

# Credit Risk Assessment Using Deep Learning

*By the Research Department*

For credit risk assessment, regulators can benefit from having a toolkit for predicting credit deterioration in the banking sector's corporate loan portfolio. This article illustrates an experimental study of how Machine Learning can be deployed to fill such a role. Compared with traditional approaches, a deep learning model (a subfield of Machine Learning) often offers better performance and eliminates the need for time-consuming feature extraction manually.

With the development of the in-house Data Science Lab under the HKMA's Digitalisation Programme, which can handle a computationally-intensive model with ease, a deep learning model is trained on a set of transaction-level corporate loan data collected under the HKMA's Granular Data Reporting programme. The trained deep learning model can consistently detect corporate credit deterioration at both the transaction level and aggregate level with a three-month lead time. Specifically, the model identifies (i) the economic sector of the borrower, (ii) the identity of the reporting authorized institutions, and (iii) the macroeconomic factors (such as the unemployment rate and Purchasing Manager Index) as three major factors driving the downgrades.

While the Machine Learning model has performed reasonably well at both transaction and aggregate levels, the model is not without its limitations. In particular, because of the rather short time span of our dataset, there is no guarantee that the trained model could perform equally well once we transition to the next stage of the business cycle. That said, the Machine Learning model's self-learning capability and flexible architecture make it a versatile tool with great potential.

## Introduction

With the emergence of an increasingly complex and large volume of data (i.e. big data), conventional analytical tools can no longer meet our surveillance needs, and therefore it is imperative to develop advanced analytical tools.

For the banking sector, loan classification (the process of assigning individual loans to groups based on the perceived credit risk) is a promising

area that is amenable to automation and reduced subjectivity by applying data mining techniques. Indeed, developing a systematic approach in predicting credit deterioration has become even more important since the COVID-19 outbreak. With structural changes accelerated by the pandemic, certain sectors may now be under more pressure. For assessment purposes, regulators can benefit from having a toolkit for predicting credit deterioration in the banking sector's corporate loan portfolio.

This article illustrates an experimental study of how Machine Learning (ML) can fill this role. Specifically, this study explores the possibility of applying ML to translate the voluminous transaction level loan data into insights facilitating risk monitoring. The information content of the granular data can be aggregated using ML to help (i) predict the trend of asset quality; and (ii) identify areas of weaknesses. While ML is often criticised by others as a “black box”<sup>1</sup>, this study goes one step further to open the “black box” and shed insight on which variables are most crucial in determining loan quality.

## Data and Methodology

Under the HKMA's Granular Data Reporting (GDR) programme, participating pilot authorized institutions (AIs) have started reporting transaction-level corporate loan data to the HKMA on a monthly basis since April 2019<sup>2</sup>. As of March 2022, the dataset comprises over 248,000 outstanding corporate loans, amounting to a total of HK\$5.1 trillion (roughly 60% of the corporate loans outstanding in Hong Kong). The data reported is also highly granular, covering more than 100 unique data fields.

As an illustration of how ML could facilitate the HKMA's surveillance work, this study focuses on loan classification reported by AIs, which is one of the most closely tracked indicators of the banking sector's health. Under the HKMA loan classification

system, loans are divided into five categories: (i) Pass; (ii) Special Mention; (iii) Substandard; (iv) Doubtful; and (v) Loss, based on AIs' assessment of borrowers' capacity to repay, and on the degree of doubt about the collectability of loan principal or interest by considering a range of factors. The loan classification process typically involves expert judgement and experience of AIs' credit analysts. During an economic downturn, AIs may need to reclassify some of their loan portfolios to reflect the increased uncertainty regarding their collectability. As such, a model that is capable of predicting changes in the credit quality of loans before they would otherwise be noticed by the loan officers, would be helpful in facilitating the HKMA's formulation of timely and targeted policy responses. A more granular analysis on corporate loans can also better inform policymakers of the costs and benefits of extending relief measures, such as moratoria on loan principal repayment. By identifying variables that are most influential on AI's downgrade decisions<sup>3</sup>, the model can systematically determine which attributes of the loans warrant closer monitoring.

Since loan degradation is normally a gradual process, an experienced credit analyst could often tell whether a loan is susceptible to downgrading through studying its characteristics and comparing them with those characteristics of downgraded loans from the past. This is similar to how a supervised ML algorithm operates in practice (Chart 1).


<sup>1</sup> In ML, models are developed by algorithms without requiring human intervention in the process. Because of this, functional forms of an ML model could sometimes get extremely complex, making it difficult for humans to comprehend how the model arrives at its predictions.

<sup>2</sup> This study covers corporate loans reported by 17 pilot banks. The dataset is expected to expand further to include 53 banks by end-2022, covering roughly 90% of corporate loans outstanding in Hong Kong.

<sup>3</sup> In this study, loan downgrades refer to the reclassifications of loans into a lower category by AIs (e.g. from Pass to Special Mention) in any one of the following three months starting from the position date.

CHART 1

## Visualising Machine Learning



In general, there are two key elements of a supervised ML algorithm: (i) output (what we want to know); and (ii) input (what we already knew). In the context of this study, the output is how likely a given set of loans would be downgraded in a three-month period ahead, and the input is their latest loan characteristics (and the macroeconomic variables, including Purchasing Manager Index (PMI), inflation, unemployment rates, etc.). A typical supervised ML process comprises three steps:

### 1. Labelling observations

Before starting the ML process, each observation is first assigned a label (*Positive / Negative*) based on whether it was downgraded in the following three-month period.

### 2. Dividing sample

All observations are then randomly split into two groups – the *train set* (80% of data) and the *test set* (20% of data). Data from the *test set* is completely withheld during the training phase and will only be used in the model evaluation stage. This is to ensure the trained model can generalise well to previously unseen data.

### 3. Training model

Similar to how humans learn from experience, an ML algorithm performs a task repeatedly, each time tweaking its approach slightly in an attempt to improve the outcome. This is typically the procedure which requires the most processing power and is not practical for a desktop computer to carry out. Now,

leveraging the power of the in-house Data Science Lab developed under the HKMA's Digitalisation Programme, which can be thought of as a supercomputer, such data processing can be completed within minutes.

At the training stage, an ML algorithm will randomly draw observations from the *train set* and build a model that best explains the output using information available from the input. The performance of the model will then be assessed using another batch of randomly drawn data from the *train set*. In each of the following iterations, the above procedure is repeated and the model parameters are further fine-tuned. This iterative process continues until it can no longer enhance the performance of the model.

There is an extensive variety of supervised ML algorithms available and each has its own virtues and shortcomings. This study uses the Deep Neural Network (DNN), a state-of-the-art ML technique frequently employed in tackling complex real-world problems (such as fraud detection, image recognition and natural language processing), to identify loans that are vulnerable to downgrades. Compared with other approaches, the training of a DNN model requires (i) a much larger volume of data; and (ii) greater computational power. In exchange, a DNN model (i) often offers better performance; and (ii) eliminates the need to carry out time-consuming feature extractions manually.


## Empirical results

After dropping those loan data with incomplete information, there are 2,913,272 observations left (from April 2019 to March 2022) in our dataset<sup>4</sup>. The dataset is further randomly divided into two parts, with 80% (2,330,617 observations) being the *train set* and 20% (582,655 observations) being the *test set*.

The performance of the fully trained model is evaluated using data from the *test set*. In the context of this study, as the model is tasked with handling a binary classification problem (whether a loan would be downgraded), its performance can be assessed using a metric called “Area Under the Curve of Receiver Operating Characteristic” (also known as AUC)<sup>5</sup>. Chart 2 provides a comparison of how our trained DNN model performs against benchmarks (Hosmer et al. (2013)).

CHART 2

Comparing model performance using AUC




Sources: Hosmer et al. (2013) and HKMA staff estimates

As shown above, the model has an AUC of **0.97**, which is considered to be outstanding by literature. Overall, the model performance has been decent, especially considering how little manual effort is required in setting it up. The model is also easily amenable to the incorporation of even more explanatory variables that may help further improve its performance.

Because of the complexity and non-linearity of the DNN model, it is difficult to comprehend how a trained model generates its predictions. For this reason, deep learning is sometimes criticised as a “black-box”. While opaqueness is often not the primary concern for ML applications like image recognition, it is necessary for a supervisory perspective to be able to pinpoint factors that drive banks’ decisions to downgrade a loan. One way to accomplish this would be through “reverse engineering”. Specifically, the relative importance of a given variable can be inferred by measuring how its presence in the model would alter the model performance, *ceteris paribus*. Chart 3 shows the top ten factors that are most influential in determining the model prediction.

CHART 3

Contribution of variables to loan downgrade predictions



Note: Variable drop-out loss refers to the drop in model performance when a variable is removed from the model. A higher drop-out loss indicates greater importance of the variable.

<sup>4</sup> Due to data quality constraints, certain GDR data fields are excluded in this study. For a complete list of the data fields used, please refer to the Appendix.

<sup>5</sup> When assessing the performance of a classifying model, we should consider both its (i) ability to detect positive cases; and (ii) likelihood of setting off false alarms. These two aspects can be measured using True Positive Rate (TPR) and False Positive Rate (FPR) respectively. The TPR and FPR of a model could, however, vary depending on how we set its classification threshold. AUC is an aggregate measure of model performance across all possible classification thresholds. For further details, please refer to <https://developers.google.com/machine-learning/crash-course/classification/roc-and-auc>.

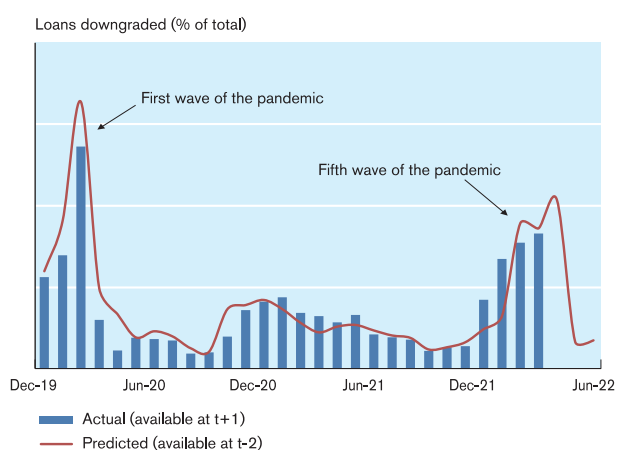
According to Chart 3, the economic sector of the borrower, the identity of the reporting AI, and the macroeconomic factors (such as unemployment rate and PMI) are the most important explanatory variables in determining how likely a given loan would be downgraded in a three-month period ahead<sup>6</sup>.

## Application to macro-prudential surveillance

Apart from identifying the characteristics of loans that could render them vulnerable to downgrades, the ML model could also be deployed to offer timely insights on the wider trend of AIs' asset quality (Chart 4).

**CHART 4**

### Comparing model forecasts and actual downgrades



Sources: GDR and HKMA staff estimates


As shown in the above chart, the ML model performs reasonably well in forecasting the portion of loans being downgraded at the aggregate level, with three-month lead time. In particular, it provides an early-warning signal of credit deterioration at the onset of the first and fifth waves of the pandemic. According to the model, the pace of loan downgrades is expected to stabilise in the second quarter of 2022.

As certain sectors may suffer more than others during a downturn, the ability to identify the hardest hit sectors early could also help the formulation of a more targeted, and hence more effective, policy response. This is another area where the ML model could potentially be deployed (Table 1).

<sup>6</sup> It is noted that fields not covered by existing returns / surveys (e.g. value of collaterals / identity of the borrower's ultimate parent) generally show lower data quality. Such a discrepancy in data quality could lead to understatement of the importance of these new variables. That said, this issue could hopefully be resolved as data quality continues to improve over time.

TABLE 1

Deep learning-based downgrade forecast on corporate loans (by economic sector)



Note: The length of the blue bars on the left is proportional to the relative amounts of corporate loans outstanding. Sectors with downgrade ratio two (four) standard deviations above the mean are highlighted in yellow (red).

Sources: GDR and HKMA staff estimates

Table 1 shows the predicted loan downgrades by economic sector since 2020<sup>7</sup>. According to the table, “Building and construction, property development and investment”, which is the single largest economic sector by loan size, has seen a transitory increase in loan downgrades during early-2020, possibly reflecting the deterioration in the market outlook for commercial properties (e.g. offices and retail shops) amid the pandemic. The sector also recorded an increase in downgrades in early-2022, partly due to renewed market concerns about the liquidity conditions of Mainland property developers. The contact-intensive sectors (such as “Hotels, boarding houses & catering”) also witnessed visible increases in downgrades during the first half

of 2020 following the tightening of social distancing measures. Finally, the model suggests that although the re-imposition of stringent containment measures to tackle the fifth wave of the pandemic could bring renewed challenges to these sectors, the negative impact is expected to be short-lived this time as their downgrade ratios have already shown signs of stabilising.

While the ML model has performed reasonably well at both transaction and aggregate levels, the model is not without its limitations. In particular, because of the rather short time span of our dataset, there is no guarantee that the trained model could perform equally well once we transit to the next stage of the business cycle. In addition, given that the 17 pilot AIs combined account for just 60% of the corporate loans outstanding in Hong Kong, the trained model, even if 100% accurate, may not be able to provide a complete picture of the credit risk faced by the Hong Kong banking sector.

<sup>7</sup> Loan downgrade ratio is calculated as (i) amount of loans reclassified into a lower category, divided by all outstanding loans. Sector classification used in this study is according to the classification criteria as defined under Table 3.5 – Loan and advances for use in Hong Kong by economic sector of the HKMA’s *Monthly Statistical Bulletin*.

## Conclusion

As data is becoming increasingly complex, advanced analytical tools are needed to help us unlock the full potential. For the banking sector, loan classification is an area that is amenable to the application of ML techniques, and given the expected rise in problematic loans once COVID-related policy support expires, regulators can benefit from having a toolkit that can accurately predict credit deterioration in the banking sector's corporate loan portfolio. Against this backdrop, this study illustrates how ML can be deployed to fill this role.

In leveraging the power of deep learning, this study trains an ML model that is capable of detecting credit deterioration at the transaction level using the GDR dataset. The model identifies (i) the economic sector of the borrower; (ii) the identity of the reporting AI; and (iii) the macroeconomic factors as three major factors driving loan downgrades.

As mentioned above, although the ML model has performed reasonably well at both transaction and aggregate levels, the model does have limitations. In particular, because of the rather short time span of our dataset, there is no guarantee that the trained model could perform equally well once we transit to the next stage of the business cycle. However, the ML model's self-learning capability and flexible architecture make it a versatile tool with great potential.

## Reference

Hosmer Jr, D. W., Lemeshow, S., & Sturdivant, R. X. (2013). *Applied logistic regression* (Vol. 398). John Wiley & Sons.

## Appendix: Explanatory variables in the model

Category	Data field
Loan characteristics	Loan amount at origination (HK\$ equivalent)
	Closing outstanding amount (HK\$ equivalent)
	Tenor at origination
	Reporting position date
	Specific provision
	Interest payment frequency
	Currency
	Governing law(s)
	Type of facility
	Loan use economic sector
	Recourse
	Seniority/Lien
	Counterparty type of Mainland exposure
	Syndicate indicator
	Loan purpose
	Loan use location
	SME Financing Guarantee indicator
	SME Loan Guarantee indicator
	Loan status
	HKMA loan classification
	Interest rate type
	Reference rate
	Reporting bank AI code
Macroeconomics	Unemployment rate
	Inflation rate
	PMI
	Retail sales (value)
	Retail sales (volume)
	Tourist arrivals
	Imports (value)
	Exports (value)
	Imports (yoy growth)
	Exports (yoy growth)
Financial market	Hang Seng Index