LAB49

FINRA Request for Comments on Artificial Intelligence (AI) in the Securities Industry

A Capability Model for AI

Submitted on August 31, 2020

August 31, 2020

Submitted electronically to pubcom@finra.org

Marcia E. Asquith
Office of the Corporate Secretary
FINRA
1735 K Street, NW
Washington, DC 20006-1506

Re: Request for Comments on Artificial Intelligence (AI) in the Securities Industry

Dear Ms. Asquith:

Lab49 Consulting Ltd. ("Lab49") appreciates the opportunity to provide our comments to the Financial Industry Regulatory Authority ("FINRA") on "Artificial Intelligence (AI) in the Securities Industry."

Lab49 is a global strategy, design, and technology consulting firm specializing in capital markets. With over 18 years of experience executing on our clients' most critical technology initiatives across the Americas, EMEA, and Asia-Pacific, we firmly believe that the future of financial services regulation is linked to the use of AI.

FINRA member firms' ambition to improve decision-making will lead to the creation of complex AI models, and supervising decisions made or assisted by these models will become increasingly complex and challenging. Explainability of the models and ensuing actions of member firms must be handled carefully to avoid an erosion of confidence.

To ensure confidence is maintained, firms need to take a holistic view of AI capabilities with a focus on Risk Management and model explainability. Accordingly, FINRA's guidance on AI capabilities is required.

Lab49's response to your request for comments offers a capability model for AI as a foundation for FINRA guidance, and draws out key considerations to enable model explainability.

Contents

A Capability Model for AI

Capability models are an established way to support top-down planning and execution within an organization, and standard capability models exist for many domains.

Lab49 believes that FINRA's guidance on AI model management capabilities will be required to establish and maintain confidence in the decisions made by the member firms using AI.

We propose leveraging a standard framework[1] to establish a baseline capability model focusing on critical areas where regulatory guidance is warranted. We further propose formalizing the capability model and making it a part of a future FINRA regulatory taxonomy[2] to aid regulatory rulebook discoverability.

Figure 1 illustrates an example of the proposed capability model.

Establishing transparency in the decision-making process is of paramount importance to instill confidence in the decisions made by member firms. To enable transparency, firms need to ensure they have sound risk management capability and robust data and AI capabilities, with a focus on explainable models.

Figure 1 highlights critical links between Risk Management, AI Model Governance, Model Explainability, Data Modelling, Data Quality and Metadata Management. We will refer to these concepts in our response, starting with Risk Management.

[1] See Data Management Body of Knowledge (DMBOK) https://dama.org/content/body-knowledge (March 2011)

[2] See Development of a Taxonomy-Based Machine-Readable Rulebook https://www.finra.org/sites/default/files/SPNotice-7-30_Lab49_Comments.pdf (October 12, 2018)

Figure 1. AI Capability Model

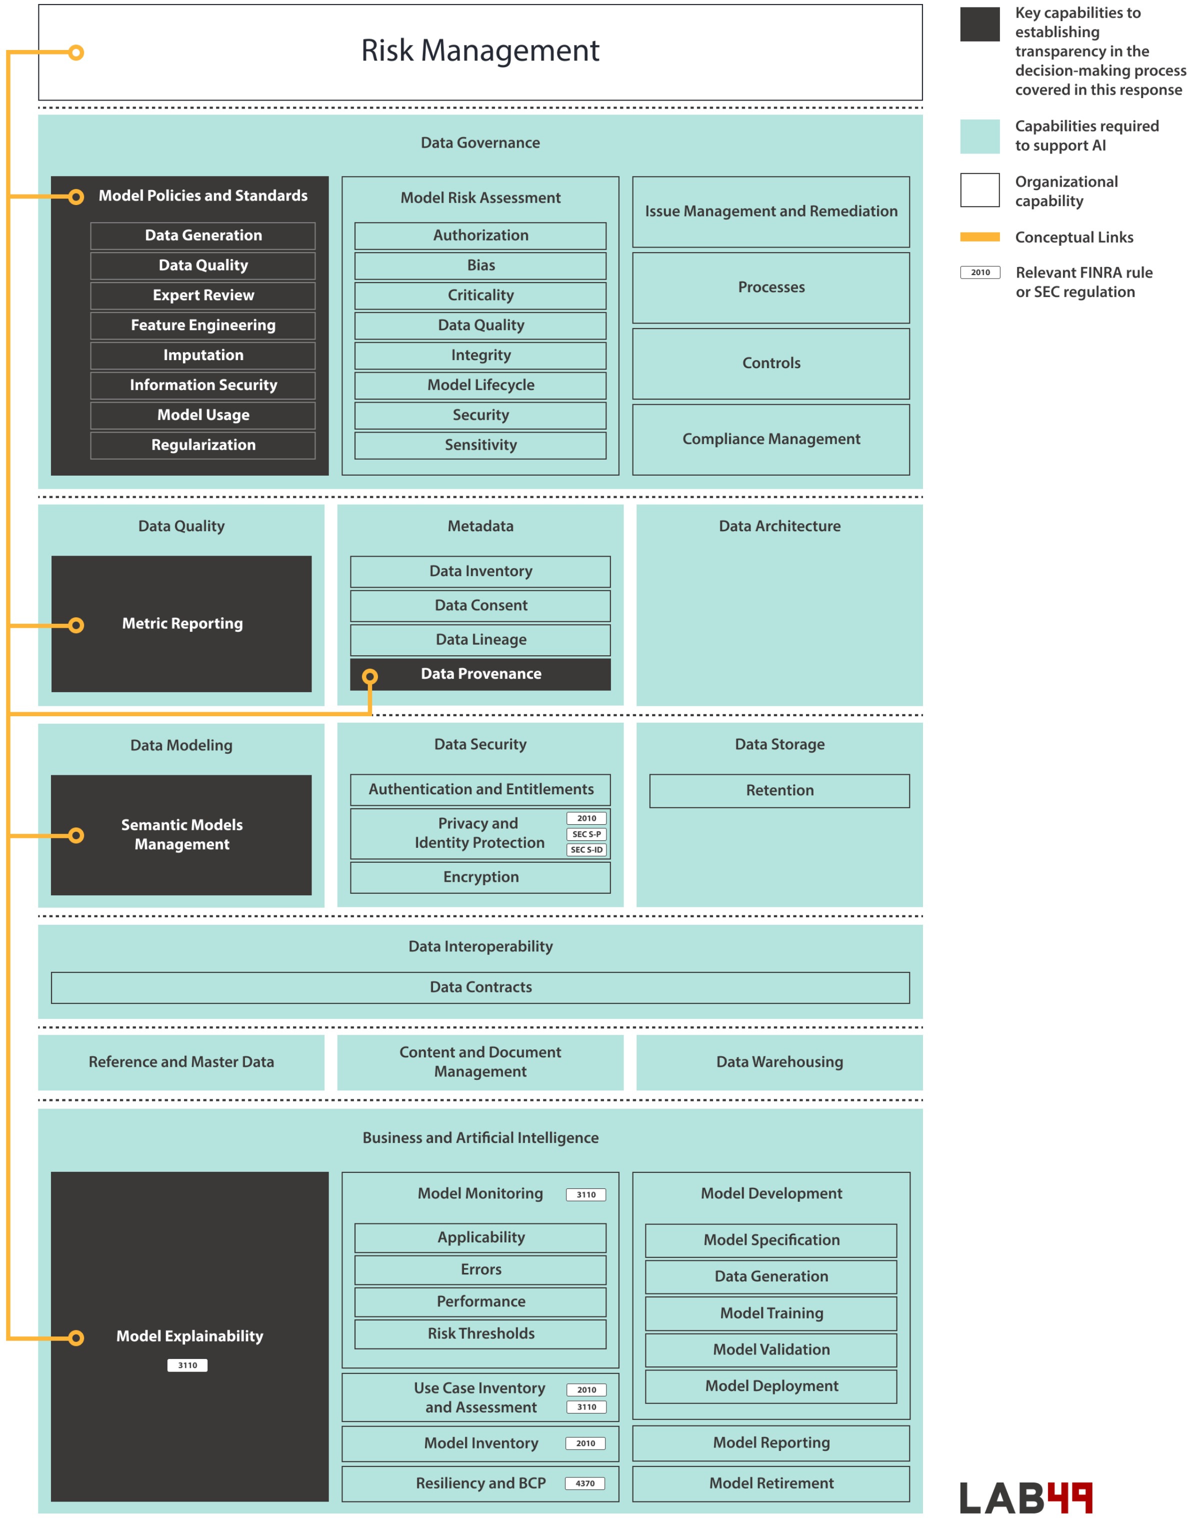

Risk Management

Data Governance
- Model Policies and Standards
  - Data Generation
  - Data Quality
  - Expert Review
  - Feature Engineering
  - Imputation
  - Information Security
  - Model Usage
  - Regularization
- Model Risk Assessment
  - Authorization
  - Bias
  - Criticality
  - Data Quality
  - Integrity
  - Model Lifecycle
  - Security
  - Sensitivity
- Issue Management and Remediation
- Processes
- Controls
- Compliance Management

Data Quality
- Metric Reporting

Metadata
- Data Inventory
- Data Consent
- Data Lineage
- Data Provenance

Data Architecture

Data Modeling
- Semantic Models Management

Data Security
- Authentication and Entitlements
- Privacy and Identity Protection 2010 SEC S-P SEC S-ID
- Encryption

Data Storage
- Retention

Data Interoperability
- Data Contracts

Reference and Master Data

Content and Document Management

Data Warehousing

Business and Artificial Intelligence
- Model Explainability 3110
- Model Monitoring 3110
  - Applicability
  - Errors
  - Performance
  - Risk Thresholds
- Use Case Inventory and Assessment 2010 3110
- Model Inventory 2010
- Resiliency and BCP 4370
- Model Development
  - Model Specification
  - Data Generation
  - Model Training
  - Model Validation
  - Model Deployment
- Model Reporting
- Model Retirement

■	Key capabilities to establishing transparency in the decision-making process covered in this response
🟦	Capabilities required to support AI
□	Organizational capability
▬	Conceptual Links
2010	Relevant FINRA rule or SEC regulation

LAB49

Risk Management

The application of AI across the organization presents significant risks. There are implications for the risk management function of FINRA member firms adopting AI as these risks require careful consideration and management. FINRA advises that member firms adopting AI should conduct their own due diligence, however member firms will benefit from guidance on the critical considerations along the dimensions of the AI capability model.

AI models may present risks to:

- Authorization: the authorization of processes running AI models to access and leverage data; the authorization to create output or trigger events for a member firm.
- Bias: the potential of the AI model to consistently learn the wrong things.
- Criticality: the importance of the model's correct operation to the firm's activities.
- Data Quality: the measure of accuracy, completeness, consistency, timeliness, availability and fitness for use of data.
- Integrity: completeness of considerations for modelling and adherence to a firm's policies and ethical code.
- Model Lifecycle: the potential for model drift, changes in the environment that impact AI models.
- Security: the potential for adversarial attacks from internal or external agents.
- Sensitivity: the use of private or material non-public information by the AI model.

Failure in any of these key risk areas may present legal, financial, operational, ethical or reputational risks to the FINRA member firms, requiring a comprehensive framework designed to address them. Table 1 highlights an indicative set of critical policy and capability considerations in each risk area.

Table 1. Indicative Set of Policy Areas

Risk Area	*Policy Areas and Capabilities*
Authorization	Feature Engineering
Bias	Feature Engineering, Data Quality, Model Validation
Criticality	Model Usage
Data Quality	Data Governance
Integrity	Imputation, Data Generation, Regularization
Model Lifecycle	Expert Review, Model Risk Assessment, Model Monitoring
Security	Information Security
Sensitivity	Privacy & Identify Protection, Expert Review, Feature Engineering

FINRA member firms need to manage risks across the model lifecycle in alignment with defined policies, but policy is not the only lever to manage risk. Member firms should complement policies with standards, processes and controls. Two critical processes for consideration are Model Risk Assessment and Use Case Risk Assessment. These processes are part of the Data Governance and Business and Artificial Intelligence capabilities within the capability model.

Model Risk Assessment

Model Risk Assessment is one of several capabilities delivering Data Governance.

FINRA member firms should consider risk identification and impact assessment of AI models at each stage of the model lifecycle. FINRA states the need for the assessment very broadly, and guidance on the required standards, processes and controls would be beneficial to member firms.

For example, understanding how a model will behave in all situations will mitigate model risk. Similar to requirements for financial risk modelling, outcomes from AI models should be validated against a range of real and generated data sets. Member firms will need to consider adopting the following into AI policies:

- Data quality controls applied to all input data;
- Data generation processes include historical and simulated market scenarios under stress conditions;
- Data generation processes consider the impact of significant but low probability events and pro-cyclical market conditions on model outcomes.

Taking a risk-based assessment approach throughout the model lifecycle enables member firms to enact controls aligned to policy. This will mitigate risks presented not only by a specific AI model, but also risks presented by changes in the environment.

Processes for risk assessment throughout the AI model lifecycle should be defined by member firms, starting with use case assessment.

Use Case Inventory and Assessment

The ideation process requires governance and oversight to ensure that the desired AI model outcomes are aligned to the member firm's risk appetite. FINRA recommends that a model inventory should be maintained. Lab49 believes that member firms should start with an inventory of use cases and link them to corresponding risks.

For example, exploration of new ideas for AI models may expose areas of a member firms' organization, data or policy that are lacking oversight. Linking use cases to risks and models will allow member firms to identify issues early in the lifecycle, reducing cost and impact.

Maintaining a use case inventory will present additional benefits for member firms and regulators by providing an auditable data set of historical decisions that can be reviewed and revisited as policy and regulatory changes occur.

Model Explainability

FINRA member firms are obligated to provide fair and ethical decisions to their customers, along with investment recommendations that are suitable and free of conflicting interests. Whether these decision processes are manual or automated, customers must be able to trust that the financial advice they receive suits their objectives and risk profile. Member firm stakeholders must be confident that their automated decision processes are ethical, fair and performant prior to placing them in production.

In order to engender this confidence, member firms should consider approaches to providing explanations as to how key decisions and recommendations are derived. Many decision models are easily interpretable, such as models based on Linear or Logistic Regression, Naïve Bayes, Decision Trees and Rules Engines. However, when these decisions are implemented via AI-driven modelling techniques, some unique challenges arise. Member firms will benefit from guidance on how to address these challenges while meeting their regulatory obligations.

The primary challenge with respect to AI explainability is rooted in the relationship between model performance and complexity – in order to improve accuracy, models necessarily get more complex. This complexity constrains both model interpretability and explainability, and may be a by-product of factors such as:

- The degree of multicollinearity among model features;
- Cognitive load created by very large feature sets;
- Engineered or abstract concepts as features;
- Black-box models and ensemble approaches.

These factors may limit explanatory precision, comprehension, or both. In addition, firms are concerned about adversarial attacks designed to trick AI models, and the risk of disclosing proprietary intellectual property when providing detailed model explanations.

In the face of these challenges, member firms that wish to leverage AI should be prepared to answer questions related to investment advice suitability and freedom from conflicts of interest in order to mitigate financial, operational and reputational risk. These explanations should be scoped to various user groups[3], including:

- Individual clients may require explanations for specific recommendations;
- Compliance may require audits of AI models to ensure that obligations are met; and
- Firm management (Business Strategy and Systems Development) may require validation that models meet business, technical and regulatory requirements.

Certain patterns emerge that address explainability in the context of model risk management. Ultimately these patterns attempt to establish a level of confidence such that AI models conform not only to performance specifications, but also meet regulatory obligations for ethical and customer-centric output.

Global Explanation

Global explanation patterns require transparency in model development. This ensures that stakeholders have visibility into the methods and techniques utilized to manage data, train models and deploy them to production. Highlighted in the capability model are several areas defined within Model Policies and Standards that should be complied with when developing AI models.

A test-driven approach that not only assesses model performance at a global level, but also utilizes targeted tests for bias and conflict of interest, will give stakeholders confidence that models meet these standards. User groups that would benefit from global explanation include compliance users and firm management.

Local Explanation

Local explanation patterns should have the capability of providing explanations for discrete model output. This capability provides individual clients with the rationale for decisions and recommendations that impact them. Accuracy in local explanation is paramount as discrepancies may lead to reputational risk with financial consequences.

[3] See Four Principles of Explainable Artificial Intelligence https://www.nist.gov/document/four-principles-explainable-artificial-intelligence-nistir-8312 (August 2020)

Implications for Model Governance

Existing governance processes should encompass AI model explainability in order to fully address risk concerns. By establishing these governance practices, member firms can meet regulatory obligations while pursuing the strategic benefits of AI. Model governance should be embedded throughout the model lifecycle with multiple touchpoints across the Software Development Life Cycle (SDLC). FINRA member firms would benefit from guidance on SDLC controls to manage model risk.

Model Specification and Development

The SDLC Requirements Analysis phase should be augmented to ensure that acceptance criteria are defined not just for model performance, but also to meet obligations related to fairness, ethics and customer-centric interests, by way of model explainability. For example,

- Protected classes should be identified for each model use case, and test cases established to ensure that training data and algorithms provide unbiased recommendations;
- Methods and tooling should be specified to enable model explanations at various levels of detail, supporting the user groups mentioned earlier; and
- Model versioning, logging and reporting requirements should be specified for audit purposes.

Once models are specified, tests should be written that exercise them across dimensions of performance and ethics.

Model Testing, Deployment and Operation

Just as models should pass strict verification criteria prior to production, they should be regularly monitored and assessed against testing thresholds. The world changes, or at least the data being fed to the model does, and model drift is inevitable. As these thresholds are approached and breached, models should be retired.

Often organizations deploy models that continue to learn from incoming data and update their parameters in an automated fashion. Governance processes should be in place to determine the appropriateness of this approach for a given use case. To be safe, each iteration of a model should be thoroughly validated prior to any change being approved.

Provenance, Data Quality and Semantic Models

Data provenance, data quality, and semantic models are important when providing rationale for decisions or recommendations leveraging AI. FINRA member firms should consider existing approaches and standards for capturing and distributing related information.

Data Provenance

Provenance describes the origins and history of a specific piece of data.

Who owns provenance information, when in the lifecycle it is captured, what granularity of provenance information is needed, where it is stored, and how it is integrated are all essential considerations for data provenance.

As the complexity of systems is growing, integration of provenance information becomes crucial. Consider the output of a sophisticated process based on data originating from multiple sources, going through numerous stages of data pipelines and performed on behalf of a human agent. Each step in this process can be using different tools and technologies, and the use of standards can help integrate the pieces of the data provenance.

FINRA member firms should consider open standards such as PROV[4] to allow integration of provenance information in a complex technology landscape, covering process, data flows, and decision responsibility.

[4] See An Overview of the PROV Family of Documents https://www.w3.org/TR/2013/NOTE-prov-overview-20130430/ (April 30, 2013)

Data Quality

Confidence is influenced by awareness of the quality of data that lead to a specific decision.

Data quality is evaluated and reported across dimensions of accuracy, completeness, consistency, timeliness, availability and fitness for use[5].

As data quality concerns become mainstream, there is a higher demand for visibility of the quality of the data that led to a specific decision. Consider a previous example of a process resulting in a decision. Its consumers, whether humans or machines, need to receive information about data quality.

FINRA member firms should consider open data quality vocabularies such as DQV[6] to report data quality in a machine-readable way that can be linked it to provenance information and definitions of data.

[5] See APRA CPG 235 Managing Data Risk https://www.apra.gov.au/sites/default/files/CPG-235-Managing-Data-Risk.pdf (September 2013)

[6] See Data on the Web Best Practices: Data Quality Vocabulary https://www.w3.org/TR/vocab-dqv/ (December 15, 206)

Semantic Models

Evaluation of provenance and quality by consumers of data depends on shared meaning.

Concepts can be defined using knowledge organization systems on a spectrum from weaker to stronger semantics. An appropriate system needs to be selected based on the use case.

Glossaries define terms to help member firms understand and share meaning of AI model inputs and outputs. Taxonomies and thesauri define hierarchies and link concepts to further help with data classification. Ontologies define classes, properties and relationships to allow for inference of new facts based on existing data.

In the earlier process example, it would be appropriate to create and publish formal definitions of related terms, process steps and roles.

FINRA members should consider the use of Semantic Web and Linked Data standards including Reference Data Framework[7], Reference Data Framework Schema[8], Web Ontology Language[9] and Simple Knowledge Organization Systems[10] to define concepts in an interoperable and reusable way.

[7] See Reference Data Framework (RDF) https://www.w3.org/RDF/ (March 14, 2014)

[8] See RDF Schema (RDFS) https://www.w3.org/TR/rdf-schema/ (February 25, 2014)

[9] See Web Ontology Language (OWL) https://www.w3.org/OWL/ (December 11, 2013)

[10] See SKOS Simple Knowledge Organization System – Home Page https://www.w3.org/2004/02/skos/ (December 13, 2012)

Closing

Lab49 appreciates the opportunity to provide comments on the "Artificial Intelligence in the Securities Industry" paper published by FINRA.

FINRA can play a leading role in the adoption of AI by the member firms by following a capability-based approach to supervisory guidance that supports FINRA's regulatory priorities.

To learn more about our perspective and how Lab49 can help, please contact this team at data@lab49.com.

Sincerely,

Kelly Attrill, Product Management, Sydney.

James Durham, User Experience, Washington, DC.

Eugene Morozov, Data and Engineering, Sydney.

Yuvraj Sidhu, Product Management, Washington, DC.