Reinforcement Learning Approaches to the Problem of Actuarial Credibility

Austin Riis-Due[1], David Landriault[1], Bin Li[1]

Actuarial Credibility has evolved from the early 1900s as a method of determining how much confidence an actuary has in a particular collection of data. It is critical to the pricing assumptions underpinning new lines of business that initially contain little information in their emerging experience. Over time, it has become commonplace in practice to use the so-called square root rule, which assigns this weight solely based upon the number of points of data recorded. While this simple idea works in practice, we are able to show that reinforcement learning (RL) based approaches, including deep reinforcement learning based approaches such as RecurrentPPO and DQN, are able to outperform this traditional weighting scheme using data generated by multiple models, having similar performance to Bayesian methods despite being restricted to a subclass of estimators and without the need for an informative Bayesian prior. Our findings demonstrate a significant improvement in accuracy and adaptability in actuarial credibility assessment using reinforcement learning methods compared to practicing methodology. This work is a joint partnership with a North American Insurer.

The aim of this research is to fit the actuarial credibility problem into the reinforcement learning framework, so that issues such as requirements on the data to be conditionally independent and identically distributed data, requirements on the loss function such as symmetry and squared loss, and parametric assumptions can all be loosened. We achieve this by modifying the loss function and using RL algorithms that don't assume anything about the data inherently. We then allow the model to learn a weight for credibility directly, fitting the problem completely within the context of RL.

The principal results of the paper include that despite restricting ourselves to a linear subclass of estimators, we achieve a 46-57% reduction of the average pricing error in pure premium compared to American and Buhlmann methodologies and a similar performance to well-informed Bayesian models on real and fabricated data. We also implement alternative loss metrics to directly load the pure premium for margin and expenses, biasing the credibility problem and using it to price directly. We test scenarios with correlation both between policyholders and temporally and see 73-81% reductions in error compared to the industry benchmark of American credibility.

Major conclusions we draw are that the credibility problem is interesting as an RL problem and RL could be leveraged with big data to produce more accurate updates in real time of pure losses, helping insurers to reprice policies in real time and understand potential losses as they emerge. Moreover, RL can be used for a variety of statistical problems in actuarial science and multi player RL may have great impacts on insurer war games.

Main references include Klugman et al. (2019), Bühlmann and Gisler (2005), Bühlmann (1967), Bellman (1954), Mnih et al. (2013), Watkins (1992). Sutton et al. (1999)

Keywords: Reinforcement Learning, Actuarial Credibility, Stochastic Control, Group Insurance, Optimal Pricing

[1]University of Waterloo, Emails: acriisdu@uwaterloo.ca, david.landriault@uwaterloo.ca, bin.li@uwaterloo.ca

References

Bellman, R. (1954). The theory of dynamic programming. *Bulletin of the American Mathematical Society*, 60(6):503 – 515.

Bühlmann, H. and Gisler, A. (2005). *A Course in Credibility Theory and its Applications*. Springer.

Bühlmann, H. (1967). Introductory report experience rating and credibility. *ASTIN Bulletin*, 4(3):199 – 207.

Klugman, S. A., Panjer, H. H., and Willmot, G. E. (2019). *Loss Models: From Data to Decisions. Fifth edition.* Society of Actuaries.

Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. A. (2013). Playing atari with deep reinforcement learning. *CoRR*.

Sutton, R. S., McAllester, D., Singh, S., and Mansour, Y. (1999). Policy gradient methods for reinforcement learning with function approximation. In Solla, S., Leen, T., and Müller, K., editors, *Advances in Neural Information Processing Systems*, volume 12. MIT Press.

Watkins, C. (1992). Technical note: Q learning. *Machine Learning*, 8:279–292.