experts in information, application, and network security. When a firm uses our cloud-based tools, Google Cloud becomes its partner in security, protecting customer data by monitoring data health, detecting anomalous behaviors, and proactively preventing security incidents by utilizing machine intelligence.

To that end, Google Cloud regularly undergoes independent verification of security, privacy, and compliance controls to ensure compliance with stringent global data protection and security standards. Further, Google Cloud is continually expanding the tools available to customers to optimize for confidentiality and control for high-sensitivity workloads. Recent examples include:

- Confidential Computing to encrypt customer data in use (encryption in transit and at rest are already provided by default).
- External Key Management options to allow customers to keep data encryption keys offsite and air gapped from Google.
- Access Transparency, Access Approval and Key Access Justifications to ensure that customers can understand why access for their data is being requested, even by Google, and deny access to their data, should they wish.
- Assured Workloads to give customers confidence that they are meeting specific compliance requirements and ensuring data is only accessible and supported by designated personnel and locations.

These and other controls have been developed to give Google Cloud customers confidence that even the most sensitive data or highly regulated workloads can be managed in a cloud environment.

V. Model Risk Management

While the Federal Reserve's Model Risk Management (MRM) Guidance was developed before there was meaningful use of AI/ML-based models in the financial services industry, many of the high-level principles of the MRM Guidance remain largely applicable and useful even today. That said, the fact that the MRM Guidance does not speak specifically to the AI/ML context may raise some challenges in terms of uncertainty that threaten to impede progress in the development and adoption of AI/ML-based models by the industry. Absent clarity in these circumstances, and with heavy consequences of failing to meet regulatory requirements and expectations, financial services institutions may be incentivized to take overly conservative approaches that could result in longer times to production or in initiatives being taken off the table.

At the same time, any effort to create clarity through prescriptive rules that are frozen in time is bound to be ineffective, particularly given the pace of technological change in the AI/ML space. Finding a way to lower regulatory uncertainty without creating new regulatory blockers in the form of prescriptive, and quickly-outdated, requirements is, and will continue to be, a central challenge for market participants and regulators alike.

A number of approaches may help to meet the challenge:

- First, the Financial Regulators should explore the full range of tools in their toolkit; some may be better suited to encouraging responsible innovation than others. For example, as discussed above in Section II, clear signals of regulatory intent to encourage innovation may be helpful for both market participants and stakeholders within the regulatory community (e.g., examiners). Similarly, because blockers to innovation are often found not in regulations *per se*, but in their application (including uncertainty about their application) creating controlled testing environments, safe harbors, or spaces for deployment of approved pilots that market participants could avail themselves of in order to work with regulators in an iterative fashion to determine how regulations should apply in new contexts.

- Second, it will be important that guidance and regulatory expectations on AI/ML-based models recognize the benefits of iteration, testing and measuring empirical outcomes, and strategic incorporation of domain understanding as key to realizing improved outcomes, consistent with the safety and soundness of the financial system.

  - *Risk Mitigation:* Risks regarding safety and soundness can potentially be mitigated or reduced through approaches such as the following (which are not intended as a comprehensive listing). If used, each of these approacheswould need to be tailored to the particular context.

  - **Upfront testing** to ensure performance is repeatable and robust from historical testing to production use (e.g., model doesn't only perform well in backtest but also translates to production use).

  - **Consistent monitoring of key metrics** to detect signs of abnormal model performance, or model performance degradation (e.g., model performance may decline for a certain risk typology over time).

  - **Production roll-out schedules** that test model performance on small representative slices (e.g., a subset of the population that is representative of the characteristics of the overall population), expanding to greater slices as great confidence is built in performance (e.g., rollout to 1% of business, expand to5%, etc). Considerations on the specifics of the rollout schedule may vary based on the use case and context

  - **Human-in-the-loop** reviews for low confidence predictions, and to gather feedback to improve model performance (e.g., Fraud analyst confirming a prediction about fraud).

  - **Fall back options.** For example, if a new model's performance starts to quickly degrade, a previous version of the model can be quickly swapped back into the production traffic.

  - **Remediation plans** for big identified issues (e.g., performance issue on a particular risk typology). An investigation could be kicked off to create an update to fix the issue and then go through the iteration loop outlined in the bullets above.

- *Rapid Iteration:* With safety and soundness risk mitigations in place, financial institutions can embrace the benefits of rapid iteration. AI/ML-based models are not static, they can be thought of as systems that constantly improve and change to meet the requirements. Thus, prioritization should be on testing and measuring improvements. Speed of iteration with robust measurement can lead to greater exploration of the problem space, which can lead to better solutions and ultimately a more robust and stable system. In fact, speed of iteration can be one of the most important aspects of ensuring safety and soundness of the financial systems in financial services, especially in changing or adversarial problem spaces. For example, fraud patterns are constantly changing, and financial institutions' ability to quickly adopt and protect customers may be dependent on the ability to quickly test, validate, and roll out new approaches to better detect the new risk patterns. This adaptation to new patterns is one of AI/ML systems greatest strengths but is dependent on the ability to move quickly where necessary (consistent with safety and soundness principles)

- *Domain Understanding:* The safety and soundness risk mitigation, and rapid iteration focuses are not substitutes for domain understanding. Domain understanding is important in all phases, including in the model building phase (e.g., deep understanding of fraud typologies to inform an effective system for detecting fraud). A key focus should be on ensuring that domain understanding drives the development of the hypotheses that can be tested andvalidated to prove robustness. Domain understanding can also be used to validate or enhance how the models adapt to new patterns and/or populations. For example, a ML technique called Active Learning could be used to suggest new types of fraud for a fraud investigator to review. If the fraud investigator confirms thenew type of fraud, that can be fed back into the AI/ML system as a new label so the system can quickly adapt over-time (this can be used in combination with other techniques that can help identify new patterns prior to having labels). This dynamic feedback loop with domain experts is a core part of robust ML systemsover-time.

VI. Supporting Community Institutions

One of the most important aspects of cloud technologyis its ability to help smaller financial institutions innovate, compete, and better serve their customers. Cloud helps to create a base layer technology stack upon which smaller institutions can build scalable business solutions and incorporate leading software and machine learning tools traditionally available only to larger competitors with large data science teams. Similarly, many low- and no-code cloud solutions like Document AI make the benefits of AI accessible tosmaller banks. Open cloud-based platforms can also help smaller institutions reconfigure their internal systems away from legacy "walled-garden" designs that can box-in an institution's ability to innovate and help avoid vendor lock-in.

In order to facilitate smaller institution adoption of promising AI/ML technologies there are a few steps that regulators can take to help level the playing field relative to large financial institutions.

First, smaller institutions and community banks may be even more sensitive to regulatory uncertainty than larger institutions. Smaller institutions may lack the means -- in terms of time, money, and resources -- to divine regulatory attitudes and requirements. Continued, clear expression of regulator receptivity to bank adoption of new technologieswill be particularly important for these institutions.

The development and recognition of industry standardsand best practices could further reduce uncertainty around technology adoption and respective responsibilities. More specifically, clear playbooks that establish an accepted standard of carewhen it comes to incorporating, managing, and overseeing AI/ML technology can reduce the burden on smaller institutions seeking to remain digitally competitive. Regulators can foster such efforts by granting established standards regulatory significance, including through inclusion in guidance or safe harbors as outlined above.

On the industry side, technology vendors and bank customers can work closely together to identify and document shared and respective responsibilities to ensure that gaps in oversight do not occur. Regulators play a role in advancing these industry efforts by articulating that evolving industry norms can help reduce risk, while permitting an adaptive framework that can adjust to technological development. Smaller institutions will substantially benefit from such an approach and more readily be able to adopt advanced technologies that can level the playing field with a range of competitors.

VII. Conclusion

Google Cloud appreciates the opportunity to provide comments in response to the RFI. We look forward to continuing to work with Financial Regulators as their work in this area continues.

Sincerely,

Behnaz L. Kibria
Senior Policy Counsel