Evaluating the Impact of Low-Light Image Enhancement Methods on Runner Re-Identification in the Wild

Oliverio J. Santana[a], Javier Lorenzo-Navarro[b], David Freire-Obregón[c], Daniel Hernández-Sosa[d] and Modesto Castrillón-Santana[e]
SIANI, Universidad de Las Palmas de Gran Canaria, Spain

Keywords: Computer Vision, Person Re-Identification, in the Wild Sporting Events, Low-Light Image Enhancement.

Abstract: Person re-identification (ReID) is a trending topic in computer vision. Significant developments have been achieved, but most rely on datasets with subjects captured statically within a short period of time in rather good lighting conditions. In the wild scenarios, such as long-distance races that involve widely varying lighting conditions, from full daylight to night, present a considerable challenge. This issue cannot be addressed by increasing the exposure time on the capture device, as the runners' motion will lead to blurred images, hampering any ReID attempts. In this paper, we survey some low-light image enhancement methods. Our results show that including an image processing step in a ReID pipeline before extracting the distinctive body appearance features from the subjects can provide significant performance improvements.

1 INTRODUCTION

Our ability to effortlessly identify all subjects of an image relies on a solid semantic understanding of the people in the scene. However, despite the human capability, this ability remains a challenge for our state-of-the-art visual recognition systems. One of the primary goals in this area is person re-identification (ReID), where the task is to match subjects captured in different spots and times.

ReID research has seen significant progress in the last few years (Penate-Sanchez et al., 2020; Ning et al., 2021). New challenging in the wild datasets have emerged due to the evolution of capture devices, going beyond short-time ReID with homogeneous illumination conditions. These new collections have shown several compelling problems in traditional ReID benchmarks. In this regard, long-term ReID copes with substantial variability in space and time. For instance, individuals may be recorded not uniquely affected by pose and occlusion variations but also strongly forced due to different resolutions, i.e., producing multi-scale detections, appearance inconsistencies due to clothing change, and many environmental and lighting variations.

State-of-the-art face recognition approaches applied in surveillance and standard ReID scenarios have evidenced low performance (Cheng et al., 2020; Dietlmeier et al., 2020) due to the need for high resolution and image quality. Lately, low-light image enhancement has attracted the community's attention on this subject to cope with illumination changes in the target images.

Poorly illuminated images suffer from low contrast and high level of noise (Rahman et al., 2021). Significant issues like over-saturation are introduced on image regions with high-intensity pixels by straightly adjusting the illumination property. Consequently, the community has proposed several image enhancement methods to address this problem.

Traditional low-light image enhancement methods rely on the Retinex Theory (Land, 1977), which focuses on an image's dynamic range and color constancy and recovers the contrast by accurately estimating illumination.

Nonetheless, some of these methods may generate color distortion or they are prone to enhance image brightness presenting unnatural effects (He et al., 2020). More recently, deep learning techniques have pushed advances on this challenging task (Zhai et al., 2021) in two directions: image quality and processing time decrease.

[a] https://orcid.org/0000-0001-7511-5783
[b] https://orcid.org/0000-0002-2834-2067
[c] https://orcid.org/0000-0003-2378-4277
[d] https://orcid.org/0000-0003-3022-7698
[e] https://orcid.org/0000-0002-8673-2725

Santana, O., Lorenzo-Navarro, J., Freire-Obregón, D., Hernández-Sosa, D. and Castrillón-Santana, M.
Evaluating the Impact of Low-Light Image Enhancement Methods on Runner Re-Identification in the Wild.
DOI: 10.5220/0011652000003411
In *Proceedings of the 12th International Conference on Pattern Recognition Applications and Methods (ICPRAM 2023)*, pages 641-648
ISBN: 978-989-758-626-2; ISSN: 2184-4313

However, despite encouraging progress on low-light image enhancement, we believe there is room for improvement. For instance, only a few recent works have been tested on video clips (Lv et al., 2018; Li et al., 2022). Unlike still images, video clips present a continuous lighting signal variation that may affect subsequent processes, such as person ReID.

This work takes a step towards image enhancement evaluation in dynamic scenarios, proposing a novel ReID pipeline to unify both tasks, image enhancement and ReID. Both are computed differently; the former exploits the complete scene to enhance the image by using as much information as possible, whereas the latter focuses on a specific region of interest of the improved image.

Additionally, we developed a ReID assessment considering different state-of-the-art methods for image enhancement. In this regard, traditional techniques and deep learning approaches have been analyzed on a sporting dataset that contains 109 different identities in a 30 hours race under variable lighting conditions, scenarios, and accessories. Our results reveal the challenge of performing a runner ReID process in these difficult conditions, especially when this process involves images captured under poor illumination. Our results also show that deep learning approaches provide better results than traditional methods, as they consider the overall context of the image instead of focusing on individual pixels. Nevertheless, there is still significant room for improvement for ReID in this scenario. This work is an interesting first step in this direction, combining different image processing techniques to boost ReID under challenging conditions.

The rest of this paper is organized as follows. Section 2 discusses the related work in the literature. Section 3 presents the proposed pipeline. Section 4 describes the considered collection, the experimental protocol, and the ReID evaluation experiments on the enhanced images. Finally, Section 5 presents our concluding remarks.

2 RELATED WORK

Many methods have been proposed to enhance images captured in low-light conditions. In general, illumination enhancement models based on the Retinex Theory (Land, 1977) decompose images into two components: reflectance and illumination. Some techniques have been proposed for enhancing images working on both components (Jobson et al., 1997; Wang et al., 2013; Fu et al., 2016; Ying et al., 2017; Li et al., 2018), but they involve a high computational cost.

Guo et al. (Guo et al., 2017) proposed a low-light image enhancement method that reduces the solution space by only estimating the illumination component. First, an illumination map is generated based on the maximum value of each RGB color channel for each image pixel. The map is then refined according to the illumination structure using an augmented lagrangian algorithm. Liu et al. (Liu et al., 2022) proposed an efficient algorithm based on a membership function and gamma correction. This method also has lower computational cost than other traditional approaches, resulting in images that do not exhibit over-enhancement or under-enhancement. Since these methods are explicitly designed for low-light imagery, they are tuned to enhance underexposed images. Zhang et al. (Zhang et al., 2019) proposed performing a dual illumination map estimation for both the original and the inverted images, generating two different maps. Underexposed and overexposed regions of the image can be corrected through these maps.

Beyond methods based on the Retinex Theory, the rise of deep learning has led to the development of image enhancement methods based on neural networks (Cai et al., 2018; Park et al., 2018; Wang et al., 2019; Kim, 2019; He et al., 2020). Recently, Hao et al. (Hao et al., 2022) proposed a decoupled two-stage neural network model which provides comparable or better results than other state-of-the-art approaches. The first neural network learns the scene structure and the illumination distribution to generate an image looking close to optimal lighting conditions. The second neural network further enhances the image, suppressing noise and color distortion.

Low illumination scenarios have been previously addressed in different computer vision problems. Liu et al. (Liu et al., 2021) presented an extensive review of low-light enhancement methods and showed the results of a face detection task. The authors also introduced the VE-LOL-L dataset collected with low-light conditions and annotated for face detection. Ma et al. (Ma et al., 2019) proposed the TMD^{2}L distance learning method based on Local Linear Models and triplet loss for video re-identification. The method results are compared with other low-illumination enhancement methods in three datasets, two simulated based on PRID 2011 and iLIDS-VID, and a newly collected LIPS dataset. Unlike Ma's approach, which seeks to obtain person descriptors directly from low-illumination images, our proposal is more similar to Liu's, having a preprocessing stage for the runner images but in a ReID context instead of face detection.

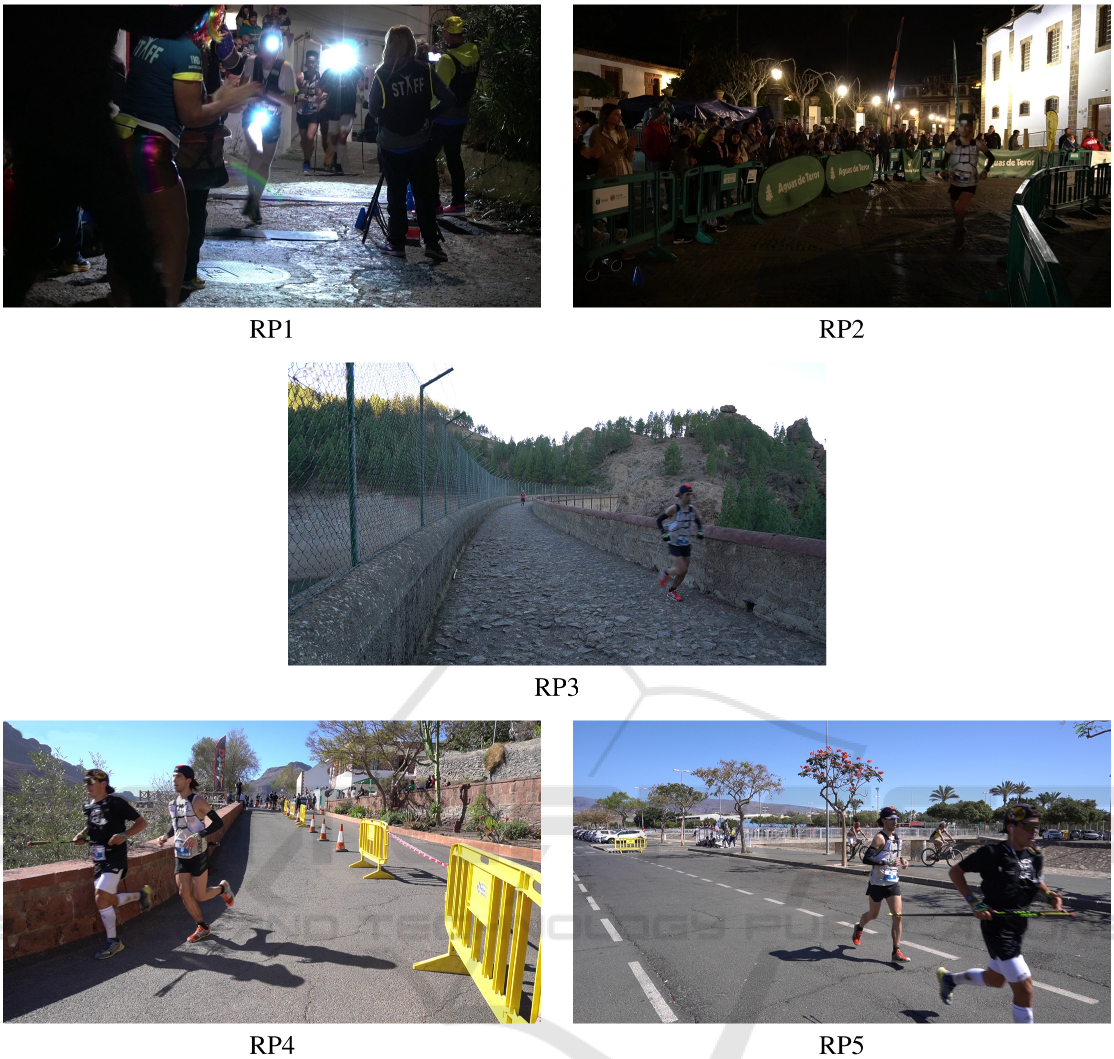

Figure 1: Example of leading runner(s) captured in every recording point (RP1 to RP5).

3 PROPOSED PIPELINE

Our experiments use the TGC2020 dataset (Penate-Sanchez et al., 2020). This dataset has been used recently to tackle several computer vision problems such as facial expression recognition (Santana et al., 2022), bib number recognition (Hernández-Carrascosa et al., 2020), and action quality assessment (Freire-Obregón et al., 2022a; Freire-Obregón et al., 2022b). It comprises a collection of runner images captured in a set of recording points (RPs) as shown in Figure 1. This work focuses on applying several enhancement techniques to sporting footage in the wild and analyzing these techniques relying on the ReID performance. However, to tackle ReID properly, scenes must be cleaned of distractors: staff, other runners, the public, and vehicles.

Figure 2 shows the devised ReID pipeline, divided into three main parts. Firstly, the image is enhanced to improve the runner's visibility by applying a low-light image enhancement method to the input footage, generating improved footage. As shown in the experiments section, several techniques have been tested to perform the enhancement process: LIME (Guo et al., 2017), the dual estimation method (Zhang et al., 2019), and the decoupled low-light image enhancement method (Hao et al., 2022).

Secondly, we are interested in specific regions where runners are located. Once bodies are detected by applying a body detector based on Faster R-CNN with Inception-V2 pre-trained with the COCO dataset (Huang et al., 2017), we have used DeepSORT to

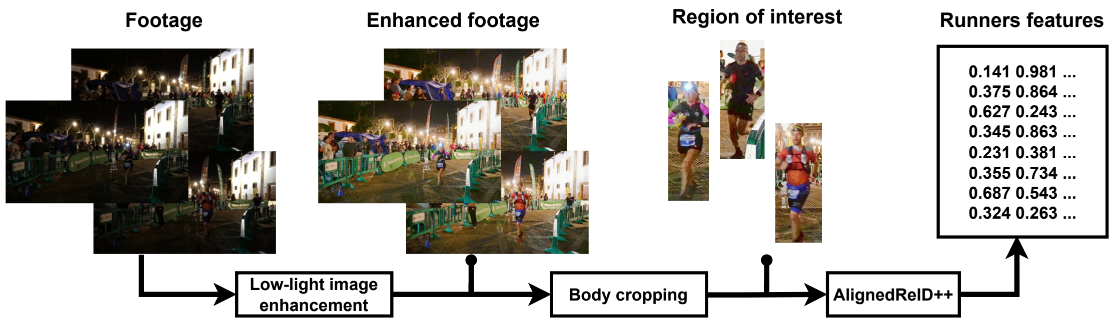

Figure 2: Proposed ReID pipeline. The devised process comprises three main parts: the footage enhancement, the region of interest cropping, and the feature extraction.

track runners in the scene (Wojke et al., 2017). This tracker is based on the SORT tracker and introduces a set of deep descriptors to integrate the appearance information along with the position information given by the Kalman filter used in SORT.

Finally, the embeddings of each runner are extracted from the resulting cropped bodies using the AlignedReID++ deep learning model (Luo et al., 2019) trained on the Market1501 dataset (Zheng et al., 2015). The resulting embeddings are adopted as features to compute ReID results. For every RP but the first one, we probe each runner against the previous RPs, which act as galleries, thus preserving the temporal progression of the runners throughout the race.

ReID performance is measured using the mean average precision score (mAP). This metric is ideal for datasets where the same subject may appear multiple times in the gallery because it considers all occurrences of each particular runner. The mAP score can be defined as:

$$mAP = \frac{\sum_{i=1}^{k} AP_i}{k} \tag{1}$$

Where AP_i stands for the area under the precision-recall curve of probe i and k is the total number of runners in the probe. In this regard, we have tried Cosine and Euclidean distances to compute the mAP values and obtained similar results. Hence we present results just for the latter.

4 EXPERIMENTS

This section describes the experimental evaluation adopted and summarize the achieved results. Firstly, we evaluate runner ReID in the original dataset, i.e., without preprocessing, using the proposed pipeline to set our comparison baseline. We then apply different low-light image enhancement methods to explore their impact on the ReID results. The objective is to check to what extent the proposed preprocessing variants allow to improve the quality of the final ReID performance.

Table 1: Location and recording starting time for each RP. The reader may observe that the first two RPs were captured during the night. Therefore, just artificial illumination was available.

	location	km	starting time
RP1	Arucas	16.5	00:06
RP2	Teror	27.9	01:08
RP3	Presa de Hornos	84.2	07:50
RP4	Ayagaures	110.5	10:20
RP5	Parquesur	124.5	11:20

4.1 Original Dataset

As described by the authors (Penate-Sanchez et al., 2020), the dataset used in this work was collected by recording participants during the 2020 edition of the TransGranCanaria (TGC) ultra-trail race, particularly those taking part in the TGC Classic, where the challenge was to cover 128 kilometers by foot in less than 30 hours.

Runners start at 11 pm, covering the initial part of the track with nightlight, roughly eight hours. Winners take approximately 13 hours to reach the finish line. The dataset comprises the recording in five RPs, two during the first eight hours, i.e., with nightlight conditions, and three more after 7 am, i.e., recorded with daylight. Given the recording environment, the images captured at each RP vary significantly in lighting conditions; see Figure 1. Table 1 shows the distance from the departure line where each point is located in the race track and the starting time for image recording at each RP.

Although runners wear a bib number and carry a tracking device, these monitoring systems do not identify the person wearing them, leaving the door open to potential cheating (e.g., several runners can share a race bib to boost the ranking of the bib owner),

Table 2: ReID results (mAP) for the original dataset, i.e. without applying any preprocessing to the images.

	Gallery			
	RP1	RP2	RP3	RP4
RP2	14.80			
RP3	9.11	7.17		
RP4	22.66	13.13	19.21	
RP5	22.41	20.32	15.09	48.29

Table 3: ReID mAP percentage variation for the enhanced dataset.

	Gallery	LIME	Dual	Net-I	Net-II
RP2	RP1	+6.75	+5.15	+10.15	**+12.18**
RP3	RP1	+4.79	+0.61	**+9.91**	+4.68
	RP2	+3.61	+1.59	**+11.13**	+5.22
RP4	RP1	-0.22	-0.89	**+3.25**	+1.04
	RP2	+3.84	+3.42	**+11.20**	+10.63
	RP3	+0.94	-0.34	**+5.65**	+3.88
RP5	RP1	+1.86	+0.42	**+4.19**	+2.85
	RP2	+6.02	+4.19	**+8.68**	+7.15
	RP3	+1.69	-0.72	**+4.00**	+2.83
	RP4	-0.28	-3.53	**+8.14**	+3.02

hence the need to apply ReID techniques in these long-distance competitions.

For this study, we used only the images of the 109 runners identified in every RP. As mentioned above, RP1 and RP2 samples were captured during the nighttime, while the images from RP4 and RP5 were captured during the daytime.

The images from RP3 provide an in-between scenario, as they started to be recorded in the early morning hours in a place where the sunlight was blocked by trees and mountains, creating a situation closer to twilight than to full sunlight. Once more, the reader may observe Figure 1, which shows images of runners in all five RPs to illustrate the meaningful variation in lighting conditions between different locations.

Table 2 shows the ReID mAP score using RP2, RP3, RP4, and RP5 as probe and separately each preceding RP as gallery. The highest value is obtained by probing RP5 against RP4 as gallery, since both RPs were shot in full daylight. On the other hand, the worst results are obtained by RP3, shot in some peculiar intermediate lighting conditions and having only night images available to use as galleries, revealing the great difficulty of runner ReID in the wild.

4.2 Enhanced Dataset

The preprocessed dataset (henceforth *enhanced*) applies a previous step to correct the original images in terms of illumination. In the experiments below, we do not distinguish between night and daylight RPs to use the preprocessing. Indeed, it is applied homogeneously to any image in the original dataset.

Table 3 provides results for the low-light image enhancement method (LIME) proposed by Guo et al. (Guo et al., 2017) and the dual estimation method (Dual) proposed by Zhang et al. (Zhang et al., 2019). We evaluated various gamma values in both cases and found the best ReID results using a value of 0.3. Although it is common in the literature to use more significant values, this is usually done with high-quality images. As a consequence of the wild scenario, the quality of the images is lower, and the runners are affected by motion blur. Raising the gamma value too high makes the images noisier and adversely affects the ReID process.

The LIME algorithm application leads to an evident improvement in the first three RPs, the most affected by the low light. The improvement is less noticeable in the better illuminated RPs, even a slight degradation in some cases. The results of the application of the Dual algorithm are, in general, poorer than those provided by the LIME algorithm. Both algorithms share the same mathematical basis, but the LIME algorithm focuses on improving the underexposed parts of the image.

In contrast, the Dual algorithm aims to improve both the underexposed and overexposed parts of the image. These results suggest that overexposure is not a significant problem for runner ReID when using this dataset. Consequently, in this scenario it is preferable to use the algorithm that enhances just the underexposed areas.

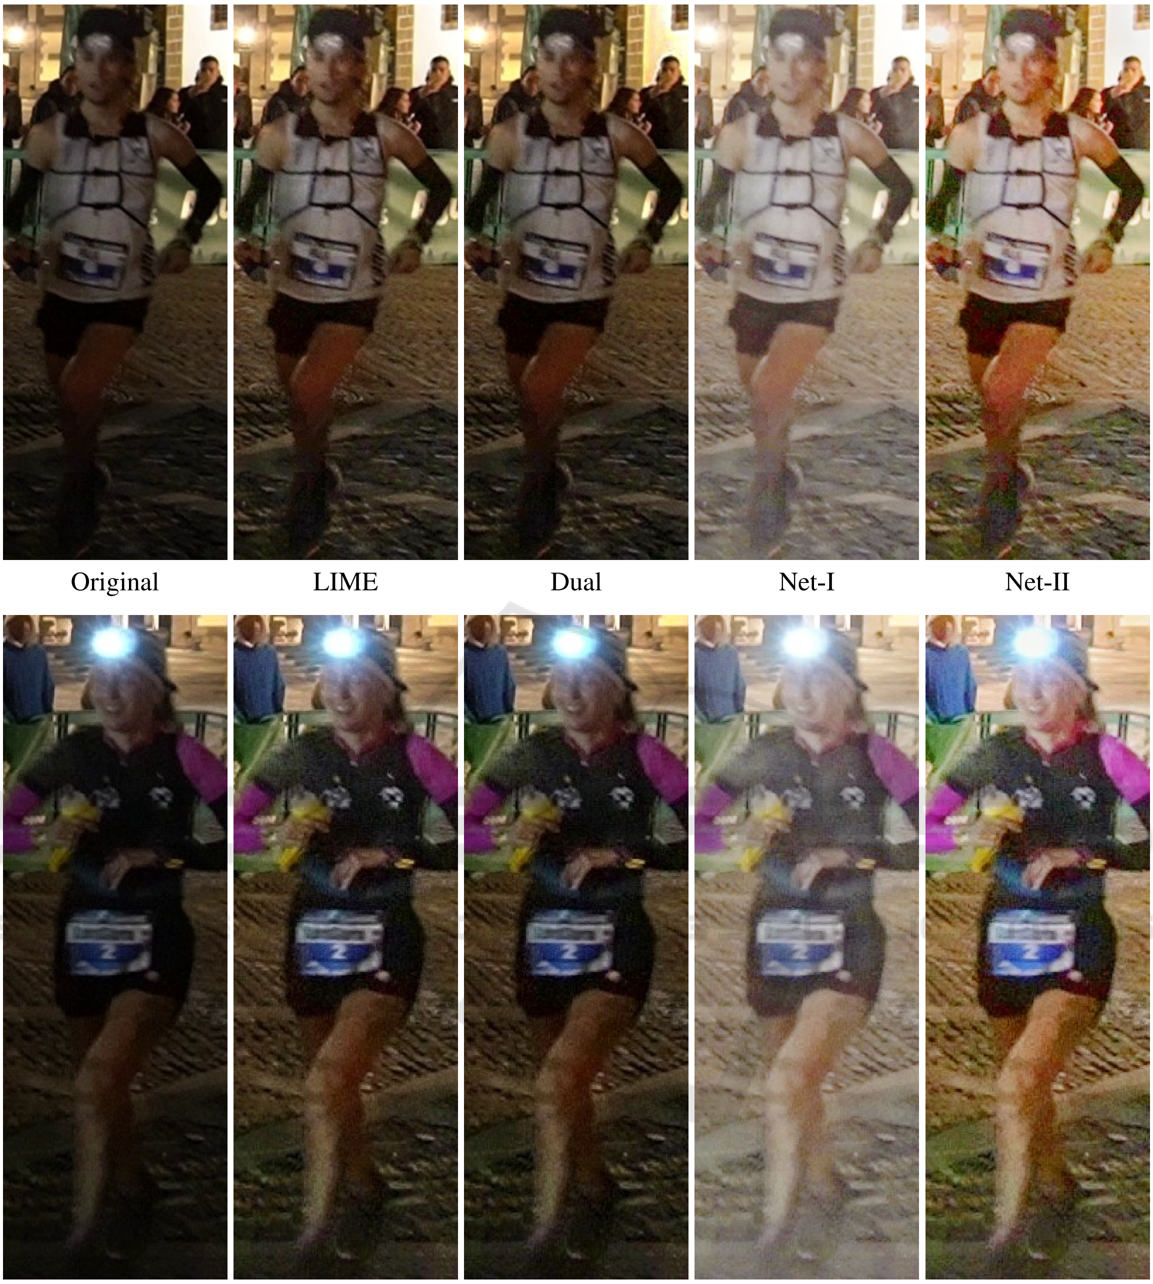

Figure 3: Images captured in RP2 that have been treated with different methods of illumination enhancement.

Opposed to these methods based on the Retinex Theory, the decoupled low-light image enhancement method proposed by Hao et al. (Hao et al., 2022) is a deep learning approach. This neural network does not apply a series of preset operations according to an algorithm but has learned to interpret the general context of the image, leading to better ReID results.

Table 3 shows separately the results of applying only the first stage (Net-I) and applying both stages (Net-II) of this decoupled neural network. In most cases, the best results are obtained using only the first stage because it is the stage trained to improve lighting conditions. The color and noise corrections applied by the second stage only provide an improvement in one case, with worse ReID results in the rest.

5 CONCLUSIONS

Runner ReID in the wild is highly complex due to the long-term component, which does not necessarily preserve clothes unchanged, and the poor lighting conditions often encountered. This latter problem cannot be easily solved by mechanical means, as increasing the exposure time to capture more light leads to degraded image quality due to motion blur.

To address these difficulties, we look for synergies between different methods proposed individually and applied in origin to different scenarios. Our results show that applying software methods for illumination enhancement is a promising approach to improve ReID results. As far as we know, the reported results are the best achieved without using temporal information (Medina et al., 2022).

It is a well-known fact that computers do not see as people do and that the enhanced image that may seem the clearest to us may not necessarily be the most useful for an automatic ReID process. Figure 3 shows the nighttime image of two runners with different techniques applied, clearly illustrating that it is not trivial to identify the best one based on human observation.

Overall, it is pretty likely that no single approach would be considered the best in all circumstances, which involves that it would be necessary to choose the method to apply according to the image features.

Determining the best strategy to achieve this goal is an interesting topic for future research. In this sense, it would be of great value for the community to establish a systematic process for choosing the best method to apply to a low-light image based on the different quality assessment metrics (Zhai et al., 2021) proposed in the literature.

ACKNOWLEDGEMENTS

This work is partially funded by the ULPGC under project ULPGC2018-08, by the Spanish Ministry of Economy and Competitiveness (MINECO) under project RTI2018-093337-B-I00, by the Spanish Ministry of Science and Innovation under projects PID2019-107228RB-I00 and PID2021-122402OB-C22, and by the ACIISI-Gobierno de Canarias and European FEDER funds under projects ProID2020010024, ProID2021010012, ULPGC Facilities Net, and Grant EIS 2021 04.

REFERENCES

Cai, J., Gu, S., and Zhang, L. (2018). Learning a deep single image contrast enhancer from multi-exposure images. *IEEE Transactions on Image Processing*, 27(4):2049–2062.

Cheng, Z., Zhu, X., and Gong, S. (2020). Face re-identification challenge: Are face recognition models good enough? *Pattern Recognition*, 107:107422.

Dietlmeier, J., Antony, J., Mcguinness, K., and O'Connor, N. E. (2020). How important are faces for person re identification? In *Proceedings International Conference on Pattern Recognition*, Milan, Italy.

Freire-Obregón, D., Lorenzo-Navarro, J., and Castrillón-Santana, M. (2022a). Decontextualized i3d convnet for ultra-distance runners performance analysis at a glance. In *International Conference on Image Analysis and Processing (ICIAP)*.

Freire-Obregón, D., Lorenzo-Navarro, J., Santana, O. J., Hernández-Sosa, D., and Castrillón-Santana, M. (2022b). Towards cumulative race time regression in sports: I3d convnet transfer learning in ultra-distance running events. In *International Conference on Pattern Recognition (ICPR)*.

Fu, X., Zeng, D., Huang, Y., Liao, Y., Ding, X., and Paisley, J. (2016). A fusion-based enhancing method for weakly illuminated images. *Signal Processing*, 129(C).

Guo, X., Li, Y., and Ling, H. (2017). Lime: Low-light image enhancement via illumination map estimation. *IEEE Transactions on Image Processing*, 26(2):982–993.

Hao, S., Han, X., Guo, Y., and Wang, M. (2022). Decoupled low-light image enhancement. *ACM Trans. Multimedia Comput. Commun. Appl.*, 18(4).

He, W., Liu, Y., Feng, J., Zhang, W., Gu, G., and Chen, Q. (2020). Low-light image enhancement combined with attention map and u-net network. In *2020 IEEE 3rd International Conference on Information Systems and Computer Aided Education (ICISCAE)*, pages 397–401.

Hernández-Carrascosa, P., Penate-Sanchez, A., Lorenzo-Navarro, J., Freire-Obregón, D., and Castrillón-Santana, M. (2020). TGCRBNW: A dataset for runner bib number detection (and recognition) in the wild. In *Proceedings International Conference on Pattern Recognition*, Milan, Italy.

Huang, J., Rathod, V., Sun, C., Zhu, M., Korattikara, A., Fathi, A., Fischer, I., Wojna, Z., Song, Y., Guadarrama, S., and Murphy, K. (2017). Speed/accuracy trade-offs for modern convolutional object detectors. In *CVPR 2017*, pages 649–656, Hawai, USA. IEEE.

Jobson, D., Rahman, Z., and Woodell, G. (1997). A multiscale retinex for bridging the gap between color images and the human observation of scenes. *IEEE Transactions on Image Processing*, 6(7):965–976.

Kim, M. (2019). Improvement of low-light image by convolutional neural network. In *2019 IEEE 62nd International Midwest Symposium on Circuits and Systems (MWSCAS)*, pages 189–192.

Land, E. H. (1977). The retinex theory of color vision. *Scientific American*, 237(6):108–129.

Li, C., Guo, C., Han, L., Jiang, J., Cheng, M.-M., Gu, J., and Loy, C. C. (2022). Low-light image and video enhancement using deep learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):9396–9416.

Li, M., Liu, J., Yang, W., Sun, X., and Guo, Z. (2018). Structure-revealing low-light image enhancement via robust retinex model. *IEEE Transactions on Image Processing*, 27(6):2828–2841.

Liu, J., Xu, D., Yang, W., Fan, M., and Huang, H. (2021). Benchmarking low-light image enhancement and beyond. *International Journal of Computer Vision*, 129(4):1153–1184.

Liu, S., Long, W., Li, Y., and Cheng, H. (2022). Low-light image enhancement based on membership function and gamma correction. *Multimedia Tools and Applications*, 81:22087–22109.

Luo, H., Jiang, W., Zhang, X., Fan, X., Qian, J., and Zhang, C. (2019). AlignedReID++: Dynamically matching local information for person re-identification. *Pattern Recognition*, 94:53 – 61.

Lv, F., Lu, F., Wu, J., and Lim, C. S. (2018). MBLLEN: Low-light image/video enhancement using cnns. In *BMVC*.

Ma, F., Zhu, X., Zhang, X., Yang, L., Zuo, M., and Jing, X.-Y. (2019). Low illumination person re-identification. *Multimedia Tools and Applications*, 78(1):337–362.

Medina, M. A., Lorenzo-Navarro, J., Freire-Obregón, D., Santana, O. J., Hernández-Sosa, D., and Santana, M. C. (2022). Boosting re-identification in the ultra-running scenario. In *Proceedings of the 11th International Conference on Pattern Recognition Applications and Methods (ICPRAM)*, pages 461–469.

Ning, X., Gong, K., Li, W., Zhang, L., Bai, X., and Tian, S. (2021). Feature refinement and filter network for person re-identification. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(9):3391–3402.

Park, S., Yu, S., Kim, M., Park, K., and Paik, J. (2018). Dual autoencoder network for retinex-based low-light image enhancement. *IEEE Access*, 6:22084–22093.

Penate-Sanchez, A., Freire-Obregón, D., Lorenzo-Melián, A., Lorenzo-Navarro, J., and Castrillón-Santana, M. (2020). TGC20ReId: A dataset for sport event re-identification in the wild. *Pattern Recognition Letters*, 138:355–361.

Rahman, Z., Pu, Y.-F., Aamir, M., and Wali, S. (2021). Structure revealing of low-light images using wavelet transform based on fractional-order denoising and multiscale decomposition. *The Visual Computer*, 37(5):865–880.

Santana, O. J., Freire-Obregón, D., Hernández-Sosa, D., Lorenzo-Navarro, J., Sánchez-Nielsen, E., and Castrillón-Santana, M. (2022). Facial expression analysis in a wild sporting environment. In *Multimedia Tools and Applications*.

Wang, R., Zhang, Q., Fu, C.-W., Shen, X., Zheng, W.-S., and Jia, J. (2019). Underexposed photo enhancement using deep illumination estimation. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6842–6850.

Wang, S., Zheng, J., Hu, H.-M., and Li, B. (2013). Naturalness preserved enhancement algorithm for non-uniform illumination images. *IEEE Transactions on Image Processing*, 22(9):3538–3548.

Wojke, N., Bewley, A., and Paulus, D. (2017). Simple online and realtime tracking with a deep association metric. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 3645–3649. IEEE.

Ying, Z., Li, G., Ren, Y., Wang, R., and Wang, W. (2017). A new low-light image enhancement algorithm using camera response model. In *2017 IEEE International Conference on Computer Vision Workshops (ICCVW)*, pages 3015–3022.

Zhai, G., Sun, W., Min, X., and Zhou, J. (2021). Perceptual quality assessment of low-light image enhancement. *ACM Trans. Multimedia Comput. Commun. Appl.*, 17(4).

Zhang, Q., Nie, Y., and Zheng, W.-S. (2019). Dual illumination estimation for robust exposure correction. *Computer Graphics Forum*, 38(7):243–252.

Zheng, L., Shen, L., Tian, L., Wang, S., Wang, J., and Tian, Q. (2015). Scalable person re-identification: A benchmark. In *Proceedings of the International Conference on Computer Vision*.