Event-based Non-Rigid Reconstruction from Contours

Yuxuan Xue^{1,2} yuxuan.xue@tuebingen.mpg.de Haolong Li¹ haolong.li@tuebingen.mpg.de Stefan Leutenegger² stefan.leutenegger@tum.de

Joerg Stueckler¹ joerg.stueckler@tuebingen.mpg.de ¹ Max Planck Institute for Intelligent Systems, Tuebingen, Germany

² Technical University of Munich, Munich, Germany

Abstract

Visual reconstruction of fast non-rigid object deformations over time is a challenge for conventional frame-based cameras. In this paper, we propose a novel approach for reconstructing such deformations using measurements from event-based cameras. Under the assumption of a static background, where all events are generated by the motion, our approach estimates the deformation of objects from events generated at the object contour in a probabilistic optimization framework. It associates events to mesh faces on the contour and maximizes the alignment of the line of sight through the event pixel with the associated face. In experiments on synthetic and real data, we demonstrate the advantages of our method over state-of-the-art optimization and learning-based approaches for reconstructing the motion of human hands.

1 Introduction

Event cameras offer a considerable number of advantages in computer vision tasks over conventional cameras, such as low latency, high dynamic range and virtually no motion blur. Unlike conventional frame-based cameras that capture images at a fixed rate, event cameras asynchronously measure per-pixel brightness change, and output a stream of events that encode the spatio-temporal coordinates of the brightness change and its polarity. While several approaches for event-based cameras have been proposed for optical flow estimation or simultaneous localization and mapping $[\Box]$, only little work has been devoted to non-rigid reconstruction $[\Box]$, \Box].

In this paper, we present a novel non-rigid reconstruction approach for event-based cameras. Our algorithm takes event streams as input and outputs the reconstructed object pose parameters, assuming a low-dimensional parameterized shape template of a deforming object (i.e. hand and body model). We propose a novel optimization-based method based on expectation maximization (EM). Our method models event measurements at contours in a probabilistic way to estimate the association likelihood of events to mesh faces and



Figure 1: Our approach reconstructs non-rigid deformation states of objects from event streams of event-based cameras within an expectation maximization (EM) framework. In the E-step, the association probability of events to contour mesh faces is estimated from the contour measurement likelihood. In the M-step, the expected value of the measurement likelihood over the the association probability is maximized for pose parameter θ^* .

maximize the measurement likelihood. The approach is evaluated on synthetic and real data sequences, and improvements over the state-of-the-art optimization and learning-based methods for hand reconstruction are demonstrated. For generating the sequences, we develop a novel event-based camera simulator for non-rigid deforming objects. Details of our proposed simulator and a comparison with the state-of-the-art event-based camera simulators [1, 2], can be found in the supplementary material.

2 Related Work

A multitude of approaches has been proposed in recent years for scene reconstruction and rigid tracking using event-based cameras [**B**, **ID**]. Only recently, approaches have been presented that track 3D object motion [**III**]. One can distinguish state-of-the-art approaches into approaches that accumulate disparity space images based on multi-view stereo [**III**], maximize contrast of reprojected events in a reference frame [**B**, **ID**], or apply a generative event measurement model derived from the brightness constancy assumption [**B**].

Reconstruction and tracking of non-rigid shapes is a challenging problem in computer vision. For monocular frame-based cameras, several approaches have been proposed. They can be classified into methods that align shape templates (e.g. [17, 24, 51]) or approaches that use regularizing assumptions such as low-rank approximations to achieve non-rigid structure from motion (e.g. [1, 2, 2, 2]). RGB-D cameras simplify the task due to the availability of dense depth for which several methods have been proposed recently (e.g. [2, 16]).

Non-rigid reconstruction and tracking with event-based cameras has only recently attained attention in the computer vision community. Nehvi et al. [1] propose a differentiable event stream simulator by subtracting renderings of parametrized hand models. The paper demonstrates the use of the simulator for non-rigid motion tracking from event streams by optimization. Rudnev et al. [23] train a deep neural network on synthetic event streams to estimate the deformation of a MANO [22] hand model. To input the event data into the neural network, they propose to represent the data in local time windows. Different to these methods, we propose geometric contour alignment in a probabilistic optimization framework.

Some research in recent years has been devoted to event stream simulation [13, 21, 23]. We propose an event data simulator which generates synthetic events and other data modalities of human body motion, especially of hand deformation by simulating events with an adaptive sampling rate.

3 Event-based Non-Rigid Reconstruction from Contours

Our event-based method tracks parameterized non-rigid objects assuming static background. Typically, for deforming objects with low texture such as hands or human bodies, the majority of events is generated at the contour between the object and the background. Hence, we formulate the reconstruction problem in a probabilistic way using a contour measurement model for the events. Assuming a known initial state, we optimize for the pose parameters of the parametric object model incrementally from the event stream.

3.1 Expectation Maximization Framework

We formulate the 4D reconstruction problem as maximum-a-posteriori estimation of the model parameters θ given the event observations **x** from the event camera

$$\boldsymbol{\theta}^* = \operatorname*{arg\,max}_{\boldsymbol{\theta}} \ln p(\mathbf{x} \mid \boldsymbol{\theta}) + \ln p(\boldsymbol{\theta}), \tag{1}$$

where $p(\theta)$ is a constant-velocity prior on θ . The analytical formulation of the likelihood is difficult because there is no observable relation between measurement **x** and model parameters θ available. Since our background is static and the objects are textureless, most of the event measurements are generated at the contour of the deforming object. We thus assume that each event corresponds to a point on the observed contour of the object. We introduce the latent variable $z_i = j$ which represents the association between the event x_i and a mesh face of the object with index j.

We use the expectation-maximization (EM) framework to find the model parameters with the latent association,

$$\boldsymbol{\theta}^* = \arg\max_{\boldsymbol{\theta}} \sum_{i=1}^N \ln \sum_{j=1}^F p(x_i, z_i = j \mid \boldsymbol{\theta}) + \ln p(\boldsymbol{\theta}), \tag{2}$$

where *N* is the number of aggregated events in an event buffer and *F* is the number of mesh faces. In practice, we aggregate a fixed number of events into an event buffer, assume that the event observations are independent from each other, and optimize over this event buffer. The optimization of the next event buffer is initialized with the parameters from the previous buffer. In the E-step, we update a probabilistic belief on the latent association variables using a variational approximation given the current estimate of parameters $\overline{\theta}$ from the previous iteration,

$$q(z_i) \leftarrow \underset{q(z_i)}{\arg\max} \sum_{j=1}^{F} q(z_i = j) \ln \frac{p(x_i, z_i = j \mid \overline{\theta})}{q(z_i = j)}.$$
(3)

The optimal solution of this step is $q(z_i) = p(z_i | x_i, \overline{\theta})$. In the M-step, the parameters are updated by maximizing the expected log posterior with the probabilistic, i.e. soft, data association from the E-step,

$$\boldsymbol{\theta} \leftarrow \arg\max_{\boldsymbol{\theta}} \sum_{i=1}^{N} \sum_{j=1}^{F} q(z_i = j) \ln p(x_i, z_i = j \mid \boldsymbol{\theta}) + \ln p(\boldsymbol{\theta}).$$
(4)

The posterior includes terms for the expected value of the measurement likelihood under the association probability distribution and a prior term on the parameters. In the following, we explain the concrete form of the EM steps in our approach in detail.

3.2 Data Association

Ideally we can use mesh rasterization to find the association of pixels to all mesh faces that intersect the line of sight through each pixel. Due to limits in the image resolution, initial inaccuracies of the shape parameters during optimization, and complex mesh topologies which allow for multiple layers being intersected by the line of sight, the rasterization often misses the correct association of contour mesh faces with event pixels. For instance, if the shape estimate is off, the contour mesh face could be observed a few pixels off the actual event location. If fingers are bent in front of the palm, events generated on the contour of the finger might hit palm mesh faces, but miss mesh faces on the finger which generate events.

The EM-framework requires to quantify the probability of associating a mesh face with an event. Intuitively, the closer the mesh face to the event's unprojection ray, the higher the probability it causes the event. Inspired by SoftRasterizer [\square], we formulate the contour measurement likelihood as

$$p(x_i \mid z_i = j, \theta) \propto \sigma\left(\delta_j^i \frac{d_{\text{lat}}^2(i, j)}{\alpha}\right) \exp\left(-\frac{d_{\text{long}}(i, j)}{\beta}\right) \exp\left(-\frac{r_{\text{ang}}(i, j)}{\gamma}\right), \quad (5)$$

with lateral distance d_{lat} , longitudinal distance d_{long} and angular error r_{ang} between the line of sight through event x_i and the mesh face f_j , and sigmoid function σ . The angular error r_{ang} measures the deviation of the direction of the line of sight from being orthogonal to the normal of the mesh face. Hyperparameters α , β , and γ are used to control the sharpness of the individual terms for the probability distribution. The lateral distance d_{lat} is the distance between the line of sight and the closest edge of the mesh face. The sign indicator is defined as $\delta_j^i := \{+1, \text{ if } x_i \in f_j; -1, \text{ otherwise }\}$. We use a maximal lateral distance threshold to reject outlier events due to noise and unmodelled effects. As the longitudinal distance d_{long} , we determine the projected distance between the event pixel and the mesh face center on the line of sight. As sketched above, the line of sight may intersect multiple mesh face closer to the camera. The longitudinal distance gives higher likelihood to the mesh face closer to the camera. The line of sight through an event caused by the object contour should be approximately orthogonal to the normal of the corresponding mesh face. The angular error r_{ang} is thus computed by the absolute dot product between the unit direction vector of the line of sight and the face normal.

3.3 E- and M-Steps

In the E-step, the association probability is calculated from the measurement likelihood,

$$p(z_{i} = j \mid x_{i}, \theta) = \frac{p(x_{i} \mid z_{i} = j, \theta)p(z_{i} = j \mid \theta)}{\sum_{j'} p(x_{i} \mid z_{i} = j', \theta)p(z_{i} = j' \mid \theta)} = \frac{p(x_{i} \mid z_{i} = j, \theta)}{\sum_{j'} p(x_{i} \mid z_{i} = j', \theta)},$$
(6)

where we assume that the prior probability of the latent variable is uniform. For the M-step, we evaluate the measurement likelihood as

$$p(x_i, z_i = j \mid \boldsymbol{\theta}) = p(x_i \mid z_i = j, \boldsymbol{\theta}) p(z_i = j \mid \boldsymbol{\theta}) \propto \sigma\left(\delta_j^i \frac{d_{\text{lat}}^2(i, j)}{\alpha}\right) \exp\left(-\frac{r_{\text{ang}}(i, j)}{\gamma}\right).$$
(7)

Here we do not include the longitudinal distance term. Ideally, this term should assign a constant probability to mesh faces on the same occlusion layer. Notably the term does not depend continuously on the shape parameters. If included in the M-step, our approximative Gaussian term for the E-step would falsely incentivize shape parameters for which the mesh intersects the line of sight closer to camera. The angular error term in M-step encourages the alignment of events with contours. For scenes with many outlier events (e.g. textured objects), we choose a larger value of γ . The prior term for the M-step is a constant velocity prior on the parameters, i.e., $\ln p(\theta) = k ||\mathbf{v} - \mathbf{v}'||_2^2$, $\mathbf{v} = \frac{\theta - \theta'}{\Delta t}$, where θ' and \mathbf{v}' are the parameters and velocity for the previous event buffer and Δt is the time difference between the two event buffers. We alternate E-step and M-step until convergence. When a new event buffer is available, we initialize θ based on the current estimate of \mathbf{v} .

4 Experiments

We evaluate and demonstrate our event-based non-rigid reconstruction approach on synthetic and real sequences using MANO and SMPL-X object models and involving random motions and various background textures. We provide qualitative and quantitative results, comparing with state-of-the-art baselines. An evaluation of the robustness against noisy events and initial poses, an ablation study for the terms in our E- and M-steps, and results for a hard-EM variant are given in the supplementary material. Please also refer to the supplemental video for qualitative results.

4.1 Experiment Setup

Implementation details To compensate between the accuracy and the efficiency, we accumulate events in event buffers and optimize a single set of shape parameters for the whole event buffer. Similar to [22], we accumulate buffers with a fixed number of events, therefore choosing their temporal length adaptively. For our real captured data sequences, we accumulate 100 events per buffer. For synthetic data generated by our simulator, we stack 300 events into each buffer. Our simulator simulates the Prophesee camera. We use the pinhole camera model for the event camera and assume the camera intrinsics are calibrated. We have several hyperparameters in our framework. Please refer to the supplementary file for the tuning process of hyperparameters. Our algorithm optimizes for the pose parameters of the MANO hand model [22] or the SMPL body model [13, 13]. In case of the hand model, the pose parameters are in PCA space and the MANO modelling approach reconstructs the vertex offsets, which are used together with the canonical pose vertices to generate the posed mesh. For the SMPL model, the pose parameters are the orientation parameters and Linear Blend Skinning (LBS) is used to recover the posed mesh.

Datasets We generate synthetic datasets of sequences with three types of different objects, namely the MANO [22] hand (Fig. 3a), the SMPL-X [13, 13] hand (Fig. 4a), and the

Scenario	Method	mean MPJPE (mm)	median MPJPE (mm)
MANO hand	Nehvi et al. [🖪]	11.61	10.85
	Ours	4.52	4.27
SMPL-X hand	Rudnev et al. [🔼]	11.88	10.73
	Ours	1.11	0.76
SMPL-X arm & hand	Ours	15.39	3.93

6 XUE, LI, LEUTENEGGER, STUECKLER: EVENT-BASED NON-RIGID RECONSTRUCTION

Table 1: Results on synthetic sequences of different objects

combined SMPL-X [I], I] arm and hand (5a). MANO hand sequences are generated with a single hand mesh at a fixed position and orientation. We vary the full 45-dimensional pose parameter space to generate varying hand poses. For SMPL-X hand sequences, the hand is attached to the whole human body which prevents observing the inside of the hand mesh. We vary the first 6 principal component pose parameters to simulate time-varying and realistic hand deformations. In the SMPL-X arm and hand sequences, we synthesize the arm motion by the 3-DoF rotation of the elbow joint and the hand motion by the 6 principal pose parameters. We use a custom event simulator with adaptive sampling rate to generate the synthetic sequences for the different object models. A detailed description of the event simulator can be found in the supplementary material. We simulate a Prophesee camera with the image size of 1280×720 . The event contrast threshold is 0.5. For each sequence, the background image is randomly chosen from a texture-rich indoor scene in the YCB video dataset [29]. To introduce noise into the event generation process, we sample the contrast threshold of each pixel from a Gaussian distribution with standard deviation 0.0004. The threshold of salt-and-pepper noise is 10^{-5} . For further details on the noise generation process in our simulator, please refer to the supplementary material.

Evaluation metrics For the synthetic data, 3D ground-truth positions for all joints and mesh vertices as well as pose parameters are known. We evaluate using the Mean Per Joint Position Error (MPJPE [23]), the percentage of correct 3D Joints (3D-PCK [13]), and the area under the PCK-curve (AUC [13]) with thresholds ranging from 0 to 50 mm. For hand sequences, we consider the 15 hand skeleton joints. For arm and hand sequences, only the forearm and the hand have motion. Thus, we consider one wrist joint and 15 hand joints.

4.2 Quantitative Evaluation

We compare our approach quantitatively with the state-of-the-art event-based non-rigid object reconstruction methods: Nehvi's optimization-based approach [13] evaluates using the MANO hand model, while Rudnev's approach [23] is designed for the SMPL-X hand model. To the best of our knowledge, previous event-based reconstruction approaches have not been demonstrated on combined arm and hand motion of a SMPL-X model. Hence, we only provide results for our method on these sequences.

For synthetic MANO hand sequences, we use the MANO [22] hand model as the parametric mesh template. In experiments, we initialize the optimized parameters with the ground-truth pose parameters and evaluate, how well the approach can keep track of the hand deformation. Our approach reconstructs the 45-dimensional pose parameter. We report quantitative results MPJPE and AUC on these sequences in Table 1. We also compare our method to the state-of-the-art event-based hand tracking approach [13] for the MANO



Figure 3: Qualitative reconstruction results on synthetic MANO hand sequences.

model. Similar to our approach, Nehvi's method is optimization-based and requires the initial parameters of the mesh template. To ensure a fair comparison, we use Optuna [II] to tune hyperparameters in Nehvi's method and our method. We observe that our approach is about 2.5-times more accurate than Nehvi's method. We also show the 3D-PCK curve of both approaches in Fig. 2a. Apparently, our method has higher AUC than Nehvi's method. Results in Table 1 and Fig. 2a demonstrate that our method outperforms Nehvi's method clearly.

In the SMPL-X hand sequences, the hand is attached to a human body model. Here, our approach reconstructs the 6-dimensional pose parameters, which is consistent with the evaluation conducted in Rudnev's method [23]. We report quantitative results in Table 1. It can be seen that our approach achieves better performance than Rudnev's method [23]



Figure 4: Qualitative reconstruction results on SMPL-X hand sequences.

8 XUE, LI, LEUTENEGGER, STUECKLER: EVENT-BASED NON-RIGID RECONSTRUCTION



(c) Ours

Figure 5: Qualitative reconstruction results on SMPL-X arm and hand sequences.

in MPJPE. Rudnev's method is learning-based and does not require the knowledge of the initial pose parameters. We use the network trained by [\square] which is limited to the resolution (240×180) of the DAVIS 240C camera. Thus, we simulate event streams of the same motion with the intrinsics provided in [\square] for Rudnev's method. We observe that since the global rotation and translation are fixed, events are only generated where the deformation occurs. Rudnev's method seems to perform less well in this case than our approach.

Finally, we evaluate the performance of our approach on sequences which combine arm and hand motion using the SMPL-X model. In the synthetic data generation process, we vary 6 principal parameters to synthesize hand poses and the 3 rotation parameters of the elbow joint. Our approach jointly optimize these hand and elbow parameters. The median MPJPE in Tab. 1 demonstrates that our approach can reconstruct the motion of the arm and hand with high accuracy. The mean MPJPE is higher than the median MPJPE due to failures in some sequences. We show failure cases and their analysis in the supplementary material. The difference in accuracy to the SMPL-X hand sequences can be explained by the fact that for the SMPL-X arm & hand sequences, also the elbow joint needs to be reconstructed. Moreover, the hand is visible on different scales in the image (the hand is smaller for SMPL-X arm & hand). Hence, the absolute error in mm becomes higher. As an incremental optimization-based approach, our approach can also drift, but it can snap the mesh silhouette to the observed events on the contour if sufficient observations are available. In the supplementary material, we provide a plot of median error over time for the MANO hand dataset.

Differently to [II], [I], our implementation is not real-time capable, due to the complete evaluation of all event measurement likelihoods for all mesh faces in the soft E-step. Depending on the motion, a sequence can be split into 50 to 500 event buffers. For each event buffer, the current average run-time of the method is 8.76 seconds on MANO hand sequences, and 50.72 seconds on SMPL-X arm and hand sequences. Our current implementation in PyTorch. In future work, the optimization process could be implemented more efficiently by associating mesh faces with a local search, tailor code with CUDA/C++, and using second-order Gauss-Newton methods instead of the current gradient-descent algorithm. In the supplementary material, we also provide evidence that hard data association with the most likely mesh face in the E-step only slightly affects accuracy. Due to the non-convexity of the problem, our approach needs a sufficiently good initial guess of the pose. In the supplementary material, we evaluate reconstruction accuracy vs. varying noise levels for the initial pose. We also evaluate the effect of varying noise in the events on accuracy in the

XUE, LI, LEUTENEGGER, STUECKLER: EVENT-BASED NON-RIGID RECONSTRUCTION 9



Figure 6: Qualitative results on a real event sequence from a DAVIS240C camera. Lin et al. [1] infer MANO pose parameters from intensity images; Rudnev et al. [2] infer 6 principal MANO pose parameters; Nehvi et al. [1] and our approach optimize 45 MANO pose parameters. Our approach recovers the deformation most similar to the ground truth.

supplementary material.

4.3 Qualitative Evaluation

Synthetic Data We show qualitative results of our approach and state-of-the-art baseline approaches [13, 23] on synthetic sequences. For each object, the ground-truth RGB images, accumulated events during the motion, and reconstruction results are shown. We crop all images with a fixed ratio to increase the view of the objects. Results on synthetic MANO hand sequences of our approach and Nehvi's approach [13] are shown in Figs. 3c and 3d, respectively. It can be observed that our approach reconstructs the deformation of the hand well, while Nehvi's approach struggles to track the hand pose accurately. Note that Nehvis method does assume black background and generatively models the specific log intensity changes induced at the optical flow at contours. Our approach only assumes that events are generated by contours without explicit dependency on the optical flow, hence, it is more robust to textured backgrounds. We compare our approach with Rudnev's method [23] in Fig. 4d on a SMPL-X hand deformation sequence. While our approach can reconstruct the hand motion well, Rudnev's approach performs less accurately. For the sequences with combined arm and hand motion of the SMPL-X model, we show qualitative results of our approach in Fig. 5c. Our proposed approach can reconstruct the motion well.

Real Data In Fig. 6, we also show qualitative results of our approach with the MANO hand model on real sequences with hand motion captured with a DAVIS240C camera. The camera also records grayscale intensity frames for reference. Since our approach requires

an initialization of the hand pose parameters, we use [I] on the first image frame and set rotation and translation manually, since the pretrained model did not yield proper poses on the DAVIS gray scale images. Further details on the initialization procedure are provided in the supplementary material. We compare our approach qualitatively with state-of-the-art image-based (MeshGraphormer [I]) and event-based [I], [] methods. MeshGraphormer is a learning-based approach which predicts MANO pose parameters from grayscale images. It has solid reconstruction performance for slower motions, but suffers from motion blur for fast motions. Furthermore, the temporal resolution of the reconstruction result is limited by the frequency of the frames. Compared to the result of MeshGraphormer [I] and event-based approaches [I], [], our approach follows the ground-truth reference more closely.

4.4 Assumptions and Limitations

Our approach uses a loose coupling of frames and events by initializing the optimization from the gray-scale frame. A possible direction of future work is to extend the method by feeding the frame-based information at a specific lower rate and use the events to estimate pose between frames in a tightly-coupled joint optimization framework. In our experiments, self-occlusions occur within the hand (for instance between fingers, or fingers and the palm, see also Fig. 3). The more self occlusions, the more unconstrained the pose parameters get due to the partial observations and low number of events. The constant velocity model and the PCA subspace of MANO can help to regularize the motion in this partially constrained setting. Our method relies on events on the contour and cannot estimate deformation if there are little contour events due to similar background color or insufficient motion. Due to the image projection, the contour information seems not sufficient yet for reconstructing shape parameters concurrently with rotation and translation of the objects with our formulation. To address challenging settings like 6D pose estimation or crossing hands in future work, one could for instance investigate including learned temporal priors, texture-based cues, or combining events with frames in a joint optimization framework.

5 Conclusion

We present a novel non-rigid reconstruction approach for event cameras. Our approach formulates the reconstruction problem as an expectation-maximization problem. Events are associated to observed contours on parametrized mesh models and an alignment objective is maximized to fit the mesh parameters with event measurements. Our method outperforms qualitatively and quantitatively state-of-the-art event-based non-rigid reconstruction approaches [13, 23]. We also demonstrate that our proposed approach is robust to noisy events and initial parameter estimates. In future work, texture-based reconstruction from events and frames could be combined with our approach or the run-time of our implementation could be improved by searching for correspondences efficiently.

Acknowledgement

This work was supported by Cyber Valley and the Max Planck Society. The authors thank the Empirical Inference Department at Max Planck Institute for Intelligent Systems for providing the DAVIS event camera.

References

- [1] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In Ankur Teredesai, Vipin Kumar, Ying Li, Rómer Rosales, Evimaria Terzi, and George Karypis, editors, *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4-8, 2019,* pages 2623–2631. ACM, 2019. doi: 10.1145/3292500.3330701. URL https: //doi.org/10.1145/3292500.3330701.
- [2] Aljaz Bozic, Pablo R. Palafox, Michael Zollhöfer, Angela Dai, Justus Thies, and Matthias Nießner. Neural non-rigid tracking. In Advances in Neural Information Processing Systems (NeurIPS), 2020.
- [3] Christoph Bregler, Aaron Hertzmann, and Henning Biermann. Recovering non-rigid 3d shape from image streams. In *Conference on Computer Vision and Pattern Recognition* (*CVPR*), pages 2690–2696, 2000.
- [4] Yuchao Dai, Hongdong Li, and Mingyi He. A simple prior-free method for non-rigid structure-from-motion factorization. *Int. J. Comput. Vis.*, 107(2):101–122, 2014.
- [5] Guillermo Gallego, Mathias Gehrig, and Davide Scaramuzza. Focus is all you need: Loss functions for event-based vision. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12280–12289, 2019.
- [6] Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew J. Davison, Jörg Conradt, Kostas Daniilidis, and Davide Scaramuzza. Event-based vision: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(1):154–180, 2022. doi: 10.1109/TPAMI.2020.3008413. URL https://doi.org/10.1109/TPAMI.2020.3008413.
- [7] Ravi Garg, Anastasios Roussos, and Lourdes Agapito. Dense variational reconstruction of non-rigid surfaces from monocular video. In *IEEE Conference on Computer Vision* and Pattern Recognition (CVPR), pages 1272–1279, 2013.
- [8] Hanme Kim, Stefan Leutenegger, and Andrew J. Davison. Real-time 3d reconstruction and 6-dof tracking with an event camera. In *European Conference on Computer Vision* (ECCV), pages 349–364, 2016.
- [9] Jose Lamarca, Shaifali Parashar, Adrien Bartoli, and J. M. M. Montiel. Defslam: Tracking and mapping of deforming scenes from monocular sequences. *IEEE Trans. Robotics*, 37(1):291–303, 2021.
- [10] Haolong Li and Joerg Stueckler. Tracking 6-dof object motion from events and frames. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 14171–14177, 2021.
- Kevin Lin, Lijuan Wang, and Zicheng Liu. Mesh graphormer. In 2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021, pages 12919–12928. IEEE, 2021. doi: 10.1109/ICCV48922. 2021.01270. URL https://doi.org/10.1109/ICCV48922.2021.01270.

12 XUE, LI, LEUTENEGGER, STUECKLER: EVENT-BASED NON-RIGID RECONSTRUCTION

- [12] Shichen Liu, Weikai Chen, Tianye Li, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pages 7707–7716, 2019. doi: 10.1109/ICCV.2019.00780.
- [13] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. ACM Trans. Graphics (Proc. SIGGRAPH Asia), 34(6):248:1–248:16, October 2015.
- [14] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. In 3D Vision (3DV), 2017 Fifth International Conference on. IEEE, 2017. doi: 10.1109/3dv.2017.00064. URL http://gvv. mpi-inf.mpg.de/3dhp_dataset.
- [15] Jalees Nehvi, Vladislav Golyanik, Franziska Mueller, Hans-Peter Seidel, Mohamed Elgharib, and Christian Theobalt. Differentiable event stream simulator for non-rigid 3d tracking. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1302–1311, 2021.
- [16] Richard A. Newcombe, Dieter Fox, and Steven M. Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 343–352, 2015.
- [17] Dat Tien Ngo, Jonas Östlund, and Pascal Fua. Template-based monocular 3d shape recovery using laplacian meshes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 38(1):172– 187, 2016.
- [18] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June* 16-20, 2019, pages 10975–10985. Computer Vision Foundation / IEEE, 2019. doi: 10.1109/CVPR.2019.01123. URL http://openaccess.thecvf.com/ content_CVPR_2019/html/Pavlakos_Expressive_Body_Capture_ 3D_Hands_Face_and_Body_From_a_CVPR_2019_paper.html.
- [19] Henri Rebecq, Timo Horstschaefer, and Davide Scaramuzza. Real-time visual-inertial odometry for event cameras using keyframe-based nonlinear optimization. In *British Machine Vision Conference (BMVC)*, 2017.
- [20] Henri Rebecq, Guillermo Gallego, Elias Mueggler, and Davide Scaramuzza. EMVS: event-based multi-view stereo - 3d reconstruction with an event camera in real-time. *Int. J. Comput. Vis.*, 126(12):1394–1414, 2018.
- [21] Henri Rebecq, Daniel Gehrig, and Davide Scaramuzza. ESIM: an open event camera simulator. In 2nd Annual Conference on Robot Learning, CoRL 2018, Zürich, Switzerland, 29-31 October 2018, Proceedings, volume 87 of Proceedings of Machine Learning Research, pages 969–982. PMLR, 2018. URL http://proceedings.mlr.press/v87/rebecq18a.html.

XUE, LI, LEUTENEGGER, STUECKLER: EVENT-BASED NON-RIGID RECONSTRUCTION 13

- [22] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: modeling and capturing hands and bodies together. ACM Trans. Graph., 36(6):245:1–245:17, 2017. doi: 10.1145/3130800.3130883. URL https://doi.org/10.1145/3130800.3130883.
- [23] Viktor Rudnev, Vladislav Golyanik, Jiayi Wang, Hans-Peter Seidel, Franziska Mueller, Mohamed Elgharib, and Christian Theobalt. Eventhands: Real-time neural 3d hand pose estimation from an event stream. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12365–12375, 2021.
- [24] Mathieu Salzmann and Pascal Fua. Reconstructing sharply folding surfaces: A convex formulation. In *IEEE Conference on Computer Vision and Pattern Recognition* (*CVPR*), pages 1054–1061, 2009.
- [25] Vikramjit Sidhu, Edgar Tretschk, Vladislav Golyanik, Antonio Agudo, and Christian Theobalt. Neural dense non-rigid structure from motion with latent space constraints. In *European Conference on Computer Vision (ECCV)*, volume 12361, pages 204–222, 2020.
- [26] Timo Stoffregen and Lindsay Kleeman. Event cameras, contrast maximization and reward functions: An analysis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12300–12308, 2019.
- [27] Antoni Rosinol Vidal, Henri Rebecq, Timo Horstschaefer, and Davide Scaramuzza. Ultimate slam? combining events, images, and IMU for robust visual SLAM in HDR and high-speed scenarios. *IEEE Robotics Autom. Lett.*, 3(2):994–1001, 2018.
- [28] Timo von Marcard, Roberto Henschel, Michael Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *European Conference on Computer Vision (ECCV)*, sep 2018.
- [29] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. In Hadas Kress-Gazit, Siddhartha S. Srinivasa, Tom Howard, and Nikolay Atanasov, editors, *Robotics: Science and Systems XIV, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, June 26-30, 2018*, 2018. doi: 10.15607/RSS.2018.XIV.019. URL http://www.roboticsproceedings.org/rss14/p19.html.
- [30] Rui Yu, Chris Russell, Neill D. F. Campbell, and Lourdes Agapito. Direct, dense, and deformable: Template-based non-rigid 3d reconstruction from RGB video. In *IEEE International Conference on Computer Vision (ICCV)*, pages 918–926, 2015.