Developing Broad-based Credit Scores to Balance Many Objectives

Dr. Gerald Fahner

FICO

Many objectives and constraints matter when developing credit scoring models for high-stakes predictions and decisions, including predictive power, robustness, interpretability, fairness and of course compliance. Nested within high-level objectives one may further identify more granular subgoals. For example, a broad-based credit score should accurately differentiate individuals' credit risk for various combinations of loan types and across the credit lifecycle (origination/account management.)

Score development can be thought of as the art and science of determining the "perfect balance" between the many objectives, subject to the non-negotiable compliance constraint. As credit score developers know, one can often improve interpretability and soundness of a model by applying sensible constraints on the functional form of the predictor function, for example by imposing monotonicity constraints, without giving up materially on predictive performance (sensible constraints can indeed improve robustness and mitigate compliance risks.) Such free-lunch opportunities are related to the "flat maximum effect" long known for linear models and subsequently also observed and described as the "Rashomon effect" for nonlinear, algorithmic models. Similarly, we find that broad-based scores can sometimes be improved on some predictive performance dimensions without materially impacting other performance dimensions.

To aid the understanding of possible tradeoffs between the many predictive performance objectives on broad-based scores, and to help stakeholders in a score development to elicit the most adequate model for their needs, we are researching and developing novel approaches rooted in multi-objective optimization to map out the multi-dimensional Pareto frontier spanned by interpretable and compliant models.

Increasing the number of objectives quickly drives up the computational cost for sampling the Pareto front—the curse of dimensionality. To mitigate this serious challenge, we developed computationally lightweight approaches to rapidly train large numbers of interpretability-constrained Generalized Additive Models (taking the form of scorecards) to form a representative sample approximating the Pareto frontier for maximizing Divergence across multiple definitions of positive and negative payment behaviors for multiple loan types and lifecycle stages.

A related challenge is to convey a human understanding of the resulting high-dimensional tradeoff. We will also demonstrate some new techniques we're developing for interactive visualization and exploration of high-dimensional Pareto surfaces.