論文の内容の要旨

論文題目

Reliable Machine Learning from Limited Data and Supervision: A Risk Modification Approach

(限られたデータと教師からの高信頼機械学習：リスク修正アプローチ)

氏名 石田 隆

We are surrounded by an immense number of automated and intelligent systems empowered by machine learning in our daily lives. We are starting to take machine learning applications such as object detection, handwriting recognition, speech recognition, and text generation for granted, which were not easily available just a decade ago.

Although machine learning seems to have been successfully applied in the real world, building a practical machine learning system is still extremely difficult. For example, collecting a correct class label can be difficult. In this case, we may instead need to learn from *weak supervision*, where the supervision of the data can be given in a weaker or alternative form than the usual supervision. Another example is learning from *limited data*. This may cause overfitting issues, especially when models with a large capacity is used. These examples are common real-world scenarios due to high data collecting and data labeling costs, and various practical constraints. In this dissertation, we focus on making machine learning more *reliable* under weak supervision or limited data. We summarize our contributions as follows.

In Chapter 3, we introduce our novel problem setting of learning a binary classifier from only positive data, without any negative data or unlabeled data. As an example, this task is important in purchase prediction. We can easily collect customer data from our own company (positive data), but not from rival companies (negative data). However, even in this challenging scenario, we still want to perform binary classification. Our finding is that if one can equip positive data with confidence (positive-confidence), one can successfully learn a binary classifier with the optimal convergence rate to solve this problem, which we call *positive-confidence classification*. The key technique is to reformulate the classification risk into a formulation that only requires the positive-confidence data, even though naively computing the classification risk requires both positive and negative data. This leads to a simple empirical risk minimization framework that is model-, loss-, and optimization-independent. We also show the consistency and an estimation error bound for positive-confidence classification, and experimental results showing the effectiveness of our method.

In Chapter 4, we discuss a new type of weak supervision called *complementary labels*, which are helpful for multi-class classification. A complementary label specifies a class that a

pattern does *not* belong to. Collecting complementary labels would be less laborious than collecting ordinary labels, since annotators/labelers do not have to carefully choose the correct class from a long list of candidate classes. However, complementary labels are less informative than ordinary labels and thus a suitable approach is needed to better learn from them. For this challenging problem, we show that an unbiased estimator to the classification risk can be obtained only from complementarily labeled data, if a loss function satisfies a particular symmetric condition. We then derive estimation error bounds for the proposed method and prove that the learned classifier achieves the optimal convergence rate. We also show that learning from complementary labels can be easily combined with learning from ordinary labels, i.e., ordinary supervised learning, resulting in even better generalization performance. We further extend the unbiased risk estimator to arbitrary losses and models, and improve it by a non-negative correction and a gradient ascent trick. We show experimental results demonstrating the usefulness of our approach.

In Chapter 5, we introduce a novel regularizer that can be used to avoid overfitting. Overparameterized deep networks have the capacity to make the empirical risk go to zero, resulting in an overconfident model with degraded test performance. While previous regularization methods indirectly cope with this problem, we propose a direct solution called *flooding* that intentionally prevents further reduction of the empirical risk when it reaches a reasonably small value, which we call the flood level. Our approach makes the flooded empirical risk float around the flood level by performing mini-batched gradient descent as usual but gradient *ascent* if the empirical risk is below the flood level. This can be implemented with one line of code and is compatible with any stochastic optimizer and other regularizers. With flooding, we will have a "random walk" with the same non-zero empirical risk, and we expect the model to go into an area with a flat loss landscape that leads to better generalization. We experimentally show that flooding improves the generalization performance and as a byproduct, induces a double descent curve of the test loss.

In Chapter 6, we summarize our contributions in this dissertation and conclude. We discuss possible extensions for positive-confidence classification, complementary-label classification, and flooding. Finally, we discuss future directions for reliable machine learning.

In summary, this dissertation is devoted to making machine learning more reliable under limited data and supervision. A common approach in all of our methods is *risk modification*. We start by observing the classification risk as the final target which we want to minimize by training our classifier. Then, we take a modification step by either rewriting or correcting the risk. Risk rewriting is used in positive-confidence learning and complementary-label learning. Since we do not have access to fully labeled data, we rewrite the classification risk in an alternative formulation that utilizes weakly supervised data. This leads to a theoretically

grounded algorithm with an unbiased estimator of the classification risk and an estimation error bound of the learned classifier achieving the optimal convergence rate. Another technique is risk correction. In our regularization method, we employ this risk correction technique, by putting a lower-bound on the empirical risk. We also use a risk correction technique in complementary-label learning to avoid severe overfitting. A notable advantage of the risk modification approach is that it usually produces a framework that can be applied to a variety of domains, optimizers, and models. To conclude, this dissertation demonstrates that risk rewriting and risk correction can be a powerful approach in building practical and useful algorithms for learning from limited data and supervision.

(973 words/2000 words)