

Twelve Considerations When Evaluating AI Vendors for Financial Services



Executive Summary

Financial services companies have unique needs when it comes to AI. Large transaction volumes create both large datasets and requirements for high throughput for real-time and batch applications. These volume & time related requirements, combined with high regulation levels, require a different approach to AI than other industries. Aligning organizational needs with platform requirements ensures that AI providers have the experience and capabilities to support AI or advanced analytics programs in financial services. The twelve critical considerations when evaluating AI and machine learning platforms are:

1. Proven Experience in Financial Services

When looking at AI vendors, look for an in-depth understanding of financial services use cases. Look for experience with real-time fraud detection, customer churn prevention, personalized x-sell and upsell, credit scoring, Know Your Customer (KYC), and Anti-money laundering (AML). Ask vendors to provide information on how they solve the issue and look for vendors who don't just describe the use case but have delivered a solution to existing customers. Ask for a list of customer references and look for results from those customers to understand how the provider might impact your business.

Consider a question like this in an RFI. *Please describe the critical use cases that you support for your financial services clients. Provide as much detail as you can to explain how your products or services help your customers in the financial services industry. Where possible, please provide specific information for how your platform or solution solves the issue and provide results metrics.*

2. Third Party Validation

The AI vendor landscape includes legacy vendors, startups, and service providers. To ensure that you are looking at the best possible vendor to fit your needs, ask for third-party

research that validates the vendor as a viable and robust player. Example responses would include significant analyst ratings like inclusion in the Gartner Magic Quadrants and Forrester Waves for Data Science and Machine Learning Platforms. Vendors that are not in these analyst reports may either be too small, not growing with the market, or may not have legitimate AI capabilities. In particular, be careful to distinguish service companies masquerading as AI platforms from actual platform providers. These companies will sell AI solutions but will not have unique AI technology of their own which could leave organizations exposed to unseen risk over the long-term. Using third party analysts is a good first cut at the providers you should consider.

Consider a question like this in an RFI. *Please describe any third-party validation for your company or products, including awards and inclusion in industry analyst reports*

3. Services and Support

To better understand the vendor's business operations, it pays to ask how they service and support customers. First, make sure they have a plan for implementation and training with details for the resources needed from your team. After a successful implementation and rollout,

make sure you understand how the vendor will support your ongoing projects and the cost. Generally, you want to educate your organization to provide maximum agility and flexibility and ensure you don't become dependent on their resources to scale.

Consider a question like this in an RFI. *Please describe how your company provides service and support to customers. What services do you offer for implementation, onboarding, training, and ongoing support?*

Key Technical Capabilities

The AI vendor landscape is broad with legacy vendors and new entrants. Having a vendor explain their fundamental capabilities will allow you to understand if that vendor aligns with your needs and resources. In particular, ask the vendor to describe their key differentiators from other players in the market. If there are key capabilities that you require, note them, and ask vendors to describe their offerings. The following are some technical considerations:

4. Is Open Source Right For You?

Open Source - You may not have considered open-source platforms in the past. With the proliferation of new algorithms and techniques, open-source providers can often provide the best balance between innovation and resource availability. The question for open source only providers comes down to their business model and viability. Look for vibrant open source communities with hundreds of thousands of members. These large groups will not only make sure to adopt the latest techniques and build them into the product where you can take advantage of them. A community also ensures that you can find the talent to run the system for your organization. With commercial only platforms, the services and roadmap are in the complete control of the vendor.

Closed Source - Anybody who doesn't mention open-source is probably a commercial-only platform. Financial services companies have extensive experience with legacy analytics platforms that have created vendor lock-in and cost millions of dollars per year in licensing and maintenance fees. Whatever technology you choose, make sure you understand the total cost of ownership over time and how it will adapt to changing business needs. With commercial only platforms, the services and roadmap are in the complete control of the vendor.

Both Open and Closed Source - Some vendors will offer a blend of open source and commercial products. The open-source component can provide access to a low-cost data science platform or algorithms. In contrast, commercial offerings can provide differentiated capabilities like automation and operations. With such a hybrid offering, it is up to your business if you want to engage with the commercial offerings or use the open-source.

Consider a question like this in an RFI. *Please describe how your product uses open-source algorithms and technology. Does your company participate in the development of open-source products, and if so, which ones? Does your company have an established community, and if so, how many people are in your community?*

5. Multi-Cloud, On-premise, or Hybrid Infrastructure

Where are you going to deploy your AI platform, and where do you plan to deploy your models? If you know you need to deploy software on-premise or in a virtual private cloud (VPC), only vendors who provide those options will meet your needs. Make sure that the vendor has complete support and not just integration with part of your preferred platform. For example, the vendor may have a data connector with cloud infrastructure, but may not support storage or compute options from that provider. The limited integration will likely have cost, performance, and scale implications if you are standardized on that platform.

Consider a question like this in an RFI. *Describe how you support <insert your environment here>. What services or products do you specifically support? Please provide examples where your customers have used these services.*

6. Automated Machine Learning (AutoML)

AutoML allows data science teams to explore more options and build more projects by automating key data science processes like feature engineering, algorithm selection, and model tuning. While AutoML was somewhat limited a few years ago, several fully-featured products exist in the market today. These products can also accelerate AI model development without having to grow your data science team. Of particular interest in financial services is the agility that AutoML provides. When market conditions change, AI

models need to be updated and tested quickly. Using automation can help discover new models rapidly and get your organization back to business faster than a purely artisanal approach.

Consider a question like this in an RFI. *Does your platform include AutoML capabilities? If so, how is your AutoML different from other automated machine learning systems. Does your AutoML have any features that are of particular interest in the financial services industry? Does your AutoML allow for model building on large datasets?*

7. Big Data

In financial services, the volume of transactions leads to enormous datasets for AI training, often terabytes in size. Training on big data is especially critical when the incident you are trying to predict is small compared to the overall data volume. For example, in payments fraud, the number of fraudulent transactions is between 15 and 50 basis points. With such sparse incidents, a sample may contain only a few fraud examples, leading to low accuracy.

Consider a question like this in an RFI. *How does your platform handle large datasets terabytes in size or larger? If you use sampling, how is that supported? Can your system build models on datasets exceeding a terabyte?*

8. Scoring speed and volume

For many use cases in financial services, like real-time fraud detection, scoring happens during the transaction in milliseconds. During the holidays, for example, the number of payment transactions can spike up dramatically. AI vendors in financial services need to provide for high-speed scoring at high volume. When talking to vendors, make sure they understand the scale of predictions required and guarantee the response times in milliseconds for predictions at scale. In this case, speed has serious implications for your top & bottom line.

Consider a question like this in an RFI. *How does your platform handle scoring for high-volume transactions? What kind of response time can customers expect from your prediction/scoring mechanism? Do you have examples you can provide of high-volume, low-latency scoring for financial services?*

9. Explainability

In financial services, models go through rigorous validation and certification for regulatory compliance. A model that is not explainable to the business or regulators cannot be used to support regulated business functions. Vendors should have explainability and interpretability techniques for a range of model types. Vendors should also provide model documentation that includes experimentation options, model profiles, and explainability tests and results. For predictions, vendors must provide reason codes for each model response to explain each prediction. Explainability is particularly vital in credit scoring, where customers now have the right to inquiry for denied credit. With reason codes, an agent can quickly pinpoint credit denial's critical factors, like missing payments in their credit history.

Consider a question like this in an RFI. *Please describe how you support explainable AI. What techniques do you offer within your platform to explain models and their outputs?*

10. Model Fairness and Bias Detection

Bias in models is a well-documented issue of particular concern to financial services companies related to fair lending and other practices. While responsible AI practices ultimately reside within each organization, technology can help find systematic bias in data and models before those models impact consumers. Techniques such as disparate impact analysis can show where models are predicting different results for groups of customers. When such groups' characteristics include location, gender, or other protected items, the analysis shows a bias inherent in the model either for or against that group.

Consider a question like this in an RFI. *Please explain how you support responsible AI practices with your customers. What technology do you provide to support bias detection and model fairness?*

11. Production Scoring

If organizations can't operationalize models in production, they cannot provide value to the business. Understanding the capabilities that vendors offer to put models into production and maintain them is essential to success with AI. Operationalizing models involves two key areas. First, how does the vendor encapsulate the scoring or runtime model for production? A simple file of the scoring model is necessary but insufficient. For scoring, the file will need to run in a production environment. The vendor may provide such an environment as part of their product, often called a scoring server. An alternative to the embedded server is a stand-alone object that customers can deploy on their infrastructure. This stand-alone option provides for more control and supports more deployment scenarios.

Consider a question like this in an RFI. *What capabilities do you offer for production scoring? Where does your platform deploy models?*

12. Application Integration

Integration of production model services into business applications where end-users interact with predictions is crucial for models to provide value. Understanding the capabilities that the vendor offers to integrate models or even create applications is vital in understanding how quickly you will generate value from your AI investments.

Consider a question like this in an RFI. *Please describe how AI models integrate into business applications. Do you provide any AI applications as part of your offering or capabilities to build AI applications?*

H2O.ai is an open source leader in AI with a mission is to democratize AI for everyone. H2O.ai is transforming the use of AI with software with its category-creating visionary open source machine learning platform, H2O. More than 20,000 companies use open-source H2O in mission-critical use cases for Finance, Insurance, Healthcare, Retail, Telco, Sales and Marketing. H2O Driverless AI uses AI to do AI in order to provide an easier, faster and cheaper means of implementing data science. H2O.ai partners with leading technology companies such as NVIDIA, IBM, Intel, AWS, Microsoft Azure and Google Cloud Platform and is proud of its growing customer base which includes Capital One, Progressive Insurance, Comcast, Walgreens and MarketAxess. For more information about how H2O.ai is driving an AI transformation, visit www.h2o.ai