Reinforcement Learning for Actuarial Credibility, with Extensions

Austin Riis-Due

University of Waterloo
200 University Ave W, Waterloo, ON N2L 3G1
Canada

Joint work with: David Landriault and Bin Li

This paper applies reinforcement learning techniques to the actuarial problem of determining a credibility weighting for claims data. We demonstrate that reinforcement learning (RL) based approaches, including both RecurrentPPO and DQN, outperform classical credibility methods, despite being restricted to a subclass of estimators and without the need for a Bayesian prior. The algorithm provides an assumption free way to determine the best credibility action given claims information and its performance over traditional methods is demonstrated in multiple simulation studies including cases with non Markovian drivers for claims.