

An open question in ensemble-based active learning is how to choose one classifier type, or appropriate combinations of multiple classifier types, to construct ensembles for a given task. While existing approaches typically choose one classifier type, this paper presents a method that trains and adapts multiple instances of multiple classifier types toward an appropriate ensemble during active learning. The method is termed Adaptive Heterogeneous Ensembles (henceforth referred to as AHE). Experimental evaluations show that AHE constructs heterogeneous ensembles that outperform homogeneous ensembles composed of any one of the classifier types, as well as bagging, boosting and the random subspace method with random sampling. We also show in this paper that the advantage of AHE over other methods is increased if (1) the overall size of the ensemble also adapts during learning; and (2) the target data set is composed of more than two class labels. Through analysis we show that the AHE outperforms other methods because it automatically discovers complementary classifiers: for each data instance in the data set, instances of the classifier type best suited for that data point vote together, while instances of the other, inappropriate classifier types disagree, thereby producing a correct overall majority vote.