

Delta Sampling R-BERT for limited data and low-light action recognition

Sanchit Hira*

sanchithira76@gmail.com

Ritwik Das*

ritwikdas54@gmail.com

Abhinav Modi

abhinav.modi888@gmail.com

Daniil Pakhomov
Johns Hopkins University
dpakhom1@jhu.edu

Abstract

We present an approach to perform supervised action recognition in the dark. In this work, we present our results on the ARID dataset[60]. Most previous works only evaluate performance on large, well illuminated datasets like Kinetics and HMDB51. We demonstrate that our work is able to achieve a very low error rate while being trained on a much smaller dataset of dark videos. We also explore a variety of training and inference strategies including domain transfer methodologies and also propose a simple but useful frame selection strategy. Our empirical results demonstrate that we beat previously published baseline models by 11%.

1. Introduction

Applications for action recognition have been on the rise with the advances in the computer vision technologies. Video action recognition finds applications in diverse areas including human-computer interaction, robotics and security surveillance. One of the important contributors to these advances is the increase in number of video datasets (Kinetics[33], HMDB51[34], DAVIS[45], EPIC-Kitchens[12], AVA[22]) which were developed for tasks such as action recognition, localization and segmentation. Most of the current work[18, 46, 2, 1, 31] in this domain is focused on well-illuminated and high contrast datasets. However, many real-world applications of computer vision involve operation under non-ideal low contrast and lighting conditions. These include low-visibility environments, smart clothing sensors, self-driving at night, ultra-thin headset cameras, in-vivo imaging, night security surveillance etc. In this work, we focus on action recognition under dark lighting conditions.

In this work, we develop our strategies on the ARID[60] dataset. For a vast majority of the videos in this dataset,

*indicates equal contribution




Figure 1. Performance of current CNN based approaches to action recognition vs. our approaches.

the illuminance is very low. This is a stark difference with other publicly available datasets like HMDB51, Kinetics-400, Charades [48] etc. where the action is clearly distinguishable to the human eye. To combat this we conduct experiments using multiple image enhancement methods.

A challenging part about this dataset is the difference in distribution of the length of the clip for different action classes as well as the frame number where the action takes place. This causes the network to correlate action classes with frame positions. This leads to significantly worse performance on out of distribution test videos. Another related challenge is that for a large number of videos, the action occurs in only a small duration of the whole video. This implies that cropping clips from the larger videos and feeding them to the network might crop out the action itself. Therefore the optimal strategy would be to feed the entire video to the network. However, because of the memory constraints




Figure 2. Sampled frames from ARID dataset, a dataset dedicated for the action recognition task in dark videos. Figure from [60]

of the hardware and varying video lengths, this is not trivial. In this work, we present a novel sampling strategy and compare it with other methods.

The third challenge is that the size of the dataset under consideration is relatively small compared to other datasets. For instance, Kinetics-400 has 247K videos, EPIC-Kitchens[12] has 28K training videos, Something-Something v1 [21] has 86K videos. In comparison, this dataset only contains 1.8K videos. An additional compounding challenge is the low number of significantly different videos. Only 4.7% of the entire dataset which is equivalent to 87 scenes consists of unique backgrounds and human subjects. The rest 95.3% of the dataset is essentially a repetition of the same subject(s) in the same background(s) doing the same corresponding action. On an average, a particular subject does the same action in the same background about 20 times with very slight variations in action performance itself.

Solving all the above problems of dark frames, lopsided action distribution across classes and extremely small size of the dataset, in this paper we propose **Delta-Sampling R(2+1)D BERT (R-BERT)**. Our results in Figure 1 demonstrate an 11% improvement over previous baselines[60].

2. Related Work

The literature analysis is divided into two different sections: (i) Advances in approaches to action recognition(2.1) and (ii) Image Enhancement techniques(2.2).

2.1. Advances in action recognition

Advances in the domain of action recognition can be attributed to the rise in both the number and complexity of datasets available for benchmarking. Datasets such as KTH [47] and Weizermann [4], which contain a relatively small number of action classes were replaced by larger and more challenging datasets like Kinetics [7], UCF101[51] and HMDB51 [35].

A wide range of solutions have been proposed for action recognition tasks using pooling[20, 66], fusion [32], and using Recurrent Networks to combine various temporal features. VideoLSTM[39] performs sequential temporal modeling on 2D CNN features extracted from video frames using convolution LSTM with spatial attention. However, these methods are restricted to 2D CNNs. Also, temporal modeling was generally achieved by an additional optical flow stream or temporal pooling layers which operated on the channel dimension. 3D CNNs (C3D[54]) were also proposed for video classification and have been used for action recognition tasks. As compared to pooling and fusion techniques 3D CNNs use 3D filters to process temporal and spatial information throughout the whole network. C3D was shallow by design and thus lacked generalizability. To overcome this issue, Inception 3D (I3D) [8] and Resnet-like versions[25] of C3D were developed. These architectures were effective but require large amounts of data for training and have huge computational and memory costs compared to their 2D counterparts. R(2+1)D [56] and S3D[59] outperformed traditional 3D CNNs by factorizing the spatial and temporal convolutions.

Other recent developments include Channel-Separated Convolutional Networks(CSN)[55], which separate channel interactions and spatio-temporal interactions in 3D CNNs. This has similarities to depth-wise separable convolutions in images. Slow-fast networks[16] consisted of a slow(lower frame rate) branch and a faster(higher frame rate) branch temporally. The intuition behind this is to capture dynamics happening at different rates within the same network. The underlying architecture is a 3D CNN. 3D CNNs show promise but still lack in an effective temporal fusion strategy at the end of the architecture. This was solved by combining R(2+1)D with Bi-directional Encoder Representations from Transformers(BERT)-based temporal modeling[31]. Recently, Bertasius *et al.* [3] have introduced a convolution-free approach to video modeling called TimeSformer[2]. It is exclusively based on space-time self-attention, suitable for long-term video modeling and achieves unprecedented results on the Kinetics-400 and Diving48[38] datasets.

2.2. Image Enhancement

Conventional histogram based methods perform image enhancement by expanding the dynamic range of an image on both global[11] and local levels[52]. Many conventional methods are also proposed based on the Retinex Theory [36] which involves decomposing an image into reflectance and illumination. Fu *et al.* [17] developed a weighted variational model for estimating the illumination and reflectance. Wang *et al.*[57] proposed an algorithm to preserve naturalness of the image while processing images with non-uniform illumination. Guo *et al.*[24] perform

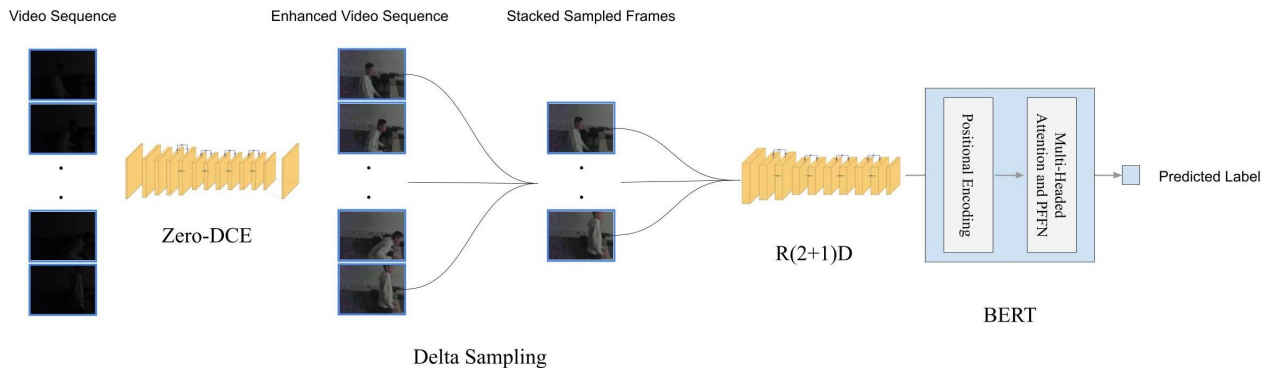


Figure 3. Overview of the proposed technique. Raw frames are first extracted from provided low-light videos. These frames are then enhanced using Zero-DCE technique. The frames are then sampled using our proposed delta sampling technique and passed to a R(2+1)D BERT network which produces the final classification.

structure-aware smoothing by estimating a coarse illumination map and then refining it by imposing a structure prior. These methods do not take into account the noise which exists in images captured in low-light. Li *et al.* proposed a Retinex model which tries to solve this issue by considering an additional noise map [37]. Yuan *et al.* perform global optimization to estimate an S-shaped curve for a given image for automatic exposure correction [65]. Most conventional methods change the distribution of image histogram or rely on potentially inaccurate physical models. This enables light enhancement on images but with the cost of creating unrealistic artifacts in the enhanced images. Advances in learning based methods were also seen in the domain of image enhancement. Bychovsky *et al.* [5] introduced the first and largest dataset for supervised image tone adjustment.

Yan *et al.* [61] presented a machine-learning based ranking approach for automatic color enhancement in an image while accounting for intermediate decisions of a human in the editing process. Yan *et al.* [62] bundled features from a pixelwise, global and contextual descriptor to achieve semantic-aware automatic photo enhancement. With advances in deep learning, a deep autoencoder based model called “LLNet” [42] was proposed for low-light image enhancement. Gharbi *et al.* [19] introduced data dependent lookup to perform real-time image enhancement on full-resolution images. Yang *et al.* [63] reformulate image correction problem as a deep High Dynamic Range (HDR) mapping problem to enhance Low Dynamic Range (LDR) images. Cai *et al.* [6] introduced a multi-exposure dataset and proposed a CNN based contrast enhancement method for a single image from multi-exposure images. However, most CNN based approaches are resource intensive as they require large quantities of training data. Recently, Chen *et al.* [10] developed an unsupervised learning-

based model using a two-way generative adversarial network (GAN) for image enhancement. Ignatov *et al.* [29] developed a weakly-supervised image-to-image GAN-based network, while, Deng *et al.* proposed EnhanceGAN [13], an approach to image correction by adversarial learning based on binary labels on aesthetic quality. Chen *et al.* [9] proposed a new dataset which includes raw sensor data to address extreme low-light image enhancement. EnlightenGAN [30] learns to enhance low-light images in an unsupervised manner by using unpaired low/normal light data. However, unsupervised GAN-based solutions usually require careful selection of unpaired training data. Reinforcement learning was also employed to enhance the image adjustment process [28, 43]. Recently, Guo *et al.* proposed a novel method to perform Zero-Reference Deep Curve Estimation (Zero-DCE) [23] thus eliminating the need for labelled or carefully selected unlabelled data.

3. Methodology

3.1. Dataset Description

Our dataset comprises of 6 different action classes from the ARID dataset, which are walking, running, standing, sitting, waving and turning. The task is to classify RGB videos into one of these 6 action classes. The dataset provided includes a total of 1837 videos. The training and the validation split consist of 1470 samples and 367 samples respectively. The video length in the dataset ranges from 33 to 225 frames.

3.2. Image Enhancement

All of the videos in the dataset are taken at night or under very low-lighting conditions. Figure 2 gives us an insight into the dataset. To address this issue we train an enhancement method called Zero-Reference Deep Curve Estimation (Zero-DCE) for low-light image enhancement. In

this method we train a model to infer dynamic range adjustments for a given image. It is equivalent to a curve estimation problem and does not require paired high and low lighting images to achieve high quality results. In this approach, each image is a distinct training sample with no paired supervision required. A weighted, linear combination of the following loss functions is used to train the DCE Model.

Spatial Consistency Loss :

$$L_{spa} = \frac{1}{K} \sum_{i=1}^K \sum_{j \in \Omega(i)} (|(Y_i - Y_j)| - |(I_i - I_j)|)^2, \quad (1)$$

where K is the number of local region, and $\Omega(i)$ is the four neighboring regions (top, down, left, right) centered at the region i . Y and I are the average intensity values of the local region in the enhanced version and input image, respectively.

Exposure Control Loss :

$$L_{exp} = \frac{1}{M} \sum_{k=1}^M |Y_k - E|, \quad (2)$$

where M represents the number of nonoverlapping local regions of size 16×16 , Y is the average intensity value of a local region in the enhanced image. E is a hyperparameter describing the well exposedness level.

Color Constancy Loss : The color constancy loss L_{col} can be expressed as:

$$L_{col} = \sum_{(p,q) \in \varepsilon} (J^p - J^q)^2, \varepsilon = \{(R, G), (R, B), (G, B)\}, \quad (3)$$

where J^p denotes the average intensity value of p channel in the enhanced image and (p, q) represents a pair of channels.

Illumination Smoothness Loss :

$$L_{tv_A} = \frac{1}{N} \sum_{n=1}^N \sum_{c \in \xi} (|\nabla_x A_n^c| + |\nabla_y A_n^c|)^2, \xi = \{R, G, B\}, \quad (4)$$

where N is the number of iteration, ∇_x and ∇_y represent the horizontal and vertical gradient operations, respectively.

Total Loss. The total loss can be expressed as:

$$L_{total} = W_{spa} L_{spa} + L_{exp} + W_{col} L_{col} + W_{tv_A} L_{tv_A}, \quad (5)$$

where W_{col} and W_{tv_A} are the weights of the losses.

3.3. Delta Sampling Strategy

The second challenge of this dataset is that the action happens only over a duration of few frames in a particular video for a large majority of the videos. A typical video consists of the following: Preparation for the action, actual action and end of action. Furthermore, the distribution of video length is different for different actions as shown in the Figure 4. This leads to overfitting, creating a bias for the

frame number at which the maximum change occurs and/or to the frame distribution instead of learning the action itself. We propose a simple but useful Delta Sampling strategy to solve this problem which beats other learning-based benchmarks like SlowFast Networks and Temporal Attention as well as other sampling heuristics.

Consider an input video $X = (X_1, \dots, X_N)$, X_i represents the i^{th} frame and N represents the total number of frames in the original video. Let T be the number of frames the network takes as input. We first calculate a base sampling rate for the video as N/T . Then we further add a δ to the calculated base sampling rate. Thus, the overall sampling rate of the video becomes $(N + T * \delta)/T$. In our case, we sample δ to be a random floating point number from a uniform distribution in the range $[\beta, \gamma]$.

$$\delta \sim Uniform[\beta, \gamma] \quad (6)$$

If β is less than zero, it is equivalent to cropping the video which leads to a much higher chance of loss of information on the important action delineation frames. Therefore in our experiments we set the lower bound of β to be 0. Furthermore we cap the maximum sampling rate to α , as a very large sampling rate might cause very little information to be captured in the sampled frames.

$$S_\delta = \min\left(\frac{N + T \cdot \delta}{T}, \alpha\right) \quad (7)$$

where S_δ is the adjusted sampling rate of the input video.

After the above sampling transformation, let the new video sequence be represented as $X^* = (X_1^*, \dots, X_{N/S_\delta}^*)$. The input to the network is padded with blank frames on both sides. If P_1 and P_2 are the number of blank frames on both sides, then P_1 is sampled as follows.

$$P_1 \sim Uniform[0, T - N/S_\delta] \quad (8)$$

Then, P_2 can be calculated as

$$P_2 = T - N/S_\delta - P_1 \quad (9)$$

Therefore, we can calculate

$$\bar{X} = (P_1 \text{ frames}, X_1^*, \dots, X_{N/S_\delta}^*, P_2 \text{ frames}) \quad (10)$$

where \bar{X} denotes the final video input sequence to the network. Our approach is summarized in Figure 5.

3.4. Network Architecture

We feed in input video sequence through image enhancement and delta sampling strategy as described in the previous sections. For the video-classification network, we use R(2+1)D - BERT with a Resnet-34 backbone. Our findings show that the temporal attention mechanism within

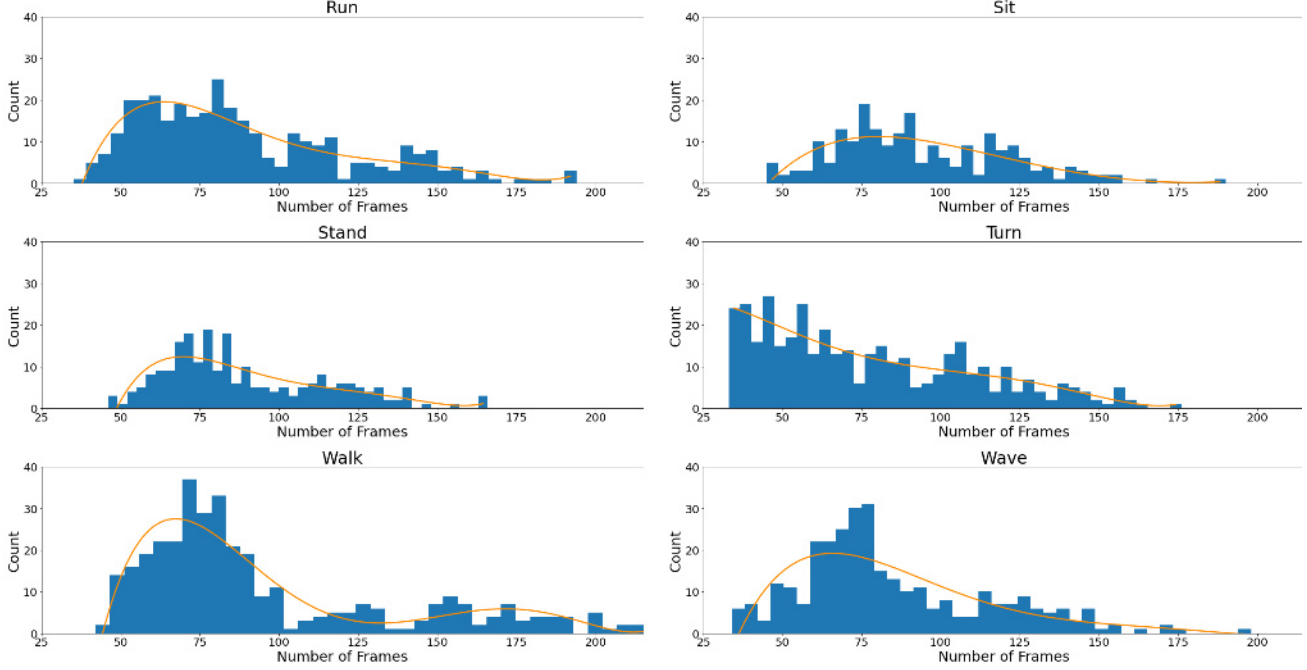


Figure 4. Distribution of frame length for each action. Differences in distribution across action classes leads to overfitting to the distribution itself instead of the action class. We solve this by proposing Delta Sampling Strategy.

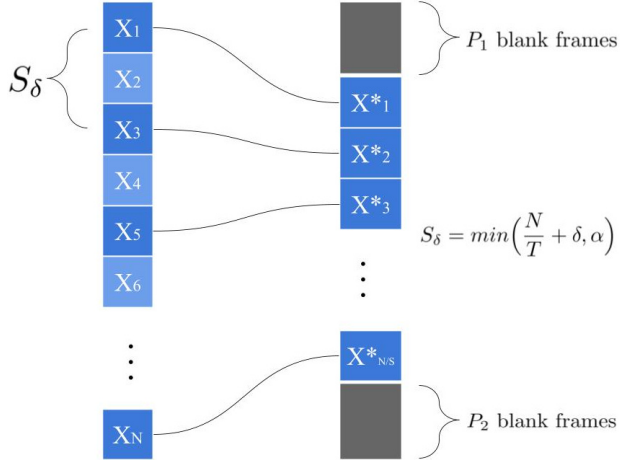


Figure 5. Frame sampling using Delta Sampling strategy. Sampling is done at rate S_δ and the sampled frames are padded with P_1 and P_2 blank frames at beginning and end respectively.

BERT gives a much better performance than temporal average pooling as shown in Section 4. The BERT model preserves the temporal positional information in the input features by adding a learned positional encoding. A classification embedding is also appended [14]. After attending the features with the self-attention mechanism of BERT, a Position wise Feed Forward Network (PFFN) is applied. Finally, a simple linear layer generates the final classification output of the model.

$$y_i = PFFN \left(\frac{1}{N(x)} \sum_{\forall j} g(\mathbf{x}_j) f(\mathbf{x}_i, \mathbf{x}_j) \right) \quad (11)$$

, where \mathbf{x}_i are the i^{th} temporal index of the input embeddings and \mathbf{x}_j are all possible combinations of input embeddings. $f(\mathbf{x}_i, \mathbf{x}_j) = \text{softmax}_j \left(\theta(\mathbf{x}_i)^T \phi(\mathbf{x}_j) \right)$ gives the similarity between \mathbf{x}_i and \mathbf{x}_j , where $\theta(\cdot)$ and $\phi(\cdot)$ are linear projections. PFFN is defined in equation 12, where $GELU(\cdot)$ represents the Gaussian Error Linear Unit activation function [26].

$$PFFN(x) = \mathbf{W}_2 GELU(\mathbf{W}_1 \mathbf{x} + \mathbf{b}_1) + \mathbf{b}_2 \quad (12)$$

We train our model end-to-end with focal loss[40]. We use horizontal flipping, random rotation and random translation for video augmenting our dataset. Figure 3 shows the end-to-end pipeline of the proposed method.

4. Implementation and Experiments

4.1. Image Enhancement

For image enhancement, we train Zero-DCE model on 2000 random images of different illumination and 2500 images from the ARID dataset. We sample the 2500 images uniformly from the videos in our training set. We use the

Method	Accuracy
No Enhancement Baseline	87.48
Gamma Image Correction	84.62
Enlighten GAN	87.50
Zero-DCE Enhanced Images	90.46

Table 1. Comparison of different methods of enhancing brightness of low-light images.

Adam optimizer and train our model for 100 epochs. We set W_{tv_A} to 200, W_{col} to 5 and W_{spa} to 10. We set the learning rate to $1e - 4$. We also explore Gamma Image Correction(GIC) and EnlightenGAN to enhance the images. However the latter two methods do not result in any improvement in the downstream action recognition task. For gamma correction we observe that a blanket gamma correction does not improve our results irrespective of the value of gamma and in-fact decreases it. We hypothesize that this is due to varying degrees of darkness across the dataset which cannot be captured properly by the gamma correction table. Upon a visual inspection, as shown in Figure 6 we notice that EnlightenGAN had a lot of image artifacts in addition to enlightening the image uniformly with little consideration of lighting conditions at existing brighter areas of the image. Zero-DCE on the other hand solves both of these problems and the generated dynamic lighting is much more suitable to the downstream tasks. We use the trained DCE model as a frozen-network for both training and inference phases of R(2+1)-D BERT. Table 1 shows the performance benefit of using Zero-DCE as compared to other baselines.

4.2. Delta Sampling Strategy

For the Delta Sampling strategy, our network takes $T = 64$ frames as input. The number of frames in the original videos vary from a minimum of 33 to a maximum of 225 as shown in the distribution curves in Figure 4. We set $\beta = 0$, $\gamma = 1.5$ and the maximum sampling rate(α) = 4. P_1 and therefore P_2 are randomly determined at train time within their respective bounds. During test-time, we average predictions across 5 different sets of (δ, P_1, P_2) values as a test-time sampling augmentation technique. We compare our method to the following alternative approaches.

1. **Slow Fast Network:** Previous works have proposed similar approaches like SlowFast network. It involves a slower and a faster network streams which operate at different temporal resolutions. Intuitonally, it tries to solve the same problem of different actions occurring at different rates across the dataset. We demonstrate that our approach is superior than the SlowFast network.
2. **Squeeze-Excitation Module:** The second method we test our strategy against is attention along the time di-

Method	Accuracy
Slow-Fast Network	44.40
Squeeze Excitation Module	70.78
Constant Sampling Rate	80.93
Length Adjusted Sampling Rate	85.29
Variable Sampling Rate	88.28
Delta Sampling Strategy (Ours)	90.46

Table 2. Comparison of different methods of sampling frames in videos.

mension in the backbone. For this method, we use Squeeze-Excitation [27] within the Resnet-34 backbone.

3. **Constant Sampling Rate:** In this third study, we compare our approach with a constant sampling rate of $S_{constant} = N_{max}/T$. Here N_{max} is the maximum number of frames across all videos in the training set. Since $T * S_{constant} > N_i$ for all videos where $N_i < N_{max}$ we pad the remaining frames on both sides with P_1 and P_2 blank frames using the same strategy as described in Section 3.3.
4. **Length Adjusted Sampling Rate:** The fourth baseline we compare against is a sampling method at a video length adjusted sampling rate of $S_{len_adj} = N_i/T$ i.e., which covers the entire T frames where N_i is the number of frames in the i_{th} training sample. In this approach, there is no need to pad any sampled training videos since we intentionally choose S_{len_adj} to cover all T frames.
5. **Variable Sampling Rate:** For the last sampling rate comparative study, we use a randomly chosen sampling rate $S_{var} \sim Uniform[N_{min}/T, N_{max}/T]$. Here N_{min} is the minimum number of frames across all videos in the training set. Similar to the method proposed in section 3.3 we pad the remaining frames on both sides with P_1 and P_2 blank frames whenever $T * S_{var} > N_i$.

According to table 2, the experiments show that our strategy is superior to all of the above methods.

4.3. Network Architecture Details

The models are trained using Ranger optimizer [58], which combines Rectified Adam[41], LookAhead optimizer [67] and Gradient Centralization [64] into one optimizer. Given the small size of the dataset we consider here, we use a small batch size of 2 which produces a regularization effect as discussed by Smith *et al.*[50]. We train our models for 300 epochs using One Cycle Learning strategy [49] and focal loss [40], as defined in the equation 13, as our classification loss.



Figure 6. Visual Enhancement results (face blurred) across different methods. Zero-DCE is the most dynamically illuminated with more enhancement on the darker areas and lesser on the over-exposed areas.

$$FL(p_t) = -\alpha_t (1 - p_t)^\gamma \log(p_t) \quad (13)$$

To limit the computational requirements, we use resized image frames of size 112×112 and sample 64 frames from each video, using our delta sampling strategy discussed in section 3.3. The output of the spatial R(2+1)D branch is a $512 \times 8 \times 7 \times 7$ feature map, which is then average pooled in the spatial domain to reduce it to a dimension of $512 \times 8 \times 1 \times 1$. The features are then normalized in the time domain before passing them on to the BERT model which generates the final classification output. The BERT model uses a hidden size of 512 with a single layer with 8 attention heads. We use PyTorch framework [44] for our experiments.

For training the R(2+1)D BERT model, we train on 6 classes of low-light videos of the ARID Dataset. We present the following comparative studies against our strategy :

1. **3D Convolutional Networks** - We compare against C3D, 3D ResNeXt-101, 3D SqueezeNet and other 3D action recognition networks in Figure 1.
2. **R(2+1)D Temporal Global Average Pooling(TGAP)** We take R(2+1)D network and use a temporal pooling layer at the end. We see that this leads to very interesting performance pitfalls. The network errors frequently include confusion between sitting and standing (which is sitting but in reverse), and also walking and running at times. We attribute this confusion to the loss of temporal relationship due to a global average pooling layer. Hence, we use BERT in our successive experiments as it has temporal positional encodings.
3. **Efficient-Net BERT** - Here, we use a modified Efficient-Net [53] as the backbone. However, we observe that this network is extremely heavy and hard to train given limited computation resources. Therefore, we revert back to the lighter R(2+1)D backbone in the next two sections.
4. **R(2+1)D BERT with Optical Flow** - In this study, we first compute dense optical flow on the frames [15] and then pass the 2 flow frames along with the 3 RGB

Method	Accuracy
R(2+1)D TGAP	82.74
Efficient-Net BERT	64.12
R(2+1)D-BERT with Optical Flow	67.89
Delta Sampling Resnet-BERT	90.46

Table 3. Comparison of different network architectures on action recognition accuracy.

channels. We notice empirically that this experiment leads the network to hugely overfit on the training set and therefore generalize poorly. We hypothesize that this is partially due to large noise in the videos which leads to noisy optical flow values, exacerbated by the small training set size (87 unique scenes) and increase in the number of partially redundant optical flow features.

Table 3 presents the performance on TGAP and BERT architectures with different backbones and input types as discussed above.

4.4. Additional Data

We also present our improved results when we add corresponding data from HMDB-51. We choose to add the videos from HMDB-51 belonging to the same 6 classes in ARID and train them in a fully supervised way. This leads to an addition of 1420 videos to our existing training set. Since the domains of both HMDB-51 and ARID are different we compare this with Multi-Representation Adaptation Network(MRAN) [68] for cross-domain adaptation. We train MRAN with unpaired labeled videos from both source and target domains. The features extracted from the R(2+1)D model are passed to the MRAN module. The MRAN module minimizes the Conditional Maximum Mean Discrepancy (CMMD) between domains, which effectively brings the conditional distributions of the source and target domain features close to each other. The output classification of the MRAN module is then appended to the classification of the BERT module, and then final classification is obtained using a linear transform. The whole combined model is trained end-to-end using the training strategy de-

Method	Accuracy
HMDB-51 + ARID - Fully Supervised	93.73
HMDB-51 + ARID - MRAN	90.37

Table 4. Comparison of Multi-Representation adaptation network and fully supervised training.

scribed above. However, we observe better performance in the fully supervised experiment as shown in Table 4.

5. Conclusion

In this work we propose Delta-sampling Resnet BERT. We show that image enhancement techniques like ZeroDCE can significantly improve the performance of machine learning models on low-light images. We also demonstrate the effectiveness of our novel efficient frame sampling strategy on skewed frame distribution across classes and low dataset volume. On the ARID dataset, we show that these techniques in combination with recent innovations in video classification architectures, improvements in deep learning optimizers and learning rate schedules can yield state-of-the-art results. We also show that simply augmenting the enhanced videos of this dataset with videos from the HMDB-51 dataset can achieve even better performance, further demonstrating generalizability of the image enhancement and sampling strategy used.

References

- [1] Irwan Bello, William Fedus, Xianzhi Du, Ekin D. Cubuk, Aravind Srinivas, Tsung-Yi Lin, Jonathon Shlens, and Barret Zoph. Revisiting resnets: Improved training and scaling strategies, 2021.
- [2] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding?, 2021.
- [3] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? *arXiv preprint arXiv:2102.05095*, 2021.
- [4] M. Blank, L. Gorelick, E. Shechtman, M. Irani, and R. Basri. Actions as space-time shapes. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 2, pages 1395–1402 Vol. 2, 2005.
- [5] V. Bychkovsky, S. Paris, E. Chan, and F. Durand. Learning photographic global tonal adjustment with a database of input/output image pairs. In *CVPR 2011*, pages 97–104, 2011.
- [6] Jianrui Cai, Shuhang Gu, and Lei Zhang. Learning a deep single image contrast enhancer from multi-exposure images. *IEEE Transactions on Image Processing*, 27(4):2049–2062, 2018.
- [7] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4724–4733, 2017.
- [8] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
- [9] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3291–3300, 2018.
- [10] Yu-Sheng Chen, Yu-Ching Wang, Man-Hsin Kao, and Yung-Yu Chuang. Deep photo enhancer: Unpaired learning for image enhancement from photographs with gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6306–6314, 2018.
- [11] Dinu Coltuc, Philippe Bolon, and J-M Chassery. Exact histogram specification. *IEEE Transactions on Image Processing*, 15(5):1143–1152, 2006.
- [12] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The epic-kitchens dataset, 2018.
- [13] Yubin Deng, Chen Change Loy, and Xiaoou Tang. Aesthetic-driven image enhancement by adversarial learning. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 870–878, 2018.
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [15] Gunnar Farnebäck. Two-frame motion estimation based on polynomial expansion. In *Scandinavian conference on Image analysis*, pages 363–370. Springer, 2003.
- [16] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition, 2019.
- [17] X. Fu, D. Zeng, Y. Huang, X. Zhang, and X. Ding. A weighted variational model for simultaneous reflectance and illumination estimation. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2782–2790, 2016.
- [18] Deepti Ghadiyaram, Matt Feiszli, Du Tran, Xueting Yan, Heng Wang, and Dhruv Mahajan. Large-scale weakly-supervised pre-training for video action recognition, 2019.
- [19] Michaël Gharbi, Jiawen Chen, Jonathan T Barron, Samuel W Hasinoff, and Frédo Durand. Deep bilateral learning for real-time image enhancement. *ACM Transactions on Graphics (TOG)*, 36(4):1–12, 2017.
- [20] Rohit Girdhar, Deva Ramanan, Abhinav Gupta, Josef Sivic, and Bryan Russell. Actionvlad: Learning spatio-temporal aggregation for action classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 971–980, 2017.
- [21] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, Florian Hoppe, Christian Thureau, Ingo Bax, and Roland Memisevic. The “something something” video database for learning and evaluating visual common sense. pages 5843–5851, 10 2017.

- [22] Chunhui Gu, Chen Sun, David A Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, et al. Ava: A video dataset of spatio-temporally localized atomic visual actions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6047–6056, 2018.
- [23] Chunle Guo, Chongyi Li, Jichang Guo, Chen Change Loy, Junhui Hou, Sam Kwong, and Runmin Cong. Zero-reference deep curve estimation for low-light image enhancement, 2020.
- [24] X. Guo, Y. Li, and H. Ling. Lime: Low-light image enhancement via illumination map estimation. *IEEE Transactions on Image Processing*, 26(2):982–993, 2017.
- [25] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6546–6555, 2018.
- [26] Dan Hendrycks and Kevin Gimpel. Bridging nonlinearities and stochastic regularizers with gaussian error linear units. 2016.
- [27] Jie Hu, Li Shen, Samuel Albanie, Gang Sun, and Enhua Wu. Squeeze-and-excitation networks, 2019.
- [28] Yuanming Hu, Hao He, Chenxi Xu, Baoyuan Wang, and Stephen Lin. Exposure: A white-box photo post-processing framework. *ACM Transactions on Graphics (TOG)*, 37(2):1–17, 2018.
- [29] Andrey Ignatov, Nikolay Kobyshev, Radu Timofte, Kenneth Vanhoey, and Luc Van Gool. Wespe: weakly supervised photo enhancer for digital cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 691–700, 2018.
- [30] 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*, 30:2340–2349, 2021.
- [31] M. Esat Kalfaoglu, Sinan Kalkan, and A. Aydin Alatan. Late temporal modeling in 3d cnn architectures with bert for action recognition, 2020.
- [32] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014.
- [33] Will Kay, João Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset. *CoRR*, abs/1705.06950, 2017.
- [34] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. HMDB: a large video database for human motion recognition. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2011.
- [35] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. Hmdb: A large video database for human motion recognition. In *2011 International Conference on Computer Vision*, pages 2556–2563, 2011.
- [36] Edwin H Land. The retinex theory of color vision. *Scientific american*, 237(6):108–129, 1977.
- [37] M. Li, J. Liu, W. Yang, X. Sun, and Z. Guo. Structure-revealing low-light image enhancement via robust retinex model. *IEEE Transactions on Image Processing*, 27(6):2828–2841, 2018.
- [38] Yingwei Li, Yi Li, and Nuno Vasconcelos. Resound: Towards action recognition without representation bias. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 513–528, 2018.
- [39] Zhenyang Li, Kirill Gavriluk, Efstratios Gavves, Mihir Jain, and Cees GM Snoek. Videolstm convolves, attends and flows for action recognition. *Computer Vision and Image Understanding*, 166:41–50, 2018.
- [40] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [41] Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. On the variance of the adaptive learning rate and beyond. *arXiv preprint arXiv:1908.03265*, 2019.
- [42] Kin Gwn Lore, Adedotun Akintayo, and Soumik Sarkar. Llnet: A deep autoencoder approach to natural low-light image enhancement. *Pattern Recognition*, 61:650–662, 2017.
- [43] Jongchan Park, Joon-Young Lee, Donggeun Yoo, and In So Kweon. Distort-and-recover: Color enhancement using deep reinforcement learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5928–5936, 2018.
- [44] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [45] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*, 2017.
- [46] Zhaofan Qiu, Ting Yao, Chong-Wah Ngo, Xinmei Tian, and Tao Mei. Learning spatio-temporal representation with local and global diffusion, 2019.
- [47] Christian Schuldt, Ivan Laptev, and Barbara Caputo. Recognizing human actions: A local svm approach. In *Proceedings of the Pattern Recognition, 17th International Conference on (ICPR’04) Volume 3 - Volume 03*, ICPR ’04, page 32–36, USA, 2004. IEEE Computer Society.
- [48] Gunnar A. Sigurdsson, Gül Varol, Xiaolong Wang, Ivan Laptev, Ali Farhadi, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. *ArXiv e-prints*, 2016.

- [49] Leslie N Smith. Cyclical learning rates for training neural networks. In *2017 IEEE winter conference on applications of computer vision (WACV)*, pages 464–472. IEEE, 2017.
- [50] Leslie N Smith. A disciplined approach to neural network hyper-parameters: Part 1—learning rate, batch size, momentum, and weight decay. *arXiv preprint arXiv:1803.09820*, 2018.
- [51] Khurram Soomro, Amir Roshan Zamir, Mubarak Shah, Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *CoRR*, page 2012.
- [52] J Alex Stark. Adaptive image contrast enhancement using generalizations of histogram equalization. *IEEE Transactions on image processing*, 9(5):889–896, 2000.
- [53] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks, 2020.
- [54] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks, 2015.
- [55] Du Tran, Heng Wang, Lorenzo Torresani, and Matt Feiszli. Video classification with channel-separated convolutional networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5552–5561, 2019.
- [56] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6450–6459, 2018.
- [57] S. Wang, J. Zheng, H. Hu, and B. Li. Naturalness preserved enhancement algorithm for non-uniform illumination images. *IEEE Transactions on Image Processing*, 22(9):3538–3548, 2013.
- [58] Less Wright. Ranger Deep Learning Optimizer, Sept. 2020.
- [59] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 305–321, 2018.
- [60] Yuecong Xu, Jianfei Yang, Haozhi Cao, Kezhi Mao, Jianxiong Yin, and Simon See. Arid: A new dataset for recognizing action in the dark, 2020.
- [61] Jianzhou Yan, Stephen Lin, Sing Bing Kang, and Xiaoou Tang. A learning-to-rank approach for image color enhancement. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2987–2994, 2014.
- [62] Zhicheng Yan, Hao Zhang, Baoyuan Wang, Sylvain Paris, and Yizhou Yu. Automatic photo adjustment using deep neural networks. *ACM Transactions on Graphics (TOG)*, 35(2):1–15, 2016.
- [63] Xin Yang, Ke Xu, Yibing Song, Qiang Zhang, Xiaopeng Wei, and Rynson WH Lau. Image correction via deep reciprocating hdr transformation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1798–1807, 2018.
- [64] Hongwei Yong, Jianqiang Huang, Xiansheng Hua, and Lei Zhang. Gradient centralization: A new optimization technique for deep neural networks. In *European Conference on Computer Vision*, pages 635–652. Springer, 2020.
- [65] Lu Yuan and Jian Sun. Automatic exposure correction of consumer photographs. In *European Conference on Computer Vision*, pages 771–785. Springer, 2012.
- [66] Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4694–4702, 2015.
- [67] Michael R Zhang, James Lucas, Geoffrey Hinton, and Jimmy Ba. Lookahead optimizer: k steps forward, 1 step back. *arXiv preprint arXiv:1907.08610*, 2019.
- [68] Yongchun Zhu, Fuzhen Zhuang, Jindong Wang, Jingwu Chen, Zhiping Shi, Wenjuan Wu, and Qing He. Multi-representation adaptation network for cross-domain image classification. *Neural Networks*, 119:214–221, 2019.