Highly Efficient Regression for Scalable Person Re-Identification

Hanxiao Wang
hanxiao.wang@qmul.ac.uk
Shaogang Gong
s.gong@qmul.ac.uk
Tao Xiang
t.xiang@qmul.ac.uk

Vision Group,
School of Electronic Engineering and
Computer Science,
Queen Mary, University of London,
London E1 4NS, UK

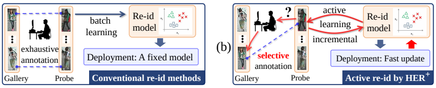

Figure 1: (a) Conventional re-id: A re-id model is trained on a fully labelled training set, then fixed for deployment; (b) Active re-id by HER$^+$: A training set is actively labelled incrementally on-the-fly as a re-id model is incrementally learned, and further updated without re-training during future deployment.

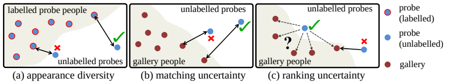

Figure 2: Joint exploration-exploitation criteria for active re-id.

This work is motivated by two very intuitive requirements for a scalable re-id system [2]: (1) Low model complexity with scalable computational cost and memory usage in model training; and (2) High model adaptability supporting fast model update to incorporate any new and increasingly larger data. A Highly Efficient Regression (HER) model is formulated by embedding the Fisher's criterion to a ridge regression model for very fast re-id model learning with scalable memory/storage usage. Importantly, this new HER model supports faster than real-time incremental model updates therefore making real-time active learning feasible in re-id with human-in-the-loop (Fig. 1).

Our Highly Efficient Regression (HER) solution for re-id has a very simple and fast closed-form solution, involved with only a set of linear equations. It is readily scalable to large data with many off-the-shelf efficient implementation available. The base HER model for adopts the form of minimising a least mean squared error:

$$\mathbf{P} = \arg\min_{\mathbf{P}} \ \tfrac{1}{2}\|\mathbf{X}^\top\mathbf{P} - \mathbf{Y}\|_F^2 + \lambda\|\mathbf{P}\|_F^2, \qquad (1)$$

where $\mathbf{X} \in \mathbb{R}^{d\times n}$ refers to the labelled data, and $\mathbf{P} \in \mathbb{R}^{d\times k}$ refers to the discriminative projection to be learned. To make the estimated subspace person identity discriminative, FDA [1] criterion are further embedded. Moreover, to incorporate new and increasingly larger data in a real-world, we further introduce an incremental learning formulation HER$^+$, enabling fast model updates without the need for re-training from scratch.

The efficient model updates achieved by HER$^+$ makes makes *active learning re-id* with human-in-the-loop feasible with reduced human labelling costs. A joint exploration-exploitation (*joint*E^2) active sampling strategy is further proposed (Fig. 2). Three criterion are considered for selecting most useful samples to maximise the re-id model's discriminative power (1) Appearance diversity exploration, (2) Matching uncertainty exploitation, and (3) Ranking uncertainty exploitation. Finally, these criterion are combined into the final active sampling strategy.

For experimental results, when evaluated under the conventional supervised re-id setting on three popular re-id benchmarks, VIPeR, CUHK01, and CUHK03, HER achieves Rank-1 rates of 45.1%, 68.3% and 60.8% respectively, outperforms all existing competitors. The computational efficiency of HER is also evaluated and it is shown that HER is the fastest in batch training over other state-of-the-art models. When evaluated under the active re-id setting where a model is trained incrementally, it is shown that: (1) HER$^+$ incremental updates is much more efficient than re-training from scratch; and (2) The proposed *joint*E^2 sampling strategy effectively reduces human labelling effort and achieves better re-id performances.

[1] Ronald A Fisher. The use of multiple measurements in taxonomic problems. *Annals of eugenics*, 7(2):179–188, 1936.

[2] Shaogang Gong, Marco Cristani, Change Loy Chen, and Timothy M. Hospedales. The re-identification challenge. In *Person Re-Identification*. Springer, 2014.