IADI Fintech Briefs provide high-level overviews and key takeaways on Fintech topics of relevance to deposit insurers.

FINTECH BRIEF

MACHINE LEARNING METHODS

POTENTIAL FOR DEPOSIT INSURANCE

NO. 3

RYAN DEFINA

September 2021

IADI Fintech Briefs are written by Members of the IADI Fintech Technical Committee (Fintech TC), and from time to time by other authors, and are published by IADI. They provide high-level overviews and key takeaways on fintech topics of relevance to deposit insurers. The views expressed are those of the author(s) and do not necessarily represent the views of IADI or the institution to which the authors are affiliated. The author is grateful to the Fintech TC and David Walker for helpful comments.

The editor of the IADI Fintech Brief series is Bert Van Roosebeke (IADI).

This publication is available on the IADI website (www.iadi.org).

International Association of Deposit Insurers (IADI), C/O Bank for International Settlements,
Centralbahnplatz 2, CH-4002 Basel, Switzerland. Tel: +41 61 280 9933 Fax: + 41 61 280 9554

Ryan Defina
IADI Research Unit
ryan.defina@iadi.org

MACHINE LEARNING METHODS POTENTIAL FOR DEPOSIT INSURANCE

Executive summary

The field of deposit insurance is yet to realise fully the potential of machine learning, and the substantial benefits that it may present to its operational and policy-oriented activities. There are practical opportunities available (some specified in this paper) that can assist in improving deposit insurers' relationship with the technology. Sharing of experiences and learnings via international engagement and collaboration is fundamental in developing global best practices in this space.

1 Introduction and Purpose

Recent innovations in the field of data science have seen the evolution of long-established machine learning algorithms. Many view this phenomenon as the next major development in data analytics with numerous potential applications for deposit insurers and other regulatory authorities.

Machine learning has emerged as a sophisticated suite of methods likely to affect deposit insurers in fundamental ways. The technical implementation of such methods has the potential to automate complex business processes and decision making for deposit insurers and their member institutions. This presents a challenge to the very essence of corporate culture and raises many questions about the future role and structure of deposit insurers and their member institutions.

This short paper introduces some key concepts and highlights potential areas for further discussion.

1.1 Motivation

Deposit insurers have access to many large administrative and transactional datasets suitable for exploration with machine learning approaches, including data on insured deposits, early warning metrics, and information applicable to promoting public awareness. Using machine learning approaches on these datasets may provide informative insights that yield productivity improvements to existing business processes and/or non-negligible cost savings. Machine learning may also enable the automation of seemingly complex expert decision-making.

Large technology firms (so called "Bigtechs") such as Facebook, Alphabet (Google), Amazon or eBay have been pioneers in realising significant commercial benefits from machine learning technology. However, the banking and financial services sectors are increasingly building technical capability in the field, and in turn reaping benefits. Research suggests that firms are likely to reduce costs by up to 22% by 2030 through utilisation of developments in statistical methodology, including machine learning, deep learning and artificial intelligence.[1] Deposit insurers should therefore expect to be further exposed to these technologies in the near future.

1.2 Background

Machine learning can be defined as the study of computer algorithms that improve automatically through experience.[2] Central banking is beginning to embrace and explore relevant applications of these algorithms. Chakraborty & Joseph (2017) present an overview of this adoption and include a number of very practical applications within the Bank of England. These include the provision of a stylised framework for bank supervision under imperfect information; improved forecasting of UK consumer price inflation; and clustering of UK firms to determine financial technology utilisation.

[1] See Autonomous Next (2019).
[2] See Mitchell (1997).

Global initiatives coordinated by international bodies such as the Bank for International Settlements have sought to highlight synergies between central banks, regulatory authorities and supervisors seeking to embrace big data technologies. Recently an Innovation Hub was established by the BIS to develop in-depth insights into critical trends in financial technology of relevance to central banks, to explore the development of public goods to enhance the functioning of the global financial system, and to serve as a focal point for a network of central bank experts on innovation.

2 Fundamentals of Machine Learning

It is important to establish a framework in which to discuss machine learning solutions and therefore better understand how they may be utilised to address business problems. The foundation for all applications of machine learning are:

Machine learning algorithms. Processes/procedures that are implemented in computer code and applied to a data source. These are used to search for relationships between data items via pre-specified criterion.

Machine learning models. Output generated by algorithms consisting of model parameters and a mapping of relationships between data items.

2.1 Building the Model

Machine learning uses algorithms to turn a dataset into a model. In high-level terms, a machine learning algorithm will take a selection of input variables and apply some optimality criterion such that relationships between variables can be identified and understood. The process is mostly data-driven, and can therefore be viewed, at times, as more robust than traditional econometric methods that can be sensitive to model assumptions. This can also limit the extent to which modelling expertise is directly incorporated, which has both pros and cons associated.

Typically, a dataset will be split into two subsets to enable model fitting to occur[3]:

Training dataset. Once the required machine learning algorithm has been selected, it is applied (or 'trained') to training data. The training dataset is chosen such that it is suitably representative of the overall dataset i.e. includes relevant modelling variables in appropriate proportions.[4] It should also provide evidence of relationships that are likely to continue into the prediction/forecast horizon. Model predictions derived from the algorithm will only be as strong its training data.

Test dataset. Model parameters estimated using the training dataset are then applied to the testing dataset. This is the remaining data not chosen to form part of the training dataset. It functions as an independent foundation for verifying the accuracy and robustness of a model. This is to guard against the possibility that model parameters are over-focussed on one sample, and hence do not accurately fit unobserved (or more practically speaking, yet-to-be observed) data points – referred to as overfitting.

In order to reduce potential overfitting and selection bias, replication methods are typically invoked. The most commonly applied in practice is cross validation. This approach fits a suite of candidate learning algorithms[5] (or different configurations of an algorithm) many times over with unique partitions of the training dataset. Results from each model run are collated and analysed to assess which algorithm performs best. Ultimately, the final model is chosen to ensure it is robust against small changes in the input data.

[3] A third subset can be used in the case of more complex cases whereby additional (intermediate) model components are sought to be estimated in a robust manner. This is referred to in practice as hyper-parameter optimisation. Such approaches are out of scope in this paper.

[4] In practice, training data is often a simple random sample (without replacement) of the full dataset. The size can vary based on subject domain, but is typically 30%-70% of all available data. It should include a broad spread of possible outcomes to minimise the risk of overfitting and the subsequent introduction of statistical bias.

[5] Various modelling options are described in Section 3.

2.2 Diagnostics for Model Validation

A standard approach used for classification exercises and demonstration is a contingency table. It represents a categorisation of model predictions based on how they compare against the actual outcome. The partitioning of datasets into training and test sets enables calculation of the TP, FP, FN and TN for a given model.[6] Adjusting the model to increase/decrease TP, FP, FN or TN will almost certainly be at the cost of another. The key message here is that machine learning is an exercise in compromise between abstracting from reality (to derive informative and interpretable insights) and letting the data speak.

Total population		Actual outcome	
		True	False
Predicted outcome	True	True positive (TP)	False positive (FP)
	False	False negative (FN)	True negative (TN)

Accuracy. The starting point for many assessments of performance is accuracy. This corresponds to a count of correct predictions as a share of all predictions i.e. (TP+TN) / (TP+TN+FP+FN). However, this metric can often yield a misrepresentative view of model performance, as unbalanced samples can increase accuracy while not achieving appropriate classification outcomes.[7]

Recall. Often referred to as the 'true positive rate' or 'sensitivity' of a machine learning algorithm. It is the ratio of true positives to the total number of actual true outcomes or TP/(TP+FN). Expressed in terms of conditional probabilities, recall can be defined as

$$\Pr(\text{True Prediction} \mid \text{Actual True})$$

Precision. Often referred to as the 'positive predicted value'. It is the ratio of true positives to the total number of true predictions or TP/(TP+FP). Expressed in terms of conditional probabilities, precision can be defined as

$$\Pr(\text{Actual True} \mid \text{True Prediction})$$

F1-score. One will often seek to trade-off between the diverging precision and recall metrics. This renders the F1-score helpful, as it is an explicit trade-off, the harmonic mean of precision and recall. As such, the F1-score is a common metric used to compare models with different features included. Use of the F1-score ultimately seeks to guard against models that overly weigh high precision to the detriment of recall (or vice versa).

$$\text{F1} - \text{score} = \frac{2.\text{Precision}.\text{Recall}}{\text{Precision} + \text{Recall}}$$

It is important to note that model selection is a field of study unto itself and that many other measures and methods are available, each at times better-tailored to optimising selection under specific algorithms, domain-specific data structures and/or industry context. Various model selection criterion[8], along with resampling approaches including the bootstrap and cross-validation are not uncommon in industry.[9] Also important is consideration of ways to present model fit to a non-technical audience.[10]

[6] Often referred to as a confusion matrix.

[7] Consider a scenario where an analyst wishes to predict when a bank failure occurs within a jurisdiction with very stable financial conditions. If data is collected frequently, most (nearly 100%) of observations will be cases of zero failures. Hence, the analyst can achieve nearly 100% accuracy by trivially predicting 'no bank failure' every single period. This however is not a helpful model, as the incidence of bank failure is of particular interest. The subsequent discussion of recall, precision and F1-score teases out this issue further.

[8] These typically include the Akaike and Bayesian Information Criterion (AIC and BIC respectively) which seek to optimise model fit subject to a penalty for additional model complexity.

[9] See Brownlee (2019) for an overview.

[10] Common tools include the confusion matrix and the AUC/ROC curves which offer an intuitive visual representation of model fit.

3 Statistical Methods

Three broad classes of machine learning algorithms are available. Each addresses a different suite of business problems.[11]

- ***Supervised learning*[12] *(or 'classification' / 'prediction')*** – Allows one to appraise the relationship between a suite of inputs and a known output. Following development of the algorithm, a new suite of inputs can be provided, and the model can be used to predict the likely category of the output. For example, a large number of credit scoring variables might accurately predict if a client is likely to default on their loan (based on historical experience).

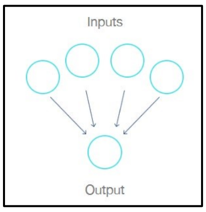

- ***Unsupervised learning*[13] *(or 'clustering' / 'segmentation')*** – Allows one to explore relationships between variables where data are bundled into groups. No explicit labels or categories denoting structure in the dataset are known. Rather, this method can identify previously unobserved dynamics in the training dataset. It is beneficial when you want to group your data into categories to simplify communication or further analysis, but you are unsure of what might determine this categorisation (or grouping). For example, a large suite of variables related to capitalisation might be used to group banks into risk-based categories, enabling tailored supervisory practices to be adopted for each bank under supervision.

- ***Reinforcement learning*** – This method seeks to learn iteratively by exploring a relatively unknown environment (limited training dataset) via a series of actions, noting feedback/reward at each stage. Over many cycles, the algorithm starts to understand optimal behaviours within the environment and note these for further analysis. For example, a knowledge-sharing platform might incorporate historical searches when deciding on which content to highlight for a given user.

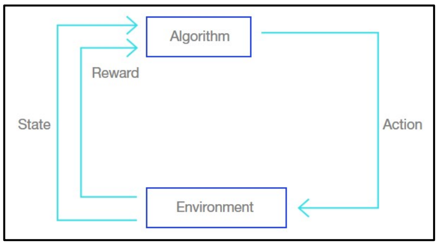

[11] Much of the content that follows is adapted from McKinsey & Company (2020).
[12] Methods for supervised learning include CART, neural networks, naïve Bayes classifiers, support vector machines, random forests.
[13] Methods for unsupervised learning include k-means, k-nearest-neighbour, centroid-based clustering, principal component analysis, Gaussian mixture models.

4 Benefits and Drawbacks for Deposit Insurers

There are three main channels to consider when articulating the benefits and drawbacks of machine learning for deposit insurers. These are:

- ***Direct interaction*** – Deposit insurers can apply machine learning to their primary oversight and assessment activities (some potential applications are specified later in the paper). This requires the DI to have a detailed technical understanding of how associated methods work, the types of problems they are designed to address, implementation issues within the context of modern IT infrastructure, configuration and maintenance workflows. The DI has full control over the extent to which these methods are utilised and will need to ensure technical outputs are mapped to clear, understandable and implementable guidance.
- ***Indirect interaction*** – Insured members of a deposit insurance system are increasingly using machine learning methods in day-to-day data analytics and business decision making. The deposit insurer may have an interest in understanding both the outputs generated by members, and the manner in which they were generated.[14] This may impact monitoring activities of the DI (and associated preparation for a potential payout) should any reported metrics generated from a machine learning algorithm not be fully comprehended. The deposit insurer's risk profile will lift substantially if the compilation methods of member reports are treated as black boxes.[15]
- ***Third-party interaction*** – This channel acknowledges that increased usage of machine learning approaches in seemingly unrelated sectors may have consequences for deposit insurers. For instance, many digital currencies use machine learning and distributed ledger technology that disrupts traditional financial service products. The deposit insurer should monitor these products to ascertain potential tertiary effects that may impact deposit taking in their jurisdiction. Understanding the tools being used will naturally inform any assessment of the relative levels of risk being introduced by these firms to the system as a whole.

Table 1 – Benefits and drawbacks of machine learning

Benefits	Drawbacks
Data-driven[16]	Different skillsets required e.g. data science skills
More effective utilisation of existing data from multiple sources	Initial investment in appropriate IT infrastructure to support computation requirements
Can assist in clarifying/verifying assumptions framing policy decision making	Interpretation of method/results can be challenging, and difficult to communicate to a non-technical audience
Reduce manual/clerical errors	Requires substantial investment in technology[17]
Productivity improvements	Need comprehensive access to granular data that is timely
Cost savings in the medium/long term	Do not always incorporate subject matter expertise to an appropriate extent i.e. becomes too data-driven

[14] The extent of interest will of course be dependent on the specific machine learning application and the deposit insurer mandate.
[15] A black box is a device, object, or system whose inner workings are unknown; only the inputs and output reactions are known.
[16] This implies that the methods are free from parameterisations that the user designs to support a pre-specified outcome.
[17] The investment may include one or more of the following: hiring suitably skilled staff such as data scientists (or similar) which typically command high wages in the labour market; retraining existing staff; building (or commissioning) advanced IT infrastructure such as distributed computing platforms. Each of these can require a substantial investment of time and money.

5 Potential applications

Likely applications of the methods and associated analysis cross a number of functional areas in deposit insurance agencies. [18] For instance, there are implications for forecasting algorithms in policy groups; document and archival management in information groups; provision of tailored advice in public awareness activities; as well as broader operational efficiency gains in service areas. Applications within the context of bank resolution are also beginning to emerge in the literature under the banner of ResTech.[19]

Table 2 – Potential applications in deposit insurance

Application	Description
Clustering	Data-driven grouping of member banks based on a potentially large number of explanatory factors. This can facilitate the provision of differential premium assessments if input variables are risk-based measures.
Classification (payout use-case)	Improve efficiency of claim processing, particularly where a standardised SCV is not available. Claims received from the public via multiple channels in non-standard formats can be triaged to ensure the most suitable follow up is conducted.
Classification (actuarial use-case)	Determine the likelihood of a given member bank failing within a specified horizon, based on historical data. Time series of historical performance metrics can be utilised.
Sentiment analysis	Tailored advice to claimants based on individual characteristics (sophisticated chat bot), plus insights on areas of misinformation/misunderstanding. This allows for tailoring public awareness efforts accordingly.

6 Improving interaction with machine learning

Effective engagement with machine learning by deposit insurers might be encouraged by considering one or more of the following:

- ***Encourage research and development.*** Foster a culture within the organisation whereby exposure to new methodologies is viewed constructively. An effective means to achieve this is through some exploration of R&D initiatives whereby the potential value proposition of such methods can be further scoped out.

- ***Commission pilot projects.*** Conducting small projects that explore new technologies can be useful when operating within the prism of a longstanding business problem. This helps to mitigate the risk that new technology may impose on the organisation's critical functions. It allows organisations to understand gaps (organisation, structural, and/or personnel-based) that may inhibit full interaction with technologies of this kind when scaled up to larger projects.

- ***Working with member banks.*** All deposit insurers understand the value of working closely with their member banks. Within the context of technology such as machine learning, the benefits are both immediate and downstream. Understanding the vision for future work, and associated pain points, can help shape planning (resourcing, timelines, stakeholder management) and allow greater appreciation for areas of emerging risk in deposit insurance.

- ***Coordinate globally.*** Engagement with international organisations and industry partners helps keep all stakeholders up to speed on the latest technical developments. It also facilitates benchmarking of performance and sharing of experiences with other deposit insurers. The formation of cross-jurisdictional working groups with well-defined terms of reference will likely be best-placed to build machine learning-oriented synergies from deposit insurers.

---

[18] The suite analysis tools now available are considerable. R packages and Python libraries are some of the most popular and are used extensively in industry applications.
[19] See Loiacono et al (2020).

- ***Invest in staff capability.*** New methods require new skills. As such, deposit insurers may consider investing in training, workshops, or conference participation to build capability among existing staff. Future recruitment may also seek to target those in the labour market with skills in data analytics and/or computer science. This includes considering selections from the private sector or from outside the jurisdiction in question to ensure the most up-to-date skillsets are captured.

7 Concluding remarks

Machine learning has great potential within deposit insurance. It is expected that associated approaches will slowly be integrated into the operations of deposit insurers over the coming years. Deposit insurers should be prepared to interact constructively with such technology into the future to maximise utility and mitigate emerging risks.

8 Appendix

8.1 Technical implementation of machine learning algorithms

The broad field of machine learning encompasses a large number of methods. A comprehensive open source C++ library for machine learning algorithms is available on Github called *tiny-dnn*: https://github.com/tiny-dnn/tiny-dnn. Others include:

- *Tensorflow* Python implementation of RNN (https://www.tensorflow.org/tutorials/recurrent).
- Python library *pytorch* (http://pytorch.org/).
- *Caffe* (http://caffe.berkeleyvision.org/) developed by an artificial intelligence student at UC Berkeley.
- LeNet (http://yann.lecun.com/exdb/lenet/) the first successful implementation of convolutional neural networks (CNN), by Yann LeCun dating back to 1989. https://arxiv.org/abs/1409.4842 provide the formal technical architecture.

8.2 A Selection of Deep Learning Methods

Deep learning methods are an extension of machine learning that is ever-evolving and an active area of academic research. The following summarises some of the more common methods used. Prerequisite understanding of neural network basics is required – see A Beginner's Guide to Neural Networks and Deep Learning. These methods often reap the most benefits when applied to large volumes of unstructured data, typically in the form of text, images and/or audio content. There is minimal marginal benefit of using deep learning on small and well-structured datasets.

8.3 Convolutional Neural Network (CNN)

The CNN method is used within the context of an image input. Hidden layers of neurons, not unlike that of a standard neural network, are arranged into a single layer of three dimensions. In terms of mechanics, the diagram below can read in a large number of two-dimensional images and classify them into one of many possible finite (scalar) categories.

A practical example of training a CNN algorithm involves using the MNIST dataset consisting of a relatively small sample of handwritten numbers. A C++ implementation is available on Github: https://github.com/tiny-dnn/tinydnn/tree/master/examples/mnist. Useful reference: Karpathy (2015)

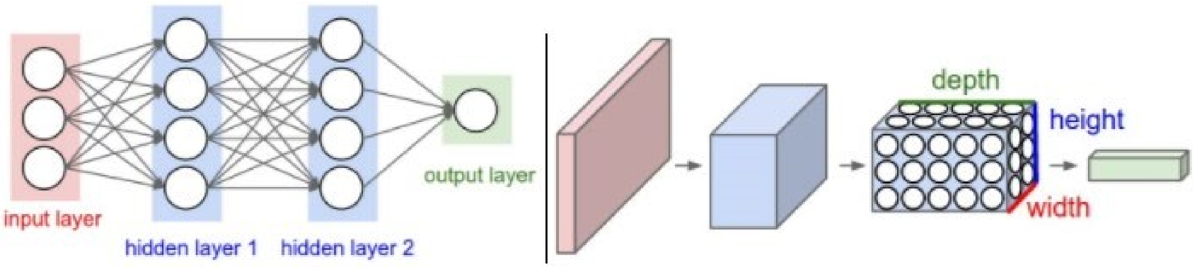

8.4 Recurrent Neural Network (RNN)

Recurrent neural networks different from classic neural networks in that they retain information more effectively, and pass this through additional steps of the algorithm. This can be thought of as multiple network copies operating in unison. Each network passes information through to others as a means to encourage more optimal overall model performance. This method is best applied to classification problems concerning handwriting and speech recognition. Useful reference: Olah (2015)

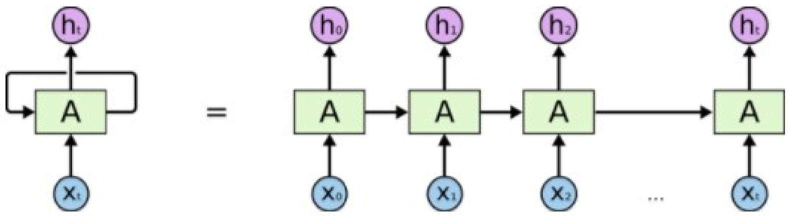

An unrolled recurrent neural network.

8.5 Other methods

Methodological development is fast-moving, with new approaches being developed all the time. Others of interest to readers may include sequence-to-sequence models; various autoencoders; and generalised adversarial networks.

9 References

Autonomous Next (2019). Augmented Finance and Machine Intelligence. *Sanford C. Bernstein & Co., LLC.*

Broeders D, J Prenio (2018), Innovative technology in financial supervision (suptech) – the experience of early users, *Financial Stability Institute Insights No. 9*

Brownlee, J. (2019), A Gentle Introduction to Model Selection for Machine Learning, *ML Mastery*

Chakraborty C, A Joseph (2017), Machine learning at central banks, *Bank of England Staff Working Paper*

LeCun Y, L Bottou, Y Bengio, P Haffner (1998), Gradient-Based Learning Applied to Document Recognition, *Proceedings of the IEEE*

Loiacono G, A Mazzullo, E Rulli (2020), ResTech: Innovative Technologies for Crisis Resolution, *Oxford Law*

McKinsey & Company (2020), An executive's guide to AI, *McKinsey Insights*

Mitchell, T (1997), Machine Learning, *McGraw Hill*

Olah, C (2015), Understanding LSTM Networks, *Academic blog*

Karpathy A (2015), CS231n Convolutional Neural Networks for Visual Recognition, *Academic blog*

Previous issues in this series

No 1 September 2021	Introductory Brief Challenges for deposit insurers	Rachel Youssef, Rose Kushmeider, and Diane Ellis (FDIC)
No 2 September 2021	Data Standardisation	Daniel Hoople (FDIC)