

Fast and Accurate: Video Enhancement Using Sparse Depth

Yu Feng

University of Rochester

yfeng28@ur.rochester.edu

Patrick Hansen

Tenstorrent Inc.

phansen@tenstorrent.com

Paul N. Whatmough

Arm Research

paul.whatmough@arm.com

Guoyu Lu

University of Georgia

gl72151@uga.edu

Yuhao Zhu

University of Rochester

yzhu@rochester.edu

Abstract

This paper presents a general framework to build fast and accurate algorithms for video enhancement tasks such as super-resolution, deblurring, and denoising. Essential to our framework is the realization that the accuracy, rather than the density, of pixel flows is what is required for high-quality video enhancement. Most of prior works take the opposite approach: they estimate dense (per-pixel)—but generally less robust—flows, mostly using computationally costly algorithms. Instead, we propose a lightweight flow estimation algorithm; it fuses the sparse point cloud data and (even sparser and less reliable) IMU data available in modern autonomous agents to estimate the flow information. Building on top of the flow estimation, we demonstrate a general framework that integrates the flows in a plug-and-play fashion with different task-specific layers. Algorithms built in our framework achieve $1.78\times$ — $187.41\times$ speedup while providing a 0.42 dB – 6.70 dB quality improvement over competing methods.

1. Introduction

Video enhancement tasks ranging from super resolution [32, 3, 29, 16], deblurring [26, 38, 21], and denoising [28, 9] are becoming increasingly important for intelligent systems such as smartphones and Augmented Reality (AR) glasses. High-quality videos are also critical to various robotics tasks, such as SLAM [25, 7], visual odometry [15], object detection [22], and surveillance [24].

Video enhancement systems today face a fundamental dilemma. High quality enhancement benefits from accurately extracting temporal flows across adjacent frames, which, however, is difficult to obtain from low-quality videos (e.g., low-resolution, noisy). As a result, video enhancement usually requires expensive optical flow algorithms, usually in the form of Deep Neural Networks

(DNNs), to extract dense flows, leading to a low execution speed. As video enhancement tasks execute on resource-limited mobile devices and potentially in real time, there is a need for high-speed and high-quality video enhancement.

We propose a method to simultaneously increase the quality and the execution speed of video enhancement tasks. Our work is based on the realization that the accuracy, rather than the density, of the flow estimation is what high-quality enhancement requires. We propose an algorithm to estimate accurate, but sparse, flows using LiDAR-generated point clouds. Coupled with the flow estimation algorithm, we demonstrate a generic framework that incorporates the flows to build video enhancement DNNs, which are lightweight by design assisted by the accurate flows.

Our flow estimation is accurate because it does not rely on the image content, which is necessarily of low-quality in video enhancement tasks. Instead, we generate flows using the accurate depth information from LiDAR point cloud assisted with the less reliable IMU information. By exploiting the spatial geometry of scene depth and the agent’s rough ego-motion (through IMU), our algorithm estimates the flows in videos using a purely analytical approach without complex feature extraction, matching, optimization, and learning used in conventional flow estimation algorithms.

Building on top of the lightweight flow estimation, we demonstrate a general framework that integrates the flows for video enhancement. The framework consists of a common temporal alignment front-end and a task-specific back-end. The front-end temporally aligns a sequence of frames by warping and concatenating frames using the estimated flows; the back-end extracts task-specific features to synthesize high-quality videos. Different from prior works that specialize the temporal alignment module for a specific task, our unified temporal alignment module applies to different tasks and, thus, empowers algorithm developers to focus energy on the task-specific back-end.

We demonstrate our framework on a range of video en-

hancement tasks including super resolution, deblurring, and denoising on the widely-used KITTI dataset [14]. Across all tasks, our system has better enhancement quality than state-of-the-art algorithms measured in common metrics such as Peak Signal-to-Noise Ratio (PSNR) and Structure Similarity Index Measure (SSIM) [33]. Meanwhile, we improve the execution speed on all tasks by a factor of 8.4 on average (up to 187.4 times). The code will be open-sourced.

2. Related Work

Video Enhancement The general theme in today’s video enhancement algorithms is to first align neighboring frames from time $t - n$ to time $t + m$ and then fuse the aligned frames to enhance the target frame t . Much of the prior innovations lie in how to better align frames.

Alignment could be done explicitly or implicitly. Explicit approaches perform an explicit flow estimation between frames [3, 28, 16]. The flows are then used to align frames either in the image space [3, 16] or in the feature space [28]. Obtaining accurate flows typically requires expensive flow estimation algorithms (e.g., dense optical flow [16] or complicated DNNs [28, 3]), which lead to low execution speed. Implicit approaches, instead, align frames in latent space using algorithms such as deformable convolution [29, 32] or recurrent neural networks [38]. Classic examples include EDVR [32], TDAN [29] and ESTRNN [38]. These algorithms tend to be more accurate than explicit approaches when the temporal correlation is not obvious in pixels pace.

Our work differs from prior works in two main ways. First, both implicit and explicit approaches are computationally-heavy, as they extract flows from purely the vision modality. We demonstrate a very fast algorithm to extra flows by fusing LiDAR and IMU data. We show that accurate flows enable a simple downstream DNN design, achieving state-of-the-art task quality while being an order of magnitude faster. Second, the alignment modules in prior works usually are specialized for specific enhancement tasks. We instead show a common alignment module based on our estimated flows broadly applies to a range of video enhancement tasks. This greatly eases development and deployment effort in practice.

LiDAR-Guided Vision Fusing point clouds and images is known to improve the quality of vision tasks such as object detection [4, 35, 36], segmentation [10, 19], and stereo matching [6, 31], but literature is scarce in LiDAR-camera fusion for video enhancement.

Fusion networks usually extract features from (LiDAR-generated) point clouds and images, and align/fuse the two sets of features before feeding them to the task-specific block. Unlike prior fusion algorithms that extract features from point clouds, we propose a different way of using point cloud data, i.e., estimating explicit pixel flow from point

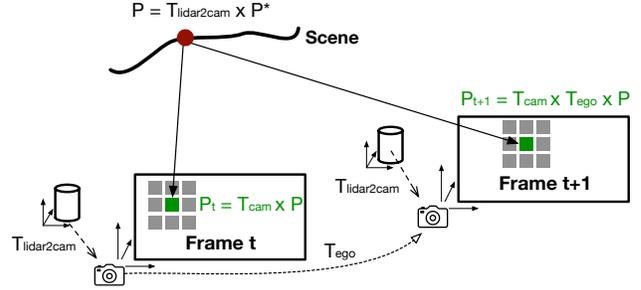


Fig. 1. LiDAR-guided flow estimation. P^* is the 3D coordinates of a point in the LiDAR coordinate system at time t . $T_{lidar2cam}$ is the transformation matrix from the LiDAR coordinate system to the camera coordinate system, which is fixed over time assuming the configuration of the LiDAR and camera is rigid. T_{cam} is the camera matrix. T_{ego} is the camera egomotion from Frame t to Frame $t + 1$.

clouds. The estimated flows are accurate and, thus, provide targeted guidance to video enhancement tasks.

Flow Estimation Estimating flows between frames is a fundamental building block. Video-based flow estimation has made great strides through DNNs [12, 23, 27]. These methods, however, are computationally intensive. When incorporated into a high-level vision task such as deblurring and denoising, the flow estimation quickly becomes a speed bottleneck. Many flow estimations algorithms use only video frames, which, while is less restrictive, also means the flow accuracy degrades when operating on low-quality videos. Our method is image content-independent and thus better estimates flows from low-quality videos. It is also fast, because it relies on simple geometric transformations.

Existing video enhancement tasks usually use dense and per-pixel flow estimation [16, 3, 28]. In contrast, our method generates sparse flows from point clouds. A key contribution of our work is to demonstrate that even a sparse flow can greatly boost the quality of video enhancement.

3. Main Idea and Optimizations

We first describe the lightweight flow estimation algorithm (Sec. 3.1), followed by a generic DNN architecture that integrates the flows for video enhancement (Sec. 3.2).

3.1. Lightweight and Accurate Flow Estimation

Overall Algorithm The key idea is to use the depth data from LiDAR to generate flows in a lightweight fashion. Fig. 1 illustrates the idea. For any point P^* in a point cloud, it is captured by two consecutive camera frames. At time t , P^* ’s coordinates in the camera coordinate system are $P = T_{lidar2cam} \times P^*$, where $T_{lidar2cam}$ is the LiDAR to camera transformation matrix, which is usually pre-calibrated. Thus, the corresponding pixel coordinates in the image at time t are $P_t = T_{cam} \times P$, where T_{cam} is the

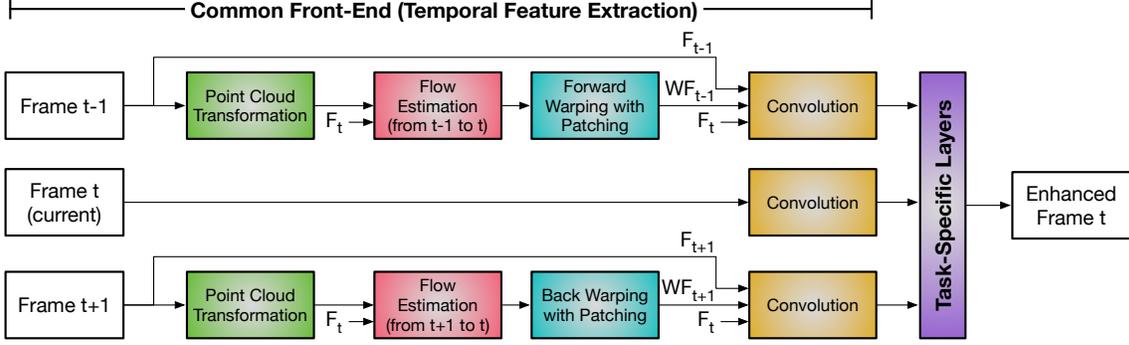


Fig. 2. Overview of our two-stage video enhancement DNN architecture. The front-end performs lightweight flow estimation to align (previous and later) frames with the current frame at time t in order to extract temporal features. The extracted features carry temporal correlations across frames and are then processed by task-specific layers to produce an enhanced frame. We merge point clouds (using the estimated ego-motion) before flow estimation and warp pixels in patches after flow estimation, both to mitigate the sparsity of LiDAR-generated point clouds.

camera matrix.

At time $t + 1$, the coordinates of the same point in the scene in the camera coordinate system are $T_{ego} \times P$, where T_{ego} is the transformation matrix of the camera egomotion. Thus, the pixel coordinates of the point at $t + 1$ are $P_{t+1} = T_{cam} \times T_{ego} \times P$. Accordingly, the pixel's motion vector can be calculated in a computationally very lightweight manner:

$$\begin{aligned} \delta_t &= P_{t+1} - P_t \\ &= T_{cam} \times T_{ego} \times T_{lidar2cam} \times P^* \\ &\quad - T_{cam} \times T_{lidar2cam} \times P^*. \end{aligned} \quad (1)$$

Egomotion The camera egomotion T_{ego} could be derived in a range of different methods. In our system, we estimate T_{ego} using the measurements from the IMU, which is widely available in virtually all intelligent devices. We note that the IMU data, while being a readily available sensor modality, is known to be a rough and imprecise estimation of the true egomotion [8]. One of our contributions is to show how the rough egomotion estimation can provide decent flow estimation for high-quality video enhancement.

The IMU provides the translational acceleration ($\hat{\mathbf{a}}$) and the angular velocity ($\hat{\omega}$). Given $\hat{\mathbf{a}}$, the translation component $T_{3 \times 1}$ in T_{ego} is calculated by:

$$T_{3 \times 1} = [\overline{\Delta x} \quad \overline{\Delta y} \quad \overline{\Delta z}] \quad (2)$$

where $\overline{\Delta x}$, $\overline{\Delta y}$, and $\overline{\Delta z}$ are the three translational displacements integrated from $\hat{\mathbf{a}}$ using Euler's method. Similarly, the rotational component $R_{3 \times 3}$ in T_{ego} is estimated from $\hat{\omega}$:

$$R_{3 \times 3} = R_{3 \times 3}^y \times R_{3 \times 3}^p \times R_{3 \times 3}^r \quad (3)$$

where $R_{3 \times 3}^y$, $R_{3 \times 3}^p$, and $R_{3 \times 3}^r$ denote the three rotational matrices, which are integrated from the three rotational displacements in $\hat{\omega}$ using Euler's method.

Alternatively, T_{ego} could be estimated through point cloud registration [2, 37], which is generally more costly but more accurate. As we will demonstrate later (Sec. 5.3), even with relatively less accurate estimations from the IMU, our method can still reliably deliver higher task accuracy while greatly reducing the execution time. When given more accurate egomotions (e.g., from point cloud registration), our framework would have a higher video enhancement quality.

A key reason why video enhancement benefits from our flow estimation is that our algorithm is purely based on 3D geometry and geometric transformation *without* relying on the image content. No pixel content participates in the flow estimation Eqn. 1. Therefore, it estimates flows accurately even when the image content is of low-quality, e.g., low resolution or noisy, which is exactly the kind of scenario video enhancement tasks target at.

3.2. A Generic DNN Architecture

Our flow estimation is a building block for simultaneously improving the quality and execution speed of video enhancement. We propose a generic DNN architecture that incorporates the estimated flows for a range of video enhancement tasks. Fig. 2 shows an overview of the architecture, which consists of two main modules: a common frame fusion front-end and a task-specific back-end.

Temporal Feature Extraction Our network uses a common front-end shared across different enhancement tasks. The goal of the front-end is to extract temporal correlations across frames in preparation for task-specific processing. Fig. 2 shows an example that extracts temporal features across three frames: the current frame F_t and the frame before (F_{t-1}) and after (F_{t+1}) the current frame, which we call the temporal frames. More temporal frames

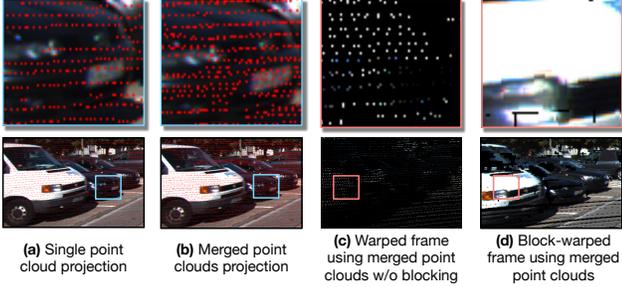


Fig. 3. Effectiveness of techniques to increase point cloud density. (a): frame overlaid with a projected single point cloud (red pixels are projected points). (b): frame overlaid with the projection of five merged point clouds. (c): warped frame using flows estimated from five merged point clouds without blocking. (d): frame warped using flows from five merged point clouds and warped in 5×5 blocks.

are possible in principle.

The front-end first calculates the flows between each temporal frame and the current frame using the algorithm described in Sec. 3.1. A critical challenge we face is that the estimated flows are necessarily sparser than the corresponding image, because LiDARs generally have lower resolutions than that of cameras. For instance, the Velodyne HDL64E LiDAR, a high-end, high-density LiDAR, generates about 130,000 points per frame, whereas an image with a 720p resolution contains about 1 million points. Fig. 3(a) shows the effect of using sparse point clouds, where only a small amount of pixels have points associated with them when projecting a single point cloud to the image.

To mitigate the sparsity of LiDAR-generated point clouds, we propose to register multiple point clouds together to form a dense point cloud. We register point clouds by simply transforming adjacent point clouds using the ego-motion T_{ego} calculated from the IMU measurements (Eqn. 2 and Eqn. 3). Fig. 3(b) shows that when projecting multiple registered point clouds, many more pixels are associated with points. Alternatively, one could leverage computationally intensive, but potentially more accurate, registration algorithms (e.g., using Iterative Closest Point [2, 5]).

Even with multiple point clouds, not every image pixel in F_{t-1} (or F_{t+1}) has a corresponding flow. As a result, when warping images using flows the warped images will have many “holes”, as illustrated in Fig. 3(c). While one could merge more point clouds to increase the point density, doing so is susceptible to mis-registration, which is especially significant when merging a long sequence of point clouds where errors can accumulate.

To address this issue, we propose blocked warping, which duplicates a pixel’s flow to its neighboring pixels (e.g., a 5×5 block) during warping. This is analogous to blocked-based motion compensation in conventional video compression. The assumption is that points corresponding

to the neighboring pixels have similar motion in the 3D space, and thus their pixel flows are similar. We warp a temporal frame (F_{t-1} or F_{t+1}) to the current frame using the blocked flows. The result is shown in Fig. 3(d), which has much dense pixels (fewer “holes”) than in Fig. 3(c).

Finally, each warped temporal frame (e.g., WF_{t-1}), along with its unwarped counterpart (e.g., F_{t-1}) and the current frame (F_t), are concatenated and go through a convolutional layer to extract the temporal correlations between the temporal frame and the current frame. The features of the current frame are extracted independently.

Task-Specific Layers The back-end of our architecture takes the extracted temporal features to perform video enhancement. The exact design of the back-end layers is task-specific. Our goal of this paper is *not* to demonstrate new task-specific layers; rather, we show that our temporal feature extraction front-end is compatible with different task layers in a plug-and-play manner.

To that end, we implement three back-end designs for three video enhancement tasks, including super-resolution, denoising, and deblurring, by directly using designs from other algorithms (with slight modifications so that the interface matches our front-end). The layers for super-resolution and deblurring connect the temporal features from the front-end in a recurrent fashion, similar to designs of RBPN [16] and ESTRNN [38], respectively. The denoising layers concatenate the temporal features, which then enter a set of convolutional layers, similar to DVDnet [28].

4. Evaluation Methodology

Applications and Baselines We evaluate three video enhancement tasks, namely super-resolution, deblurring and denoising.

- Super-resolution: we compare with two baselines: RBPN [16] and VESPCN [3]. RBPN uses a recurrent encoder-decoder to learn temporal correlations; VESPCN warps images in the pixel space and fuses multiple warped frames through a CNN to upsample.
- Deblurring: we compare with ESTRNN [38], which uses RNN to learn the temporal features; we also compare with DEEPGYRO [20], which fuses IMU with image data for single-image deblur.
- Denoising: we compare with DVDNET [28], which uses CNN to extract explicit motion and warp frames.

In addition, we also designed a simple LiDAR-camera fusion baseline for *each* task. This baseline, which we call VEFUSION, resembles many LiDAR/camera fusion DNNs [13]: it first concatenates the projected point cloud and the image; the concatenated data then enters the task-specific layers. Our proposed method also leverages point

clouds for video enhancement, but uses point clouds in a different way: instead of fusing points with pixels, we use point clouds to generate flows. This baseline allows us to assess the effectiveness of this way of using point cloud for video enhancement. We make sure VEFUSION has roughly the same amount of parameters as our proposed method such that the performance difference is due to the algorithm.

Variants We evaluate two variants of our methods: OURS-S uses a single point cloud for flow estimation, and OURS-M uses five point clouds for flow estimation.

Dataset We use the KITTI dataset [14], which provides sequences of synchronized LiDAR, camera, and IMU data. Following the common practices, we preprocess the dataset for different tasks. For super-resolution we downsize the videos by $4\times$ in both dimensions using bicubic interpolation, similar to VESPCN [3]; for deblurring we add Gaussian blur to the videos, similar to EDVR [32]; for denoising we apply random noises to the videos, similar to DVDnet [28].

Evaluation Metrics To evaluate the efficacy of our method, we use two metrics, PSNR and SSIM, to qualitatively evaluate the results. We also show the runtime performance of different methods by measuring the execution time of different methods on two platforms, one is the Nvidia RTX 2080 GPU; the other is the mobile Volta GPU on Nvidia’s recent Jetson Xavier platform [1]. Each execution time is averaged over 1000 runs.

Design Parameters Unless otherwise noted, we use a block size of 3×3 in super resolution, and a block size of 7×7 in deblurring and denoising tasks. Five point clouds are registered for flow estimation. We will study the sensitivity to these two design parameters (Sec. 5.4).

5. Evaluation

We show that the execution speed of our method is on average an order of magnitude faster than existing methods while at the same time delivering higher task quality, both objectively and subjectively (Sec. 5.1). We study the accuracy of our flow estimation (Sec. 5.2) and the sensitivity of our method on key design parameters (Sec. 5.4).

5.1. Overall Evaluation

Results Overview OURS-M and OURS-S consistently outperform the baselines in both quality and speed. OURS-M is slightly better than OURS-S by using multiple point clouds for flow estimation. A naive fusion of point cloud and images, as done by VEFUSION, has significantly lower quality than our methods, albeit with a similar speed.

Super-resolution Tbl. 1 compares different super-resolution algorithms. We also show the execution time of different methods normalized to that of OURS-M.

Overall, OURS-M achieves the highest visual quality both in terms of PSNR and SSIM among all methods.

Table 1. Super-resolution comparison. Execution times are normalized to that on OURS-M; H and M denote the high-end 2080 Ti GPU and the mobile Volta GPU, respectively.

	<u>RBPN</u>	<u>VESPCN</u>	<u>VEFUSION</u>	<u>OURS-S</u>	<u>OURS-M</u>
PSNR (dB)	27.08	24.78	26.95	27.43	27.50
SSIM	0.860	0.787	0.854	0.873	0.872
Time (H)	36.10	0.55	1.00	1.00	1.00
Time (M)	7.24	0.13	1.00	1.00	1.00

OURS-S has similar SSIM but lower PSNR. OURS-M achieves a $36.10\times$ speedup against RBPN on 2080 Ti and $7.24\times$ speedup on the mobile GPU, showing the effectiveness of our lightweight flow estimation algorithm, which executes in about $10\mu s$ on GPUs. OURS-M and OURS-S have virtually the same speed, because transforming point clouds into one frame has negligible overhead. VEFUSION has the same speed as our methods with lower quality. VESPCN is the fastest, but has a much lower super-resolution quality due to a simpler CNN.

Table 2. Deblurring comparison.

	<u>ESTRNN</u>	<u>DEEPGYRO</u>	<u>VEFUSION</u>	<u>OURS-S</u>	<u>OURS-M</u>
PSNR (dB)	34.78	31.20	35.22	35.50	36.61
SSIM	0.945	0.806	0.949	0.950	0.957
Time (H)	1.78	6.20	1.00	1.00	1.00
Time (M)	1.08	11.96	1.00	1.00	1.00

Deblurring Tbl. 2 compares different methods on video deblurring. Our method, OURS-M, achieves the highest quality both in terms of PSNR and SSIM. Compared to ESTRNN, OURS-M achieves 1.83 higher in PSNR and 0.012 higher in SSIM. Our methods are also faster than the baselines on both GPUs. The speedup on ESTRNN is not significant, because the flow estimation in ESTRNN is small to begin with (7.7% on the mobile GPU). DEEPGYRO has the lowest task quality and the slowest speed. Its low quality is attributed to the fact that it deblurs using a single image, while other methods use temporal information.

Table 3. Denoising comparison.

	<u>DVDNET</u>	<u>VEFUSION</u>	<u>OURS-S</u>	<u>OURS-M</u>
PSNR (dB)	27.19	31.60	33.34	33.89
SSIM	0.838	0.951	0.953	0.961
Time (H)	187.41	1.00	0.99	1.00
Time (M)	68.97	1.00	0.99	1.00

Denoising For video denoising, OURS-M achieves the highest quality both in PSNR and SSIM, as shown in Tbl. 3. OURS-M improves upon VEFUSION and DVDNET by a large margin — 2.29 dB and 6.70 dB in PSNR, respectively. Meanwhile, OURS-M has a $187.4\times$ speedup compared to DVDNET on 2080 Ti and $69.0\times$ speedup on the mobile GPU. The speedup comes from avoiding the expensive flow estimation algorithm DeepFlow [34] used in DVDNET.



Fig. 4. Visual comparison of different methods on various visual enhancement tasks.

Subjective Comparison Our approach is also visually better than the baselines upon subjective comparisons. Fig. 4 shows the visual comparisons on different tasks. The improvements from the baselines to Ours-M are the most significant. Ours-M is best at revealing details, such as the roads and bushes, because of its dense motion obtained from merging point clouds.

5.2. Flow Estimation Accuracy and Speed

Our lightweight flow estimation algorithm provides accurate flow information. To demonstrate the effectiveness of the estimated flows, we warp frames in the dataset using

the estimated flows and calculate the PSNR. Tbl. 4 shows the results across different flow estimation algorithms used in different networks. We also show the speed of different flow estimation algorithms normalized to that of ours.

Table 4. Flow estimation comparison. Execution time is normalized to that of ours.

	DVDNET	VESPCN	RBPN	Ours
PSNR (dB)	14.71	16.64	22.68	18.74
Time (H)	4147.5	1420.0	98694.0	1.0

Judged by the quality of warped images, our flow esti-

mation method is better than the estimation methods used in DVDNET and VESPCN, as shown in Tbl. 4. This also explains the task quality difference. Interestingly, while the frames warped using our flow estimation have a lower PSNR compared to those in RBPN, we are able to achieve a better super-resolution quality than RBPN. The reason is that our method uses the warped frames to extract temporal features (Fig. 2) while RBPN uses the actual flow values.

Our flow estimation is at least three orders of magnitude faster than other methods used in baselines. This explains the overall speed difference shown earlier, since our task-specific layers are similar to those used in the baselines.

5.3. Egomotion from Point Cloud Registration

In this section, we show that more accurate egomotion from point cloud registration can further improve task accuracy. Here, we compare two variants that both use single point cloud for flow estimation:

- OURS-S (I): obtains the egomotion from IMU.
- OURS-S (R): obtains the egomotion from point cloud registration.

Table 5. Comparison of different sources of egomotion. I and R denote obtaining egomotion from IMU and point cloud registration, respectively.

PSNR (dB)	Super-resolution	Deblurring	Denosing
<u>OURS-S (I)</u>	27.43	35.69	33.34
<u>OURS-S (R)</u>	27.32	35.91	33.45

Tbl. 5 shows the accuracy comparison between OURS-S (I) and OURS-S (R) on three vision tasks. OURS-S (R) generally has higher quality than OURS-S (I), because the former uses registration to generate a more accurate camera egomotion (compared to using the IMU). Meanwhile, limited accuracy improvements also show that a rough egomotion from IMU is sufficient for DNN model to learn the correlation between input pixel values and estimated flows.

5.4. Sensitivity Study

We use super-resolution as an example to study how the block size used in blocked warping and the number of merged point clouds used in flow estimation influence the task quality. Other tasks have a similar trend.

Block Size Larger blocks initially improve the task quality. Tbl. 6 shows how the super-resolution quality varies with the block size. When the block size initially increases from 1×1 to 3×3 , the PSNR improves because the flow density increases. Increasing the block size further degrades the quality. This is because with large blocks more

Table 6. Sensitivity of the block size and the number of merged point clouds on super-resolution on OURS-S.

Patch size	1×1	3×3	5×5	7×7
PSNR (dB)	27.02	27.43	27.29	27.26
# of point clouds	1	3	5	7
PSNR (dB)	27.43	27.47	27.50	27.52

flows are duplicated from neighbor pixels rather than calculated using depth information, reducing the flow accuracy.

Number of Merged Point Clouds Merging more point clouds leads to denser and more accurate flow estimation and thus a higher the task quality. This is evident in Tbl. 6, which shows that the PSNR of increases as the number of merged point clouds increases.

6. Conclusion

We demonstrate a general framework to build fast and accurate video enhancement algorithms. The key is to assist video enhancement with an accurate depth-driven flow estimation algorithm. Our flow estimation is accurate because it leverages the accurate depth information generated from LiDARs based on a physically-plausible scene model. We show strategies to overcome the sparsity of LiDAR point clouds. Our flow estimation is lightweight because it relies on only simple geometric transformations, enabling lean end-to-end algorithms. We propose a generic framework that integrates the flow estimation with task-specific layers in a plug-and-play manner. We achieve over an order of magnitude speedup while improving task quality over competing methods. While fusing point clouds with images has been extensively studied lately in vision tasks, we show that using point clouds for flow estimation, rather than simply fusing them with images, achieves better performance.

An implication of our framework is that the point cloud data must be attached to the video content, which could potentially increase the storage and transmission overhead. However, the overhead is likely small, because the size of point cloud data is smaller than that of images. For instance, one point cloud frame obtained from a high-end Velodyne HDL-64E LiDAR [30] is about 1.5 MB, whereas one 1080p image is about 6.0 MB in size. The overhead will become even smaller in the future as point cloud compression techniques become more mature [11, 17, 18].

7. Acknowledgement

The work was supported, in part, by NSF award #2126642 and #2044963 as well as a Meta research grant.

References

- [1] Jetson Xavier. <https://www.nvidia.com/en-us/>

autonomous-machines/embedded-systems/
jetson-agx-xavier/.

- [2] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Sensor Fusion IV: Control Paradigms and Data Structures*, volume 1611, pages 586–607. International Society for Optics and Photonics, 1992.
- [3] Jose Caballero, Christian Ledig, Andrew Aitken, Alejandro Acosta, Johannes Totz, Zehan Wang, and Wenzhe Shi. Real-time video super-resolution with spatio-temporal networks and motion compensation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4778–4787, 2017.
- [4] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1907–1915, 2017.
- [5] Yang Chen and Gérard Medioni. Object modelling by registration of multiple range images. *Image and vision computing*, 10(3):145–155, 1992.
- [6] Xuelian Cheng, Yiran Zhong, Yuchao Dai, Pan Ji, and Hongdong Li. Noise-aware unsupervised deep lidar-stereo fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6339–6348, 2019.
- [7] Younggun Cho and Ayoung Kim. Visibility enhancement for underwater visual slam based on underwater light scattering model. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 710–717, 2017.
- [8] Young-Shin Cho, Seong-Ho Jang, Jae-Sung Cho, Mi-Jung Kim, Hyeok Dong Lee, Sung Young Lee, and Sang-Bok Moon. Evaluation of validity and reliability of inertial measurement unit-based gait analysis systems. *Annals of rehabilitation medicine*, 42(6):872, 2018.
- [9] Michele Claus and Jan van Gemert. Videnn: Deep blind video denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [10] Khaled El Madawi, Hazem Rashed, Ahmad El Sallab, Omar Nasr, Hanan Kamel, and Senthil Yogamani. Rgb and lidar fusion based 3d semantic segmentation for autonomous driving. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 7–12. IEEE, 2019.
- [11] Yu Feng, Shaoshan Liu, and Yuhao Zhu. Real-time spatio-temporal lidar point cloud compression. In *IROS*, 2020.
- [12] Philipp Fischer, Alexey Dosovitskiy, Eddy Ilg, Philip Häusser, Caner Hazirbas, Vladimir Golkov, Patrick van der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *Proceedings of the 15th IEEE International Conference on Computer Vision*, 2015.
- [13] Chen Fu, Christoph Mertz, and John M Dolan. Lidar and monocular camera fusion: On-road depth completion for autonomous driving. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 273–278. IEEE, 2019.
- [14] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)*, 2013.
- [15] Ruben Gomez-Ojeda, Zichao Zhang, Javier Gonzalez-Jimenez, and Davide Scaramuzza. Learning-based image enhancement for visual odometry in challenging hdr environments. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 805–811, 2018.
- [16] Muhammad Haris, Gregory Shakhnarovich, and Norimichi Ukita. Recurrent back-projection network for video super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3897–3906, 2019.
- [17] E. S. Jang, M. Preda, K. Mammou, A. M. Tourapis, J. Kim, D. B. Graziosi, S. Rhyu, and M. Budagavi. Video-based point-cloud-compression standard in mpeg: From evidence collection to committee draft [standards in a nutshell]. *IEEE Signal Processing Magazine*, 2019.
- [18] Sébastien Lasserre, David Flynn, and Shouxing Qu. Using neighbouring nodes for the compression of octrees representing the geometry of point clouds. In *Proceedings of the 10th ACM Multimedia Systems Conference*, 2019.
- [19] Gregory P Meyer, Jake Charland, Darshan Hegde, Ankit Laddha, and Carlos Vallespi-Gonzalez. Sensor fusion for joint 3d object detection and semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [20] Janne Mustaniemi, Juho Kannala, Simo Särkkä, Jiri Matas, and Janne Heikkila. Gyroscope-aided motion deblurring with deep networks. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1914–1922. IEEE, 2019.
- [21] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [22] Jinay Parekh, Poojan Turakhia, Hussain Bhinderwala, and Sudhir N Dhage. A survey of image enhancement and object detection methods. *Advances in Computer, Communication and Computational Sciences*, pages 1035–1047, 2021.
- [23] Anurag Ranjan and Michael J Black. Optical flow estimation using a spatial pyramid network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4161–4170, 2017.
- [24] Yunbo Rao, Wei Yao Lin, and Leiting Chen. Image-based fusion for video enhancement of night-time surveillance. *Optical Engineering*, 49(12):120501, 2010.
- [25] Monika Roznere and A Quattrini Li. On the mutual relation between slam and image enhancement in underwater environments. In *ICRA Underwater Robotics Perception Workshop*, 2019.
- [26] Hyeonjun Sim and Munchurl Kim. A deep motion deblurring network based on per-pixel adaptive kernels with residual down-up and up-down modules. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [27] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8934–8943, 2018.

- [28] Matias Tassano, Julie Delon, and Thomas Veit. Dvdnet: A fast network for deep video denoising. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 1805–1809. IEEE, 2019.
- [29] Yapeng Tian, Yulun Zhang, Yun Fu, and Chenliang Xu. Tdan: Temporally-deformable alignment network for video super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3360–3369, 2020.
- [30] Inc. Velodyne LiDAR. Hdl-64e data sheet, 2018.
- [31] Tsun-Hsuan Wang, Hou-Ning Hu, Chieh Hubert Lin, Yi-Hsuan Tsai, Wei-Chen Chiu, and Min Sun. 3d lidar and stereo fusion using stereo matching network with conditional cost volume normalization. *arXiv preprint arXiv:1904.02917*, 2019.
- [32] Xintao Wang, Kelvin CK Chan, Ke Yu, Chao Dong, and Chen Change Loy. Edvr: Video restoration with enhanced deformable convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [33] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [34] Philippe Weinzaepfel, Jerome Revaud, Zaid Harchaoui, and Cordelia Schmid. Deepflow: Large displacement optical flow with deep matching. In *Proceedings of the IEEE international conference on computer vision*, pages 1385–1392, 2013.
- [35] Danfei Xu, Dragomir Anguelov, and Ashesh Jain. Pointfusion: Deep sensor fusion for 3d bounding box estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 244–253, 2018.
- [36] Jin Hyeok Yoo, Yecheol Kim, Ji Song Kim, and Jun Won Choi. 3d-cvf: Generating joint camera and lidar features using cross-view spatial feature fusion for 3d object detection. *arXiv preprint arXiv:2004.12636*, 3, 2020.
- [37] Zhengyou Zhang. Iterative point matching for registration of free-form curves and surfaces. *International journal of computer vision*, 13(2):119–152, 1994.
- [38] Zhihang Zhong, Ye Gao, Yinqiang Zheng, and Bo Zheng. Efficient spatio-temporal recurrent neural network for video deblurring. In *European Conference on Computer Vision*, pages 191–207. Springer, 2020.