Is heteroscedastic regression a sound way to approximate model uncertainty?

Kiarash Aghakasiri, Farnaz Kohankhaki, Amirmohsen Sattarifard, Bedir Tapkan, Weicong Yuan

Instructor: Martha White[1]

[1]Department of Computing Science, University of Alberta

October 2020

Model-based reinforcement learning methods are much more sample efficient than model-free methods and have the same asymptotic performance given the true model of the environment. Although their performance drops significantly when the model of the environment is even slightly inaccurate. Since model learning is a supervised learning task, we have three sorts of uncertainty : 1) Aleatoric Uncertainty (Irreducible Error), 2) Structural Uncertainty (Bias), and 3) Parametric Uncertainty (Variance).

There are many ways to capture the parametric uncertainty (variance) such as Monte-Carlo dropout, model ensembles, and random prior functions, but there aren't as many for capturing structural and aleatoric uncertainties. A promising idea for capturing aleatoric and structural uncertainties is heteroscedastic regression (Nix and Weigend, 1994) which is a method for learning the mean and the variance of a probability distribution. Although there have been empirical studies showing the effectiveness of heteroscedastic regression to capture structural and aleatoric uncertainties (Zaheer et al., 2020), the theoretical part is missing. Thus, we plan to perform more concrete experiments to see what exactly can be captured by this learned variance and then investigate the theoretical side of this idea.

Références

[1] Nix, D. A., Weigend, A. S. (1994, June). Estimating the mean and variance of the target probability distribution. In Proceedings of 1994 ieee international conference on neural networks (ICNN'94) (Vol. 1, pp. 55-60). IEEE.

[2] Zaheer, M., Sokota, S., Talvitie, E. J., White, M. (2020). Selective Dyna-style Planning Under Limited Model Capacity. arXiv preprint arXiv :2007.02418.