

Positive unlabeled data and prior shift estimation

Jan Mielniczuk^{1,2}, Wojciech Rejchel³, and
Paweł Teisseyre^{1,2}

¹ Polish Academy of Sciences, Poland

² Warsaw University of Technology, Poland

³ Nicolaus Copernicus University, Toruń, Poland

Abstract

We study estimation of a class prior for unlabeled target samples which possibly differs from that of the source population. Moreover, it is assumed that the source data is partially observable: only samples from the positive class and from the whole population are available (PU learning scenario) [1]. We introduce a novel direct estimator of the class prior which avoids estimation of posterior probability in both populations and has a simple geometric interpretation. It is based on a distribution matching technique together with kernel embedding in a Reproducing Kernel Hilbert Space and is obtained as an explicit solution to an optimisation task.

In [2], we establish its asymptotic consistency, as well as an explicit, nonasymptotic bound on its deviation from the unknown prior, which is computable in practice. We also study finite sample behaviour for synthetic and real data and show that the proposal works consistently on par or better than its competitors.

Keywords

Binary classification, Partially observed data, Distribution shift.

References

- [1] Bekker, J. and Davis, J. (2020). Learning from positive and unlabeled data: a survey. *Machine Learning* 109, 719–760.
- [2] Mielniczuk, J., Rejchel, W. and Teisseyre, P. (2026). Prior shift estimation for positive unlabeled data through the lens of Kernel embedding. In: *Proceedings of the 29th International Conference on Artificial Intelligence and Statistics (AISTATS)*.