05 Anomaly detection with low false positive rate

Semi-supervised learning for maximizing partial AUC

Abstract

The partial area under a receiver operating characteristic curve (pAUC) is a performance measurement for binary classification problems that summarizes the true positive rate with the specific range of the false positive rate. Obtaining classifiers that achieve high pAUC is important in a wide variety of applications, such as anomaly detection and medical diagnosis. Although many methods have been proposed for maximizing the pAUC, existing methods require many labeled data for training. We propose a semi-supervised learning method for maximizing the pAUC, which trains a classifier with a small amount of labeled data and a large amount of unlabeled data. To exploit the unlabeled data, we derive two approximations of the pAUC: the first is calculated from positive and unlabeled data, and the second is calculated from negative and unlabeled data. A classifier is trained by maximizing the weighted sum of the two approximations of the pAUC and the pAUC that is calculated from positive and negative data.

Partial AUC

Partial Area Under the Receiver Operating Characteristic Curve

PAUC : Area Under the ROC Curve when $\alpha \leq$ False positive rate $\leq \beta$

Evaluation measurement for binary classifiers when positive and negative sample size is unbalanced

Applications: Anomaly detection and medical diagnosis: reduce costs by reducing false alarm

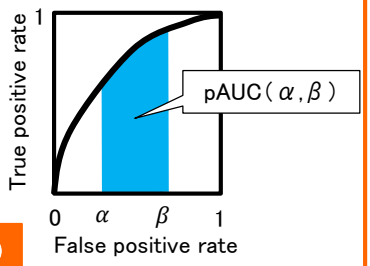

Existing method (supervised)

Inputs labeled data (positive and negative data), and learns score function s by maximizing the pAUC

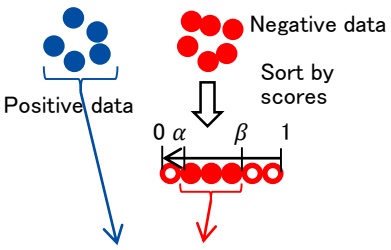

$$\mathbb{E}[I(s(\mathbf{x}^{\mathrm{P}}) > s(\mathbf{x}^{\mathrm{N}}))] = \mathrm{pAUC1}(\alpha, \beta)$$

pAUC calculated by positive and negative data

Objective function

$$L = \mathrm{pAUC1}(\alpha, \beta)$$

Proposed method (semi-supervised)

Derives two approximations of pAUC using unlabeled data, and uses them for learning the score function

Assumes data with high (low) scores as positive (negative) with the boundary of positive rate θ_{p}

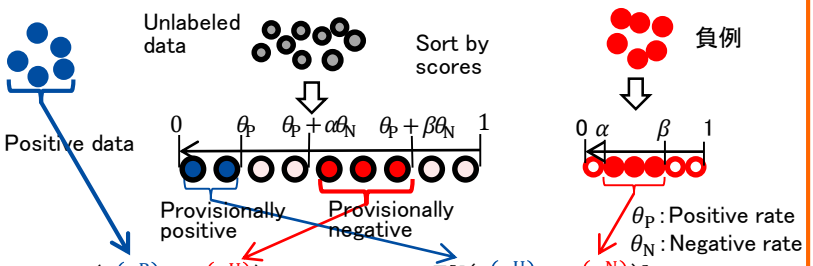

$$\mathbb{E}[I(s(\mathbf{x}^{\mathrm{P}}) > s(\mathbf{x}^{\mathrm{U}}))] \approx \mathrm{pAUC2}(\alpha, \beta)$$

Approximated pAUC calculated by positive and unlabeled data

$$\mathbb{E}[I(s(\mathbf{x}^{\mathrm{U}}) > s(\mathbf{x}^{\mathrm{N}}))] \approx \mathrm{pAUC3}(\alpha, \beta)$$

Approximated pAUC calculated by negative and unlabeled data

Objective function: weighted sum of three pAUCs

$$L = \lambda_1 \mathrm{pAUC1}(\alpha, \beta) + \lambda_2 \mathrm{pAUC2}(\alpha, \beta) + \lambda_3 \mathrm{pAUC3}(\alpha, \beta)$$

References

[1] T. Iwata, A. Fujino, N. Ueda, "Semi-supervised learning for maximizing the partial AUC," *Proc AAAI Conference on Artificial Intelligence (AAAI)*, 2020.

Contact

Tomoharu Iwata Email: cs-openhouse-ml@hco.ntt.co.jp
Ueda Research Laboratory

オープンハウス 2020