



Data Science White Paper

An AI Approach for Consumer Lending Approval & Pricing

An AI Approach for Consumer Lending Approval & Pricing



Machine Learning

Mosaic data scientists collaborate with customers, digging deep into the data to inform design and deployment of custom ML tools that make a difference.



Artificial Intelligence

Mosaic integrates powerful AI tools into clients' existing technology stack to solve complex business challenges



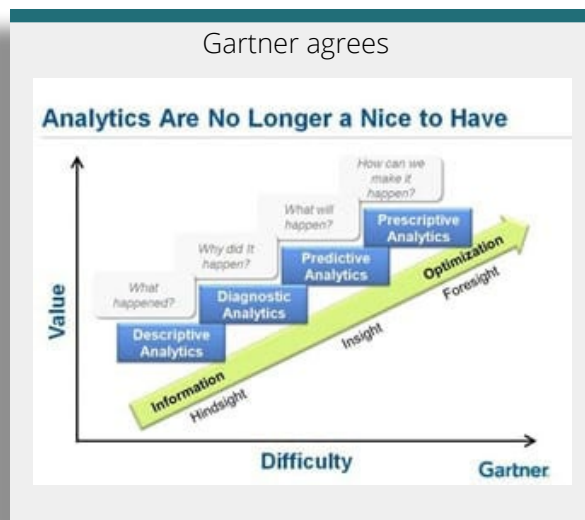
Business Analytics

Mosaic helps corporations of all shapes and sizes take advantage of their data, transforming their decision-making processes.

BACKGROUND

Artificial Intelligence (AI) has transformed how many make in the corporate world make decisions. Traditional lending practices are a prime candidate for machine learning (ML) improvements. Lenders can make more accurate and faster decisions by shifting decision-making from analysis of individuals to analysis of trends and patterns. This leads to more repeat business and lower operational costs.

According to the [Lending Times](#)¹, if you are not informing these decisions with ML & AI, your competition surely is.



HOW ML CAN HELP

Data scientists can build automated pricing and decision credit models which ingest data and decline applications when risk factors hit a certain threshold, and recommend terms based on the loan risk factors. Most current decisioning models require substantial use of manual inputs and decisions. The promise of AI-based automation takes specific funding requests from a business and makes an approve/decline decision based on attributes and historical data. This model should also recommend loan terms such as pricing and payback period.

DEVELOPING AN AUTOMATED DECISIONING MODEL

When making lending decisions, lenders face an inherent tradeoff between confidence in the decision being made and speed of the decision. Another tradeoff that is closely aligned with decision speed is the burden placed on the client in the form of stipulations and information requests. The most secure

decision can be made if all possible data – including tax returns, bank statements, and financial statements – are requested up front. However, this approach would repel most clients and likely lead to lost business to competitors due to the delay in a decision. In addition, this would often waste lender time and resources, as many clients could have been poor credit risks with only a few basic data points. Any model development strategy must address these tradeoffs to find the sweet spot between risk and reward to grow the lender's profits.

Just as in chess, lending decision makers need to understand the tradeoffs affecting their ML-generated decisions



For a typical lending scenario, Mosaic breaks down the lending process into three separate models:

1. **Decline Model**
2. **Lending Model**
3. **Pricing Model**

Decline Model

The Decline Model consists of heuristics, or rules, that are automatically applied to quickly gathered, basic information such as applicant information, balance checks, etc. These heuristics incorporate the policies of the current manual processes of underwriters, but they are automated to save time and enforce consistency. In addition, some additional heuristics are pushed “upstream” during the development and evaluation of the lending decision model.

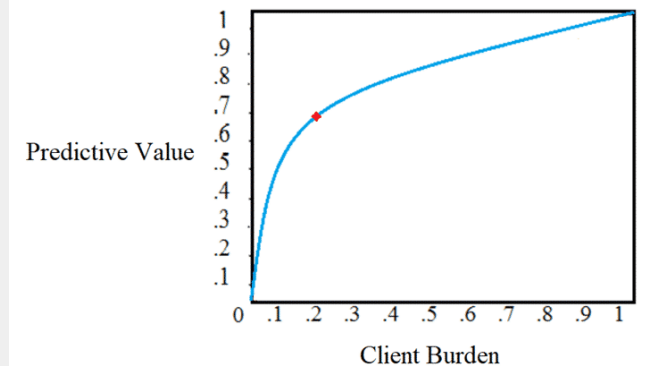
Lending Model

To address those funding requests that survive the Decline Model, the Lending Model uses historical data and classification machine learning algorithms to learn from past decisions and systematize the nuanced decisions of scarce underwriting resources. To discover the sweet spot in the tradeoff mentioned above, the model is initially built using historical cases with full datasets to establish a benchmark for model performance with the availability of “perfect” data. Groups of deal attributes according to the source of the data (e.g., tax return data) are then removed from the model one-at-a-time, continuing until only the most basic data is used in the model. The falloff in performance (see below) at each stage is recorded, thus providing an estimate of the relative importance of each group of features/attributes and their contribution to model performance.

These value estimates are then compared to subjective client burden estimates for providing the information (in consultation with lender subject matter experts) incorporating historical stipulation response rates and times. The goal of these comparisons are

to find the feature set with the best balance of predictive value and client burden. As seen below, this is the point at which the additional predictive value from obtaining more information is not commensurate with the additional burden placed on the client along with the accompanying impact on decision speed. The model feature set is further stratified according to the size of the deal, such that different funding amounts with different amounts at risk dictate different tradeoff points along the curve. Note that this process results in consistent stipulation policies and thresholds. As suggested above, the predictive value analysis typically also reveals additional heuristics that can be fulfilled with basic data, and any such findings are incorporated into the Decline Model.

An example of a hypothetical tradeoff curve between predictive value and client burden



Pricing Model

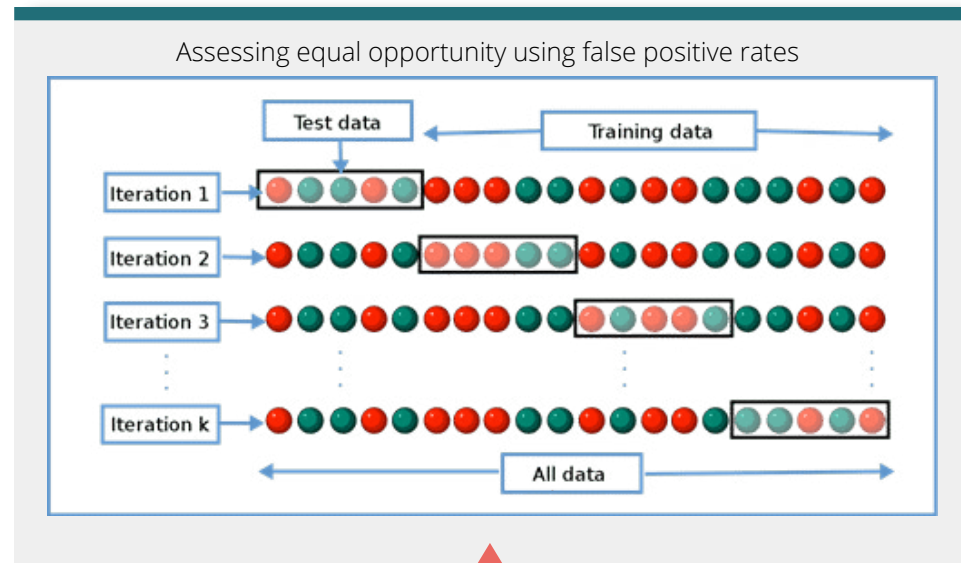
Those funding requests that are approved by the Lending Model have their deal terms set by the Pricing Model to create loan offers. The Pricing Model is possibly a set of models, one for each aspect of pricing: maximum term, rate, and maximum amount. These regression machine learning models again leverage the historic data of past offers and incorporate the knowledge of scarce underwriting resources. They likely leverage the same feature sets determined by the tradeoff analysis for the Lending Model, although Mosaic experiments to ensure that the classification feature sets are compatible with a regression target.

MODEL TESTING AND VALIDATION

The initial process described here applies to both the Lending Model and the Pricing Model. Once the feature set is determined, Mosaic evaluates multiple machine learning algorithms to determine the one with best performance. Additional testing is performed to find the optimal model parameter settings. Performance metrics used for this testing and evaluation are according to the performance section below.

Regardless of the metric, the models are initially trained with only a random subset of historical data and validated against a separate subset of data. This process of cross validation is repeated multiple times to ensure that the algorithm and parameter setting selections are repeatable and not biased by the chance selection of training sets. The final model is trained with all the relevant historical data (egregious outliers are removed).

For the Decline Model, evaluation also consists of cross-validation testing of the heuristics against past decisions, although the heuristics themselves are not trained using historical data.



Cross Validation is critical to filtering out false positives/negatives

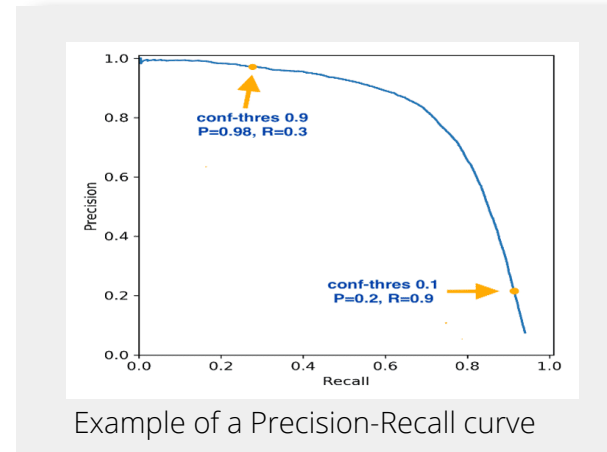


PERFORMANCE

The ultimate performance metrics for the ML solution are loss rate and default rate. To evaluate performance, it is vital to establish a baseline of current operations with respect to the governing metrics. Mosaic begins by evaluating the historical data to establish benchmarks incorporating current trends.

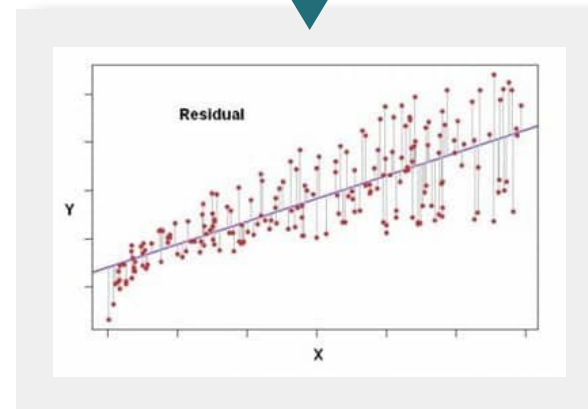
However, each model type has intermediate metrics that are used to tune the models to best achieve the higher-level objectives. For the Decline Model and the Lending Model, the intermediate metrics used are precision and recall (see the chart below). These metrics are commonly used in binary (i.e., approve/decline) classification models to balance false positives and false negatives. An ideal classification model has neither of these when tested against past decisions, but all classification models in practice will produce some decisions where the recommendation is Accept but it should have been Decline (a false negative if default is what we are trying to predict) or Decline when it should have been Accept (a false positive).

Going for high recall would mean we would rarely miss turning down a bad deal, but it would come at the expense of good deals that should have been accepted. Going for high precision means that we would rarely turn down a good deal, but it would come at the expense of accepting bad deals that should have been declined. Mosaic typically works with the client to determine the appropriate balancing of these two bad outcomes according to their impact on the business.



For the regression Pricing Model(s), several performance metrics are evaluated and considered. The basic concept is to minimize the difference/error between the prediction and the actual value, but some metrics handle outliers differently, and some reveal the direction of the error (i.e., model #1 has less average error than model #2, but it consistently recommends too high of an advance amount). The exact set used is selected in consultation with the client team.

This chart shows the difference between the actual values and the predicted values for historic data

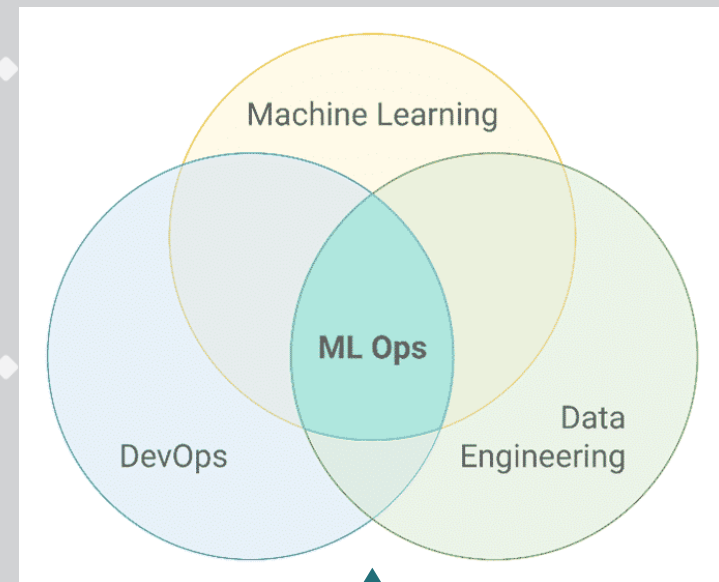


MODEL DEPLOYMENT AND INTEGRATION

Upon completion of the required machine learning models, Mosaic typically gates a decision on integrating the models into the lender's infrastructure. Mosaic waits for acceptance of model performance and mechanics before starting work on implementation.

Models are developed using open-source software such as R or Python. These can be deployed on standard hardware running on-premise or in the cloud as desired. Mosaic typically works with a client's tech team to determine the best integration points and methods, such as reads (deal attributes) and writes (recommendations) to a lender database. The methods for triggering model predictions and retraining runs can be collaboratively selected from several options.

End-to-end testing of the models and their interactions with existing systems is critical. This testing is done in a development environment in conjunction with the lender tech team, and promotion to production is done in a controlled manner according to client dev ops policies.



Machine Learning Operations (ML Ops) is a rapidly rising area of emphasis in AI

MONITORING ML MODEL DRIFT

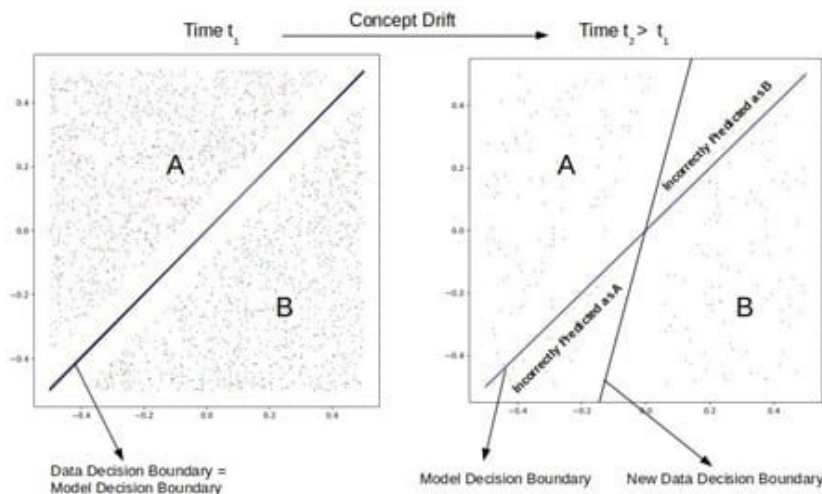
It is crucial to ensure that an automated decision process such as this continues to perform as expected over time. While periodic retraining of the models with the most recent data is standard procedure (perhaps 2-4 times a year), the decisions should be continuously monitored to ensure that customer behavior is not changing in such a way as to require acceleration of the retraining schedule. AI practitioners call this concept drift, and an example is shown in the diagram below. Significant deviation from the model performance benchmarks requires immediate action. Performance that is within acceptable limits but is trending away should also be investigated.



There are several ways of implementing continuous monitoring, and Mosaic selects the exact method in partnership with the lender tech team. A recommended approach is to leverage existing dashboarding/reporting tools to visualize the performance for the appropriate owner. That individual could then manually trigger a retraining run upon review of any deviations. A fail-safe method would also be implemented to automatically cease making advance recommendations in the case of some catastrophic failure of the system (e.g., a change is made to the application that results in an inadvertent transformation of the model input data). These approaches fall under the broad umbrella of the emerging field of Machine Learning Operations (ML Ops).

FUTURE DEVELOPMENT AND EXTENSIONS

There are several enhancements that can build upon the initial modeling foundation. A natural extension would be to create industry-specific stipulations and thresholds so that the process is more stratified and nuanced for different customer types. An additional model could be built to predict likelihood of customer acceptance based upon offer terms and pricing. Forecasting models could be built to provide future revenue visibility to lender management and stakeholders. And, while the foundation models will largely reflect existing subject matter expertise, they could be enhanced by leveraging expanded training data obtained from external parties such as brokers to include outcomes from offers that were not made or accepted (and perhaps went instead to competition). This would enable better decisions by incorporating outcomes currently invisible to lenders (e.g., a deal is declined by the lender, but that prospect goes on to pay back in full a profitable loan from a competitor).



Customer behavior can change over time – concept drift – requiring models to be retrained



CONCLUSION

Hopefully, as a reader you are enlightened as to how AI can fortify lending decisions. Based on similar engagements, Mosaic estimates this project would last for 2-2.5 months following a project kickoff. Investing in AI leads to significant automation savings, more profitable relationships, and an overall boost to any lender's bottom line.

Endnotes

1. <https://lending-times.com/2019/09/27/digital-lending-how-artificial-intelligence-machine-learning-are-making-a-difference/>

