

Exploiting Prunability for Person Re-identification

Hugo Masson

Ecole de technologie superieure

Md Amran Hossen Bhuiyan (✉ amran.apece@gmail.com)

École de technologie supérieure: Ecole de technologie superieure <https://orcid.org/0000-0002-2069-0753>

Le Than Nguyen-Meidine

Ecole de technologie superieure

Mehrsan Javan

Sportlogiq Inc.

Parthipan Siva

Sportlogiq Inc.

Ismail Ben Ayed

Ecole de technologie superieure

Eric Granger

Ecole de technologie superieure

Research

Keywords: Deep Learning, Convolutional Neural Networks, Complexity, Pruning, Domain Adaptation, Person Re-identification

Posted Date: October 7th, 2020

DOI: <https://doi.org/10.21203/rs.3.rs-84743/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Version of Record: A version of this preprint was published at EURASIP Journal on Image and Video Processing on June 25th, 2021. See the published version at <https://doi.org/10.1186/s13640-021-00562-6>.

Abstract

Recent years have witnessed a substantial increase in the deep learning (DL) architectures proposed for visual recognition tasks like person re-identification, where individuals must be recognized over multiple distributed cameras. Although these architectures have greatly improved the state-of-the-art accuracy, the computational complexity of the CNNs commonly used for feature extraction remains an issue, hindering their deployment on platforms with limited resources, or in applications with real-time constraints. There is an obvious advantage to accelerating and compressing DL models without significantly decreasing their accuracy. However, the source (pruning) domain differs from operational (target) domains, and the domain shift between image data captured with different non-overlapping camera viewpoints leads to lower recognition accuracy. In this paper, we investigate the prunability of these architectures under different design scenarios. This paper first revisits pruning techniques that are suitable for reducing the computational complexity of deep CNN networks applied to person re-identification. Then, these techniques are analysed according to their pruning criteria and strategy, and according to different scenarios for exploiting pruning

methods to re-tuning networks to target domains. Experimental results obtained using DL models with ResNet feature extractors, and multiple benchmarks re-identification datasets, indicate that pruning can considerably reduce network complexity while maintaining a high level of accuracy. In scenarios where pruning is performed with large pre-training or re-tuning datasets, the number of FLOPS required by ResNet architectures is reduced by half, while maintaining a comparable rank-1 accuracy (within 1% of the original model). Pruning while training a larger CNNs can also provide a significantly better performance than re-tuning smaller ones.

Full Text

This preprint is available for [download as a PDF](#).

Figures

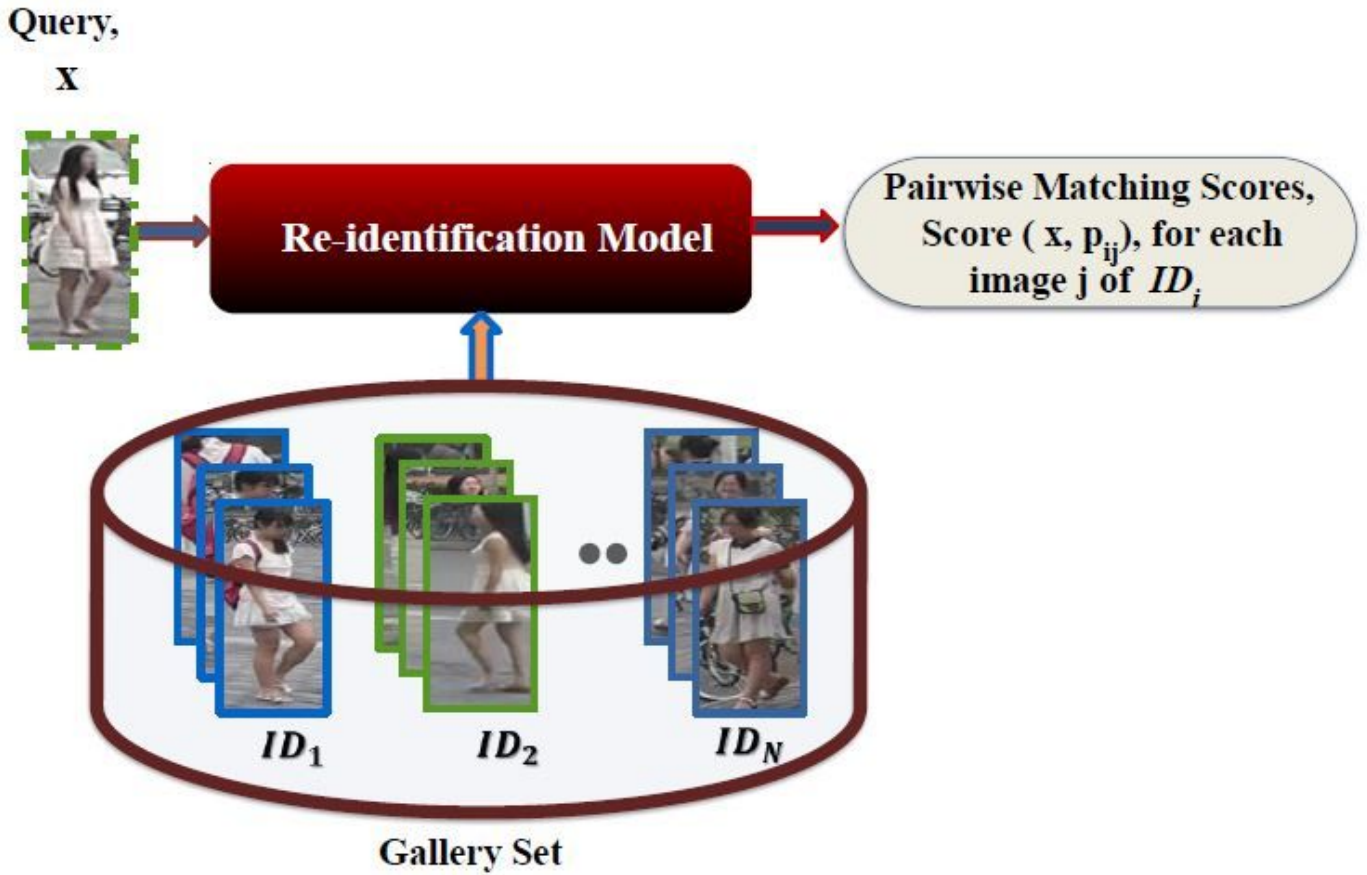


Figure 1

Illustration of a typical person re-identification system.

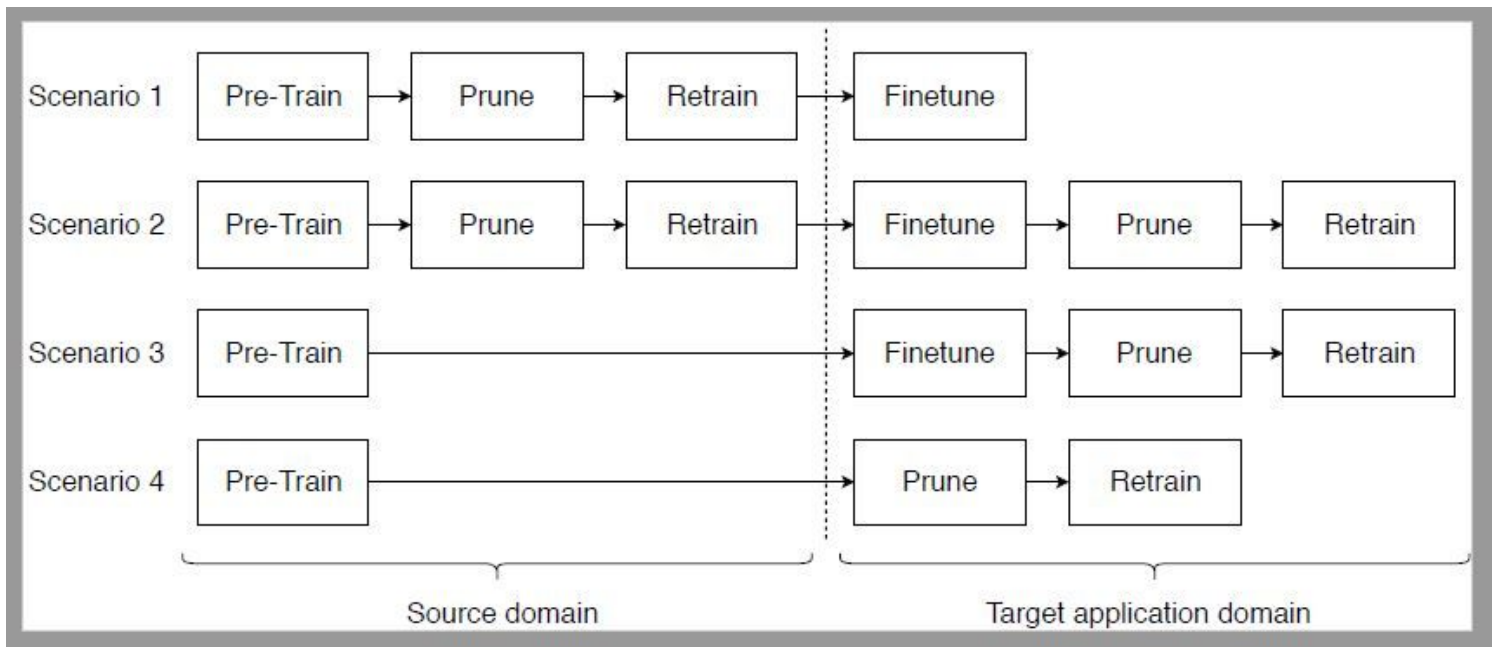


Figure 2

Scenarios for pruning and training a CNN.

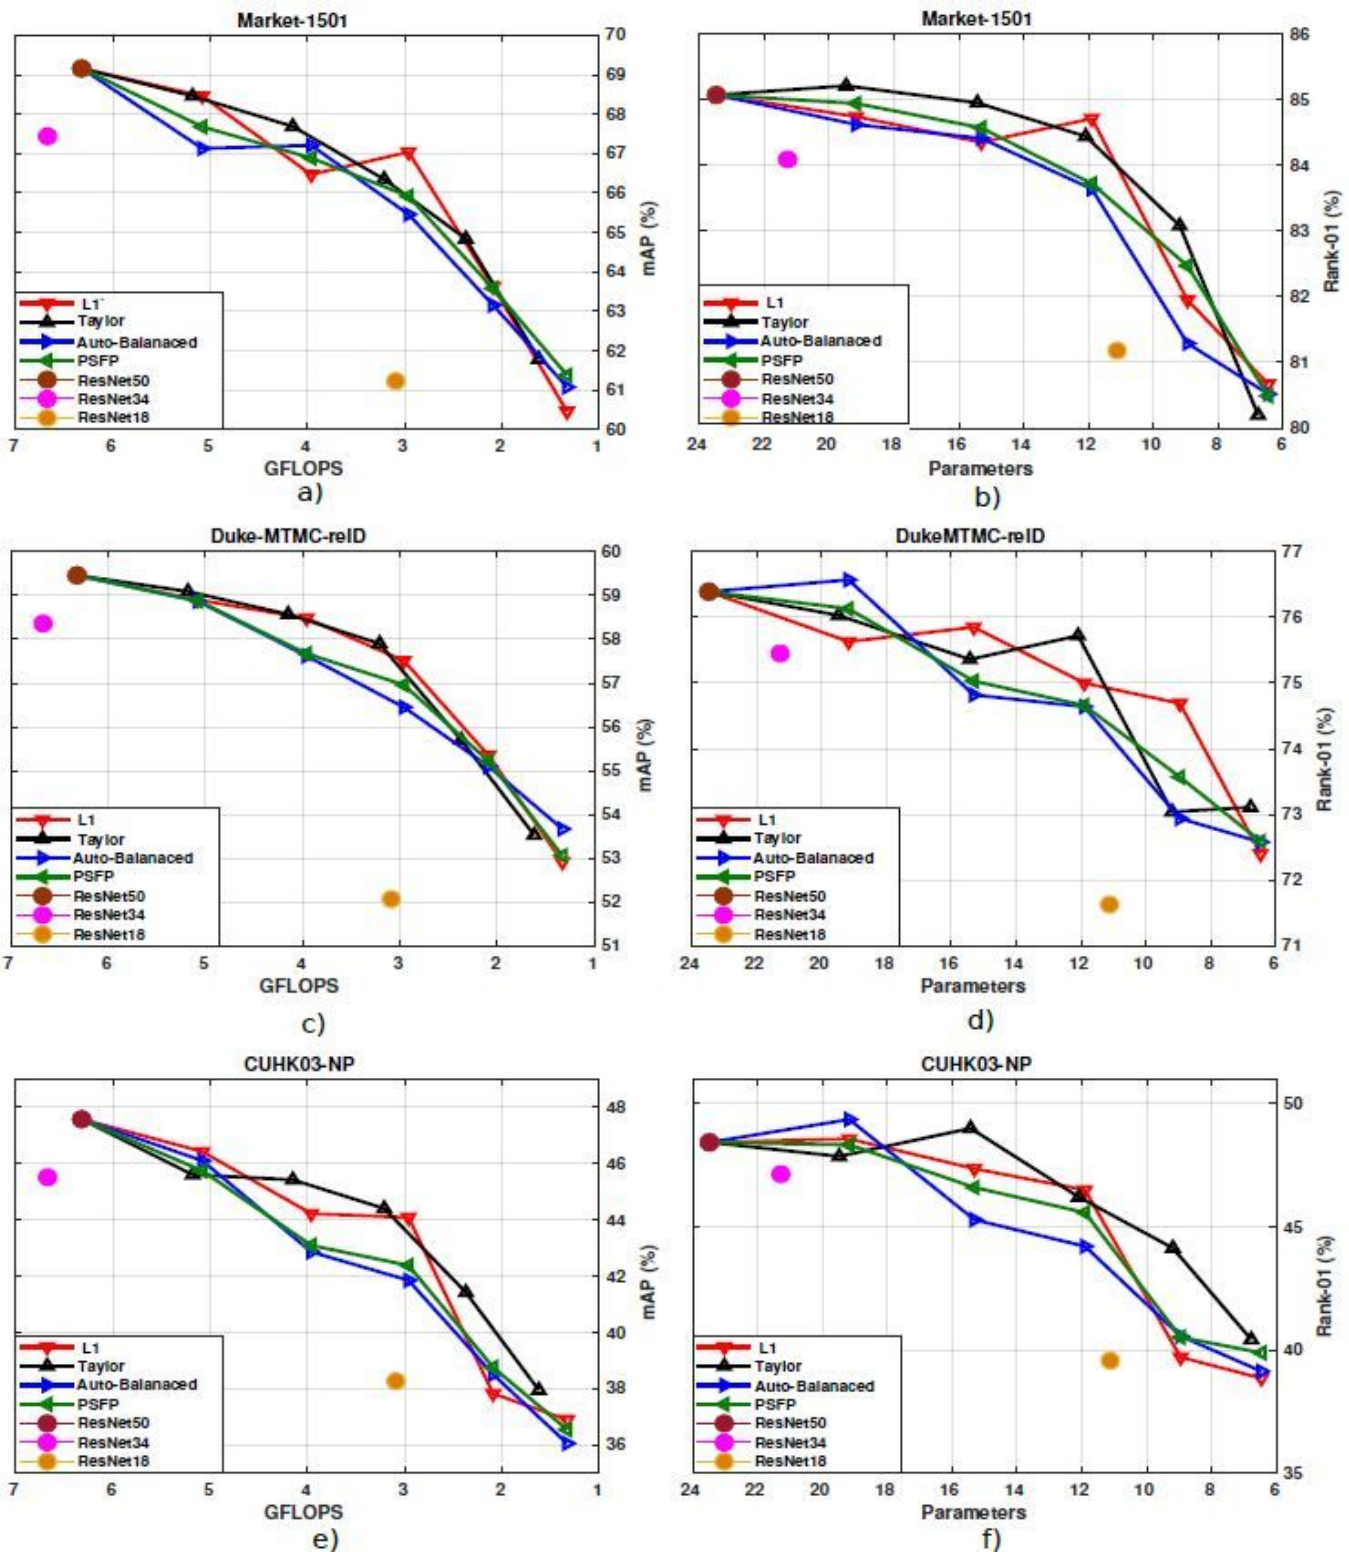


Figure 3

Comparative ReID performance analysis of the pruning methods for all the ReID datasets: (a,c,e) mAP vs GFLOPS and (b,d,f) Parameters vs Rank-01.