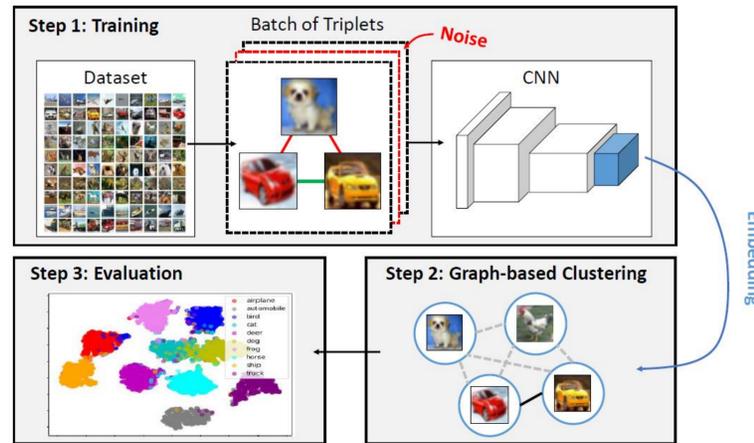


Abstract

In this work, we evaluate two different image clustering objectives, k-means clustering and correlation clustering, in the context of Triplet Loss induced feature space embeddings. Specifically, we train a convolutional neural network to learn discriminative features by optimizing two popular versions of the Triplet Loss in order to study their clustering properties under the assumption of noisy labels.



Study Setup

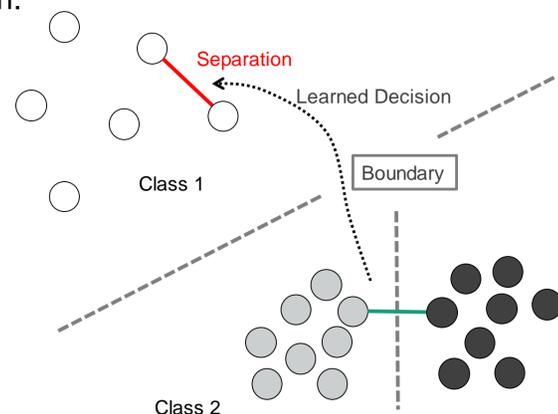
Our approach is based on the assumption that embedding features, learned from the regular Triplet Loss (1) can produce high variances in inter- and intra-cluster distances, because it only considers relative differences between the distances of positive and negative pairs. This objective is suitable for k-means clustering. Yet, the attempt to learn whether two data points should belong to the same or to a different class from their pairwise distances might fail, when the intra- and inter cluster samples are equally far away. This is shown in the following Figure where the correct decision boundaries are marked by green lines. In contrast, the red line, at the same Euclidean distance as the green lines, indicates a false separation of data. This motivates us to consider losses that preserve the distance equally between the positive pairs during the optimization.

Triplet Losses

$$L_{\text{triplet}} = \sum_{i=1}^n [\|f(x_i^a) - f(x_i^p)\|^2 - \|f(x_i^a) - f(x_i^n)\|^2 + \alpha]_+$$

$$L_{\text{triplet}_2} = L_{\text{triplet}} + [\|f(x_i^a) - f(x_i^p)\|^2 - \beta]_+$$

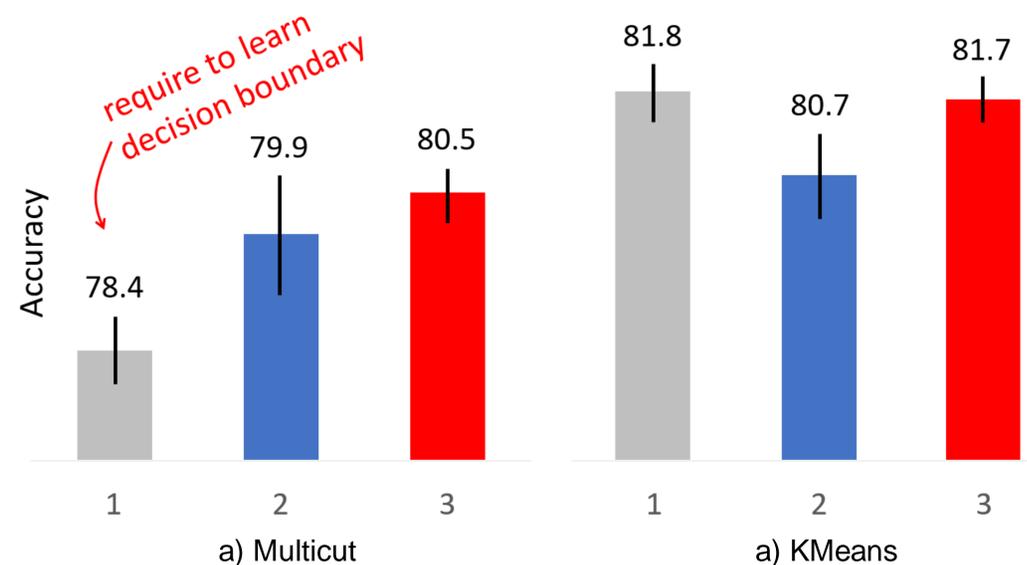
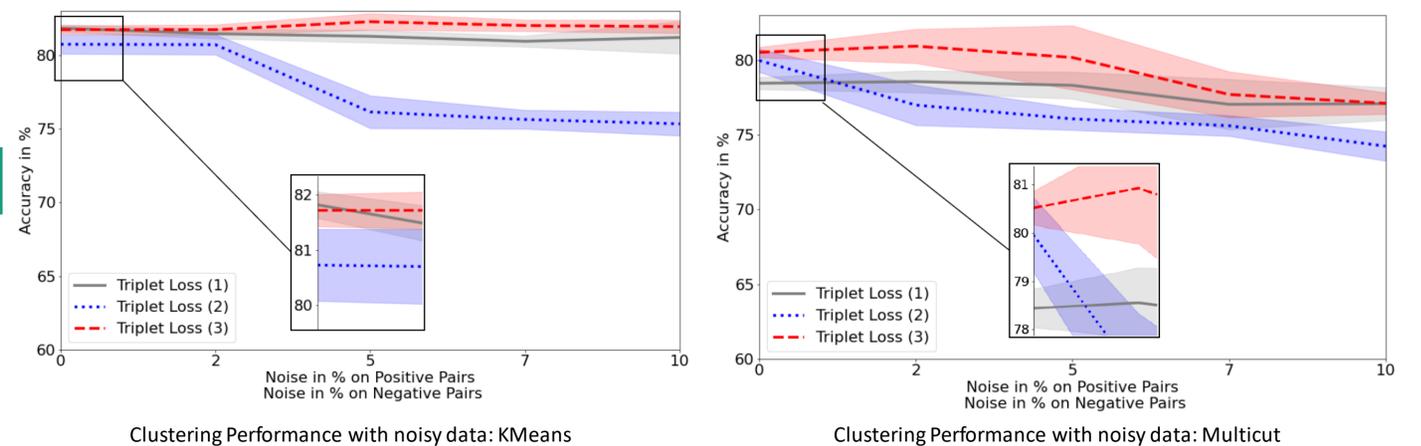
$$L_{\text{triplet}_3} = \sum_{i=1}^n [\alpha - \|f(x_i^a) - f(x_j^n)\|^2]_+ + [\|f(x_i^a) - f(x_k^p)\|^2 - \beta]_+$$



Contribution

1) We conduct a thorough study of the clustering behavior of two clustering approaches, k-means and minimum cost multicut, applied to learnt embedding spaces from three Triplet Loss formulations. 2) Our study reveals that, while the traditional Triplet Loss is well suited for k-means clustering, its performance drops under the looser assumptions made by minimum cost multicut. 3) We propose a simplification of the Triplet Loss from (3), which allows to directly compute the probability of two data points for belonging to disjoint components and is robust against noise in both clustering scenarios.

Results



Clustering Performance: Multicut vs. KMeans