Abstract

In this thesis, we conduct a comparative analysis of how various modern machine learning techniques perform when employed to asset return prediction on a relatively small sample. We consider a broad selection of machine learning methods, including e.g. elastic nets, random forests or recently highly popularized neural networks. We find that these methods fail to outperform a simple linear model containing only 5 factors and estimated via ordinary least squares. Our conclusion is that applications of machine learning in finance should be conducted carefully, because the techniques may not actually be as powerful as one might think when they are applied under unfavorable circumstances.