

Human Pose Estimation in Extremely Low-Light Conditions

Sohyun Lee^{1*}

Jaesung Rim^{1*}

Boseung Jeong²

Geonu Kim²

Byungju Woo²

Haechan Lee¹


Sunghyun Cho^{1,2†}

Suha Kwak^{1,2†}

¹Graduate School of AI, POSTECH

²Dept. of CSE, POSTECH


<http://cg.postech.ac.kr/research/ExLPose>



Well-lit image



Low-light image & its scaled version



Our result

Abstract

We study human pose estimation in extremely low-light images. This task is challenging due to the difficulty of collecting real low-light images with accurate labels, and severely corrupted inputs that degrade prediction quality significantly. To address the first issue, we develop a dedicated camera system and build a new dataset of real low-light images with accurate pose labels. Thanks to our camera system, each low-light image in our dataset is coupled with an aligned well-lit image, which enables accurate pose labeling and is used as privileged information during training. We also propose a new model and a new training strategy that fully exploit the privileged information to learn representation insensitive to lighting conditions. Our method demonstrates outstanding performance on real extremely low-light images, and extensive analyses validate that both of our model and dataset contribute to the success.

1. Introduction

Deep neural networks [6, 55, 56, 64, 66] trained with large-scale datasets [1, 18, 30, 35, 38] have driven dramatic advances in human pose estimation recently. However, their success demands high-quality inputs taken in controlled environments while in real-world applications images are often corrupted by low-light conditions, adverse weather conditions, sensor noises, motion blur, *etc.* Indeed, a precondition for human pose estimation *in the wild* is robustness against such adverse conditions.

Motivated by this, we study pose estimation under ex-

remely low-light conditions using a single sRGB image, in which humans can barely see anything. The task is highly practical as its solution enables nighttime applications of pose estimation without raw-RGB data or additional devices like IR cameras. It is at the same time challenging due to the following two reasons. The first is the difficulty of data collection. Manual annotation of human poses in low-light images is often troublesome due to their limited visibility. The second is the difficulty of pose estimation on low-light images. The poor quality of low-light images in terms of visibility and signal-to-noise ratio largely degrades prediction accuracy of common pose estimation models. A naïve way to mitigate the second issue is to apply low-light image enhancement [5, 28, 41, 42, 62] to input images. However, image enhancement is in general highly expensive in both computation and memory. Also, it is not aware of downstream recognition tasks and thus could be sub-optimal for pose estimation in low-light conditions.

To tackle this challenging problem, we first present a new dataset of real extremely low-light images with ground-truth pose labels. The key feature of our dataset is that each low-light image is coupled with a well-lit image of the same content. The advantage of using the well-lit images is two-fold. First, they enable accurate labeling for their low-light counterparts thanks to their substantially better visibility. Second, they can be utilized as *privileged information* [13, 32, 40, 57, 58], *i.e.*, additional input data that are more informative than the original ones (low-light images in our case) but available only in training, to further improve performance on low-light images. Such benefits of paired training images have also been validated in other robust recognition tasks [9, 34, 46–49]. The beauty of our dataset is that pairs of low-light and well-lit images are all

* Equal contribution. † Corresponding authors.

real and *aligned*, unlike existing datasets that provide pairs of synthetic-real images [9, 47, 48] or those of largely misaligned real images [46, 49]. Since it is practically impossible to capture such paired images using common cameras, we build a dedicated camera system for data collection.

We also propose an effective method based on learning using privileged information (LUPI) [58] to fully exploit our dataset. The proposed method considers a model taking low-light inputs as a *student* and a model dealing with corresponding well-lit images as a *teacher*. Both of the teacher and student are trained by a common pose estimation loss, and the student further utilizes knowledge of the teacher as additional supervision. Specifically, our method employs neural styles of intermediate feature maps as the knowledge and forces neural styles of low-light images to approximate those of well-lit images by an additional loss. As will be demonstrated, this LUPI approach allows the learned representation to be insensitive to lighting conditions. Moreover, we design a new network architecture that unifies the teacher and student through lighting-condition specific batch normalization (LSBN). LSBN consists of two batch normalization (BN) layers, each of which serves images of each lighting condition, *i.e.*, ‘well-lit’ or ‘low-light’. We replace BNs of an existing network with LSBNs so that images of different lighting conditions are processed by different BNs. Hence, in our architecture, the teacher and student share all the parameters except for those of their corresponding BNs, which allows the student to enjoy the strong representation learned using well-lit images.

The efficacy of our method is evaluated on real low-light images we collected for testing. Our method outperforms its reduced versions and relevant approaches such as lighting-condition adversarial learning and a combination of image enhancement and pose estimation. These results clearly demonstrate the advantages of our dataset and method. In short, our major contribution is three-fold:

- We propose a novel approach to human pose estimation in extremely low-light conditions using a single sRGB image. To the best of our knowledge, we are the first to tackle this challenging but highly practical problem.
- We build a new dataset that provides real and aligned low-light and well-lit images with accurate pose labels.
- We present a strong baseline method that fully exploits the low-light and well-lit image pairs of our dataset.

2. Related Work

Low-light Image Enhancement. Classical methods for low-light image enhancement have been developed based on histogram equalization and the Retinex theory [19, 25, 29, 33]. Recently, learning-based methods have driven remarkable advances in this field; examples include an auto-encoder [41], a multi-branch architecture [42], and a U-Net

architecture [5]. Also, Jiang *et al.* [28] proposed a GAN using unpaired low-light and well-lit images, and Wang *et al.* [62] learned mapping from low-light to well-lit images via normalizing flow. They usually require heavy computation and are learned without considering downstream recognition tasks. In contrast, we focus on learning features insensitive to lighting-condition for pose estimation while bypassing low-light image enhancement.

Low-light Datasets. SID [5] and LOL [63] provide paired low-light and well-lit images. Wang *et al.* [60] collected paired low-light and well-lit videos by playing motion using an electric slide system. Due to the difficulty of capturing paired images simultaneously, they only provide images of static objects. Meanwhile, Jiang *et al.* [27] built a dual-camera system to capture paired low-light and well-lit videos at once. Inspired by this, we construct a dedicated camera system for collecting paired low-light and well-lit images of humans with motion. Low-light datasets have also been proposed for other tasks, *e.g.*, ARID [65] for human action recognition, NOD [43] and ExDark [39] for object detection. While these datasets provide only low-light images for training recognition models, our dataset provides low-light images along with their well-lit counterparts as well as accurate human pose labels.

Pose Estimation in Low-light Conditions. Crescitelli *et al.* [7, 8] presented a human pose dataset containing 1,800 sRGB and 2,400 infrared (IR) images captured at night, and integrated sRGB and IR features for human pose estimation. Our work is clearly distinct from this in two aspects. First, we focus on human pose estimation in *extremely* low-light images, in which humans barely see anything. On the other hand, Crescitelli *et al.* [7, 8] captured images of night scenes with light sources, in which most human objects are sufficiently visible. Second, our model takes a single sRGB image captured in a low-light environment as input for pose estimation without demanding extra observations like IR images. To the best of our knowledge, our work is the first to tackle this challenging but highly practical task.

Robust Visual Recognition. The performance of conventional recognition models often gets degraded in adverse conditions [37, 39, 43, 61]. To address this issue, the robustness of visual recognition has been actively studied [17, 21, 22, 34, 49, 51, 53, 54]. Low-light visual recognition is one such direction that aims at learning features robust to limited visibility of low-light images. Sasagawa *et al.* [50] utilized a U-Net architecture to restore well-lit images from raw-RGB low-light images for object detection. Morawski *et al.* [43] incorporated an image enhancement module into an object detector and proposed lighting variation augmentations for nighttime recognition. However, the enhancement module is often substantially heavy to be integrated with recognition models. We instead learn a pose estimation model insensitive to lighting conditions so that it




Figure 1. Our dual-camera system that consists of two camera modules with a beam splitter and an ND filter.

does not need an extra module for image enhancement.

Learning Using Privileged Information. LUPI aims at exploiting privileged information available only in training to improve target models in terms of accuracy, label-efficiency, and convergence speed. Vapnik *et al.* [58] first introduced LUPI for support vector machine classifiers. The idea has been extended to tackle various tasks beyond classification, *e.g.*, object localization [13], metric learning [14], ranking [52], and clustering [12]. LUPI has also been studied for training deep neural networks: Lopez-Paz *et al.* [40] investigated the relation between LUPI and knowledge distillation [23], Lambert *et al.* [32] developed a new dropout operation controlled by privileged information, and Hoffman *et al.* [24] employed an auxiliary model that approximates the teacher using ordinary input data. Our work is clearly distinct from these in terms of the model architecture and the way of teacher-student interaction as well as the target task: Our model maximizes parameters shared by the teacher and student so that privileged information helps improve representation quality of the student, and the student is trained to approximate internal behavior of the teacher as well as predicting human poses.

3. ExLPose Dataset

We propose the first *Extremely Low-light* image dataset for human Pose estimation, coined ExLPose. The ExLPose dataset provides pairs of low-light and well-lit images that are real and aligned, as well as their ground-truth human pose labels. We believe that the accurate pose labels and well-lit counterparts for extremely low-light images will open a new and promising research direction towards human pose estimation in low-light conditions.

Since it is practically impossible to simultaneously capture paired low-light and well-lit images using common cameras, we construct a dual-camera system [27, 44] dedicated to the purpose. Our camera system is depicted in Fig. 1. It consists of two camera modules with a beam splitter that distributes light from a lens equally to the two camera modules. One of the modules captures well-lit images, and the other captures low-light images through a 1% neutral density (ND) filter that optically reduces the amount of light by 100 times. The two camera modules simultane-




Figure 2. Pairs of well-lit (*top*) and low-light (*bottom*) images captured in the same scene. The low-light images are scaled by 30 times for visualization.

ously capture a pair of low-light and well-lit images of the same scene with a synchronized shutter. As done in [45], the images are geometrically aligned using a reference image pair captured in a static scene. Details of the geometric alignment are given in Sec. C of the supplement.

We collected images of various indoor and daytime outdoor scenes in the sRGB format. To cover low-light scenes of diverse brightness levels, we collected low-light images with various exposure times. Specifically, for each scene, we first manually found a proper exposure time and a gain value to capture a well-lit image without losing much information in the highlights or in the shadows, and also without motion blur. Then, we sequentially captured low-light images reducing the exposure time by 1, 2, 4, 8, and 12 times with the same gain value as shown at the bottom of Fig. 2. At the same time, for the well-lit images, we also changed the exposure time in the same manner to synchronize the exposure time for each frame, but inverse-proportionally increased the gain value to maintain the same brightness level as presented at the top of Fig. 2. For the indoor scenes, we reduced the exposure time by 1, 2, 3, 4, and 6 times. Note that low-light and well-lit images of a pair are captured with the same exposure time.

Finally, we collected 2,556 pairs of low-light and well-lit images of 251 scenes; 2,065 pairs of 201 scenes are used for training, and the remaining 491 pairs of 50 scenes are kept for testing. We manually annotated a bounding box and 14 body joints for each person using well-lit images following the CrowdPose dataset [36], and collected annotations for 14,215 human instances. The annotations are used as the ground-truth labels for both low-light and well-lit images as the images are spatially aligned.

While the dedicated camera system allows capturing real low-light images paired with well-lit images, it may introduce two limitations regarding the generalization ability of the pose estimation method. First, the dataset does not cover diverse cameras as the system is designed with specific camera modules. Second, low-light images captured with an ND filter may have different characteristics than low-light images captured at night. Thus, for the evaluation of the generalization ability, we also propose another dataset,

		#Scenes	#Images	#Instances	Camera	Mean Intensity	Paired Well-lit	Resolution
ExLPose	Train	201	2,065	11,405	daA1920-160uc	2.0 (low-light) / 90.5 (well-lit)	✓	1920×1200
	Test	50	491	2,810				
ExLPose-OCN	Test	-	180	466	A7M3	3.8		6000×4000
			180	524	RICOH3	5.6		

Table 1. Statistics of the ExLPose and ExLPose-OCN datasets.



Figure 3. Example images of the ExLPose dataset. WL and LL indicate well-lit and low-light images, respectively. Scaled LL denotes low-light images with intensities scaled up for visualization.

named ExLPose-OCN, that provides extremely low-light images captured by Other Cameras at Night. Specifically, the ExLPose-OCN dataset provides images captured using a DSLR camera (A7M3) and a compact camera (RICOH3) and manually annotated ground-truth labels, but no well-lit images. The images are in JPEG format. For each camera, 180 images are provided. Table 1 summarizes statistics of the ExLPose and ExLPose-OCN datasets, and Fig. 3 and Fig. 4 show example images of the two datasets.

4. Proposed Method

As our target model deals with extremely low-light images, it suffers from the significantly low quality of inputs. To alleviate this, we propose a new method for learning the target model using the paired well-lit images as privileged information [58], additional high-quality input data accessible only in training. Our method introduces another model called *teacher* that takes the privileged information as input and provides rich supervision to the target model called *student*. This method for learning using privileged information (LUPI) allows the student to simulate the internal behavior of the teacher as well as learn to predict human poses.

To further exploit the privileged information, we design a single concise architecture that integrates the teacher and the student. The key idea is to let them use separate batch normalization (BN) layers while sharing all the other parameters of the network; following the domain-specific batch normalization [4] we call such a set of separate BNs lighting-condition specific BN (LSBN). This design choice allows the student to enjoy the strong representation learned




Figure 4. Example images of the ExLPose-OCN dataset. LL and Scaled LL indicate low-light images with the original intensities and intensities scaled up for visualization, respectively.

using the well-lit images (*i.e.*, the privileged information) while capturing specific characteristics of low-light images through the separate BN parameters.

Our model architecture and LUPI strategy are depicted in Fig. 5. Note that, before being fed to the student, low-light images are scaled automatically by adjusting their average pixel intensity value to a predefined constant. On the other hand, the teacher takes as input well-lit images as-is. Both of the teacher and the student are trained by a common pose estimation loss:

$$\mathcal{L}_{\text{pose}}(\mathbf{P}, \mathbf{Y}) = \frac{1}{K} \sum_{i=1}^K \|\mathbf{P}_i - \mathbf{Y}_i\|_2^2, \quad (1)$$

where P_i and Y_i denote the predicted heatmap and ground-truth heatmap for the i -th joint, and K refers to the number of joints. In addition to the above loss, the student takes another supervision based on the privileged information through the teacher. The remaining part of this section elaborates on LSBN and LUPI.

4.1. LSBN

For each iteration of training, a low-light image I^{low} and its well-lit counterpart I^{well} are given together as input and processed by different BNs in LSBNs according to their light conditions. Each LSBN layer contains two BNs, each of which has its own affine transform parameters, $(\gamma^{\text{low}}, \beta^{\text{low}})$ for low-light and $(\gamma^{\text{well}}, \beta^{\text{well}})$ for well-lit. Within a mini-batch of N samples, an LSBN layer whitens input activations and transforms them using the lighting-




Figure 5. The proposed model architecture and training strategy. Both of teacher (*bottom*) and student (*top*) are trained by the same pose estimation loss, and student takes additional supervision from teacher through LUPI. The loss for LUPI is applied to the feature maps of the first convolutional layer (*i.e.*, C1) and the following four residual blocks (*i.e.*, R1–R4) of a ResNet backbone. Teacher and student share all the parameters except LSBNs. Details of the pose estimation module are presented in Fig. a3 of the supplementary material.

condition specific affine transform parameters in a channel-wise manner. Let $\mathbf{x} \in \mathbb{R}^{N \times H \times W}$ denote a channel of activations computed from N images of a specific lighting condition, and λ be an indicator that returns 1 if the lighting condition of \mathbf{x} is ‘low-light’ and 0 otherwise. The output of the LSBN layer taking \mathbf{x} and λ as inputs is given by

$$\text{LSBN}(\mathbf{x}, \lambda) = \lambda \left(\gamma^{\text{low}} \cdot \frac{\mathbf{x} - \mu}{\sqrt{\sigma^2 + \epsilon}} + \beta^{\text{low}} \right) + (1 - \lambda) \left(\gamma^{\text{well}} \cdot \frac{\mathbf{x} - \mu}{\sqrt{\sigma^2 + \epsilon}} + \beta^{\text{well}} \right), \quad (2)$$

where μ and σ^2 denote mean and variance of the activations in \mathbf{x} , respectively, and ϵ is a small constant adopted for numerical stability.

4.2. LUPI

Since low-light and well-lit images of a pair share the same content in our dataset, we argue that the gap between their predictions will be largely affected by their style difference.¹ In our case, the style of an image is determined by high-frequency noise patterns and its overall intensity. Hence, in our LUPI strategy, the teacher provides neural styles of its intermediate feature maps as additional supervision to the student so that the student learns to fill the style gap between low-light and well-lit images in feature spaces; this approach eventually leads to a learned representation insensitive to varying lighting conditions.

To implement the above idea, we adopt as a neural style representation the Gram matrix [16], denoted by $\mathbf{G} \in$

¹It has been known that an image is separated into content and style [16, 31]. Since a pair of low-light and well-lit images in ExLPose have the same content due to our camera system, they are different only in style.

$\mathbb{R}^{C \times C}$, that captures correlations between C channels of a feature map \mathbf{F} . Specifically, it is computed by $\mathbf{G}_{i,j} = \mathbf{f}_i^\top \mathbf{f}_j$ where \mathbf{f}_i is the vector form of the i^{th} channel of \mathbf{F} . Let $\mathbf{G}^{\text{low},l}$ denote the neural style of a low-light image computed from the feature map of the l^{th} layer of the student, and similarly, $\mathbf{G}^{\text{well},l}$ be the neural style of the coupled well-lit image computed from the feature map of the l^{th} layer of the teacher. Then our LUPI strategy minimizes the following loss with respect to $\mathbf{G}^{\text{low},l}$ of all predefined layers so that the neural style of the student approximates that of the teacher:

$$\mathcal{L}_{\text{LUPI}} = \sum_l \frac{1}{4C_l^2 N_l^2} \sum_{i=1}^{C_l} \sum_{j=1}^{C_l} \left(\mathbf{G}_{i,j}^{\text{low},l} - \mathbf{G}_{i,j}^{\text{well},l} \right)^2, \quad (3)$$

where C_l and N_l are the number of channels and the spatial size of the l^{th} feature map, respectively.

4.3. Empirical Justification

To investigate the impact of LUPI, we first demonstrate that it reduces the style gaps between different lighting conditions. To this end, we compute the average Hausdorff distance [11] between the sets of Gram matrices of different lighting conditions before and after applying LUPI. Fig. 6(a) shows that the style gaps between different lighting conditions are effectively reduced by LUPI as intended.

It is also empirically examined if LSBN and LUPI of our method eventually lead to a model insensitive to lighting conditions. For this purpose, each lighting condition is represented by the set of features computed from associated images, and discrepancies between different lighting conditions are estimated by the average Hausdorff distances between such sets. As shown in Fig. 6(b), LSBN matches the




Figure 6. Quantitative analysis on (a) style gaps and (b) feature gaps measured by the average Hausdorff distance [11] between low-light and well-lit conditions. The gaps are measured on each level of the four residual blocks (*i.e.*, R1-R4) of the ResNet backbone while considering each lighting condition as a set of styles of associated images in (a) and their features in (b).

feature distributions of different lighting conditions, compared to the model trained with only pose estimation loss for both low-light and well-lit images, denoted by Baseline-all. LUPI even further closes the gaps, leading to the representation that well aligns two different lighting conditions.

5. Experiments

5.1. Single-person Pose Estimation

To ease the difficulty of the problem and focus solely on pose estimation, we first tackle single-person pose estimation with the assumption that ground-truth bounding boxes are given for individuals.

Implementation Details. We adopt Cascaded Pyramid Network (CPN) [6] with ResNet-50 [20] backbone as our pose estimation network, which is pre-trained on the ImageNet dataset [10]. All BNs in the backbone are replaced with LSBNs, and LUPI is applied to the outputs of the first convolution layer and the following four residual blocks. The average channel intensity of an input low-light image is automatically adjusted to 0.4 before being fed to the student network following Zheng *et al.* [67]. More details are given in Sec. D of the supplement.

Evaluation Protocol. We report the standard average precision (AP) scores based on object keypoint similarity following CrowdPose [35]. Our method and competitors are evaluated on both low-light (LL) and well-lit (WL) images. Low-light images for testing are further divided into three subsets according to their relative difficulty, low-light normal (LL-N), low-light hard (LL-H), and low-light extreme (LL-E), by the gain values of the coupled well-lit images. Specifically, they are split into 167 LL-N images, 169 LL-H images, and 155 LL-E images by applying two thresholds, 15 and 24, to their gain values. The mean pixel intensities of LL-N, LL-H, and LL-E images are 3.2, 1.4, and 0.9, respectively. Note that images of all three subsets are captured in extremely low-light conditions and hard to be recognized by humans. The union of the three low-light test splits is

denoted as low-light all (LL-A).

5.1.1 Quantitative Results on the ExLPose Dataset

Our method is compared with potential solutions to the target task on the five test splits of the ExLPose dataset, *i.e.*, LL-N, LL-H, LL-E, LL-A, and WL. The solutions include baselines that train the pose estimation model directly with the ExLPose dataset, those incorporating low-light image enhancement techniques as pre-processing, and domain adaptation methods that consider different lighting conditions as different domains. Specifically, the baselines are CPNs trained on low-light images, well-lit images, or both of them using the pose estimation loss only, which are denoted by Baseline-low, Baseline-well, and Baseline-all, respectively. Also, when a low-light image enhancement technique is incorporated, low-light training images are first enhanced and both of the enhanced images and well-lit images are used to train CPN, for which low-light test images are also enhanced by the same technique; we found that this strategy maximizes the advantage of image enhancement as shown in Table. a2 of the supplement. Pose estimation performance of ours and these solutions is given in Table 2.

Our method clearly outperforms all the three baselines in the four low-light splits, and is even on par with Baseline-well in the well-lit split. As expected, Baseline-low and Baseline-well are significantly inferior in the well-lit and low-light splits, respectively, due to the substantial gap between training and testing images in lighting conditions. Baseline-all achieves the best among the baselines, but its performance is still limited compared with other approaches including ours. These results suggest that it is not straightforward to learn a common model working under both low-light and well-lit conditions, as also reported in [39, 43]. In contrast to these early findings, we successfully manage to train a single model that performs well under both of the two lighting conditions through LSBN and LUPI.

To evaluate the efficacy of low-light image enhancement, we adopt two enhancement methods: LLFlow [62] as a learning-based method and LIME [19] as a traditional method based on the Retinex theory. The combinations of these enhancement techniques and CPN are denoted as ‘LLFlow + Baseline-all’ and ‘LIME + Baseline-all’. As shown in the table, low-light image enhancement helps improve performance in two evaluation settings, LL-N and LL-H, but it rather degrades performance in LL-E. Furthermore, it additionally imposes immense inference latency, and demands a large amount of memory footprint when adopting learning-based methods like LLFlow. On the other hand, our method outperforms them by large margins in all the five splits with only a small number of additional parameters and no additional inference latency.

Finally, our method is compared with two domain adaptation (DA) methods: DANN [15] for feature-level DA and

	Training data			AP@0.5:0.95					Param. (M)	Latency (sec)
	LL	WL	Enhanced-LL	LL-N	LL-H	LL-E	LL-A	WL		
Baseline-low	✓			32.6	25.1	13.8	24.6	1.6	27.37	1.07
Baseline-well		✓		23.5	7.5	1.1	11.5	68.8	27.37	1.07
Baseline-all	✓	✓		33.8	25.4	<u>14.3</u>	25.4	57.9	27.37	1.07
LLFlow + Baseline-all		✓	✓	35.2	20.1	8.3	22.1	65.1	66.23	3.34
LIME + Baseline-all		✓	✓	<u>38.3</u>	<u>25.6</u>	12.5	<u>26.6</u>	63.0	27.37	1.65
DANN	✓	✓		34.9	24.9	13.3	25.4	58.6	27.37	1.07
AdvEnt	✓	✓		35.6	23.5	8.8	23.8	62.4	27.37	1.07
Ours	✓	✓		42.3	34.0	18.6	32.7	<u>68.5</u>	27.53	1.07

Table 2. Pose estimation accuracy in AP@0.5:0.95 on the ExLPose dataset. In the case of training data, LL, WL, and Enhanced-LL indicate low-light, well-lit, and enhanced low-light images, respectively. In the case of evaluation splits, LL-E, LL-N, LL-H, LL-A, WL stand for low-light-easy, Low-light-normal, low-light-hard, low-light-all, and well-lit splits, respectively. The number of parameters and prediction latency of each method are reported along with the accuracy. Baseline-low, Baseline-well, and Baseline-all are the base pose estimation models (*i.e.*, CPN [6]) trained on low-light, well-lit, and both, respectively.

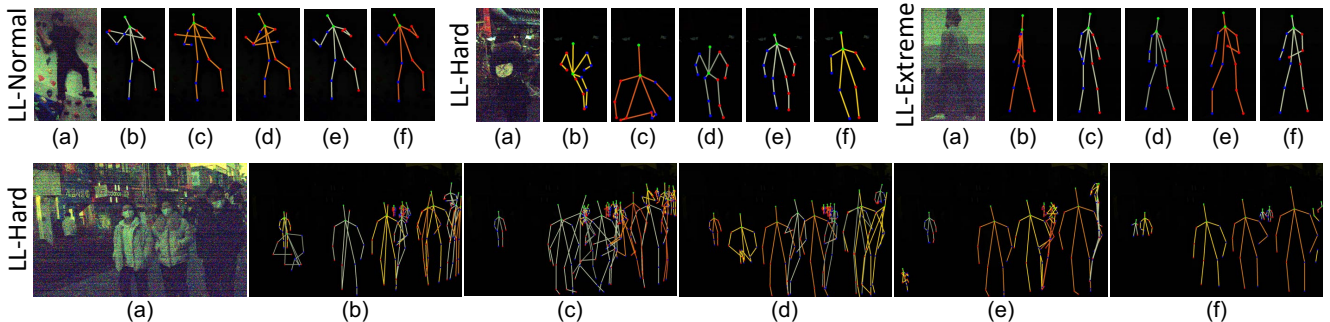


Figure 7. Qualitative results of single-person (*top*) and multi-person (*bottom*) pose estimation on the ExLPose dataset. Predicted poses and labels are visualized on corresponding low-light images. (a) Scaled low-light images. (b) Baseline-all. (c) DANN. (d) LIME + Baseline-all. (e) Ours. (f) Ground-truth. More results can be found in Sec. I of the supplementary material.

AdvEnt [59] for output-level DA. Note that existing DA methods for 2D human pose estimation [3, 26] are not compared since they are essentially based on DANN [15]. For training CPN with these methods, we assign well-lit images to a source domain and low-light images to a target domain, and utilize pose labels of both domains. The results in the table show that the direct adaptation between low-light and well-lit conditions is not effective due to the large domain gap, which suggests the necessity of LSBN. Thanks to LSBN and LUPI, our method clearly surpasses the two methods in every split. Further analysis on the impact of LSBN and LUPI is presented in Sec. G of the supplement.

5.1.2 Qualitative Results on the ExLPose Dataset

Our method is qualitatively compared with Baseline-all, DANN, LIME + Baseline-all in Fig. 7. As shown in the figure, Baseline-all and DANN fail to predict poses frequently. LIME + Baseline-all performs best among the competitors, but often fails to capture details of poses, in particular under more difficult low-light conditions. Our method clearly exhibits the best results; it estimates human poses accurately even under the LL-E condition with severe noises.

	Label		Paired WL-LL	LL-N	LL-H	LL-E	LL-A	WL
	WL	LL						
(a)	✓	✓	✓	42.3	34.0	18.5	32.7	<u>68.5</u>
(b)	✓	✓		<u>41.1</u>	<u>30.7</u>	<u>15.5</u>	<u>30.2</u>	65.9
(c)	✓			5.0	4.7	4.5	4.6	52.8

Table 3. Analysis on the significance of the paired low-light and well-lit images of the ExLPose dataset. (a) LSBN + LUPI w/pair (Ours) (b) LSBN + LUPI w/o pair. (c) Ours-UDA w/o pair.

5.1.3 Significance of the ExLPose Dataset

The ExLPose dataset has two significant properties: It provides pairs of aligned low-light and well-lit images, and it thus enables accurate labeling of low-light images. The impact of these properties is investigated by additional experiments, whose results are summarized in Table 3.

First, we study the importance of pairing low-light and well-lit images. To this end, we compare our method (LSBN + LUPI w/pair) with its variant (LSBN+LUPI w/o pair) disregarding the pair relations, trained with unrelated low-light and well-lit images. As demonstrated in the table, our method clearly outperforms the variant, which suggests the contribution of pairing low-light and well-lit im-

AP@0.5:0.95	LL-N	LL-H	LL-E	LL-A	WL
Baseline-all	33.8	25.4	14.3	25.4	57.9
LUPI	34.2	23.1	11.2	24.0	61.7
LSBN	39.0	30.2	<u>18.3</u>	30.1	<u>67.2</u>
LSBN + LUPI- <i>feat</i>	<u>39.7</u>	<u>30.8</u>	17.4	30.4	65.2
LSBN + LUPI (Ours)	42.3	34.0	18.5	32.7	68.5

Table 4. Analysis on the impact of LUPI and LSBN.

AP@0.5:0.95	A7M3	RICOH3	Avg.
Base-low	23.7	23.9	23.8
Base-well	15.2	15.6	15.4
Base-all	32.8	<u>31.7</u>	<u>32.2</u>
LLFlow + Base-all	25.6	28.2	27.0
LIME + Base-all	<u>33.2</u>	28.4	30.7
DANN	27.9	30.6	29.3
AdvEnt	28.2	29.0	28.6
Ours	35.3	35.1	35.2

Table 5. Quantitative results on the ExLPose-OCN dataset. Base denotes Baseline (*i.e.*, CPN [6]).

ages. Next, we investigate the impact of accurate pose labels for low-light images by comparing our method with its another variant trained in an unsupervised domain adaptation (UDA) setting. To be specific, the second variant (Ours-UDA w/o pair) is trained with both LSBN and LUPI using labeled well-lit images and unlabeled low-light images, which are unpaired, following the common problem setting of UDA. As demonstrated in the table, this variant performs worst, which justifies the significance of pose labeling on low-light images as well as that of pairing low-light and well-lit images in our dataset.

5.1.4 Ablation Study

We investigate the impact of LSBN and LUPI by the ablation study in Table 4. In the table, LSBN and LUPI denote variants of ours trained with either LSBN or LUPI, respectively. The inferior performance of LUPI implies that LSBN is essential to bridge the large gap between low-light and well-lit conditions effectively. For the same reason, LSBN significantly improves performance over Baseline-all. LUPI further contributes to the outstanding performance of our method: It improves LSBN substantially in every split when integrated. For an in-depth analysis on the impact of LUPI, our method is also compared with another variant (LSBN + LUPI-*feat*) that directly approximates features of the teacher instead of its neural styles. The performance gap between our model and this variant empirically justifies the use of neural styles for LUPI.

5.1.5 Results on the ExLPose-OCN Dataset

To demonstrate the generalization capability of our method, we evaluate our method and the other solutions on the




Figure 8. Qualitative results on the ExLPose-OCN dataset. Results are visualized on low-light images.

AP@0.5:0.95	LL-N	LL-H	LL-E	LL-A	WL
Baseline-low	25.4	18.2	6.0	17.2	1.2
Baseline-well	15.3	2.7	0.4	6.7	59.9
Baseline-all	26.4	18.2	6.1	17.6	52.3
LLFlow + Baseline-all	28.7	15.7	5.3	17.4	<u>60.7</u>
LIME + Baseline-all	<u>31.9</u>	<u>21.2</u>	<u>7.6</u>	<u>21.1</u>	<u>57.7</u>
DANN	28.0	17.5	5.3	17.8	52.0
Ours	35.6	25.0	11.6	25.0	61.5

Table 6. Quantitative results of multi-person pose estimation.

ExLPose-OCN dataset. As shown in Table 5, all enhancement and DA methods do not perform well on the ExLPose-OCN dataset; their performance is even inferior to that of Baseline-all (Base-all). In contrast, our method achieves the best on the ExLPose-OCN dataset also, which suggests that it improves the generalization ability as well as pose estimation accuracy. Fig. 8 also demonstrates that our method qualitatively outperforms Baseline-all (Base-all). Additional qualitative results are given in Sec. I.1 of the supplement.

5.2. Multi-person Pose Estimation

Our method and most of the potential solutions can be extended to multi-person pose estimation since they are applicable to person detection as well as pose estimation. To this end, we train Cascade R-CNN [2] for person detection with the ExLPose dataset, and then utilize bounding boxes predicted by the detector instead of ground-truths; technical details and results of person detection in low-light conditions are presented in Sec. D.2 and Sec. H of the supplement. In Table 6, each method is applied to both person detection and pose estimation models, and our method outperforms all the others in every split. As shown in the bottom row of Fig. 7, our method qualitatively outperforms other solutions.

6. Conclusion

We have introduced the first attempt to estimate human poses in extremely low-light images. To this end, we have first presented the ExLPose dataset that provides paired low-light and well-lit images with accurate human pose labels. The novel model and training strategy also have been proposed for learning representations insensitive to lighting conditions using the paired well-lit images as privileged information. Our method has been evaluated on real images taken under various low-light conditions, where it clearly outperforms domain adaptation and image enhancement methods.

Acknowledgement. This work was supported by Samsung Research Funding & Incubation Center of Samsung Electronics under Project Number SRFC-IT1801-52.

References

- [1] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [2] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [3] Jinkun Cao, Hongyang Tang, Hao-Shu Fang, Xiaoyong Shen, Cewu Lu, and Yu-Wing Tai. Cross-domain adaptation for animal pose estimation. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [4] Woong-Gi Chang, Tackgeun You, Seonguk Seo, Suha Kwak, and Bohyung Han. Domain-specific batch normalization for unsupervised domain adaptation. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [5] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [6] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded pyramid network for multi-person pose estimation. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [7] Viviana Crescitelli, Atsutake Kosuge, and Takashi Oshima. An rgb/infra-red camera fusion approach for multi-person pose estimation in low light environments. In *Proc. IEEE Sensors Applications Symposium (SAS)*, 2020.
- [8] Viviana Crescitelli, Atsutake Kosuge, and Takashi Oshima. Poison: Human pose estimation in insufficient lighting conditions using sensor fusion. *IEEE Transactions on Instrumentation and Measurement (TIM)*, 2021.
- [9] Dengxin Dai, Christos Sakaridis, Simon Hecker, and Luc Van Gool. Curriculum model adaptation with synthetic and real data for semantic foggy scene understanding. *International Journal of Computer Vision (IJCV)*, 2020.
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: a large-scale hierarchical image database. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [11] M-P Dubuisson and Anil K Jain. A modified hausdorff distance for object matching. In *Proc. International Conference on Pattern Recognition (ICPR)*, 1994.
- [12] J. Feyereisl and U. Aickelin. Privileged information for data clustering. *Information Sciences*, 2012.
- [13] Jan Feyereisl, Suha Kwak, Jeany Son, and Bohyung Han. Object localization based on structural svm using privileged information. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2014.
- [14] Shereen Fouad, Peter Tino, Somak Raychaudhury, and Petra Schneider. Incorporating privileged information through metric learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2013.
- [15] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *Journal of Machine Learning Research (JMLR)*, 2016.
- [16] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [17] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *Proc. International Conference on Learning Representations (ICLR)*, 2015.
- [18] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [19] Xiaojie Guo, Yu Li, and Haibin Ling. Lime: Low-light image enhancement via illumination map estimation. *IEEE Transactions on Image Processing (TIP)*, 2017.
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [21] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *Proc. International Conference on Learning Representations (ICLR)*, 2019.
- [22] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. In *Proc. International Conference on Learning Representations (ICLR)*, 2018.
- [23] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [24] Judy Hoffman, Saurabh Gupta, and Trevor Darrell. Learning with side information through modality hallucination. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [25] Haidi Ibrahim and Nicholas Sia Pik Kong. Brightness preserving dynamic histogram equalization for image contrast enhancement. *IEEE Transactions on Consumer Electronics (TCM)*, 2007.
- [26] Tomas Jakab, Ankush Gupta, Hakan Bilen, and Andrea Vedaldi. Self-supervised learning of interpretable keypoints from unlabelled videos. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [27] Haiyang Jiang and Yinqiang Zheng. Learning to see moving objects in the dark. In *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- [28] Yifan Jiang, Xinyu Gong, Ding Liu, Yu Cheng, Chen Fang, Xiaohui Shen, Jianchao Yang, Pan Zhou, and Zhangyang Wang. Enlightengan: Deep light enhancement without paired supervision. *IEEE Transactions on Image Processing (TIP)*, 2021.
- [29] Daniel J Jobson, Zia-ur Rahman, and Glenn A Woodell. Properties and performance of a center/surround retinex. *IEEE Transactions on Image Processing (TIP)*, 1997.

- [30] Sam Johnson and Mark Everingham. Clustered pose and nonlinear appearance models for human pose estimation. In *Proc. British Machine Vision Conference (BMVC)*, 2010.
- [31] Dmytro Kotovenko, Artsiom Sanakoyeu, Sabine Lang, and Bjorn Ommer. Content and style disentanglement for artistic style transfer. In *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- [32] John Lambert, Ozan Sener, and Silvio Savarese. Deep learning under privileged information using heteroscedastic dropout. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [33] Chulwoo Lee, Chul Lee, and Chang-Su Kim. Contrast enhancement based on layered difference representation of 2d histograms. *IEEE Transactions on Image Processing (TIP)*, 2013.
- [34] Sohyun Lee, Taeyoung Son, and Suha Kwak. Fifo: Learning fog-invariant features for foggy scene segmentation. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [35] Jiefeng Li, Can Wang, Hao Zhu, Yihuan Mao, Hao-Shu Fang, and Cewu Lu. Crowdpose: Efficient crowded scenes pose estimation and a new benchmark. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [36] Siyuan Li, Iago Breno Araujo, Wenqi Ren, Zhangyang Wang, Eric K. Tokuda, Roberto Hirata Junior, Roberto Cesar-Junior, Jiawan Zhang, Xiaojie Guo, and Xiaochun Cao. Single image deraining: A comprehensive benchmark analysis. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [37] Jinxiu Liang, Jingwen Wang, Yuhui Quan, Tianyi Chen, Jiaying Liu, Haibin Ling, and Yong Xu. Recurrent exposure generation for low-light face detection. *IEEE Transactions on Multimedia (TMM)*, 2021.
- [38] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: common objects in context. In *Proc. European Conference on Computer Vision (ECCV)*, 2014.
- [39] Yuen Peng Loh and Chee Seng Chan. Getting to know low-light images with the exclusively dark dataset. *Computer Vision and Image Understanding (CVIU)*, 2019.
- [40] D. Lopez-Paz, B. Schölkopf, L. Bottou, and V. Vapnik. Unifying distillation and privileged information. In *Proc. International Conference on Learning Representations (ICLR)*, 2016.
- [41] Kin Gwn Lore, Adedotun Akintayo, and Soumik Sarkar. L1-net: A deep autoencoder approach to natural low-light image enhancement. *Pattern Recognition Letter (PRL)*, 2017.
- [42] Feifan Lv, Feng Lu, Jianhua Wu, and Chongsoon Lim. Mblen: Low-light image/video enhancement using cnns. In *Proc. British Machine Vision Conference (BMVC)*, 2018.
- [43] Igor Morawski, Yu-An Chen, Yu-Sheng Lin, and Winston H. Hsu. Nod: Taking a closer look at detection under extreme low-light conditions with night object detection dataset. In *Proc. British Machine Vision Conference (BMVC)*, 2021.
- [44] Jaesung Rim, Geonung Kim, Jungeon Kim, Junyong Lee, Seungyong Lee, and Sunghyun Cho. Realistic blur synthesis for learning image deblurring. In *Proc. European Conference on Computer Vision (ECCV)*, 2022.
- [45] Jaesung Rim, Haeyun Lee, Jucheol Won, and Sunghyun Cho. Real-world blur dataset for learning and benchmarking deblurring algorithms. In *Proc. European Conference on Computer Vision (ECCV)*, 2020.
- [46] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Guided curriculum model adaptation and uncertainty-aware evaluation for semantic nighttime image segmentation. In *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- [47] Christos Sakaridis, Dengxin Dai, Simon Hecker, and Luc Van Gool. Model adaptation with synthetic and real data for semantic dense foggy scene understanding. In *Proc. European Conference on Computer Vision (ECCV)*, 2018.
- [48] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision (IJCV)*, 2018.
- [49] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. ACDC: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [50] Yukihiro Sasagawa and Hajime Nagahara. Yolo in the dark-domain adaptation method for merging multiple models. In *Proc. European Conference on Computer Vision (ECCV)*, 2020.
- [51] Steffen Schneider, Evgenia Rusak, Luisa Eck, Oliver Bringmann, Wieland Brendel, and Matthias Bethge. Improving robustness against common corruptions by covariate shift adaptation. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2020.
- [52] Viktoriia Sharmanska, Novi Quadrianto, and Christoph H. Lampert. Learning to rank using privileged information. In *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2013.
- [53] Baifeng Shi, Dinghuai Zhang, Qi Dai, Zhanxing Zhu, Yadong Mu, and Jingdong Wang. Informative dropout for robust representation learning: A shape-bias perspective. In *Proc. International Conference on Machine Learning (ICML)*, 2020.
- [54] Taeyoung Son, Juwon Kang, Namyup Kim, Sunghyun Cho, and Suha Kwak. Urie: Universal image enhancement for visual recognition in the wild. In *Proc. European Conference on Computer Vision (ECCV)*, 2020.
- [55] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [56] Alexander Toshev and Christian Szegedy. Deeppose: Human pose estimation via deep neural networks. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [57] Vladimir Vapnik and Rauf Izmailov. Learning using privileged information: Similarity control and knowledge transfer. *Journal of Machine Learning Research (JMLR)*, 2015.
- [58] Vladimir Vapnik and Akshay Vashist. A new learning paradigm: Learning using privileged information. *Neural Networks*, 2009.

- [59] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [60] Ruixing Wang, Xiaogang Xu, Chi-Wing Fu, Jiangbo Lu, Bei Yu, and Jiaya Jia. Seeing dynamic scene in the dark: A high-quality video dataset with mechatronic alignment. In *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [61] Wenjing Wang, Wenhan Yang, and Jiaying Liu. Hla-face: Joint high-low adaptation for low light face detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.
- [62] Yufei Wang, Renjie Wan, Wenhan Yang, Haoliang Li, Lap-Pui Chau, and Alex C Kot. Low-light image enhancement with normalizing flow. In *Proc. AAAI Conference on Artificial Intelligence (AAAI)*, 2022.
- [63] Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. In *Proc. British Machine Vision Conference (BMVC)*, 2018.
- [64] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *Proc. European Conference on Computer Vision (ECCV)*, 2018.
- [65] Yuecong Xu, Jianfei Yang, Haozhi Cao, Kezhi Mao, Jianxiong Yin, and Simon See. Arid: A new dataset for recognizing action in the dark. In *Proc. International Workshop on Deep Learning for Human Activity Recognition*, 2021.
- [66] Feng Zhang, Xiatian Zhu, and Mao Ye. Fast human pose estimation. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [67] Yinqiang Zheng, Mingfang Zhang, and Feng Lu. Optical flow in the dark. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.