An energy-based model approach to the estimation of rare event probabilities

L. Friedli[†,1], N. Linde[§,1], D. Ginsbourger[△,2], A. Doucet[3]

[†] PhD student (presenting author) [§] PhD supervisor [△] PhD co-supervisor

PhD expected duration: Nov. 2019 – Oct. 2023

[1] Institute of Earth Sciences, University of Lausanne, Switzerland
lea.friedli@unil.ch niklas.linde@unil.ch

[2] Institute of Mathematical Statistics and Actuarial Science, University of Bern, Switzerland
david.ginsbourger@stat.unibe.ch

[3] Department of Statistics, Oxford University, United Kingdom
doucet@stats.ox.ac.uk

Abstract

We consider non-linear Bayesian inversion aiming to infer unknown property fields $\boldsymbol{\theta}$ given measurements $\boldsymbol{y}$. In practice, one is often not interested in the field itself, but in the distribution of a quantity depending on the field through a non-linear relationship $\boldsymbol{\theta} \mapsto R(\boldsymbol{\theta})$. If one is interested in the probability of this quantity exceeding a critical threshold, $\mathbb{P}(R(\boldsymbol{\theta}) \geq T)$, this is related to the problem of evaluating the risk of failure of a system. Examples include the probability of hazardous contaminants arriving at a groundwater well or a satellite colliding with an object floating in space. For a small probability of failure, a traditional Monte Carlo approach is to be avoided as it requires an extensively large number of samples, which is unfeasible especially for expensive-to-simulate models.

We consider energy-based models (EBM) and write the marginal posterior distribution of the risk quantity $R = R(\boldsymbol{\theta})$ as an energy density function with free energy $r \mapsto F(r)$. To estimate $F(r)$ efficiently, the bias potential $r \mapsto V(r)$ and the corresponding PDF $r \mapsto p_V(r)$,

$$p_V(r) = \frac{\exp(-(F(r) + V(r)))}{\int \exp(-(F(s) + V(s)))ds}, \tag{1}$$

are introduced. Our approach relies on optimizing the potential $V(r)$ such that $p_V(r)$ approximates a pre-defined PDF $r \mapsto p(r)$. Finding the minimizing bias potential $V(r)$ is equivalent to finding the free energy $F(r)$, whereby, an accurate estimation in the region targeted within $p(r)$ is emphasized. This formulation reduces the potentially high-dimensional problem of estimating $p(\boldsymbol{\theta}|\boldsymbol{y})$ and subsequently $R(\boldsymbol{\theta})$ to an optimization of a one-dimensional function $V : \mathbb{R} \to \mathbb{R}$. This approach is based on the variational method of Valsson and Parrinello (2014) [4], which is related to metadynamics, a popular method developed in physics to perform free energy estimation.

In practice, the potential $V(r)$ is parameterized using methods such as neural networks, splines or radial basis functions. Subsequently, $V_{opt}(r)$ is approximated by minimizing the divergence between $p(r)$ and $p_V(r)$ using stochastic gradient descent. Traditionally, the variational method

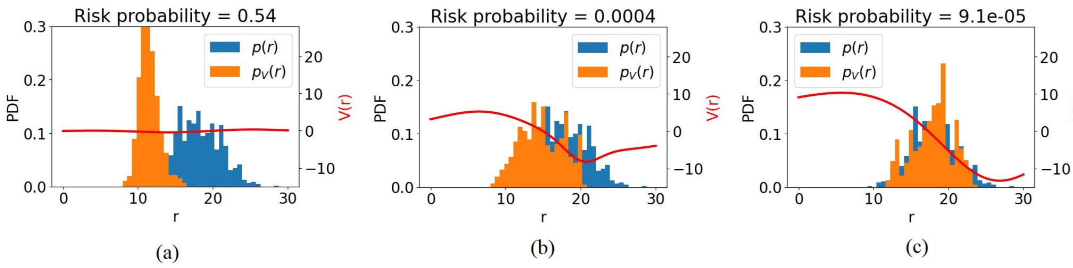

Figure 1: Illustration of the energy-based model approach for an exemplary setting targeting $\mathbb{P}(R(\boldsymbol{\theta}) \geq 18) = 8.8 \times 10^{-5}$ with $\boldsymbol{\theta} \in \mathbb{R}^9$, a multivariate Gaussian posterior $p(\boldsymbol{\theta}|\boldsymbol{y})$, $R(\boldsymbol{\theta})$ being the sum of squares, $p(r) = \mathcal{N}(18, 3)$ and $V(r)$ parameterized using radial basis functions. (a)-(c): Histograms depicting samples of $p_V(r)$ and $p(r)$ for an evolving potential $V(r)$. The corresponding estimates of $\mathbb{P}(R(\boldsymbol{\theta}) \geq 18)$ are denoted on top of the figures. While for the final estimate (c) we performed 40'000 model evaluations with the EBM approach, the traditional Markov chain Monte Carlo approach requires at least 5 million evaluations for a reasonable estimation of the risk.

employs a distance related to the Kullback-Leibler divergence [4][3]. We are presently investigating the possibility of using other distance measures such as the Kernel-Stein discrepancy [2]. Furthermore, as we know the pre-defined $p(r)$ analytically, ongoing work focuses on resulting possibilities for diagnostic tools and further enhancement of the approach.

We are currently testing the proposed EBM approach with applications from the geosciences and other domains. Thereby, we compare its performance with state-of-the-art methods and a new sequential application of the Sequential Monte Carlo method [1], which is explored in another project of the same research team. For different test cases, we show that the EBM approach ensures stable estimates of the risk even in situations when the probability of occurrence is less than one in a million. In the simple analytical test case of Figure 1, it enables accurate estimates with one hundred times less model evaluations than traditional Monte Carlo.

Short biography (PhD student)

Lea Friedli graduated with a MSc in Statistics from the University of Bern. Her PhD in Environmental Science takes place within the framework of the Swiss National Science Foundation project number 184574 on "GEOFACES: GEOphysics-based FAlsification and Corroboration in the Earth Sciences".

References

[1] Frédéric Cérou, Pierre Del Moral, Teddy Furon, and Arnaud Guyader. Sequential monte carlo for rare event estimation. *Statistics and computing*, 22(3):795–808, 2012.

[2] Will Grathwohl, Kuan-Chieh Wang, Jörn-Henrik Jacobsen, David Duvenaud, and Richard Zemel. Learning the stein discrepancy for training and evaluating energy-based models without sampling. In *International Conference on Machine Learning*, pages 3732–3747. PMLR, 2020.

[3] Michele Invernizzi, Omar Valsson, and Michele Parrinello. Coarse graining from variationally enhanced sampling applied to the ginzburg–landau model. *Proceedings of the National Academy of Sciences*, 114(13):3370–3374, 2017.

[4] Omar Valsson and Michele Parrinello. Variational approach to enhanced sampling and free energy calculations. *Physical Review Letters*, 113(9):090601, 2014.