Partially Factorized Variational Inference for overdispersed Poisson claim counts

Ruben Viduli *[1], George Tzougas †[2], and Giovanni Rabitti ‡[1]

[1]University of Trieste
[2]Heriot-Watt University

Abstract

Overdispersion, which occurs when the variability of claim counts exceeds that predicted by a Poisson model, is typically linked to unobserved heterogeneity across policyholders, reflecting latent differences in underlying risk not captured by observed covariates, and can be naturally accommodated by standard Bayesian mixed Poisson models by introducing latent random effects that induce extra-Poisson variation, providing a coherent probabilistic representation of risk heterogeneity and individual risk profiles. In this framework, the latent claim intensity can be interpreted as a policyholder-specific risk profile, whose posterior distribution summarizes both expected claim frequency and the uncertainty surrounding that risk assessment. Moreover, to account for heterogeneous variability across observed risk classes, we allow the dispersion structure itself to vary with covariates, so that both the mean and the degree of overdispersion may depend on policyholder characteristics. While these hierarchical models are highly flexible and capable of representing complex data structures, posterior inference using Markov Chain Monte Carlo (MCMC) methods can become very computationally demanding, particularly when applied to high-dimensional datasets with a large number of covariates. These computational demands often limit the practical applicability of fully Bayesian approaches in large-scale modeling, as the required time can make such analyses infeasible.

To address these challenges, we develop Variational Inference (VI) alternatives for a sequence of Poisson mixture models, with particular focus on the Poisson-Gamma (Negative Binomial), Poisson-Lognormal, and Poisson-Inverse Gaussian specifications. We implement both a standard Mean-Field VI (MFVI) approach and a Partially Factorized VI (PFVI) framework. MFVI assumes full independence among all model parameters, allowing for extremely fast posterior approximation but often at the expense of both the accuracy and the ability to capture posterior dependencies inherent in hierarchical models.

Our primary contribution lies in the PFVI formulation, which relaxes the full-independence assumption by preserving dependence structures within carefully selected parameter groups. This structured approximation more faithfully represents the hierarchical nature of mixed Poisson

*E-mail address: ruben.viduli@phd.units.it
†E-mail address: george.tzougas@hw.ac.uk
‡E-mail address: giovanni.rabitti@deams.units.it

models while retaining the scalability and efficiency advantages of VI. PFVI enables better uncertainty quantification by capturing posterior correlations between the key model parameters, thereby improving the approximation of posterior risk profiles and bridging the gap between computational tractability and inferential fidelity.

We conduct multiple simulation studies to evaluate both the posterior approximation accuracy and the computational efficiency of the proposed methods as the complexity increases, using MCMC as a benchmark to evaluate the results. Finally, the proposed methods are illustrated through a real-world case study using an open-access insurance claims dataset, where the resulting posterior distributions are used to characterize and compare policyholder risk profiles under covariate-dependent overdispersion.

Keywords: Variational Inference, Bayesian Hierarchical model, Claim Frequency, Individual Risk Profiling.

References

[1] Goplerud, M., Papaspiliopoulos, O. and Zanella, G. (2024), "Partially factorized variational inference for high-dimensional mixed models." *Biometrika*, vol. **112**(2).

[2] Blei, D. M., Kucukelbir, A. and McAuliffe, J. D. (2017), "Variational Inference: A Review for Statisticians." *Journal of the American Statistical Association*, vol. **112**(518), pp. 859–877.

[3] Fuzi, M. F. M., Jemain, A. A. and Ismail, N. (2016), "Bayesian quantile regression model for claim count data." *Insurance: Mathematics and Economics*, vol. **66**, pp. 124–137.

[4] Kim, S., Chen, Z., Zhang, Z., Simons-Morton, B. G. and Albert, P. S. (2013), "Bayesian Hierarchical Poisson Regression Models: An Application to a Driving Study With Kinematic Events." *Journal of the American Statistical Association*, vol. **108**(502), pp. 494–503.