MLBoost Revisited: A Faster Metric Learning Algorithm for Identity-Based Face Retrieval

Romain Negrel romain.negrel@unicaen.fr Alexis Lechervy alexis.lechervy@unicaen.fr Frederic Jurie frederic.jurie@unicaen.fr Normandie Univ, UNICAEN, ENSICAEN, CNRS France

This paper focuses on the problem of identitybased face retrieval [2], a problem heavily depending on the quality of the similarity function used to compare the images. Instead of using standard or handcrafted similarity functions, one of the most popular ways to address this problem is to learn adapted metrics, from sets of similar and dissimilar example pairs. This is generally equivalent to projecting the face signatures into an adapted (possibly low-dimensional) space in which the similarity can be measured by the Euclidean distance. For large scale applications, the dimension of this subspace should be as small as possible to limit the storage requirements, and the projections should be fast to compute. Since the Euclidean distance fulfill the second requirement, producing face representations adapted to the Euclidean metric is interesting. However, such representations usually have very large sizes. Several methods have been proposed to learn projections that are capable of reducing the size of the signatures while preserving their performance. Most of these approaches are based on metric leaning algorithms [1] used to learn Mahalanobis-like distances:

$$D_{\mathbf{W}}(\mathbf{x}_i, \mathbf{x}_j) = (\mathbf{x}_i - \mathbf{x}_j)^\top \mathbf{W}(\mathbf{x}_i - \mathbf{x}_j), \quad (1)$$

with **W** a positive semi-definite matrix. To guarantee this property and to reduce the size of the signatures, these methods use the factorization $\mathbf{W} = \mathbf{L}\mathbf{L}^{\top}$ with $\mathbf{L} \in \mathcal{M}_{D \times d}$ as projection matrix: $\mathbf{y}_i = \mathbf{L}^{\top} \mathbf{x}_i$. It is important to control the rank of **W** so that the dimension of the reduced signature is as small as possible.

In this paper, we focus on a particular metric learning algorithm so-called MLBoost [3], a supervised method based on boosting. MLBoost learns the metric incrementally by aggregating several weak metrics:

$$D_{\mathbf{W}}(\mathbf{x}_i, \mathbf{x}_j) = \sum_t \alpha^{(t)} \mathbf{D}_{\mathbf{z}^{(t)}} \left(\mathbf{x}_i, \mathbf{x}_j \right)$$
(2)

-		-	1	-		-
Sign.	Final	n=	n=	n=	n=	n=
	Dim.	1	10	20	50	100
$\begin{array}{c} \text{LBP} \\ D = 9860 \end{array}$	FULL	31.9	53.7	60.5	68.8	74.7
	16	18.7	43.5	52.7	64.3	74.0
	32	31.4	57.0	63.1	72.1	77.3
	128	36.4	54.8	62.9	71.6	77.5
	512	38.5	58.6	63.6	74.0	79.2
VGG-Face $D = 4096$	FULL	89.6	96.9	97.4	98.1	98.3
	16	82.0	94.1	96.7	97.6	98.6
	32	89.4	96.2	97.4	98.1	98.6
	128	90.8	95.7	97.2	98.1	98.8
	512	92.4	96.7	97.6	98.3	98.6

Table 1: Performance of MLBoost with low-cost weak metrics ($\tau = 5\%$) and rank constraints ($R \in \{16, 32, 128, 512\}$) for face retrieval on the LFW Dataset [2] ("FULL" rows correspond to the performance of the non-reduced original signatures).

with $\alpha^{(t)}$ the weights of the weak metrics and $\mathbf{z}^{(t)}$ the projector vectors of these weak metrics. Here, we propose two improvements over ML-Boost [3]: (i) A new method for **computing weak metrics at a lower computational cost**; (ii) An explicit way **to control the rank of the produced metric** during the learning phase, allowing to fix the size of the low-dimensional space in which the images are represented.

Our experiments demonstrated a more than $10 \times$ speedup in addition to a significant improvement in the performance of the signatures even when their dimensions are strongly reduced (Table 1).

- [1] A. Bellet, A. Habrard, and M. Sebban. *Metric Learning*. Morgan & Claypool Publishers, 2015.
- [2] B. Bhattarai, G. Sharma, F. Jurie, and P. Pérez. Some faces are more equal than others: Hierarchical organization for accurate and efficient largescale identity-based face retrieval. In ECCV Workshops, pages 1–13, 2014.
- [3] R. Negrel, A. Lechervy, and F. Jurie. Boosted metric learning for efficient identity-based face retrieval. In *BMVC*, volume 13, pages 1007–1036, 2015.