FI·CONSULTING

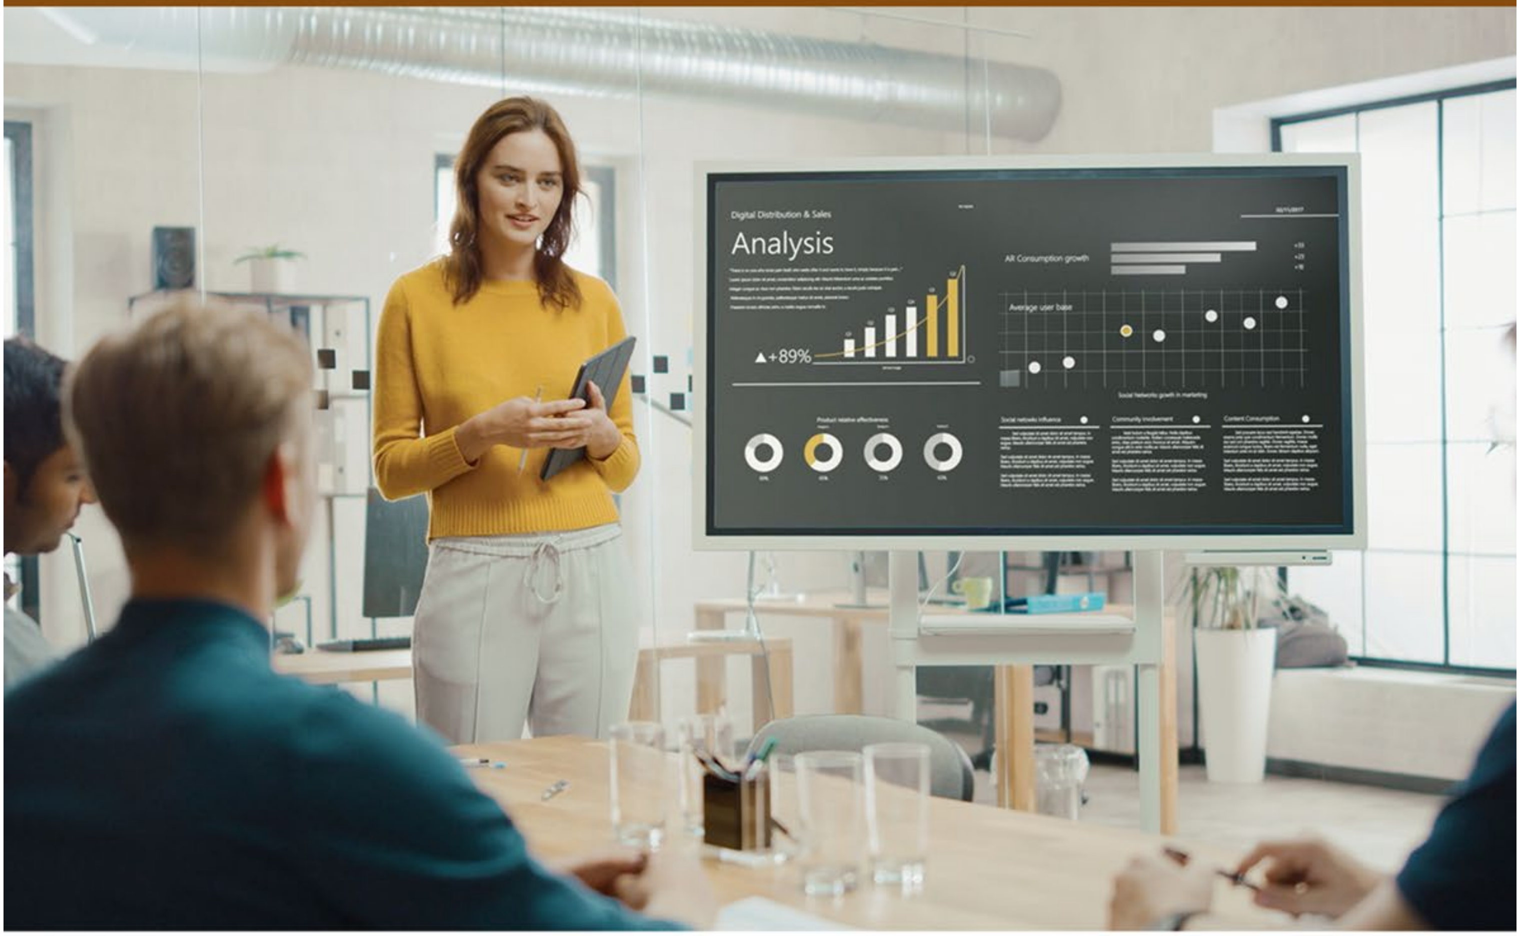

Fintech in Housing Finance: Request for Information

Thursday, October 13, 2022

Federal Housing Finance Agency
Office of Financial Technology
400 7th Street SW, 5th Floor
Washington, D.C. 20219

POC: Andrew T. Eastman, Chief of Staff
Phone: 571.255.6771
eastman@ficonsulting.com

UEI: KA1KEJ7MNSY1
CAGE: 4MEL7
TIN: 73-1666529
GSA MAS: GS-00F-346CA

Submitted electronically via
https://www.fhfa.gov/AboutUs/Contact/Pages/Request-for-Information-Form.aspx

Contents

1 Executive Summary

FI Consulting, Inc. (FI) is a government and commercial prime contractor with a successful history of helping the Federal Government, Mortgage Associations, and Commercial Lenders manage their complex portfolios. For 20 years, we have supported customers such as the Federal Housing Finance Agency (FHFA), US Department of Agriculture (USDA) Rural Development (RD), the Federal National Mortgage Association (Fannie Mae), Federal Home Loan Mortgage Corporation (Freddie Mac), Government National Mortgage Association (Ginnie Mae), and the US Treasury improve their financial technology processes, tools, and environments.

We provide these agencies and organizations with data analytics, model development, policy analysis, and technologies that improve their loan, grant, and financial portfolio equity and performance. For example, our work significantly reduces the time analysts spend creating, updating, and producing reports by automating import and export procedures, and modernizing legacy models. We differentiate ourselves from other services providers in the following ways:

- We provide delivery teams where every member brings expertise from data, statistical programming, and technology domains. This allows our teams to peer review each other's work before formal testing or delivery which leads to fewer defects.
- We are experts in helping program, procurement, finance, and IT staff collaborate effectively—this saves time and money throughout our engagements by ensuring all parties understand requirements, timing, and dependencies.
- Our work consistently stands up to scrutiny from internal and external stakeholders, including market participants, Office of Management and Budget (OMB), Office of the Inspector General (OIG), and financial statement auditors providing peace of mind that government or market participants will not find errors with their reports and data products.
- We deliver solutions and training so that our customers can own and update deliverables after our contract ends saving money by reducing or eliminating the need for future operation and maintenance contracts.

We believe these qualifications apply specifically to FHFA's commitment to increase the liquidity of mortgage investments and improve the distribution of investment capital available for mortgage financing. Our ability to apply financial technologies can be focused on simplifying, streamlining, and spurring efficiency in accessing financing.

Our response to FHFA's request for information focuses on our suggestions to the questions in Section C: Equitable Access to Mortgage Credit. As you will read, we apply the use of graph intelligence to the FHFA challenges in Section C. We are also providing a White Paper developed in collaboration with one of our strategic partners, Katana Graph (www.katanagraph.com). Further information on our credentials and qualifications can be found at www.ficonsulting.com. If you have any questions or concerns with the information provided in our response, please feel free to contact me at 571.255.6771 or eastman@ficonsulting.com.

With regards,

Andrew T. Eastman, Chief of Staff

1.1 Equitable Access to Mortgage Credit

General Questions on Equitable Access		
Question C.1	What new fintech tools and techniques are emerging that could further equitable access to mortgage credit and sustainable homeownership? Which offer the most promise? What risks do the new technologies present?	**A growing number of financial institutions that employ AI/ML techniques in credit risk models are finding that using graph intelligence platforms to accelerate feature development can materially increase the accuracy of their models.** **Graph intelligence platforms are an evolutionary step forward from graph databases. Graph platforms must be scalable and able to handle data pipelines in their entirety to drive significant efficiency increases in data model creation and production. A robust graph platform unifies the following capabilities:** ▪ **Graph databases to store and manage data** ▪ **Graph neural network compute engines to train models** ▪ **Model serving and inferencing services to deploy models** **Graph intelligence platforms have proven to be useful in automatically developing features for use in AI/ML models. Traditionally, consumers with high credit scores have received multiple low-rate credit options, while underscored-individuals face limited and expensive options, despite being equally creditworthy.** **Because graph platforms can store connections between entities, a graph of a given set of data contains much more information than the same data stored in tabular form. Graph algorithms can then be used to automatically develop a much larger set of features from a given dataset. This larger set of features then makes the AI/ML models more accurate and reliable.**
Question C.2	What emerging techniques are available to facilitate or evaluate fintech compliance with fair lending laws? What documentation, archiving, and explainability requirements are needed to monitor compliance and to facilitate understanding of algorithmic decision-making?	**Most AI/ML models are difficult to explain because they are computing quantities of data and finding connections between those data points in a way no human could ever do on their own. Common ways in which regulators think about explainability include:** ▪ **Global interpretability - Understanding how a model works holistically, for example, asking which features contributed the most to the model's decisions.** ▪ **Local interpretability - Why did the model make this specific prediction? For example, lenders will send adverse action notices to explain why a particular applicant was denied credit.** ▪ **Contrasts and Counterfactuals - If an applicant is denied credit, stakeholders might want to know what needs to change for underwriters to approve the applicant.** ▪ **Partial Dependence - Isolating the specific contribution of a feature (or set of features), by perturbing the feature and holding the other features constant.** **We use the following model-agnostic techniques for explainability that can be used on any machine learning model, no matter the complexity:**

		▪ **PDP (Partial Dependence Plots) -PDP explains the global behavior of a model by showing the relationship of the marginal effect of one or two predictors on the response variable. PDP plots are very intuitive, and easy for stakeholders to understand as they provide a visual representation of how one or two features influence the predicted outcome of the model, while marginalizing over the remaining features.** ▪ **ICE (Individual Conditional Expectation) - ICE is a more granular view of a PDP plot, while PDP plots show the average prediction across all instances, ICE plots show the prediction for each instance individually. A PDP is the average of the lines of an ICE plot. As ICE plots display more information, they can provide more insights.** ▪ **SHAP (SHapley Additive exPlanations) - Shapley values are a concept from cooperative game theory to split the reward from a game fairly among different players. The analogy for model interpretation is that Shapley values can be used to split the output of a model fairly among different features. Shapley values are additive, allowing us to decompose any prediction as a sum of the effects of each individual feature value.** ▪ **ALE (Accumulated Local Effects) - ALE is similar to Partial Dependence in that they both aim to understand how a feature influences the response variable on average. However, ALE helps mitigate the problem of correlated variables that leads to incorrect interpretation of Partial Dependence Plots. ALE mitigates this by using only the conditional distribution of a feature and mitigates the problem of correlated variables by using differences in predictions instead of averages.**
Question C.3	Are there effective ways to identify and reduce the risk of discrimination, whether during development, validation, revision, and/or use fintech models or algorithms? Please provide examples if available.	**A key benefit of AI/ML graph-based feature development is that the models ignore systemic biases; they are only looking for connections that relate and can be analyzed, which leads financial institutions to be better able to find people who have good risk based on others who have similar features. Relying primarily on traditional metrics, such as credit scores, excludes a large fraction of Americans from credit markets altogether.** **Graphs are uniquely positioned to find meaning in certain types of data that other models would find too irregular to be useful. Graph platforms also capture data sets that contain sparse amounts of data better than other databases and platforms. This is very pertinent for credit risk modeling, as companies want to be able to look back over the entire credit history of an individual, but there may be gaps in this data. Or, they may be examining an applicant with a limited credit history. It is difficult for other platforms and traditional feature development approaches to make sense of such limited data sets, whereas graphs allow users to get more signal, and therefore more value, out of all available data** **Another key benefit of employing graph technology for credit data is that much more data can be added to graphs than when conventional feature development methods are used. For instance, with graphs, utility, phone, and TV bill payments or deposit account information, can be their own features, benefitting applicants with non-traditional credit histories.**

		Graph feature development allows models to have greater context and reach more informed conclusions about the credit risk of borrowers. Graphs could extend opportunities to recent college graduates starting their first job, or recent immigrants, who despite being potentially creditworthy, might exhibit low credit scores, and help them obtain their first credit.

Introduction

A growing number of financial institutions that employ AI/ML techniques in credit risk models are finding that using graph intelligence platforms to accelerate feature development can materially increase the accuracy of their models.

Graph intelligence platforms are an evolutionary step forward from graph databases. Graph platforms must be scalable and able to handle data pipelines in their entirety to drive significant efficiency increases in data model creation and production. A robust graph platform unifies the following capabilities:

- Graph databases to store and manage data
- Graph neural network compute engines to train models
- Model serving and inferencing services to deploy models

With such a platform, data scientists are more efficient working on specific graph neural network models and can use the tools they are already comfortable with (such as Jupyter Notebooks and Python libraries). As a result, data scientists spend far less time to productionalize graph-based approaches. This is highly beneficial in the world of credit risk modeling.

Graph Platforms and Credit Risk Modeling

Creating effective credit risk models is one of the most important tasks that data scientists perform for financial institutions. This sort of modeling has taken place for generations and generally is performed with advanced, state of the art techniques.

Credit risk modeling gets so much attention and investment because the stakes are extremely high — it is a core part of the foundation of how financial organizations operate. If a financial institution (like a bank, a credit card company, or a commercial lender) can better understand which borrowers are lower risk and which are higher, their lending businesses become more profitable. Just a small improvement in the effectiveness of credit models can improve financial performance by large amounts. Depending on the scale of the business, this can mean millions or even tens of millions of dollars in increased revenue, as well as significant savings from not making bad or non-performing loans.

Over the recent decades, because of the need to accurately predict risk, financial institutions have been at the forefront of deploying AI/ML models in production. Models such as XGBoost, LightGBM, and other techniques have been improving performance substantially over the techniques used in previous generations.

Graph intelligence platforms have proven to be useful in automatically developing features for use in AI/ML models. Because graph platforms can store connections between entities, a graph of a given set of data contains much more information than the same data stored in tabular form. Graph algorithms can then be used to automatically develop a much larger set of features from a given dataset. This larger set of features then makes the AI/ML models more accurate and reliable.

To perform this automatic feature development, graph platforms must be able to achieve levels of scalability that eluded the first generation of graph technology.

This paper will explore this story in more detail by answering the following questions:

- Why are small increases in credit risk modeling so important?

- Why are graph platforms able to create more features out of a given dataset than traditional methods?
- Why does a larger set of features improve performance of AI/ML models?
- How does graph-based feature development change the process of creating AI/ML models?
- What are the requirements of a graph intelligence platform to do automatic feature development?

A Little Boost Goes a Long Way: The High Stakes Game of Credit Risk Modeling

Given what is involved in assessing whether a potential borrower is a worthy credit risk, it should be no surprise that credit risk modeling is extremely difficult. Models have to take into account an astonishing number of factors, from credit history to borrowing behavior to debt levels to future potential earnings of similar applicants, when determining the credit risk of an individual or business.

Due to the onerous nature of identifying credit risk accurately, financial institutions are constantly looking for ways to gain even a slight advantage over their competitors. But even small improvements matter: a model that is even 3-4% more accurate than another model can result in millions of dollars of additional revenue or avoided losses when it is applied to a significant array of applications. Financial institutions are always looking for ways to make their models even marginally better because of the fact that this can lead to outsized results.

Using graph-based feature development can produce these types of significant advancements in AI/ML models. This is especially true for companies that are just starting with AI/ML models — graphs can allow them to leapfrog their competition. Here are some of the key reasons graph platforms are specifically beneficial for credit risk modeling:

- The speed with which graph-based feature development can provide insights. Speed and efficiency are of paramount importance when it comes to credit risk modeling. Often, decisions on whether to lend money must be made in a matter of minutes, or even seconds. For instance, if a customer wants to apply for a store-backed credit card while making a purchase, his or her credit worthiness must be assessed immediately — while they're standing at the register. Any latency that occurs during the transaction could lead to the business losing out on offering a credit card to a valuable customer with a dependable credit history. Graph platforms meet the need of improving the modeling, so that decisions can be made faster down the line.
- Quicker results. Generally, when someone applies for a credit card, models can incorporate 18 to 24 months of their credit history. With a graph, this history will be built in 18 separate graphs. Using tabular models to create and analyze this array of data is cumbersome and time-consuming. With graphs, it's surprisingly easy and efficient. Graph technology can assess how an applicant's creditworthiness and history has evolved over that one- to two-year period. Then, once the model is trained, insights can be generated with alacrity, which can make all the difference in such transactions.
- Graph platforms provide far greater depth of analysis than is possible with traditional modeling. Essentially, all modeling is a form of pattern recognition. AI/ML credit risk modeling is examining patterns of the individual applicant, as well as the behaviors of applicants like them, to produce a judgment on the person's credit risk. With graph-based feature development, based on the very

nature of the way graphs work, connections are made between data points that would likely remain invisible with tabular data stores.

- Graph platforms drive incremental improvements to models over time. While incremental improvements might sound modest, as mentioned earlier in this section, in the credit risk modeling business, small gains can produce exponential progress, offering companies a notable return on investment. Uninformed financial institutions may think that the differences between a simple, inexpensive model that is 75% accurate, and a more expensive, complex model that is 80% accurate are insignificant. But the credit space is hyper-competitive, and thus, over time, this 5% increase can lead to massive financial yields.

The Power of Graph-Based Feature Development

Given the claims made in the previous section about the power of graph platforms to train AI/ML modeling for credit risk, it's worth expanding on how and why graphs produce more accurate models. They do so primarily through graph-based feature development.

Simply put, a feature is a dimension or aspect of a given entity. For instance, the features of a 40-year old middle manager who lives in Nebraska would be his age, occupation, income, and geographic location, but also his credit and loan history, whether he owns his house, and if so, how many years he has left on his mortgage, as well as other factors like whether he has student debt and his purchasing behavior. These are just a small glimpse of the features that need to be accounted for when determining if he is credit worthy or not.

Of course, it's not enough to know just this man's features — these have to be coupled with features of other people like him and the reliability of the financial institutions where he already has lines of credit. Graph platforms can show the links in these types of datasets by:

- Showing connections: Graphs help to connect features and show the relationships between people and the data. In essence, feature development allows the models to have greater and greater context to reach more informed conclusions about the credit risk of borrowers.
- Discovering new relationships: What makes graph-based feature development uniquely powerful is that users can go beyond just these features and use graph algorithms to identify new, previously undiscovered relationships which then create additional features that help to separate the risk of one person from the next.
- Revealing relationships automatically: Graph-based feature development reveals significant relationships automatically, unlike traditional feature development. In graph-based feature development, algorithms analyze the graph and then come up with features that are encoded into the vector representing each entity being scored by the model, in this case borrowers. Because the development occurs automatically, graphs can create hundreds of features, whereas when data scientists work without an algorithm, they face human limitations and could only produce a small fraction of the insights of graph technology. Humans simply cannot make sense of the amount of data available on their own when it comes to credit risk modeling. Graph platforms make the work of data scientists more efficient.
- Ignoring biases: Another key benefit of AI/ML graph-based feature development is that the models do not care if features are hand-crafted; they are only looking for connections that relate and can be analyzed, which leads financial institutions to be better able to find people who have

good risk based on others who have similar features. Thus, graphs can make traditional feature development more powerful and graph queries can reveal more of the data. With automatic feature development, graphs allow companies to increase the dimensionality of their AI/ML models.

- Effectively handling small data sets: Graph platforms also capture data sets that contain sparse amounts of data better than other databases and platforms. This is very pertinent for credit risk modeling, as companies want to be able to look back over the entire credit history of an individual, but there may be gaps in this data. Or, they may be examining an applicant with a limited credit history. It is difficult for other platforms and traditional feature development approaches to make sense of such limited data sets, whereas graphs allow users to get more signal, and therefore more value, out of all available data.
- Finding meaning in a wider range of data: Graphs also are uniquely positioned to find meaning in certain types of data that other models would find too irregular to be useful. An example of this is with tradeline data, which is often too haphazard for standard modeling techniques. Only by putting tradeline data in the graph and then generating embedding vectors can companies actually leverage that data.
- Increasing the amount of data that can be used: Another key benefit of employing graph technology for credit data is that much more data can be added to graphs than when conventional feature development methods are used. For instance, with graphs, time can become a dimension. Or, every line of credit that a consumer has can be its own feature. To give an example, this would be relevant when categorizing auto tradelines that have been more than 30 days delinquent in the past year. In this case, there may be a massive amount of data on those tradelines and payment history, but a graph will learn iteratively how to represent that information in a vector. The model figures out the best way to encode that information and how to represent it in a graph, then showing it as a single feature vector.
- Making automation easier: The automation of this type of analysis is vastly easier than traditional methods in which analysts would have to write queries or a program to understand the calculations necessary to come up with results. With a graph platform, there's a node or a connection for each piece of data. The connections graphs store help to show the maximum signal that differentiates the credit risk of people from each other.

All of these factors contribute to graph-based features that are more effective than conventional approaches because they offer a true 360 degree, unified view of the consumer. With graph platforms, users get more information because they can create relationships between the data and also have more automation of feature creation. The data is also highly reusable across other models, meaning that one encoding of a vector could be used by more than one model.

That's not to suggest there are no drawbacks to graph-based feature development. The most notable challenge to the approach is that the complexity of the algorithms that are used and the automatic feature development can raise questions of explainability. At times it may not be clear what a feature represents so data scientists might not be able to decipher the rationale behind the conclusions the models are reaching. Actually constructing and thinking in graphs can also be challenging for some data scientists to adjust to.

Yet, ultimately, graphs provide tremendous value to businesses. While companies can get a lot out of their data without running any type of AI/ML modeling, they gain more leverage and power out of their graph platforms by using the graph algorithms that are inherent to the data. Graphs help to produce powerful models and standardize data in an intuitive way users can work with — even if that data is irregular, like with credit tradeline data. The result is that users and companies extract more value from their data when using graphs than they would otherwise and extend their insights to alternative data sets, such as utility payments, to achieve ever greater insights into the behavior and credit worthiness of individuals.

AI/ML Models Thrive on Larger Feature Sets

Another main reason that graph platform-based models are ideal for credit risk modeling is that they can handle huge amounts of data in a way that other modeling approaches cannot. This matters because AI/ML models in general perform better and are more accurate when they have large sets of data to analyze.

The greater the variety and quantity of data that can be fed into an AI/ML model, the better its predictions can become. The reason for this is that more data provides ever richer context, and covers ever more potentialities for the models to examine.

Graph intelligence platform modeling thus fits perfectly into this paradigm as the increased dimensionality of graph-based feature development is ideal for high-quality AI/ML performance. Some of the reasons for this are as follows:

- Graphs use large vectors that provide deeper context for the models. The result is that these models work better with graphs and can generate incisive predictions about the future. For instance, when it comes to credit risk, because graphs show connections between people, if many people suddenly start defaulting on their loans, the graph-based vectors can then flag as potential risks other individuals who have similar characteristics to those going into default. This is only possible with graphs because of the way it surfaces connections between people and data points.
- Graph platforms are uniquely situated to address the specific issues inherent to credit risk modeling. One of these issues is dealing with the extensive number of credit lines a single individual can have — like a person with two mortgages, 12 credit cards, and three auto loans. Graphs create a customized fingerprint for individuals with these specific types of tradelines that can be differentiated from others — such as a person with two credit cards and one student loan, but who is only renting at that point in his life. Every person gets mapped uniquely with graphs and then placed nearest those other people with the most similar credit histories, to provide a more accurate view of their creditworthiness.
- Graphs are particularly useful for fraud detection, a major component of credit risk modeling. When used for fraud detection, modeling predictions are mission critical for lenders. Even a small improvement in fraud detection accuracy can lead to significant value for a business. Graphs are ideal for fraud detection modeling because they can handle huge volumes of data so well. More accurate fraud detection leads to long-term ROI and cost savings, as fewer bad candidates receive loans, leading to less future defaults. But additionally, the model will also be better able to identify good bets, leading to approvals of more trustworthy borrowers and greater revenue over the long-term. As an example, we have found that just a four point accuracy improvement in credit risk modeling can lead to a 30% decrease in defaults.

How Graph-based Feature Development Changes the AI/ML Modeling Process

Another core benefit of graph-based feature development is how it changes the AI/ML modeling process. Compared to other forms of modeling, with graphs, data scientists focus less on discovering features directly, and more on actually using the features the models have generated to make accurate predictions. Data scientists are thus freed up to help the business get more out of its data.

Graph Platforms Aid Data Scientists in Improving Efficiency

With graph intelligence platforms, data scientists have to change their approach to their work. The construction of graphs becomes a crucial skill because of the connections that graphs unlock. Data sets must also be selected for their value in creating these connections.

Fortunately for the data scientist, more data than ever before can be added to the models because the data scientist doesn't have to make complete sense of the additional data on her own to create new features — instead the automatic graph algorithms will discover the features. Ultimately, graph-based feature development allows data scientists to improve their understanding of the value of the data, making the models ever more powerful. This type of feature development is only truly available through a graph intelligence platform, rather than just a graph database.

Compare that to XGBoost modeling, where data scientists are simply trying to measure performance. Graphs provide far greater depth than this type of modeling, and, as mentioned earlier, are uniquely situated to create accurate models even in environments with sparse data.

Handling Raw, Incomplete, Messy Data

In tabular approaches, data scientists have to spend far more time cleaning and imputing missing values in order to work with the data. In graphs, though, they can better work with the data in its raw form rather than doing that kind of complex summarization because the graphs can handle the missing data.

Graph Platforms Help to Improve Compliance

The one major drawback of modeling in general is that data scientists do lose the ability to completely explain the reasoning behind the findings the models produce. Most AI/ML models are a bit of a black box because they are computing quantities of data and finding connections between those data points in a way no human could ever do on their own. Graph-based feature development makes explainability harder at times when it is not clear what the automatically discovered features mean.

Explainability is incredibly important in the credit space because credit risk models are so important to the lives of people, that they should be explainable in some way — as well as because of the regulations and laws preventing lenders from not providing credit lines to people based on discriminatory rationales.

However, graphs can help companies comply with these regulatory and moral requirements. Explainability is an issue not just with graph modeling, but any form of credit risk modeling. Graphs essentially make playing this difficult game of chess easier — but it's still a difficult game. Ultimately, what companies should keep in mind is that all models are black boxes in some ways, and graph models do not suffer from a unique lack of explainability compared to these other modeling approaches.

Why New Levels of Scalability Are Required to Perform Automatic Feature Development

It should be clear that AI/ML models operate most effectively with vast amounts of data. But that also requires any database that is using these models to be able to scale and do so rapidly.

Not all graph intelligence platforms are made the same. Many graph platforms simply cannot scale and therefore cannot handle the avalanche of data necessary to have AI/ML models become as accurate as companies in credit risk modeling need them to be. Additionally, because the datasets are so large, companies need a graph platform that can do feature development automatically. Only scalable graph intelligence platforms can handle the volume necessary to make credit risk AI/ML modeling work efficiently.

Scalable graph intelligence platforms should aid the user and especially data scientists in being able to get data into the graph format without unnecessary logistical hurdles. One very effective way to do this is to use data frames. These data frames are incredibly straightforward and easy for data scientists to work with. It's simple to turn a data frame into a graph and a remarkably easy way to transform tabular data into a graph structure.

Conclusion

Modern graph platforms should improve the efficiency of the data science team to create models and push them into production. It is important to combine the necessary capabilities for effective and efficient modern credit risk modeling into a single, unified platform.

As mentioned earlier, with that platform, data scientists save considerable time iterating on the specific graph neural network models. Graph platforms are hugely beneficial to the work of credit risk modeling specifically and data scientists in general. Graph intelligence platform capabilities can lead to a dramatic reduction in the time data scientists spend creating and productionizing models, data scientists can often eliminate nearly 80% of the time they've spent on creating models in the past. Additionally, because of the graph intelligence platform's range of capabilities, the insights and models data scientists can produce are far more reliable and valuable than ever before.