Beyond cross-entropy: learning highly separable feature distributions for robust and accurate classification

Arslan Ali, Andrea Migliorati, Tiziano Bianchi, Enrico Magli
Politecnico di Torino

GCCS Architecture

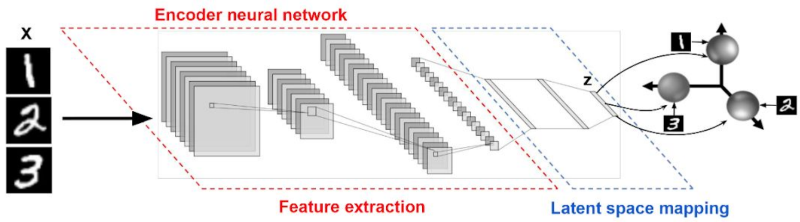

Fig. 1: The GCCS architecture takes input data and learns discriminative features that are mapped onto Gaussian target distributions in the latent space.

- GCCS: Gaussian class-conditional simplex loss
- High Separability between classes
- Extract discriminative features from the input data
- Map the features to well behaved, target Gaussian distributions

GCCS loss

Proposed loss

Authorized users loss

$$\mathcal{L}_i = \frac{1}{i}\left[\log\frac{|\Sigma_{Ti}|}{|\Sigma_{Oi}|} - D + \mathrm{tr}(\Sigma_{Ti}^{-1}\Sigma_{Oi}) + (\boldsymbol{\mu}_{Ti} - \boldsymbol{\mu}_{Oi})^{\intercal}\Sigma_{Ti}^{-1}(\boldsymbol{\mu}_{Ti} - \boldsymbol{\mu}_{Oi})\right]$$

$$\mathcal{K}_i = \left(\frac{x-\boldsymbol{\mu}_{Oi}}{\sigma_{Oi}}\right)^4$$

- target statistics
- batch statistics

Total loss: $\mathcal{L}^{\mathrm{GCCS}} = \sum_{i=1}^{D}\left[\mathcal{L}_i + \lambda(\mathcal{K}_i - 3)\right],$

GCCS decision rule

- Partition the decision space into Voronoi regions
- Compute the distance from all the centers - choose minimum

$$\widehat{y} = \arg\max_i z_i,$$

- Index of the predicted class for the test image