decision boundary

Author and editors

Alexander Jung*, Department of Computer Science, Aalto University, Espoo, Finland (ORCID 0000-0001-7538-0990)
Editors: Konstantina Olioumtsevits (ORCID 0000-0003-2697-8659), Ekkehard Schnoor (ORCID 0000-0001-8252-1499), Aalto University, Espoo, Finland
*Corresponding author: alex.jung@aalto.fi

Abstract

Definition

Consider a hypothesis $h\colon \mathcal{X} \to \mathcal{Y}$ whose finite label space $\mathcal{Y}$ represents different classes. Its decision boundary consists of those feature vectors that are arbitrarily close to two or more different decision regions. More precisely, $\mathbf{x} \in \mathcal{X}$ lies on the decision boundary if every neighborhood of $\mathbf{x}$ contains feature vectors $\mathbf{x}'$ and $\mathbf{x}''$ such that $h(\mathbf{x}') \neq h(\mathbf{x}'')$. Thus, arbitrarily small perturbations near the decision boundary can result in different predicted labels (or classes).

A common construction of a classifier is by thresholding a real-valued scoring function $g\colon \mathcal{X} \to \mathbb{R}$. The classifier assigns different labels according to whether $g(\mathbf{x})$ lies below or above a threshold c. If g is continuous, its decision boundary is the level set $\{\mathbf{x} : g(\mathbf{x}) = c\}$, provided that g takes values on both sides of c arbitrarily close to every point in this set. For a linear classifier with $g(\mathbf{x}) = \mathbf{w}^{\top}\mathbf{x}$ and $\mathbf{w} \neq \mathbf{0}$, this level set is the hyperplane $\mathbf{w}^{\top}\mathbf{x} = c$, whose normal vector is $\mathbf{w}$ (Bishop, 2006, Sect. 4.1; Hastie et al., 2009, Ch. 4).

Fig. 1 shows that the geometry of a decision boundary depends on the model. A linear classifier produces a straight boundary, a decision tree produces an axis-aligned, piecewise constant boundary, and logistic regression applied to polynomial features can produce a curved boundary.

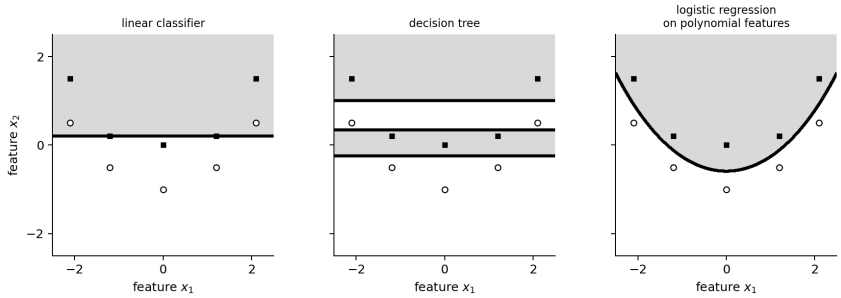


Figure 1: Decision regions learned from the same ten-point training set by three classifiers. Gray marks the region assigned the label +1; the thick black lines are the learned decision boundaries. Open circles have label −1, and filled squares have label +1. The experiment is generated by `pythondemos/decisionboundary.py`

Cross References

hypothesis, feature vector, label space, boundary, decision region, neighborhood

References

1. Bishop (2006). Pattern Recognition and Machine Learning. Springer Science+Business Media, New York, NY, USA.

2. Hastie et al. (2009). The Elements of Statistical Learning: Data Mining, Inference, and Prediction. 2nd ed., Springer Science+Business Media, New York, NY, USA.