Abstract

Despite intensive research in the last years, the biometric authentication problem is still not fully solved; thus, it is not deployed in critical sectors. In this dissertation, we first consider a generic biometric authentication scenario in which the framework can be applied to any biometric trait, i.e., face, fingerprint, iris, and so on. The features extracted from biometric traits are mapped onto a latent space such that authorized and unauthorized users follow simple and well-behaved distributions. We show that, by learning a regularized mapping instead of a classification boundary, higher performance and improved robustness is achieved. Secondly, a deep unconstrained face verification scenario is considered. In the proposed approach, no specific metric on facial features is imposed; instead, the decision space is shaped by learning a latent representation in which matching and non-matching pairs are mapped onto clearly separated and well-behaved target distributions.

The second part of the dissertation focuses on a robust and accurate classification system using deep learning. In recent years, deep learning has shown outstanding performance in several applications, including image classification. However, deep classifiers are known to be highly vulnerable to adversarial attacks, in that a minor perturbation of the input can easily lead to an error. Providing robustness to adversarial attacks is a very challenging task, especially in problems involving a large number of classes, as it typically comes at the expense of a reduced accuracy. We propose a novel approach for training deep, robust multiclass classifiers that provide adversarial robustness while at the same time achieving or even surpassing the classification accuracy of state-of-the-art methods.