

Sparse and Privacy-enhanced Representation for Human Pose Estimation

Ting-Ying Lin^{1*}

tingyinglin@m111.nthu.edu.tw

Lin-Yung Hsieh^{1*}

linyunghsieh@gapp.nthu.edu.tw

Fu-En Wang¹

fulton84717@gapp.nthu.edu.tw

Wen-Shen Wuen²

vincent_wuen@novatek.com.tw

Min Sun¹

sunmin@ee.nthu.edu.tw

¹ Vision Science Lab

National Tsing Hua University

Hsinchu, Taiwan

² Novatek Microelectronics Corp.

Hsinchu, Taiwan

Abstract

We propose a sparse and privacy-enhanced representation for Human Pose Estimation (HPE). Given a perspective camera, we use a proprietary motion vector sensor (MVS) to extract an edge image and a two-directional motion vector image at each time frame. Both edge and motion vector images are sparse and contain much less information (*i.e.*, enhancing human privacy). We advocate that edge information is essential for HPE, and motion vectors complement edge information during fast movements. We propose a fusion network leveraging recent advances in sparse convolution used typically for 3D voxels to efficiently process our proposed sparse representation, which achieves about 13x speed-up and 96% reduction in FLOPs. We collect an in-house edge and motion vector dataset with 16 types of actions by 40 users using the proprietary MVS. Our method outperforms individual modalities using only edge or motion vector images. We also demonstrate the generalizability of our approach on MMH-PSD and HumanEva datasets. Finally, we validate the privacy-enhanced quality of our sparse representation through face recognition on CelebA (a large face dataset) and a user study on our in-house dataset. The code and dataset are available on the project page: <https://lyhsieh.github.io/sph/>.

1 Introduction

With the advance of deep learning [1] for visual perception and the lower cost of connected camera systems, human society is at the beginning of the smart camera systems era [2] facilitating our daily life. However, there are two key challenges to be addressed. Firstly, the dilemma between convenience and privacy. We want the systems (*e.g.*, in our home) to recognize our behavior and assist us, but we also need to ensure they protect our privacy. Secondly, applying deep models on the cloud's visual data is costly and introduces delays.




Figure 1: Examples from SPHP. Our motion vector sensor (MVS) can extract an edge image and a two-directional motion vector image (visualized as horizontal MV_x and vertical MV_y , respectively) at each time frame. We also provide annotated body joints for Human Poses (HP) and corresponding grayscale images.

We aim to run these models efficiently on the edge device, which supports real-time response and protects our privacy. Therefore, a sparse and privacy-enhanced representation for human pose estimation is desired.

An event-based camera is a sparse representation approach that only records the intensity variances between active pixels. Hence, privacy can be enhanced, and recorded identities cannot be directly recognized compared to traditional cameras. Human pose estimation from event data has been studied in recent years [4, 5]. However, networks trained with event data usually have a significantly lower performance in HPE due to no event data recorded on still body parts. In order to achieve better HPE performance by sparse representation, a new approach is necessary for privacy-enhanced smart cameras on edge devices.

Inspired by the seminal edge (*i.e.*, boundary) [28] and motion [4] literature before deep learning, we advocate that edge information is essential for HPE, and motion vectors complement edge information during fast movements. Besides, both edge and motion vectors are sparse representations that can be efficiently processed and enhance privacy. Hence, we propose to use edge to capture the contour of all body parts and replace event data with two-directional motion vectors. Specifically, we use a proprietary motion vector sensor (MVS) to extract an edge image and a two-directional motion vector image at each time frame. Thus, our method differs from pure software solutions that may already contain privacy-sensitive data. We propose a fusion network leveraging these complementary sparse representations (*i.e.*, edge and motion vectors) to perform better than each representation. Moreover, we exploit the sparsity of our data to replace standard convolution with submanifold sparse convolutions [3] typically used for 3D voxels to speed up the inference time by 13 times.

Unlike RGB images, our proposed edge and motion vector representation lacks labeled data. Therefore, we collect an in-house dataset called Sparse and Privacy-enhanced Dataset for Human Pose Estimation (**SPHP**), as shown in Figure 1. This dataset includes recordings of 40 individuals performing 16 exercise and fitness-related actions, such as stretching or jogging, captured from 4 distinct viewing angles. We also collect grayscale images synchronized with edge and motion vectors to speed up joint position annotation. First, we apply a pre-trained natural-image-based HPE network on the grayscale images to obtain high-quality "initial" joint labels. Next, we ask human annotators to check and correct the errors if needed. This approach is estimated to save about 2,000 hours of manual joint labeling effort.

To demonstrate the applicability of our proposed framework, we begin by conducting extensive experiments on our **SPHP** dataset. Our fusion method surpasses the performance of individual modalities (*i.e.*, edge or motion vectors) for human pose estimation, especially for fast-moving joints, with a maximum relative improvement of 11% over using only the edge modality. We also compare traditional and submanifold sparse convolutions to demonstrate the advantage of our sparse data. The results show that sparse convolution networks can reduce the number of FLOPs of HPE by 96% while maintaining an acceptable error rate.

In addition, we achieve a 13x acceleration in inference time, which is particularly beneficial for real-time applications. Besides, we confirm the generalizability of our approach on additional HPE datasets, such as MMHPSD[65] and HumanEva[40]. The results on MMHPSD highlight that motion vectors capture more valuable information compared to traditional event cameras. To verify the privacy-enhanced attribute of our data, we use ArcFace [41] to evaluate the cosine similarity between various modalities of faces on the large-scale face attributes dataset, CelebA [26]. Using edge as input results in a 10.2% recall drop compared to RGB images, indicating that privacy is enhanced by converting RGB images to edge images. Moreover, we conduct a user study to check if humans can identify a specific individual from 10 leaked edge images, given a grayscale reference face. The significantly low accuracy of 18.8% demonstrates the limited ability of humans to recognize faces in edge modality.

Our main contributions are summarized as follows:

1. We introduce the Sparse and Privacy-enhanced Dataset for Human Pose Estimation (**SPHP**), which consists of synchronized, complementary images of edge and motion vectors along with ground truth labels for 13 joints.
2. Our fusion method outperforms individual modalities (*i.e.*, edge or motion vectors) for human pose estimation, particularly for fast-moving joints, with a maximum relative improvement of 11% on our dataset when compared to using only the edge modality. Additionally, we further demonstrate the generalizability on other HPE datasets.
3. The high sparsity of **SPHP** enables us to achieve a maximum 13x acceleration in inference time and decrease FLOPs by 96% after applying sparse convolution, significantly reducing computational costs compared to traditional convolutional neural networks.
4. We demonstrate the privacy-enhanced nature of our sparse representation through face recognition experiments. We utilize edge images as input and validate a recall drop of over 10.2% in comparison to RGB images on the CelebA dataset. Furthermore, our user study shows that humans have limited ability to recognize faces in edge modality.

2 Related work

2.1 Human Pose Estimation

Human Pose Estimation (HPE) is a rapidly developing field of research in computer vision, intending to predict 2D/3D positions of joints from various types of input, including RGB, grayscale, or even depth image [27, 39, 40]. HPE has numerous potential applications, such as action recognition, health monitoring [7], and physical education. Currently, CNN-based methods [13, 42, 47, 51, 54] achieve the state-of-the-art performance. Some focus on single-person pose estimation [29, 45] while others can perform HPE on multi-person [8, 14, 53]. These methods are typically trained and evaluated on RGB-based image datasets such as MS-COCO [25], MPII [0], and CrowdPose [24].

Top-down and bottom-up approaches are two common strategies in HPE. A top-down approach is a two-stage method that performs human detection at each input image to find the region of interest before predicting joint positions. The accuracy of a top-down approach may rely on the quality of the human detection process. Representative networks for top-down paradigms include HRNet [43], Mask-RCNN [08], CFN [20], CPN [00], G-RMI [50] and SimpleBaseline [50]. A bottom-up approach detects each joint position and

assembles them into groups. Representative works are DeepCut [63], DeeperCut [22] and OpenPose [8]. Bottom-up approaches have shown greater efficiency in HPE for multiple humans compared to top-down approaches and can effectively handle complex poses involving self-occlusion. In this study, we adopt a bottom-up approach for our backbone networks.

2.2 Sparse Convolution

Convolution Neural Networks (CNNs) have been proven effective for many visual recognition tasks, including human pose estimation. However, CNNs’ high computational cost makes it challenging to use on resource-limited embedded systems. The cost significantly increases when higher dimensional convolutions, such as 3D CNN on 3D point clouds, are needed. Despite the challenge, several sparse convolution techniques have been proposed to leverage the sparsity property within the data. For instance, [12, 15] propose sparse convolutional layers for processing 3D data efficiently. [9] proposes submanifold sparse convolution, which leads to computational saving but remains state-of-the-art performance. [46] designs a sparse data structure and realizes fast computation for 3D shape analysis. [36, 42] use sparse convolutional layers to maintain the efficiency for processing point clouds. Other than 3D data, handwritten characters are a 2D sparse data example demonstrating sparse convolution’s effectiveness for character recognition [16]. Nevertheless, fewer studies [32] have been conducted to leverage sparse convolution for human pose estimation.

2.3 Privacy Enhancing

Using RGB images to recognize human behavior has many useful applications but raises privacy concerns. Hence, enhancing privacy using software or hardware-level techniques becomes an essential research direction. At the software level, early methods [50] involved algorithms such as encryption, image filtering, and object/people removal. Besides, [38] introduces inverse super-resolution, generating multiple low-resolution training videos (*e.g.*, 16×12) from high-resolution videos for human activity recognition. [35, 49] adopt adversarial learning to remove privacy-sensitive information from images while balancing the trade-off between privacy and task. At the hardware level, privacy-enhanced optics are designed to filter out private information while still retaining functionality such as human pose estimation and action recognition. [52] applies privacy-enhanced optics to block sensitive information from the incident light field before sensor measurements. [2, 19, 20] use optimal encoders to protect privacy, and use convolutional neural networks to extract features of specific tasks. In this work, we use a proprietary motion vector sensor (MVS) to extract an edge image and a two-directional motion vector image at the hardware level. We show that MVS significantly enhances privacy through our designed face recognition experiments.

3 Approach

We describe the detailed process to obtain our sparse representation (Section 3.1). Then, we introduce a fusion model (Section 3.2), which benefits from both edge and motion vectors. To enhance the efficiency of our fusion model, we utilize sparse convolution (Section 3.3).

3.1 Sparse Representation

We use a proprietary Motion Vector Sensor (MVS) to extract an edge image and a two-directional motion vector image at each time frame. Our MVS provides sparse and privacy-

enhanced representation, as detailed below, compared to traditional RGB/grayscale cameras.

Edge Image. MVS uses an efficient hardware implementation of edge detection, similar to Canny edge detection [4], to generate edge images. Each pixel in the edge image has a value within the range of $\{0, 255\}$. A higher value indicates a stronger intensity of the edge.

Motion Vector. Inspired by the motion detection of the *Drosophila* visual system [5] and designed with patent-pending pixel-sensing technology, MVS detects vertical and horizontal motion vectors, denoted as MV_X and MV_Y shown in Figure 1, by analyzing changes in illumination. Each value falls within the range of $\{-128, 128\}$. The magnitude and sign of a value represent the strength and direction of motion, respectively.

Since MVS only produces non-zero values when detecting edge or motion vectors, the resulting data can be very sparse. This greatly reduces resource usage during computation and storage. Additionally, the output data from MVS only focuses on changes in illumination over time and space, reducing the risk of privacy exposure (*i.e.*, privacy-enhanced).

3.2 Fusion Model

Edge and motion vector information complement each other. While edge is sufficient for detecting clear and non-blurred body joints, incorporating motion vectors into our model can effectively address the challenges posed by fast movements and overlapping boundaries, which may confuse edge-based HPE models. Hence, we aim to combine the complementary information of edge and motion vectors while keeping our model compact and efficient. We directly concatenate an edge image and a two-directional motion vector image, proposing the early fusion model (referred to as FS) as illustrated in Figure 2. Our FS model can leverage various single-branch network architectures designed for compactness and efficiency.





Figure 2: Fusion model. We employ early fusion on different data modalities, including edge and two-directional motion vectors (MV_X and MV_Y). The resulting output channel N corresponds to the number of joints.

Figure 3: Environment Setting. In an area of $6 \times 5 \text{m}^2$, two MVSs are located 3 meters away from the subject, with a viewing angle of 45 degrees centered on the subject's position.

3.3 Sparse Convolution

Sparse convolution methods, studied mainly on 3D point clouds [6] and 2D hand-written characters data [7], have shown on-par performance with dense methods while requiring substantially less computation. Our proposed edge and motion vector representation is 96.23% and 99.13% sparse on average, respectively. Hence, we leverage the submanifold sparse convolution network introduced in [8] to train our compact fusion model. This approach offers a

powerful tool for processing sparse data, allowing us to effectively exploit the sparsity of our representation while capturing the necessary information for accurate predictions. We also achieve on-par performance with 96% fewer FLOPs compared to traditional convolution.

4 Our SPHP Dataset

We use our proprietary MVS to collect an in-house edge and motion vector dataset with 16 types of actions performed by 40 users, as shown in Figure 1.

Data Format. Motion Vector Sensor (MVS) captures three channels of 8-bit information, conveying edge, horizontal, and vertical motion vectors at each time frame. MVS has a resolution of 640×480 and a field-of-view of 60 degrees. When collecting data, MVS records videos for 10 seconds per clip, with 30 frames per second.

Environment Setting. We collect the data in a room with an area of $6 \times 5 \text{m}^2$. Two MVSs are located 3 meters away from the subject, with a viewing angle of 45 degrees centered on the subject’s position as shown in Figure 3. This setup allows us to collect data from two different viewing angles simultaneously, which increases the data collection efficiency.

4.1 Participants and Actions

We collected data from 40 subjects, having a balanced distribution of 20 male and 20 female participants. The average ages of male and female participants are 32 years and 22 years, respectively. Besides, the average heights of male and female participants are 172 and 161 centimeters, respectively. Participants performed 16 fitness-related actions, which are listed in Table 1 and categorized into four classes based on the type of movement: C1 for upper-body movements, C2 for lower-body movements, C3 for slow whole-body movements, and C4 for fast whole-body movements.

To diversify the viewing angles of our dataset, we apply a novel strategy to capture each action from multiple perspectives. Firstly, we place two cameras within an interval of 45 degrees. Then, we instruct the participants to face various directions (*i.e.*, 0, 15, 30, and 45 degrees, respectively) while capturing their actions. In each direction, every participant will perform four actions, totaling 16 actions, as listed in Table 1.

Table 1: The 16 actions in our SPHP dataset are categorized into four classes: C1 for upper-body movements, C2 for lower-body movements, C3 for slow whole-body movements, and C4 for fast whole-body movements.

C1	C2	C3	C4
1. Arm abduction	5. Leg knee lift	8. Squat	12. Elbow-to-knee
2. Arm bicep curl	6. Leg abduction	9. Walk in place	13. Jump in place
3. Wave hello	7. Leg pulling	10. Standing side bend	14. Jumping jack
4. Punch up forward		11. Roll wrists & ankles	15. Hop on one foot
			16. Jog in place

4.2 Efficient Annotation

Our dataset is annotated with the positions of 13 joints: nose, left/right shoulder, left/right elbow, left/right hand, left/right hip, left/right knee and left/right foot. We use a well-trained grayscale HPE model to speed up the annotation process to mark the initial joint positions.

Dataset	Sub#	Seq#	Frame#	Dir.	MM
DHP19 [9]	17	33	87K		
MMHPSD [65]	15	12	240K		✓
SPHP (ours)	40	16	384K	✓	✓

Table 2: A comparison of our SPHP dataset with event-based datasets in terms of the number of subjects (Sub#), action sequences (Seq#) per subject, frames (Frame#), directional motion (Dir), and multi-modality (MM).

Then, we manually confirm the initial position as acceptable or make fine-grained adjustments on each frame. All videos in our dataset are fully annotated.

4.3 Dataset Comparison

We compare our **SPHP** dataset and two event-based datasets, namely DHP19 [9] and MMHPSD [65], in Table 2. Our dataset outperforms these datasets in various aspects. Firstly, our multi-modality data includes two-directional motion vectors, providing more detailed and nuanced information about actions. Secondly, SPHP boasts the highest number of subjects and frames, and exhibits greater diversity in different unique action types compared to DHP19, which has a large proportion of symmetric actions. Additionally, unlike DHP19 and MMHPSD, SPHP captures 10-second videos for each action rather than recording a fixed number of repetitions, thus providing more comprehensive information.

5 Experimental results

5.1 Human Pose Estimation

To showcase the superiority of our sparse data in HPE, we conduct the experiments on three datasets: **SPHP**, MMHPSD [65] and HumanEva [40]. Within our SPHP dataset, we compare the performance across various input modalities, including edge, motion vectors, a fusion of edge and motion vectors, and grayscale images. Besides traditional convolution, we also employ submanifold sparse convolution to assess the computational efficiency gained from exploiting the sparsity in input data. Additionally, we assess the generalization capability of our method on MMHPSD and HumanEva. Notably, within the MMHPSD dataset, we incorporate the provided event data for further comparison with our motion vectors.

Implementation. We test on 3 CNN backbones, including the DHP19 [9] proposed model (218K), U-Net-Small (1.9M), and U-Net-Large (7.7M). U-Net-Small and U-Net-Large are built based on the architecture proposed by [57], incorporating three downsampling and upsampling operations. The output dimensions of the 3×3 convolutions in the U-Net-Large are twice as large as those of the U-Net-Small. The input frames are resized to 288×384 for the SPHP dataset, whereas the MMHPSD dataset retains its original size of 256×256 . For each joint, the model outputs a heatmap that indicates the likelihood of the joint position at each pixel. To generate a target heatmap for a joint, we initialize an all-zero map of the same size as the input frame and set a value of 1 to the pixel corresponding to the annotated joint position. The heatmap is then smoothed using Gaussian blur with $\sigma = 4$. Mean Squared Error (MSE) is employed as the loss function.

Evaluation Metrics. Mean Per Joint Position Error (**MPJPE**) is chosen for evaluation. $MPJPE = \frac{1}{N} \sum_i \|y_i - \hat{y}_i\|$ calculates the Euclidean distance between predicted positions \hat{y}_i and ground truth positions y_i for each joint, where N is the number of joints.

Table 3: MPJPE (lower is better) on SPHP with different input types. Different input modalities, "GR", "ED", "MV", and "FS", are grayscale, edge, motion vector, and fusion (edge and motion vector), respectively. C1, C2, C3, C4 are four classes listed in Table 1. C and SC are traditional and sparse convolutions, respectively.

Backbone	# of Params	Input	C					SC				
			C1	C2	C3	C4	Mean	C1	C2	C3	C4	Mean
DHP19 [B]	218K	GR	2.62	3.08	3.33	3.62	3.20	-	-	-	-	-
		MV	16.96	6.50	6.43	5.11	8.66	56.96	27.12	25.86	14.12	30.20
		ED	3.14	3.64	3.71	4.03	3.65	5.18	6.47	5.94	6.76	6.10
		FS	3.36	3.32	3.56	3.88	3.56	5.00	6.52	6.10	6.45	6.01
U-Net-Small	1.9M	GR	1.82	2.08	2.13	2.48	2.15	-	-	-	-	-
		MV	19.56	5.41	5.69	3.82	8.52	54.29	20.93	19.25	8.11	24.85
		ED	3.20	3.49	3.19	3.49	3.35	3.35	3.78	3.48	3.95	3.65
		FS	3.32	2.91	2.79	3.18	3.07	3.42	3.61	3.41	3.69	3.54
U-Net-Large	7.7M	GR	1.73	2.14	2.00	2.33	2.06	-	-	-	-	-
		MV	18.76	5.27	5.54	3.79	8.25	52.08	20.00	18.65	7.77	23.86
		ED	2.95	3.08	2.86	3.27	3.05	3.32	3.70	3.47	3.86	3.60
		FS	2.68	2.94	2.79	3.14	2.90	3.32	3.64	3.41	3.64	3.50

Table 4: MPJPE at different speed levels. Each joint is classified into three levels, slow, medium, and fast. "ED," "MV," and "FS" stand for edge, motion vector, and fusion (edge and motion vectors).

Backbone	Input	Slow	Medium	Fast
DHP19 [B]	MV	36.66	9.66	9.65
	ED	5.66	6.90	8.22
	FS	5.69	6.76	7.34
U-Net-Small	MV	31.07	5.33	4.77
	ED	3.51	3.95	4.33
	FS	3.48	3.62	3.85
U-Net-Large	MV	29.84	5.10	4.57
	ED	3.47	3.82	4.22
	FS	3.44	3.62	3.84

Table 5: Comparison of GFLOPs and FPS (frames per second). All of the experiments are conducted for the fusion input modality on SPHP. "C" and "SC" are traditional and sparse convolutions.

Backbone	Params	Conv.	GFLOPs	FPS
DHP19 [B]	218K	C	275.45	26.89
		SC	33.25 (↓87%)	38.88 (1.5x)
U-Net-Small	1.9M	C	1135	11.82
		SC	46.74 (↓96%)	36.13 (3x)
U-Net-Large	7.7M	C	4510	1.07
		SC	186.80 (↓96%)	13.89 (13x)

HPE Performance on SPHP. The results in Table 3 show that our fusion (FS) model achieves the best results compared to other sparse inputs (*i.e.*, Edge (ED) and Motion Vector (MV)) for both traditional and sparse convolution networks. Besides, the gap between traditional and sparse convolutions is small in the U-Net-Small/U-Net-Large backbone, whereas the gap becomes more significant when using the smallest model (DHP19).

In Table 4, we evaluate MPJPE with sparse convolution networks based on the speed of each joint, categorized into three levels: slow, medium, and fast. For a 640×480 image, a joint that moves less than 4 pixels compared to the previous frame is considered slow; between 4-6 pixels is medium; more than 6 pixels is fast. The results show that the gap between FS and ED is the largest for fast-moving joints. U-Net-Small shows the most noteworthy relative improvement of 11% (from 4.33 to 3.85), validating that motion vectors complement edges in detecting fast movements. See supplementary materials for more performance analysis of backbones and joint speed.

The qualitative results are shown in Figure 4. According to the results, the edge images alone can lead to motion blur and decreased accuracy. For instance, the joint predictions in

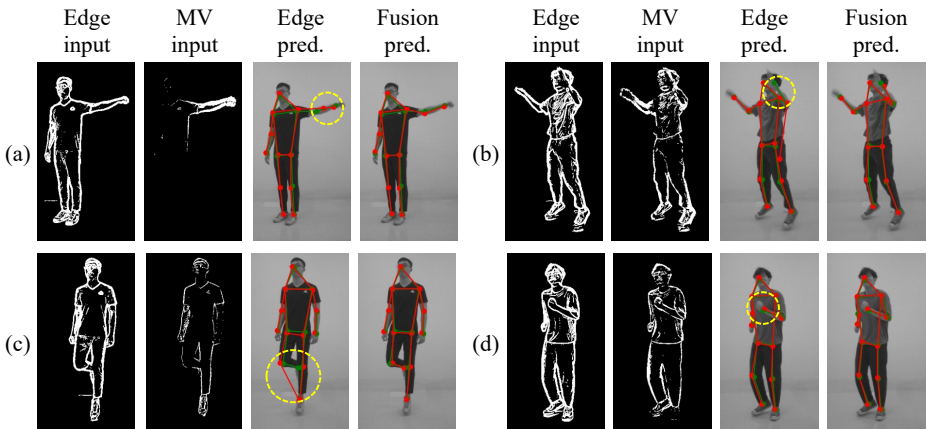


Figure 4: Qualitative results. Our fusion (*i.e.*, edge and MV) method rectifies the inaccurately predicted joints caused by motion blur and unclear contour when using only edge inputs, as shown in yellow circles. Green dots denote the ground truth, and red ones represent the predictions. Note that the grayscale images in the figure are for better visualization only and are not used as input for both predictions.

Figure 4(a)(b) are incorrect due to motion blur. Moreover, the unclear contour may also lead to inaccurate prediction, as shown in Figure 4(c)(d). Performing early fusion on edge and motion data enables models to rectify pose estimation.

Computational Efficiency. As shown in Table 5, sparse convolution can achieve a noteworthy 96% reduction in FLOPs. We also evaluate frames per second (FPS) on an Intel Core i9-7940 3.1GHz CPU and observe a 3x and 13x acceleration in U-Nets. With sparse convolution, U-Net-Small strikes the best balance between accuracy (MPJPE) and speed (FPS).

Experiments on more complex datasets. We evaluate our method on the MMHPSD [55], which provides event frames. On MMHPSD, we generate MV data using an off-the-shelf approach. The experimental results are presented in Table 6. Notably, our "MV + edge" fusion approach exhibits superior performance compared to using edge or MV modalities alone. This shows the generalizability of our method across various datasets. Furthermore, using "MV" as input demonstrates superior performance compared to using the "event" provided by MMHPSD, both with and without edges. This serves as evidence that MV captures better information than traditional event cameras. Besides MMHPSD, we include the experimental results from HumanEva [41] dataset in the supplementary material.

Table 6: MPJPE on MMHPSD [55] with different input types. Conv. stands for convolutions.

	Edge	Event	MV	Event + Edge	MV + Edge
Conv.	3.50	4.66	3.30	3.26	2.95
Sparse Conv.	3.84	9.75	6.91	3.28	3.06

5.2 Privacy-enhanced Representation

Face Recognition. To demonstrate the privacy-enhanced nature of our sparse data, we conduct face recognition experiments on CelebA [26] featuring over 10,000 individuals. We

Input	Acc.	Drop	Recall	Drop
RGB	88.9	-	84.1	-
Grayscale	88.7	0.2	84.2	-0.1
Edge	84.8	4.1	73.9	10.2

Table 7: Face recognition performance for various input types on CelebA [26]. We train the ResNet-50 [17] backbone with Arcface [18] loss to perform face recognition. The recall of edge input drops by 10.2.

convert CelebA’s RGB images into grayscale and edge images for comparison. Note that facial motion is typically small or absent, so it is not included. We perform face recognition on each type of image using the ResNet-50 [17] backbone with Arcface [18] loss.

Table 7 shows that face recognition performances are similar for RGB and grayscale images. However, when using edge input, the accuracy drops by 4.1% compared to RGB, while the recall drops even more significantly by 10.2%. This indicates that using the edge format can help reduce individuals’ privacy exposure compared to the other two formats.

Human Ability for Cross-modality Face Matching. We further test the cross-modality face-matching ability of humans on our SPHP dataset. To simulate people identifying different faces, we design a survey to analyze the ability of 100 participants to match leaked edge faces with grayscale faces from different angles and appearances. Each participant will be asked ten questions. In each question, the participant should identify the person from ten choices of edge faces given a grayscale reference face. In order to make the test more complicated, faces in the choices are presented from various angles and differ from those in the questions. The average test accuracy is only 18.8%, whereas the accuracy increases to 87.3% if we change the choices to grayscale images. These results suggest that people have a restricted ability to recognize daily faces from leaked edge images.

6 Conclusion

We introduce the novel SPHP dataset, which contains sparse edge images and two-directional motion vectors. Our proposed fusion model exhibits enhanced performance in human pose estimation compared to individual modalities using only edge or motion vectors. By leveraging our sparse data with submanifold sparse convolutions, we further reduce the FLOPs by 96% and achieve a 13x speed-up compared to traditional convolutional neural networks. In addition, we showcase the generalization capability of our method through experiments on MMHPSD and HumanEva datasets. Furthermore, to verify the privacy-enhanced nature of our representation, we conduct a face recognition experiment using edge images as inputs, which results in a recall drop of 10.2% compared to the use of RGB images. Finally, we carry out a user study to demonstrate the limited ability of humans to recognize faces across leaked edge and daily grayscale images.

7 Acknowledgements

This work is supported in part by Ministry of Science and Technology of Taiwan (NSTC 111-2634-F-002-022) and Novatek Microelectronics Corp. We thank National Center for High-performance Computing (NCHC) for computational and storage resource. Besides, we appreciate Novatek’s MVS technology in efficiently collecting edge and motion data for this project.

References

- [1] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3686–3693, 2014. doi: 10.1109/CVPR.2014.471.
- [2] Paula Arguello, Jhon Lopez, Carlos Hinojosa, and Henry Arguello. Optics lens design for privacy-preserving scene captioning. In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 3551–3555. IEEE, 2022.
- [3] Laurens van der Maaten Benjamin Graham. Submanifold sparse convolutional networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [4] Thomas Brox and Jitendra Malik. Large displacement optical flow: descriptor matching in variational motion estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2010.
- [5] Enrico Calabrese, Gemma Taverni, Christopher Awai Easthope, Sophie Skriabine, Federico Corradi, Luca Longinotti, Kynan Eng, and Tobi Delbruck. Dh19: Dynamic vision sensor 3d human pose dataset. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2019.
- [6] John Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, PAMI-8(6):679–698, 1986. doi: 10.1109/TPAMI.1986.4767851.
- [7] Ting Cao, Mohammad Ali Armin, Simon Denman, Lars Petersson, and David Ahmedt-Aristizabal. In-bed human pose estimation from unseen and privacy-preserving image domains. In *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2021.
- [8] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7291–7299, 2017.
- [9] Junyi Chai, Hao Zeng, Anming Li, and Eric WT Ngai. Deep learning in computer vision: A critical review of emerging techniques and application scenarios. *Machine Learning with Applications*, 6:100134, 2021.
- [10] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded pyramid network for multi-person pose estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [11] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [12] Martin Engelcke, Dushyant Rao, Dominic Zeng Wang, Chi Hay Tong, and Ingmar Posner. Vote3deep: Fast object detection in 3d point clouds using efficient convolutional neural networks. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.

- [13] Xiaochuan Fan, Kang Zheng, Yuewei Lin, and Song Wang. Combining local appearance and holistic view: Dual-source deep neural networks for human pose estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [14] Hao-Shu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. Rmpe: Regional multi-person pose estimation. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2334–2343, 2017.
- [15] Ben Graham. Sparse 3d convolutional neural networks. In *British Machine Vision Conference (BMVC)*, pages 150.1–150.9, 2015.
- [16] Benjamin Graham. Spatially-sparse convolutional neural networks, 2014.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [18] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [19] Carlos Hinojosa, Juan Carlos Niebles, and Henry Arguello. Learning privacy-preserving optics for human pose estimation. In *IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [20] Carlos Hinojosa, Miguel Marquez, Henry Arguello, Ehsan Adeli, Li Fei-Fei, and Juan Carlos Niebles. Privhar: Recognizing human actions from privacy-preserving lens. In *European Conference on Computer Vision (ECCV)*, pages 314–332. Springer, 2022.
- [21] Shaoli Huang, Mingming Gong, and Dacheng Tao. A coarse-fine network for keypoint localization. In *IEEE International Conference on Computer Vision (ICCV)*, pages 3047–3056, 2017.
- [22] Eldar Insafutdinov, Leonid Pishchulin, Bjoern Andres, Mykhaylo Andriluka, and Bernt Schiele. Deeppercut: A deeper, stronger, and faster multi-person pose estimation model. In *European Conference on Computer Vision (ECCV)*, 2016.
- [23] Ranjay Krishna, Mitchell Gordon, Li Fei-Fei, and Michael Bernstein. Visual intelligence through human interaction. *Artificial Intelligence for Human Computer Interaction: A Modern Approach*, 2021.
- [24] 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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [25] 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 *European Conference on Computer Vision (ECCV)*, 2014.
- [26] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015.

- [27] Manuel J. Marin-Jimenez, Francisco J. Romero-Ramirez, Rafael Muñoz-Salinas, and Rafael Medina-Carnicer. 3d human pose estimation from depth maps using a deep combination of poses. *Journal of Visual Communication and Image Representation*, 2018.
- [28] David R Martin, Charless C Fowlkes, and Jitendra Malik. Learning to detect natural image boundaries using local brightness, color, and texture cues. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2004.
- [29] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European Conference on Computer Vision (ECCV)*, 2016.
- [30] José Ramón Padilla-López, Alexandros Andre Charaoui, and Francisco Flórez-Revuelta. Visual privacy protection methods: A survey. *Expert Systems with Applications*, 42(9):4177–4195, 2015.
- [31] George Papandreou, Tyler Zhu, Nori Kanazawa, Alexander Toshev, Jonathan Tompson, Chris Bregler, and Kevin Murphy. Towards accurate multi-person pose estimation in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4903–4911, 2017.
- [32] Mathias Parger, Chengcheng Tang, Christopher D. Twigg, Cem Keskin, Robert Wang, and Markus Steinberger. Deltacnn: End-to-end cnn inference of sparse frame differences in videos. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [33] Leonid Pishchulin, Eldar Insafutdinov, Siyu Tang, Bjoern Andres, Mykhaylo Andriluka, Peter Gehler, and Bernt Schiele. Deepcut: Joint subset partition and labeling for multi person pose estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [34] Francesco Pittaluga and Sanjeev J Koppal. Privacy preserving optics for miniature vision sensors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 314–324, 2015.
- [35] Zhongzheng Ren, Yong Jae Lee, and Michael S Ryoo. Learning to anonymize faces for privacy preserving action detection. In *Proceedings of the european conference on computer vision (ECCV)*, pages 620–636, 2018.
- [36] Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. Octnet: Learning deep 3d representations at high resolutions. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [37] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [38] Michael Ryoo, Brandon Rothrock, Charles Fleming, and Hyun Jong Yang. Privacy-preserving human activity recognition from extreme low resolution. In *AAAI Conference on Artificial Intelligence (AAAI)*, volume 31, 2017.

- [39] Jamie Shotton, Andrew Fitzgibbon, Mat Cook, Toby Sharp, Mark Finocchio, Richard Moore, Alex Kipman, and Andrew Blake. Real-time human pose recognition in parts from single depth images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011.
- [40] Jamie Shotton, Ross Girshick, Andrew Fitzgibbon, Toby Sharp, Mat Cook, Mark Finocchio, Richard Moore, Pushmeet Kohli, Antonio Criminisi, Alex Kipman, and Andrew Blake. Efficient human pose estimation from single depth images. In *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2013.
- [41] Leonid Sigal, Alexandru Balan, and Michael Black. Humaneva: Synchronized video and motion capture dataset and baseline algorithm for evaluation of articulated human motion. *International Journal of Computer Vision*, 87:4–27, 03 2010. doi: 10.1007/s11263-009-0273-6.
- [42] Hang Su, Varun Jampani, Deqing Sun, Subhransu Maji, Evangelos Kalogerakis, Ming-Hsuan Yang, and Jan Kautz. Splatnet: Sparse lattice networks for point cloud processing. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [43] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5693–5703, 2019.
- [44] Yuandong Tian, C. Lawrence Zitnick, and Srinivasa G. Narasimhan. Exploring the spatial hierarchy of mixture models for human pose estimation. In *European Conference on Computer Vision (ECCV)*, pages 256–269, 2012.
- [45] Alexander Toshev and Christian Szegedy. Deeppose: Human pose estimation via deep neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1653–1660, 2014.
- [46] Peng-Shuai Wang, Yang Liu, Yu-Xiao Guo, Chun-Yu Sun, and Xin Tong. O-cnn: Octree-based convolutional neural networks for 3d shape analysis. *ACM Transactions on Graphics (SIGGRAPH)*, 36, 2017.
- [47] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional pose machines. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4724–4732, 2016.
- [48] Wenxuan Wu, Zhongang Qi, and Li Fuxin. Pointconv: Deep convolutional networks on 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [49] Zhenyu Wu, Zhangyang Wang, Zhaowen Wang, and Hailin Jin. Towards privacy-preserving visual recognition via adversarial training: A pilot study. In *Proceedings of the European conference on computer vision (ECCV)*, pages 606–624, 2018.
- [50] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *European Conference on Computer Vision (ECCV)*, pages 466–481, 2018.

- [51] Xixia Xu, Yingguo Gao, Ke Yan, Xue Lin, and Qi Zou. Location-free human pose estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13137–13146, 2022.
- [52] Helen H Yang and Thomas R Clandinin. Elementary motion detection in drosophila: algorithms and mechanisms. *Annual Review of Vision Science*, 4:143–163, 2018.
- [53] Jie Yang, Ailing Zeng, Shilong Liu, Feng Li, Ruimao Zhang, and Lei Zhang. Explicit box detection unifies end-to-end multi-person pose estimation. In *International Conference on Learning Representations (ICLR)*, 2023.
- [54] Wei Yang, Shuang Li, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang. Learning feature pyramids for human pose estimation. In *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [55] Shihao Zou, Chuan Guo, Xinxin Zuo, Sen Wang, Pengyu Wang, Xiaoqin Hu, Shoushun Chen, Minglun Gong, and Li Cheng. Eventhpe: Event-based 3d human pose and shape estimation. In *IEEE International Conference on Computer Vision (ICCV)*, 2021.