

# Privacy-Preserving Action Recognition via Motion Difference Quantization

Sudhakar Kumawat<sup>✉</sup> and Hajime Nagahara<sup>✉</sup>

Osaka University, Japan  
{sudhakar,nagahara}@ids.osaka-u.ac.jp

**Abstract.** The widespread use of smart computer vision systems in our personal spaces has led to an increased consciousness about the privacy and security risks that these systems pose. On the one hand, we want these systems to assist in our daily lives by understanding our surroundings, but on the other hand, we want them to do so without capturing any sensitive information. Towards this direction, this paper proposes a simple, yet robust privacy-preserving encoder called BDQ for the task of privacy-preserving human action recognition that is composed of three modules: *Blur*, *Difference*, and *Quantization*. First, the input scene is passed to the *Blur* module to smoothen the edges. This is followed by the *Difference* module to apply a pixel-wise intensity subtraction between consecutive frames to highlight motion features and suppress high-level privacy attributes. Finally, the *Quantization* module is applied to the motion difference frames to remove the low-level privacy attributes. The BDQ parameters are optimized in an end-to-end fashion via adversarial training such that it learns to allow action recognition attributes while inhibiting privacy attributes. Our experiments on three benchmark datasets show that the proposed encoder design can achieve state-of-the-art trade-off when compared with previous works. Furthermore, we show that the trade-off achieved is at par with the DVS sensor-based event cameras. Code available at: [https://github.com/suakaw/BDQ\\_PrivacyAR](https://github.com/suakaw/BDQ_PrivacyAR)

**Keywords:** Action recognition, Privacy, Motion difference, Quantization, Adversarial training

## 1 Introduction

For many decades, people have been fascinated with the idea of creating computer vision (CV) systems that can see and interpret the world around them. In today’s world, as this dream turns into reality and such systems begin to be deployed in our personal spaces, there is an increased consciousness about “what” these systems see and “how” they interpret it. Nowadays, we want CV systems that can protect our visual privacy without compromising the user experience. Therefore, there is a growing interest in developing such CV systems that can prevent the camera system from obtaining detailed visual data that may contain




Fig. 1: Left- An example where downsampling is effective for privacy-preserving action recognition. Right- An example where down-sampling is detrimental for action recognition since the information about the number of fingers is lost. Unlike downsampling, the BDQ preserves both spatial and temporal resolutions.

any sensitive information, but allow it to capture useful information to successfully perform the CV task [37,30,33,14,26]. For a satisfactory user experience and strong privacy protection, a CV system must satisfy the following properties:

- **Good target task accuracy.** This is necessary for maintaining a good user experience. For example, a privacy-preserving face detection model must detect faces with high precision without revealing facial identity [33], a privacy-preserving pose estimation model must detect body key-points without revealing the person identity [14], and an action recognition model must recognize human actions without revealing their identity information [37,30].
- **Strong privacy protection.** Any privacy-preserving model, irrespective of the target task must preserve common visual privacy attributes such as identity, gender, race, color, gait, etc. [22]. Note that the definition of privacy attributes for a privacy-preserving model may vary depending on its application. Furthermore, strong privacy protection is guaranteed when the model is applied at the point of capture and it is impossible for any adversary to learn or reconstruct the privacy attributes.
- **Cost-effective and low space-time complexity.** The right to privacy is considered as one of the fundamental rights in the most countries. However, with the increasing availability of low-budget consumer cameras and smartphones, it cannot be fully guaranteed unless the privacy-preserving models are affordable to everyone. This calls for a focus on implementing such models, whether in software or hardware, in a cost-effective manner. Additionally, since the privacy-preserving model needs to be applied at the point of capture in consumer cameras and smartphones with low memory, computation, and power budgets, it must be of low space-time complexity.

In this paper, we are interested in developing a privacy-preserving encoder for the task of human action recognition while keeping the above properties in mind.

In particular, we develop a simple and robust privacy-preserving encoder called BDQ that is composed of three modules: *Blur*, *Difference*, and *Quantization*. The modules are applied in a sequential manner such that *Blur* module smoothen the edges in the input frames. The *Difference* module computes pixel-wise intensity subtraction between consecutive frames to highlight motion features and suppress high-level privacy attributes. The *Quantization* module removes the low-level privacy attributes from the motion difference frames. The parameters of BDQ are optimized in an end-to-end manner via adversarial training such that it learns to facilitate action recognition while inhibiting privacy attributes. Note that, any successful action recognition model relies on the presence of both spatial and temporal cues in its input for good performance. However, in most cases, applying a privacy-preserving encoder often destroys the spatial resolution which is sometimes detrimental for action recognition as shown in Figure 1. The design of our BDQ encoder is motivated to alleviate such issues such that it preserves both spatial and temporal cues that are essential for action recognition while discarding attributes that may reveal the privacy information. In summary, the contributions of this work are as follows.

- We propose a simple yet robust privacy-preserving encoder called BDQ for the task of human action recognition. The BDQ encoder allows important spatio-temporal cues for action recognition while preserving privacy attributes at a very low space-time complexity.
- We show that the BDQ encoder achieves state-of-the-art trade-off between action recognition and privacy preservation on three benchmark datasets SBU, KTH, and IPN, when compared with other privacy-preserving models. Moreover, we show that the trade-off achieved is at par with the DVS sensor-based event cameras.
- We provide an extensive analysis of the BDQ encoder. The analysis includes an ablation study on the components of BDQ, learning and reconstructing privacy attributes by different adversaries, and a subjective evaluation.
- Finally, we also discuss the feasibility of implementing the BDQ modules using existing hardware (please refer supplementary).

## 2 Related Work

In recent years, there is a growing interest in developing privacy-preserving vision systems for various computer vision tasks such as action recognition [37,30], face detection [33], pose estimation [14,32], fall detection [3], and posture classification [12]. Here, we provide a brief overview of privacy-preserving frameworks for various computer vision tasks, especially human action recognition.

**Privacy-preserving computer vision.** Early privacy-preserving models used hand-crafted features such as blurring, down-sampling, pixelation, and face/object replacement for protecting sensitive information [2,7,23]. Unfortunately, such approaches require extensive domain knowledge about the problem setting which may not be feasible in practice. Modern privacy-preserving frameworks take a data-driven approach to hide sensitive information. Such frameworks learn

to preserve privacy via adversarial training [28] using Deep Neural Networks (DNNs) that train the parameters of an encoder to actively inhibit the sensitive attributes in an visual data against an adversarial DNN whose task is to learn privacy attributes while allowing attributes that are essential for the computer vision task [5,24,33,37,16,21]. Note that the encoder can be a hardware module [14,33] or a software module [37,28]. Besides these works, there are other imaging systems that use optical operations to hide sensitive attributes. For example, in [25], the authors design camera systems that perform blurring and  $k$ -same face de-identification in hardware. Furthermore, in [35,6], authors explore coded aperture masks to enhance privacy.

**Privacy-preserving action recognition.** Early works in this area proposed to learn human actions from low-resolution videos [9,30,29]. Ryoo *et al.* [30,29] proposed to learn image transformations to down-sample frames into low-resolution for action recognition. Wang *et al.* [35] proposed a lens-free coded aperture camera system for action recognition that is privacy-preserving. Note that the above frameworks are limited to providing visual privacy and it is not clear, to what extent they provide protection against adversaries like DNNs that may try to learn or reconstruct the privacy attributes. To solve this issue, Ren *et al.* [27] proposed to learn a video face anonymizer using adversarial training. The training uses a video anonymizer that modifies the original video to remove sensitive information while trying to maximize the action recognition performance and a discriminator that tries to extract sensitive information from the anonymized videos. Later, Wu *et al.* [37,36] proposed and compared multiple adversarial training frameworks for optimizing the parameters of the encoder function. They used a UNet-like encoder from [18] which can be seen as a 2D conv-based frame-level filter. The encoder is trained to allow important spatio-temporal attributes for action recognition, measured by a DNN, and to inhibit the privacy attributes in frames against an ensemble of DNN adversaries whose task is to learn the sensitive information. An important drawback of this framework is that the training requires an ensemble of adversaries in order to provide strong privacy protection. Concurrent to our work, Dave *et al.* [10] proposed a self-supervised framework for training a UNet-based privacy-preserving encoder.

### 3 Proposed Framework

In this section, we present the design of our BDQ encoder which is composed of a series of modules. We also discuss an adversarial training scheme to optimally train the parameters of these modules for the two seemingly contradictory tasks: protecting privacy and enabling action recognition.

#### 3.1 BDQ: Blur Difference Quantization

The BDQ encoder, as the name suggests is composed of three modules: *Blur*, *Difference*, and *Quantization*. Given a scene, the three modules are applied to it in a sequential manner as shown in Figure 2. Here, for each module, we provide

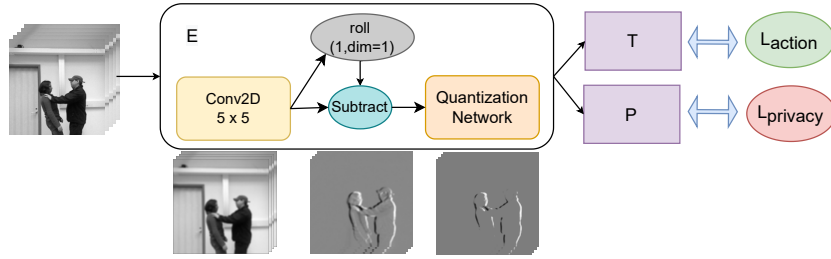


Fig. 2: The BDQ encoder architecture and adversarial training framework for privacy-preserving action recognition.  $roll(1, dim = 1)$  operation shifts frames by one along the temporal dimension.

a detailed description of its implementation and role in preserving privacy and enabling action recognition.

**Blur.** The goal of the *Blur* module is to blur the spatial edges while preserving important spatial features for action recognition. More importantly, its task is to suppress the obvious privacy features that may leak at the spatial edges in motion difference frames that will be produced by the *Difference* module. Given a frame  $v_i$ , we define a video as  $V = \{v_i | i = 1, 2, \dots, t\}$  where  $t$  is the number of frames. The blurred frame  $B_{v_i}$  is convolution of video frame  $v_i$  and a 2D Gaussian kernel  $G_\sigma$  of  $\sigma$  standard deviation and defined as  $B_{v_i} = G_\sigma v_i$ , where  $G_\sigma = \frac{1}{2\pi\sigma^2} \exp(-\frac{x^2+y^2}{2\sigma^2})$ . Window-size of the kernel is kept as  $5 \times 5$  and  $\sigma$  is learned during the adversarial training. A small window-size is chosen since it stabilizes training and avoids losing important spatial features.

**Difference.** Given two consecutive frames from an action video after passing through the *Blur* module, this module performs pixel-wise numerical subtraction between their intensity values and outputs a single frame as  $D(B_{v_i}, B_{v_j}) = B_{v_i} - B_{v_j}$ . It serves two important purposes: First, it highlights the motion features between the two frames by bringing out the direction of motion which improves action recognition accuracy. This is evident from the fact that many current state-of-the-art action recognition methods such as [17,19,20,34] use such temporal-difference modules in the feature space to develop action recognition models. Second, it suppresses obvious high-level spatial privacy cues which helps in preserving privacy. Note that this module does not contain any learnable parameters. Furthermore, when implemented for online applications where frames are continuously being captured, it needs to store a copy of the previous frame for pixel-wise numerical subtraction.

**Quantization.** Although, the *Blur* and the *Difference* modules contribute to suppressing high-level spatial privacy cues, they cannot fully protect against adversaries (see Section 5.1). This is because they still allow low-level spatial privacy cues that are good enough for learning and reconstructing the privacy attributes (see Section 5.4). To alleviate this issue, the task of the *Quantization* module is to remove such information by applying a pixel-wise quantization func-

tion on the motion difference frames that are output by the *Difference module*. Conventionally, a quantization function is defined by  $y = \sum_{n=1}^{N-1} \mathcal{U}(x - b_n)$ , where  $y$  is the discrete output,  $x$  is a continuous input,  $b_n = \{0.5, 1.5, 2.5, \dots, N - 1.5\}$ ,  $N = 2^k$ ,  $k$  is the number of bits, and  $\mathcal{U}$  is the Heaviside function. Unfortunately, such a formulation is not differentiable and therefore not suitable for back-propagation. Thus, following [38,33], we approximate the above quantization function with a differentiable version by replacing the Heaviside function with the sigmoid  $\sigma(\cdot)$  function, which results in the following formulation:  $\sum_{n=1}^{N-1} \sigma(H(x - b_n))$ , where  $H$  is a scalar hardness term. Here, the parameters learned are the  $b_n$  values. In this work, we fix the number of  $b_n$  values to be 15 and initialize them to the values 0.5, 1.5,  $\dots$ , 14.5. Furthermore, the input to the *Quantization* module is normalized to have values between 0 to 15. Note that, here we allow quantization output to be non-integer since no hardware constraint is imposed on BDQ.

### 3.2 Training BDQ Encoder

The BDQ encoder contains several parameters such as the standard deviation of the *Blur* module and the intervals of the *Quantization* module. Our goal is to set these parameters such that: (1) the privacy attributes cannot be learned from the BDQ output by any adversary, (2) recognizing action attributes must be feasible with high precision. Although, we can achieve this goal by setting these parameters heuristically, however, it has been shown that a better performance can be achieved if they are learned in a data-driven fashion [33,37,14].

Our training framework consists of three components: (1) the BDQ encoder denoted as  $E$ , (2) a 3D CNN for predicting target action attribute denoted as  $T$ , (3) and a 2D CNN for predicting privacy attribute denoted as  $P$ . The three components are connected as illustrated in Figure 2 such that the output of  $E$  is simultaneously passed to the networks  $T$  and  $P$ . Following [28], the parameter optimization of  $E$ ,  $T$ , and  $P$  is formulated as a three-player non-zero sum game where the goal of  $E$  is to maximize the likelihood of the target action attributes, measured by  $E$ , while maximizing the uncertainty in the privacy attributes, measured by  $P$ . Such a training procedure consists of two steps that are iterated until the privacy attributes are sufficiently preserved without a significant compromise in action recognition. In the first step,  $P$  is fixed and  $E$  and  $T$  are trained together using the following loss function.

$$\mathcal{L}(V, \theta_E, \theta_T) = \mathcal{X}\mathcal{E}(T(E(V)), L_{action}) - \alpha\mathcal{E}(P(E(V))) \quad (1)$$

Here,  $\mathcal{X}\mathcal{E}$  and  $\mathcal{E}$  refer to the cross-entropy loss and entropy function, respectively.  $\theta_E$  and  $\theta_T$  denote parameters of  $E$  and  $T$ , respectively.  $L_{action}$  is the ground-truth action label and  $\alpha$  is the adversarial weight that allows a trade-off between action and privacy recognition. In the second step,  $E$  and  $T$  are fixed, and  $P$  is trained using the following loss function. Here,  $\theta_P$  denote the parameters of  $P$  and  $L_{privacy}$  are the ground-truth privacy labels.

$$\mathcal{L}(V, \theta_P) = \mathcal{X}\mathcal{E}(P(E(V)), L_{privacy}) \quad (2)$$

## 4 Experiments

### 4.1 Datasets

**SBU.** The SBU Kinect Interaction Dataset [39] is a two-person interaction dataset for video-based action recognition, recorded at 15 fps. It consists of seven actors interacting in pairs in the following eight ways: approaching, departing, pushing, kicking, punching, exchanging objects, hugging, and shaking hands. Originally, the dataset comes divided into 21 sets such that each set corresponds to a pair of actors performing all the eight interactions. Furthermore, the 21 sets are created such that the same two actors may appear in two different sets. In such a case, in the first set, one actor acts and the other reacts and vice versa in the second set. For example, in set s01s02, actor 1 is acting and actor 2 is reacting; similarly, in set s02s01, actor 2 is acting and actor 1 is reacting. Since, such sets contain the same pair of actors, they can be combined into one class. Following this procedure, we reduce the number of 21 original sets to 13 different distinct actor-pair sets. For our setting, given a video, the target task is to classify it into one of the eight interaction/action classes while the privacy label prediction task is to recognize the actor-pair among the 13 actor-pairs. Note the above method is identical to the one followed in [37,36] for developing their privacy-preserving action recognition framework.

**KTH.** The KTH dataset [31] is a video-based action recognition dataset, recorded at 25 fps. It consists of 25 actors, each performing the following six actions: walk, jog, run, box, hand-wave, and hand clap. The different actions are recorded in different settings and variations including outdoor, outdoor with scale variation, outdoor with different clothes, and indoor. In our experiments, we use the six action classes for the action recognition task and the 25 actor identities for the privacy label prediction task.

**IPN.** The IPN hand gesture dataset [4] is a video-based hand gesture dataset, recorded at 30 fps. It consists of 50 actors, each performing the following 13 hand gestures that are common in interacting with a touch-less screen: pointing with one finger, pointing with two fingers, click with one finger, click with two fingers, throw up, throw down, throw left, throw right, open twice, double click with one finger, double click with two fingers, zoom in, and zoom out. In our experiments, we use the 13 hand gesture classes for the action recognition task and the gender (male/female, 2 classes) of the actors for the privacy label prediction task.

### 4.2 Implementation

**Adversarial Training.** Our adversarial training framework consists of three components: (1) the BDQ encoder, (2) an action recognition model which is set to a 3D ResNet-50 network, (3) and a privacy attribute prediction model which is set to a 2D ResNet-50 network. Furthermore, the 3D ResNet-50 and the 2D ResNet-50 networks are initialized with Kinetics-400 and ImageNet pre-trained weights, respectively. For training, we densely sample  $t$  consecutive frames ( $t = 16$  for SBU,  $t = 32$  for KTH and IPN) from the input video to form an input

sequence. For spatial data augmentation, we randomly choose for each input sequence, a spatial position and a scale to perform a multi-scale cropping where the scale is picked from the set  $\{1, \frac{1}{2^{1/4}}, \frac{1}{2^{3/4}}, \frac{1}{2}\}$ . The final output is an input sequence with size  $224 \times 224$ . As shown in Figure 2, the input sequence is then passed through the BDQ encoder whose output is then passed to the 3D ResNet-50 model for action recognition and 2D ResNet-50 model for predicting privacy labels. The optimization of parameters of the three networks is done according to the adversarial training framework discussed in Section 3.2 with  $\alpha = 2, 1,$  and  $8$ , for SBU, KTH, and IPN, respectively. Scaler hardness  $H$  is set to  $5$  for all datasets. The adversarial training is performed for 50 epochs with SGD optimizer,  $lr = 0.001$ , cosine annealing scheduler, and batch size of 16.

**Validation.** We freeze the trained BDQ encoder and newly instantiate a 3D ResNet-50 model for action recognition and a 2D ResNet-50 model for privacy label prediction. Note that the 3D ResNet-50 and the 2D ResNet-50 models are initialized with Kinetics-400 and ImageNet pre-trained weights, respectively. We use the BDQ encoder output on the train set videos to train the 3D ResNet-50 model for action recognition and the 2D ResNet-50 model for predicting privacy labels. Both the networks are trained for 50 epochs with SGD optimizer,  $lr = 0.001$ , cosine annealing scheduler, and batch size 16. For validation, we sample consecutive  $t$  frames ( $t = 16$  for SBU,  $t = 32$  for KTH and IPN) from each input video without any random shift, producing an input sequence. We then center crop (without scaling) each frame in the sequence with a square region of size  $224 \times 224$ . For action recognition, we use the generated sequence on the 3D ResNet model to report the clip-1 crop-1 accuracy. For privacy prediction, we average the softmax outputs by the 2D ResNet-50 model over  $t$  frames and report the average accuracy.

### 4.3 Results

Figure 3 presents our results on the three datasets. We use the visualization proposed in [37] to illustrate the trade-off between the action/gesture recognition accuracy and the privacy label prediction accuracy. We compare our method with two different methods for preserving privacy in action videos: Ryoo *et al.* [30] and Wu *et al.* [36].

In Ryoo *et al.* [30], the degradation encoder is a down-sampling module. Here, a high-resolution action video is down-sampled into multiple videos of a fixed low resolution by applying different image transformations that are optimized for the action recognition task. These transformations include sub-pixel translation, scaling, rotation, and other affine transformations emulating possible camera motion. The low-resolution videos can then be used for training action recognition models. In our experiments, we consider the following low spatial resolutions:  $112 \times 112$ ,  $56 \times 56$ ,  $28 \times 28$ ,  $14 \times 14$ ,  $7 \times 7$ , and  $4 \times 4$ . For each resolution, we generate corresponding low-resolution videos by applying the learned transformations and train a 3D ResNet-50 model for action recognition and a 2D ResNet-50 model for predicting privacy labels. Figure 3 row 1 reports the results of this experiment where a bigger marker depicts a larger



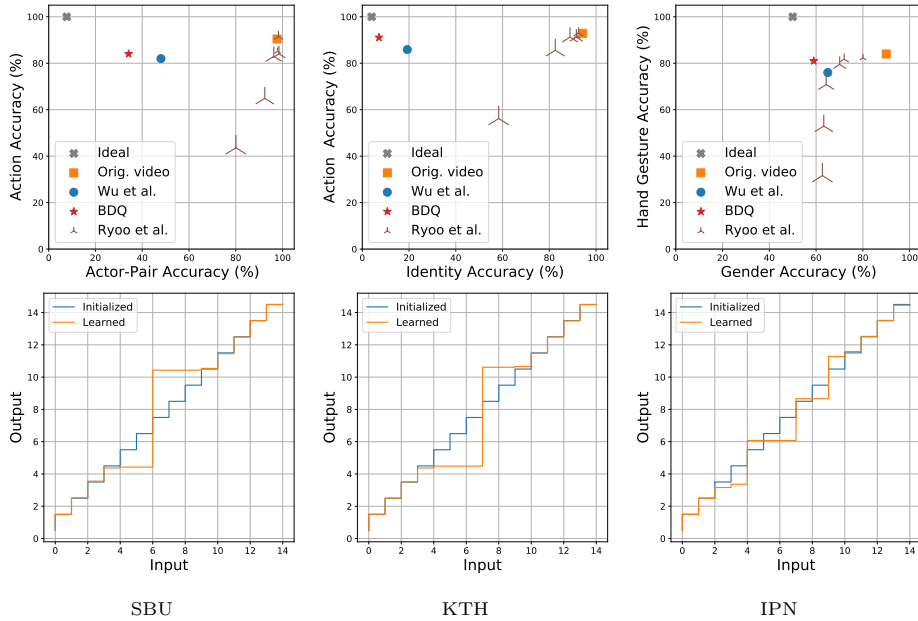


Fig. 3: Performance trade-off (row 1) and learned quantization steps (row 2) on the three datasets: SBU, KTH, and IPN.

down-sampling rate. We observe that, as the down-sampling rate increases, the privacy label prediction accuracy drops for all the datasets. Unfortunately, the action recognition accuracy drops at a faster rate which is contradictory to the desired behavior.

In Wu *et al* [36], the degradation encoder is a UNet-like network that is implemented using the model from [18]. It can be seen as a 2D convolution-based frame-level filter and converts each frame into a feature map of the same shape as the original frame. For training the encoder, we use the same adversarial training method as described in Section 3.2 which is same as the original work. However, the privacy model is reset after certain iterations to improve performance as done in the original work. For a fair comparison, the adversarial training is performed using a 3D ResNet-50 model for action recognition and a 2D ResNet-50 model for predicting privacy labels. Furthermore, the validation method followed is identical to the one mentioned in Section 4.2. Figure 3 row 1 reports our evaluation results of these experiments. We observe that, in comparison to Ryoo *et al.* [30], Wu *et al.* [36] performs signif-

Method	Params.	Size	FLOPs
Wu <i>et al</i>	1.3M	3.8Mb	166.4G
BDQ	16	3.4Kb	120.4M

Table 1: Comparison of space-time complexity of BDQ and Wu *et al.*

icantly better across all the datasets in terms of preserving privacy and enabling action recognition.

Our proposed BDQ encoder surpasses both Ryoo *et al.* [30] and Wu *et al.* [36] by a significant margin across all the datasets in preserving privacy and enabling action recognition. As seen in Figure 3 row 1, it is closer to the ideal trade-off than any other method. Finally, Table 1 compares the space-time complexity of the BDQ and Wu *et al* encoders. We observe that BDQ uses significantly less parameters and computation in comparison to the Wu *et al* encoder.

## 5 Analysis

### 5.1 Ablation study

As described in Section 3, the BDQ encoder consists of three modules: *Blur* (B), *Difference* (D), and *Quantization* (Q). Here, we study the role of each of these modules in preserving privacy and enabling action recognition. For this, we take the pre-trained BDQ encoder that was learned on the SBU dataset in Section 4 and select various combinations of its modules to study their contribution to action recognition and

actor-pair recognition. For each combination, we freeze the parameters of its module(s) and use its output to train a 3D ResNet-50 (pre-trained on Kinetics-400) for action recognition and a 2D ResNet-50 (pre-trained on ImageNet) for actor-pair recognition. Note that, for both the networks, the training, and the validation procedures are identical to the one used in Section 4. Figure 4 provides results of this study on the SBU dataset. We observe that the combinations ‘B’, ‘D’, ‘Q’, ‘B+D’, and ‘B+Q’, have very little effect in preserving privacy information and achieve results close to the case when original video is used. Interestingly, the combination ‘D’ achieves a higher action recognition accuracy among all the combinations, signifying its ability to produce better temporal features. Furthermore, a good drop in privacy accuracy is observed when ‘D’ and ‘Q’ are used together. Moreover, this accuracy further drops drastically when all ‘B’, ‘D’, and ‘Q’ are used together. Finally, we also study the effect of the adversarial parameter  $\alpha$  on these modules as shown in Figure 4. We observe that when no adversarial training is performed, i.e  $\alpha = 0$ , there is a very little drop

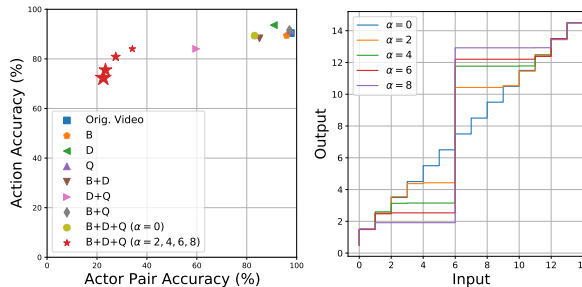


Fig. 4: Left- Results of the ablation study. Here, a bigger  $\star$  corresponds to a higher value of  $\alpha$ . Right- Effect of the adversarial parameter  $\alpha$  on the quantization steps.

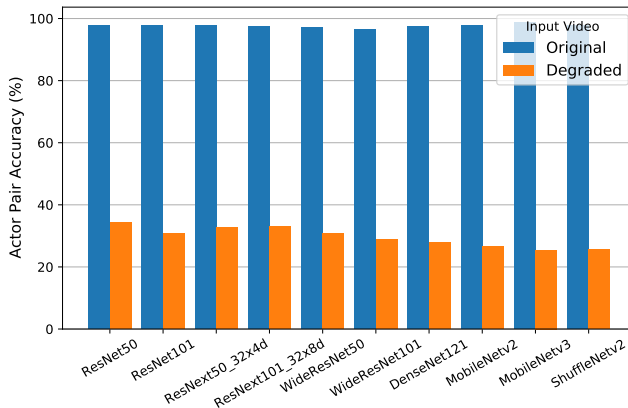


Fig. 5: Actor-pair accuracy on various image classification networks.

in privacy accuracy. However, with the increase in the value of  $\alpha$ , both action recognition and actor-pair accuracy begin to fall, with the action recognition accuracy falling more sharply. Figure 4 presents the learned quantization values corresponding to each value of  $\alpha$ . We observe that with the increase in  $\alpha$  the amount of quantization increases which leads to drop in action and actor-pair accuracies. Please refer to supplementary for more studies.

## 5.2 Strong Privacy Protection

A significant challenge for any privacy-preserving model is to provide protection against any possible, seen and unseen adversary, that may try to learn the privacy information. In order to show that our proposed framework provides such strong privacy protection, we prepare a list of ten state-of-the-art image classification networks (adversaries) as shown in Figure 5. We take the pre-trained BDQ encoder from Section 4 and use its output (degraded video) to train the above networks for predicting the actor-pair labels on the SBU dataset. Furthermore, we also train these networks on the original videos to prepare corresponding baselines for comparison. Note that all the networks are initialized with ImageNet pre-trained weights, and the training and inference procedures are identical to the one used for the adversary in Section 4. From Figure 5, we observe that the BDQ encoder consistently protects privacy information against all the networks with ResNet-50 [13] performing the best at 34.18% and MobileNet-v3 [15] performing the worst at 25.46%. Note that, among all the adversaries, the BDQ encoder had only seen ResNet-50 during its training.

## 5.3 Generalized Spatio-temporal Features

In addition to strong privacy protection, a privacy-preserving model must allow task-specific features to be learned by any network that is designed for

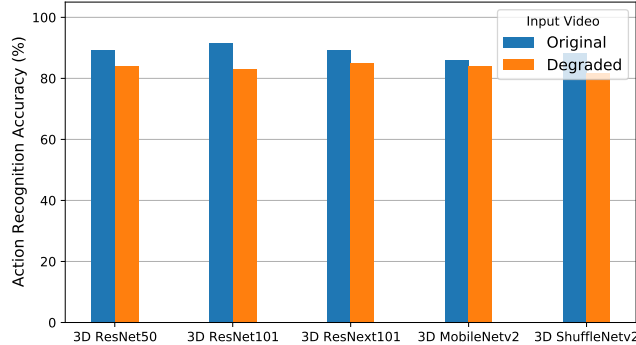


Fig. 6: Action recognition accuracy on various action recognition networks.

that task. In order to show that our proposed framework allows useful spatio-temporal cues for action recognition, we take the pre-trained BDQ encoder from Section 4 and use its output to separately train five 3D CNNs for predicting action classes on the SBU dataset, as shown in Figure 6. Furthermore, similar to Section 5.2, we also train these networks on the original videos to prepare corresponding baselines for comparison. Note that all the networks are initialized with Kinetics-400 pre-trained weights. From Figure 6, we observe that the BDQ encoder consistently allows spatio-temporal information to be learned with 3D ResNext-101 performing the best at 85.1% and 3D ShuffleNet-v2 performing the worst at 81.91%. Furthermore, all the networks achieve action recognition accuracy marginally lower than their corresponding baselines.

#### 5.4 Robustness to Reconstruction Attack

In this section, we explore a scenario where an attacker has access to the BDQ encoder such that he/she can produce a large training set containing degraded videos along with their corresponding original videos. In such a case, the attacker can train an encoder-decoder network and try to reverse the effect of BDQ, recovering the privacy information. In order to show that our proposed framework is resistant to such an attack, we train a 3D UNet [8] model for 200 epochs on the SBU dataset with input as degraded videos from the pre-trained BDQ encoder of Section 4 and output original videos as ground-truth labels. Additionally, we also train the 3D UNet model with input as degraded videos from an untrained BDQ encoder. Figure 7 visualizes some examples of reconstruction when videos from untrained (column 4) and trained (column 5) BDQ encoder are used for training. We observe that the reconstruction network can successfully reconstruct the original video with satisfactory accuracy when the input is from an untrained BDQ encoder. However, when the input is from a trained BDQ encoder, reconstruction is significantly poor and privacy information is still preserved.

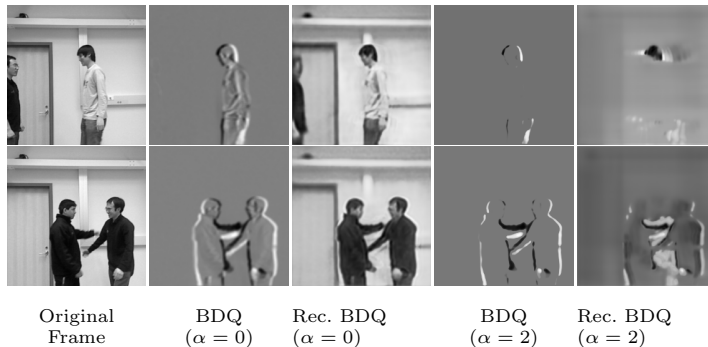


Fig. 7: Visualization of reconstruction results for  $\alpha = 0$  and 2.

### 5.5 Subjective Evaluation

In Section 5.2 and 5.4, we show that our proposed framework is robust against adversaries that may try to learn or reconstruct the privacy attributes. However, such protection is of no use unless it provides visual privacy against the human visual system. In order to show that the BDQ encoder produces encodings that provide visual privacy, we conduct a user study on the videos from the SBU dataset. The user study is composed of 60 questions where each question consists of a video sampled from the SBU validation set and applied with the BDQ encoder learned in Section 4. Furthermore, each question has seven options showing cropped faces of actors from the SBU dataset. Given a BDQ output video where two persons are interacting, the task of the user is to select the identities of both the actors from the seven options. A total of 26 participants took part in the study. Note that the random chance of selecting two actors and both of them being correct is 4.76%. Similarly, the random chance of selecting two actors and at least one of them being correct is 52.38%. In the first case, the results of our user study reveal that the participants were able to recognize both the actors correctly with an accuracy of 8.65%. Similarly, the users' accuracy for the second case was 65.64% (more details in supplementary).

### 5.6 Comparison with Event Camera

In recent years, Dynamic Vision Sensor (DVS) based cameras are being proposed as an in-home monitoring solution for privacy-preserving action detection and recognition [1]. Unlike traditional cameras that capture high-resolution videos and images, a DVS sensor detects the temporal changes in the pixel intensity at a pixel location. If the pixel intensity rises beyond a fixed threshold at a pixel, it is registered as a positive event. However, if it drops below a fixed threshold, it is registered as a negative event. Here, positive and negative events describe the direction of motion which can be encoded as outlines. At an abstract level, our BDQ encoder can be seen as a digital approximation of the DVS sensor that can be implemented in any traditional camera. In order to compare our framework

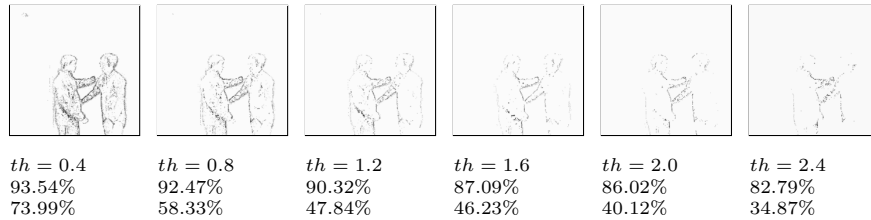


Fig. 8: Example event frames (Row 1), event threshold (Row 2), action recognition accuracy (Row 3) and actor-pair recognition accuracy (Row 4) on SBU.

with the DVS sensor, we first convert the videos from the SBU dataset into events using [11] which proposes a method for converting any video into synthetic events such that they can be simulated as events from a DVS sensor. Such a method is useful since training DNNs require huge amount of data and collecting such amount is not feasible using event cameras. Furthermore, the above method enables us to vary the pixel level threshold which decides if the intensity change (event) will be registered or not. Note that a high threshold leads to less events while a low threshold leads to more events being registered. Figure 8 displays the effect of threshold on an event frame. The events from the converted SBU dataset are first converted into event frames using [11] which are then used for training a 3D ResNet-50 model for action recognition and a 2D ResNet-50 model for actor-pair recognition. The initialization and training settings is identical to Section 4. Figure 8 reports the trade-off for each threshold. We observe that with the increase in threshold, both action recognition and actor-pair accuracy drops. Furthermore, the trade-off at threshold value 2.4 is close to the trade-off achieved by the BDQ encoder with  $\alpha = 2$  (refer Section 4.3).

## 6 Conclusions

This paper proposes a novel encoder called BDQ for the task of privacy-preserving human action recognition. The BDQ encoder is composed of three modules: *Blur*, *Difference*, and *Quantization* whose parameters are learned in an end-to-end fashion via an adversarial training framework such that it learns to allow important spatio-temporal attributes for action recognition and inhibit spatial privacy attributes. We show that the proposed encoder achieves state-of-the-art trade-off on three benchmark datasets in comparison to previous works. Furthermore, the trade-off achieved is at par with the DVS sensor-based event cameras. Finally, we also provide an extensive analysis of the BDQ encoder including an ablation study on its components, robustness to various adversaries, and a subjective evaluation.

**Limitation:** Due to its design, our proposed framework for privacy preservation cannot work in cases when the subject or the camera does not move.

**Acknowledgement:** This work was supported by JSPS KAKENHI Grant Number JP20K20628.

## References

1. <https://www.samsung.com/au/smart-home/smartthings-vision-u999/GP-U999GTEEAAC/>
2. Agrawal, P., Narayanan, P.: Person de-identification in videos. *IEEE Transactions on Circuits and Systems for Video Technology* **21**(3), 299–310 (2011)
3. Asif, U., Mashford, B., Von Cavallar, S., Yohanandan, S., Roy, S., Tang, J., Harrer, S.: Privacy preserving human fall detection using video data. In: *Machine Learning for Health Workshop*. pp. 39–51. PMLR (2020)
4. Benitez-Garcia, G., Olivares-Mercado, J., Sanchez-Perez, G., Yanai, K.: Ipn hand: A video dataset and benchmark for real-time continuous hand gesture recognition. In: *2020 25th International Conference on Pattern Recognition (ICPR)*. pp. 4340–4347. IEEE (2021)
5. Brkic, K., Sikiric, I., Hrkac, T., Kalafatic, Z.: I know that person: Generative full body and face de-identification of people in images. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. pp. 1319–1328. IEEE (2017)
6. Canh, T.N., Nagahara, H.: Deep compressive sensing for visual privacy protection in flatcam imaging. In: *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*. pp. 3978–3986. IEEE (2019)
7. Chen, D., Chang, Y., Yan, R., Yang, J.: Tools for protecting the privacy of specific individuals in video. *EURASIP Journal on Advances in Signal Processing* **2007**, 1–9 (2007)
8. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: *International conference on medical image computing and computer-assisted intervention*. pp. 424–432. Springer (2016)
9. Dai, J., Wu, J., Saghafi, B., Konrad, J., Ishwar, P.: Towards privacy-preserving activity recognition using extremely low temporal and spatial resolution cameras. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. pp. 68–76 (2015)
10. Dave, I.R., Chen, C., Shah, M.: Spact: Self-supervised privacy preservation for action recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 20164–20173 (2022)
11. Gehrig, D., Gehrig, M., Hidalgo-Carrió, J., Scaramuzza, D.: Video to events: Recycling video datasets for event cameras. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3586–3595 (2020)
12. Gochoo, M., Tan, T.H., Alnajjar, F., Hsieh, J.W., Chen, P.Y.: Lownet: Privacy preserved ultra-low resolution posture image classification. In: *2020 IEEE International Conference on Image Processing (ICIP)*. pp. 663–667. IEEE (2020)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
14. Hinojosa, C., Niebles, J.C., Arguello, H.: Learning privacy-preserving optics for human pose estimation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 2573–2582 (2021)
15. Howard, A., Sandler, M., Chu, G., Chen, L.C., Chen, B., Tan, M., Wang, W., Zhu, Y., Pang, R., Vasudevan, V., et al.: Searching for mobilenetv3. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 1314–1324 (2019)

16. Huang, C., Kairouz, P., Sankar, L.: Generative adversarial privacy: A data-driven approach to information-theoretic privacy. In: 2018 52nd Asilomar Conference on Signals, Systems, and Computers. pp. 2162–2166. IEEE (2018)
17. Jiang, B., Wang, M., Gan, W., Wu, W., Yan, J.: Stm: Spatiotemporal and motion encoding for action recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2000–2009 (2019)
18. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: European conference on computer vision. pp. 694–711. Springer (2016)
19. Li, Y., Ji, B., Shi, X., Zhang, J., Kang, B., Wang, L.: Tea: Temporal excitation and aggregation for action recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 909–918 (2020)
20. Liu, Z., Luo, D., Wang, Y., Wang, L., Tai, Y., Wang, C., Li, J., Huang, F., Lu, T.: Teinet: Towards an efficient architecture for video recognition. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 11669–11676 (2020)
21. Mirjalili, V., Raschka, S., Ross, A.: Flowsan: Privacy-enhancing semi-adversarial networks to confound arbitrary face-based gender classifiers. *IEEE Access* **7**, 99735–99745 (2019)
22. Orekondy, T., Schiele, B., Fritz, M.: Towards a visual privacy advisor: Understanding and predicting privacy risks in images. In: Proceedings of the IEEE international conference on computer vision. pp. 3686–3695 (2017)
23. Padilla-López, J.R., Chaaoui, A.A., Flórez-Revuelta, F.: Visual privacy protection methods: A survey. *Expert Systems with Applications* **42**(9), 4177–4195 (2015)
24. Pittaluga, F., Koppal, S., Chakrabarti, A.: Learning privacy preserving encodings through adversarial training. In: 2019 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 791–799. IEEE (2019)
25. Pittaluga, F., Koppal, S.J.: Pre-capture privacy for small vision sensors. *IEEE transactions on pattern analysis and machine intelligence* **39**(11), 2215–2226 (2016)
26. Raval, N., Machanavajjhala, A., Cox, L.P.: Protecting visual secrets using adversarial nets. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 1329–1332. IEEE (2017)
27. Ren, Z., Lee, Y.J., Ryoo, M.S.: Learning to anonymize faces for privacy preserving action detection. In: Proceedings of the european conference on computer vision (ECCV). pp. 620–636 (2018)
28. Roy, P.C., Boddeti, V.N.: Mitigating information leakage in image representations: A maximum entropy approach. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2586–2594 (2019)
29. Ryoo, M., Kim, K., Yang, H.: Extreme low resolution activity recognition with multi-siamese embedding learning. In: Proceedings of the AAAI conference on artificial intelligence. vol. 32 (2018)
30. Ryoo, M.S., Rothrock, B., Fleming, C., Yang, H.J.: Privacy-preserving human activity recognition from extreme low resolution. In: Thirty-First AAAI Conference on Artificial Intelligence (2017)
31. Schuldt, C., Laptev, I., Caputo, B.: Recognizing human actions: a local svm approach. In: Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004. vol. 3, pp. 32–36. IEEE (2004)
32. Srivastav, V., Gangi, A., Padoy, N.: Human pose estimation on privacy-preserving low-resolution depth images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 583–591. Springer (2019)



33. Tan, J., Khan, S.S., Boominathan, V., Byrne, J., Baraniuk, R., Mitra, K., Veer-  
araghavan, A.: Canopic: pre-digital privacy-enhancing encodings for computer vi-  
sion. In: 2020 IEEE International Conference on Multimedia and Expo (ICME).  
pp. 1–6. IEEE (2020)
34. Wang, L., Tong, Z., Ji, B., Wu, G.: Tdn: Temporal difference networks for efficient  
action recognition. In: Proceedings of the IEEE/CVF Conference on Computer  
Vision and Pattern Recognition. pp. 1895–1904 (2021)
35. Wang, Z.W., Vineet, V., Pittaluga, F., Sinha, S.N., Cossairt, O., Bing Kang, S.:  
Privacy-preserving action recognition using coded aperture videos. In: Proceed-  
ings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition  
Workshops. pp. 0–0 (2019)
36. Wu, Z., Wang, H., Wang, Z., Jin, H., Wang, Z.: Privacy-preserving deep action  
recognition: An adversarial learning framework and a new dataset. *IEEE Transac-  
tions on Pattern Analysis and Machine Intelligence* (2020)
37. Wu, Z., Wang, Z., Wang, Z., Jin, H.: Towards privacy-preserving visual recognition  
via adversarial training: A pilot study. In: Proceedings of the European Conference  
on Computer Vision (ECCV). pp. 606–624 (2018)
38. Yang, J., Shen, X., Xing, J., Tian, X., Li, H., Deng, B., Huang, J., Hua, X.s.:  
Quantization networks. In: Proceedings of the IEEE/CVF Conference on Com-  
puter Vision and Pattern Recognition. pp. 7308–7316 (2019)
39. Yun, K., Honorio, J., Chattopadhyay, D., Berg, T.L., Samaras, D.: Two-person  
interaction detection using body-pose features and multiple instance learning. In:  
2012 IEEE Computer Society Conference on Computer Vision and Pattern Recog-  
nition Workshops. pp. 28–35. IEEE (2012)