

In the past few years, complex neural networks have achieved state of the art results in image classification. However, training these models requires large amounts of labelled data. Whereas unlabelled images are often readily available in large quantities, obtaining labels takes considerable human effort. Active learning reduces the required labelling effort by selecting the most informative instances to label. The most popular active learning query strategy framework, uncertainty sampling, uses uncertainty estimates of the model being trained to select instances for labelling. However, modern classification neural networks often do not provide good uncertainty estimates. Bayesian neural networks model uncertainties over model parameters, which can be used to obtain uncertainties over model predictions. Exact Bayesian inference is intractable for neural networks, however several approximate methods have been proposed. We experiment with three such methods using various uncertainty sampling active learning query strategies.