Score matching in machine learning

Eduardo García-Portugués • Department of Statistics, UC3M

Machine learning models often rely on accurately estimating probability distributions from data, especially in unsupervised learning tasks. Score matching is an important method for fitting distribution models, offering a way to estimate parameters without needing to compute the partition function or normalizing constant – a common challenge in high-dimensional models and a requirement for applying maximum likelihood estimation. Originally introduced by Aapo Hyvärinen in 2005 (Hyvärinen, 2005), score matching has since evolved, with many improvements and extensions aimed at making it more robust and scalable (e.g., Song et al., 2019). Score matching has also been extended to handle data coming from non-Euclidean supports, such as data comprised by directions or proportions (Mardia et al., 2016; Scealy & Wood, 2023), where normalizing constants are especially complicated even in low-dimensional situations.

This thesis will focus on reviewing key developments in score matching, digesting their contributions, and proposing new applications. The goal is to understand the advantages and limitations of these methods and explore how score matching can be applied to modern machine learning problems. The thesis will also include numerical experiments to demonstrate the potential of score matching in real-world scenarios.

References

Hyvärinen, A. (2005). Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, *6*(Apr), 695–709. https://jmlr.org/papers/volume6/hyvarinen05a/hyvarinen05a.pdf

Mardia, K. V., Kent, J. T., & Laha, A. K. (2016). Score matching estimators for directional distributions. *arXiv:1604.08470*. https://doi.org/10.48550/arXiv.1604.08470

Scealy, J. L., & Wood, A. T. A. (2023). Score matching for compositional distributions. *Journal of the American Statistical Association*, *118*(543), 1811–1823. https://doi.org/10.1080/01621459.2021.2016422

Song, Y., Garg, S., Shi, J., & Ermon, S. (2019). Sliced score matching: A scalable approach to density and score estimation. *Conference on Uncertainty in Artificial Intelligence*. https://auai.org/uai2019/proceedings/papers/204.pdf