CS 349f: Technology for Financial Systems Final Report
Topic Summary: Machine Intelligence for Housing Finance (Kay Giesecke)

Abstract

The primary focus of Prof. Giesecke's research was to develop a deep-learning based model to predict the status of mortgages based on loan-level information (i.e. credit score, income, and demographic metadata). Using a large dataset of 120 million mortgages, researchers were able to train a neural network to generate a probability distribution over the space of possible loan states over time. Investment portfolios designed based on the developed models indicate that the neural network approach outperforms logistic regression both in the prevalence of portfolio-selected loans that are current after 12 months and the rarity of the number of loans that are delinquent after 12 months. The approaches presented have high potential to be applied to other areas as well, such as valuing securities built on student loans, credit card loans, or auto loans.

Introduction

Historical models to predict mortgage and prepayment risk have been assumed to be linearly dependent on predictive factors (i.e. credit score, income, and loan-to-value ratio). However, empirical results have demonstrated nonlinear effects in the relationship between such predictive factors and mortgage risk. This finding motivates the study of nonlinear models (i.e. artificial neural networks) as a technique to predict mortgage risk.

The research utilizes a dataset provided by CoreLogic that provides both loan-level and borrower-level information on 120 million mortgages, recording 3.5 billion monthly observations over at time frame of roughly two decades. This core dataset, coupled with additional metadata (from various private companies and government sources) on individual income profiles, credit scores, demographic information, social media activity, and bank transaction data, provides a comprehensive set of features upon which to construct models that are predictive of the observed nonlinearities. The dataset also provides temporal data on the mortgage status (i.e. number of days of delinquency and foreclosure status).

Methodology

The ultimate goal is that the model predicts, based on the provided features, a discrete variable representing the loan state. The state u is discretized such that it must be one of the following seven options: current, 30 days delinquent, 60 days delinquent, 90+ days delinquent, foreclosed, REO (real estate owned), or paid off. Let U_{t+1}^{n} represent the state of the n^{th} loan at time horizon $t+1$, and let X_t^n represent the loan features at time t combined with the distribution over the previous timestep's loan states. Then, the model h (parametrized by θ) is a function of X_t^n that provides a probability distribution over the loan states at timestep $t+1$.

$$P(U_{t+1}^{n} = u) = h_{\theta}(u, X_t^n)$$

A deep neural network was chosen for the model h_θ, where θ represents the weights and biases of the L layers $(W_1, \ldots, W_L, b_1, \ldots, b_L)$. The benefits of the deep learning approach to state transition prediction are that it can leverage the vast amount of data present and capture non-linear effects. In order to determine the extent to which the added complexity of the neural network model is truly necessary for accurate predictions, logistic regression was also tested.

Results

One key metric for evaluation was the performance of an investment portfolio based on the best performing 20,000 loans (according to predicted probability of being current). Portfolios constructed using both neural network and logistic regression models were analyzed with respect to (unseen) ground truth empirical data on a 12-month time horizon. When the loans were selected using the neural network model, 93.28% of the selected loans were current after the 12-month period. By contrast, using the logistic repression model, 89.09% of the selected loans were current. Furthermore, the neural network model also had fewer loans that were either foreclosed or 90+ days delinquent after the 12-month period as compared to the logistic regression model, indicating that the neural network was indeed effective.

Another key consideration for these models is explainability. In order to measure this, the researchers verified whether model-predicted trends match common sense intuitions. For instance, the model confirms the intuition that the rate of delinquency and the rate of unemployment within a given state are positively correlated, and the intuition that mortgages with high interest rate and low FICO score are more likely to be 30 days delinquent.

Conclusions

This research has shown that ML, and in particular deep learning, can successfully be applied to predicting the value of securities (i.e. mortgage backed securities) over various time horizons. Furthermore, it is clear that the approaches presented in the research can generalize and have high potential to be applied to the valuation process on securities backed by other types of loans (i.e. student loans or auto loans). Finally, the researchers conclude that tests for both explainability and statistical significance in neural networks are crucial, both from a performance analysis perspective, and to ensure trust for widespread adoption.

References

Giesecke, Kay. "Machine Intelligence for Housing Finance." *CS 349f Seminar Presentation*, Advanced Financial Technologies Laboratory, Stanford University, 6 May 2019.

Sirignano, Justin, et al. "Deep Learning for Mortgage Risk." *SSRN Electronic Journal*, 10 Mar 2018, doi:10.2139/ssrn.2799443. https://arxiv.org/abs/1607.02470

Kagan, Julia. "Mortgage-Backed Security (MBS)." *Investopedia*, 12 Mar. 2019, www.investopedia.com/terms/m/mbs.asp.