Cubic Ridge Regression for Scaling Machine Learning Credit Risk Scores

Abstract

Machine learning (ML) for bespoke credit risk analytics is well trodden ground. Both FinTech's and larger financial institutions use ML methods for in-house credit risk models. Model development initiatives vary by institution. Smaller teams typically adopt open-source software. Whereas larger institutions often leverage licensed software that can be more "point and click" for development. However, most ML binary classification algorithms provide a score value from zero to one with high decimal precision. Scaling of this raw score to match the score range and odds of an existing score, or a bureau score, is usually a business requirement. Oftentimes, this scaling process entails rounding to the nearest integer value to generate a user-friendly score. Whether writing custom code to scale or a licensed solution, one common way is a two-step linear transformation from unscaled scores, first to implied log odds, then finally to scaled values. However, the literature on scaling of credit scores is rather sparse, especially that of scaling high decimal precision ML analytics. To that end, the authors outline a novel approach to better scale a bespoke score to match the range and odds of an existing score. This involves leveraging a third order polynomial Ridge regression. We combine this technique with a known, albeit less documented, approach to match both the odds at a base score value and the points to double the odds (PDO). On an empirical data set the authors observe at least four main benefits to this novel framework. First, we can achieve greater risk differentiation in terms of unique score values, even after rounding the scaled score to the nearest integer. On our data set our cubic fit garners a roughly 54% increase in unique score values over the traditional linear method. This is an obvious benefit to business end users for setting credit strategies and policies given the increase in fidelity in risk rank ordering and risk differentiation. Second, the scaled score is linear in the log odds space, whereas the linear transformed scaled score demonstrates non-linearity in the log odds space. Third, the odds at a base score value are better matched. Fourth, the points to double the odds are preserved better under the third order polynomial transformation. We will outline in detail our novel approach and share empirical results on a proprietary data set.

Authors & Affiliations

Dr. Marick Sinay[1], Dr. Peng Jiang[1], Dr. Andrew Jennings[1]
[1]Octane Lending, New York, USA