Supplementary Material
Ev-TTA: Test-Time Adaptation for Event-Based Object Recognition

Junho Kim[1], Inwoo Hwang[1], and Young Min Kim[1,2,*]
[1]Department of Electrical and Computer Engineering, Seoul National University
[2]Interdisciplinary Program in Artificial Intelligence and INMC, Seoul National University

A. Hypothesis Testing

We handle the extreme noise in event data under low-light conditions with conditional denoising, which considers spatial consistency, as explained in Section 3.2. In this section, we further elaborate the hypothesis testing procedure. Given a batch of size N containing events in the target domain, we obtain the transformed event ratios $T(R_i)$ for $i = 1, 2, \ldots, N$ using the source domain statistics, as described in Equation (6) in the main paper. If the transformed ratios follow a standard Gaussian distribution, we can assume that the event measurement is free from noise burst.

To this end, we first calculate the batch-wise mean $\hat{\mu}$ and standard deviation $\hat{\sigma}$ of the transformed ratios $T(R_i)$. We then determine if the event ratio $R = N_{\text{pos}}/N_{\text{neg}}$ is either too large (noise burst in positive channel) or too small (noise burst in negative channel). Specifically, we apply the standard one-tailed z-test procedure [19], and label the batch as containing noise burst in the positive channel if the following inequality holds,

$$\Phi(\frac{\sqrt{N}|\hat{\mu} - \mu_{\text{thres}}|}{\hat{\sigma}}) > 0.9, \tag{1}$$

where $\Phi(\cdot)$ is the cumulative distribution function (CDF) of the standard Gaussian. Here, μ_{thres} is the threshold value for separating noise bursts, which we set to 0.25 in all our experiments. However, the choice of μ_{thres} does not have a significant impact in performance. Table A.1 verifies that the accuracy of the timestamp image [11] in validation #6, 7 from N-ImageNet is stable for various values of μ_{thres}. The criterion for determining noise burst in the negative channel is similarly defined as follows,

$$\Phi(\frac{\sqrt{N}|\hat{\mu} + \mu_{\text{thres}}|}{\hat{\sigma}}) < 0.1, \tag{2}$$

where the signs of variables in the inequality are reversed.

*Young Min Kim is the corresponding author.

μ_{thres}	0.25	0.5	0.75	1.00
Validation 6	29.20	29.39	29.13	27.57
Validation 7	38.46	37.36	37.19	37.21

Table A.1. Effect of the threshold value μ_{thres} used in hypothesis testing on the classifier performance in N-ImageNet [6].

Day	Recording
Source	2019-06-19
Day 1	2019-02-22
Day 2	2019-06-11
Day 3	2019-06-14
Day 4	2019-02-21
Day 5	2019-06-26

Table C.1. Conversion table for the Prophesee Megapixel Dataset [12] on days used in the main paper and the corresponding recordings.

B. Hyperparameter Setup

In this section we report the hyperparameters used for Ev-TTA. We mostly follow the hyperparameter setup of Tent [20] and avoid tuning the algorithm on the test set. For all experiments, we use the Adam optimizer [7]. In N-ImageNet experiments, we use a learning rate of 0.00025 with a batch size of 64, while for other datasets with smaller number of labels, we use a learning rate of 0.001 with a batch size of 128. For steering angle prediction, we use a learning rate of 0.000025 with a batch size of 64, as larger learning rates failed to converge. We employ the identical hyperparameter setup for baselines used throughout our experiments.

C. Dataset Preparation

In this section we explain the preprocessing pipelines used in datasets for our experiments.

Prophesee Megapixel Dataset For evaluating Ev-TTA in real-world environments, we use the Prophesee Megapixel

Day	Scene Type	Time	Recording
Source	City	Day	`rec1487779465`
Day 1	Freeway	Evening	`rec1487608147`
Day 2	City	Night	`rec1487355090`
Day 3	Town	Day	`rec1487856408`
Day 4	City	Day	`rec1487842276`

Table C.2. Conversion table for the DDD17 Dataset [2] on days used in the main paper and the corresponding recordings.

Method	No Adaptation	Min Entropy	Majority Vote	Random (Ours)
Accuracy	33.37	43.77	43.74	43.47

Table D.1. Ablation study on anchor event selection. We report the average accuracy of the timestamp image [11] on the N-ImageNet variants.

Dataset [12] in Section 4.1. Due to the immense size of the dataset, we select six recordings for our experiments, where the exact filename of each recording is specified in Table C.1. We further use the Prophesee Automotive Dataset Toolbox [12] to parse the bounding boxes and collect approximately 9000 bounding boxes for three classes (car, truck, bus). We discard other four classes (twowheeler, pedestrian, traffic light, traffic sign) in the dataset because the object bounding boxes are often too small and the class labels are not as frequent.

SimN-ImageNet To evaluate Ev-TTA for reducing sim2real gap, we generate SimN-ImageNet, which is a simulated version of N-ImageNet [6]. We use the event camera simulator Vid2E [4, 14] to generate synthetic events from a virtual event camera moving around images from ImageNet [15]. The event camera resolution was set to 480×640, to match the resolution of the Samsung DVS camera [16] used for creating N-ImageNet. Due to the large size of ImageNet [15], generating SimN-ImageNet using Vid2E [4, 14] takes approximately nine days on a configuration of eight 2080Ti GPUs.

DDD17 Dataset For assessing the extension of Ev-TTA to regression tasks, we use the DDD17 dataset [2] which is a dataset targeted for steering angle prediction. We select five recordings for our experiments, where the exact filenmae of each recording is specified in Table C.2. We further use the preprocessing toolkit provided by the authors [2,5] to obtain event histograms [9] from raw event data.

D. Additional Ablation Study

Anchor Event Selection We report the impact of choosing the anchor event for optimizing the prediction similarity loss and selective entropy loss in Section 3.1. Recall that in Section 3.1 we choose the anchor event as a random event slice. To validate our design choice, we implement two additional variants of Ev-TTA where the anchor is chosen more deliberately. The first variant (Min Entropy) uses the event slice with the smallest prediction entropy as the anchor. The second variant (Majority Vote) uses the event slice whose predicted class label is equal to the majority vote of the K event slices. We report the average performance of the timestamp image [11] on the N-ImageNet [6] variants under the various anchor selection schemes. As shown in Table D.1, only a small amount of performance gain exists from using deliberate anchor selection schemes. Therefore, the random selection scheme suffices for successful adaptation.

Method	No Adaptation	Ev-TTA	Augmentation
Accuracy	33.37	43.47	41.09

Table D.2. Ablation study on using event slices. We report the average accuracy of the timestamp image [11] on the N-ImageNet variants.

Using Event Slices for Adaptation We validate the use of multiple event slices for the prediction similarity loss and selective entropy loss in Section 3.1. To this end, we implement a variant of Ev-TTA that applies data augmentation to a single event slice, similar to SENTRY [13]. Instead of enforcing consistency on predictions among event slices, this variant applies the same loss formulation among augmented events. We employ three augmentation schemes: horizontal flipping, polarity flipping, and temporal flipping. Horizontal flipping is where the input event is flipped along the spatial dimension horizontally, and polarity flipping is where the event polarities are inverted. Temporal flipping is where the timestamps of the input event are reversed, similar to Tulyakov *et al.* [18]. The performance comparison between Ev-TTA and the augmentation-based variant is made on N-ImageNet [6] using the timestamp image [11] as input. As shown in Table D.2, the average accuracy is higher for Ev-TTA that uses event slices to impose temporal consistency. The design choice of using event slices instead of data augmentation leads to effective adaptation.

E. Full Evaluation Results in N-ImageNet

In this section, we report the full evaluation results of various event representations on N-ImageNet [6]. Ev-TTA shows large amount of performance improvement compared to the baselines [10, 13, 17] in all tested representations both online and offline. The results in Table E.1∼12 is the accuracy for six event representations, namely: binary event image [3], event histogram [9], timestamp image [11], time surface [8], sorted time surface [1], and DiST [6]. We provide the individual accuracy for each representation.

Change	None	Trajectory					Brightness				Average
Validation Dataset	Orig.	1	2	3	4	5	6	7	8	9	All
No Adaptation	45.86	43.01	33.62	39.47	25.39	36.23	21.16	30.02	36.52	34.92	33.37
Mummadi *et al.* [10]	-	44.90	45.25	45.45	42.66	43.95	24.27	33.84	45.00	44.52	41.09
URIE [17]	-	41.68	39.77	42.28	38.30	39.42	17.68	30.95	39.38	41.90	36.82
SENTRY [13]	-	45.90	45.10	45.72	41.93	43.96	20.06	33.94	44.87	44.44	40.66
Tent [20]	-	42.36	43.93	43.94	41.01	41.73	25.21	34.62	43.40	42.97	39.91
Ev-TTA	-	**47.15**	**46.94**	**46.58**	**44.03**	**45.66**	**29.20**	**38.45**	**47.12**	**46.12**	**43.47**

Table E.1. Offline evaluation results of timestamp image [11] on N-ImageNet [6] and its variants.

Change	None	Trajectory					Brightness				Average
Validation Dataset	Orig.	1	2	3	4	5	6	7	8	9	All
No Adaptation	45.86	43.01	33.62	39.47	25.39	36.23	21.16	30.02	36.52	34.92	33.37
Mummadi *et al.* [10]	-	42.60	43.06	43.39	40.36	41.29	24.93	33.55	42.39	42.27	39.32
URIE [17]	-	39.12	38.10	39.74	36.69	37.48	18.54	28.32	38.27	39.07	35.04
SENTRY [13]	-	42.22	42.45	43.21	39.44	40.96	20.48	31.38	41.42	41.71	38.14
Tent [20]	-	41.30	42.41	42.55	39.50	40.26	24.07	33.21	41.65	41.34	38.48
Ev-TTA	-	**43.86**	**43.91**	**44.33**	**41.16**	**42.45**	**25.86**	**34.78**	**43.84**	**43.37**	**40.40**

Table E.2. Online evaluation results of timestamp image [11] on N-ImageNet [6] and its variants.

Change	None	Trajectory					Brightness				Average
Validation Dataset	Orig.	1	2	3	4	5	6	7	8	9	All
No Adaptation	47.73	43.73	33.72	37.69	24.56	35.24	20.89	29.68	36.33	34.56	32.93
Mummadi *et al.* [10]	-	46.99	46.38	45.71	42.92	44.79	28.26	36.54	45.35	45.12	42.45
URIE [17]	-	45.08	44.36	44.18	40.40	42.48	23.71	34.48	43.77	42.99	40.16
SENTRY [13]	-	47.06	48.01	45.75	41.97	45.06	24.60	35.48	45.06	44.91	41.99
Tent [20]	-	44.88	45.00	44.20	41.31	43.11	26.94	34.65	43.75	43.57	40.82
Ev-TTA	-	**48.64**	**48.01**	**47.24**	**44.49**	**47.06**	**30.08**	**38.34**	**47.37**	**46.58**	**44.20**

Table E.3. Offline evaluation results of event histogram [9] on N-ImageNet [6] and its variants.

Change	None	Trajectory					Brightness				Average
Validation Dataset	Orig.	1	2	3	4	5	6	7	8	9	All
No Adaptation	47.73	43.73	33.72	37.69	24.56	35.24	20.89	29.68	36.33	34.56	32.93
Mummadi *et al.* [10]	-	43.71	43.67	43.20	40.33	42.54	25.65	33.66	42.55	42.76	39.79
URIE [17]	-	41.94	42.16	42.10	38.67	41.10	23.21	31.90	40.97	41.20	38.14
SENTRY [13]	-	43.31	42.77	42.78	39.33	41.68	23.20	32.36	41.91	41.86	38.80
Tent [20]	-	42.69	42.93	42.56	39.61	41.79	25.07	32.83	41.68	41.82	39.00
Ev-TTA	-	**44.94**	**44.63**	**43.31**	**41.48**	**43.46**	**26.89**	**34.71**	**43.86**	**43.42**	**40.86**

Table E.4. Online evaluation results of event histogram [9] on N-ImageNet [6] and its variants.

Change	None	Trajectory					Brightness				Average
Validation Dataset	Orig.	1	2	3	4	5	6	7	8	9	All
No Adaptation	46.36	42.68	30.68	37.74	22.99	34.74	19.00	27.85	34.03	32.08	31.31
Mummadi *et al.* [10]	-	46.07	45.02	44.94	42.35	43.95	22.90	31.58	44.66	45.50	40.77
URIE [17]	-	42.63	39.30	42.74	37.28	41.30	14.58	25.76	42.23	42.53	36.48
SENTRY [13]	-	46.43	44.27	44.39	40.20	43.56	18.54	31.94	43.69	43.52	39.62
Tent [20]	-	43.16	43.51	43.11	40.47	42.21	25.33	33.28	42.91	43.90	39.76
Ev-TTA	-	**48.51**	**46.46**	**47.01**	**43.48**	**47.10**	**29.08**	**38.39**	**46.72**	**46.76**	**43.72**

Table E.5. Offline evaluation results of binary event image [3] on N-ImageNet [6] and its variants.

Change	None	Trajectory					Brightness				Average
Validation Dataset	Orig.	1	2	3	4	5	6	7	8	9	All
No Adaptation	46.36	42.68	30.68	37.74	22.99	34.74	19.00	27.85	34.03	32.08	31.31
Mummadi *et al.* [10]	-	43.61	42.63	42.65	40.14	41.80	23.63	32.45	42.27	42.77	39.11
URIE [17]	-	41.20	39.49	41.15	37.01	40.01	19.89	28.12	40.89	40.54	36.48
SENTRY [13]	-	42.99	41.75	42.05	38.14	41.05	19.52	30.39	41.00	41.58	37.61
Tent [20]	-	42.07	41.78	41.64	39.12	40.89	24.05	31.97	41.44	41.94	38.32
Ev-TTA	-	**44.97**	**43.73**	**43.89**	**40.85**	**43.34**	**25.42**	**34.65**	**43.68**	**43.80**	**40.48**

Table E.6. Online evaluation results of binary event image [3] on N-ImageNet [6] and its variants.

Change	None	Trajectory					Brightness				Average
Validation Dataset	Orig.	1	2	3	4	5	6	7	8	9	All
No Adaptation	44.32	41.01	34.63	40.00	25.48	34.89	22.12	31.27	37.12	35.36	33.54
Mummadi *et al.* [10]	-	44.40	44.85	46.56	43.05	42.96	24.05	34.18	45.56	44.76	41.15
URIE [17]	-	36.21	38.20	36.76	34.42	37.85	10.74	24.44	38.37	38.25	32.80
SENTRY [13]	-	44.42	46.63	47.02	42.27	42.51	21.00	35.13	45.90	45.34	41.14
Tent [20]	-	41.77	45.23	45.26	41.69	41.36	26.03	34.64	43.97	43.71	40.41
Ev-TTA	-	**45.50**	**47.42**	**47.24**	**44.27**	**43.87**	**27.28**	**37.06**	**47.05**	**46.54**	**42.91**

Table E.7. Offline evaluation results of time surface [8] on N-ImageNet [6] and its variants.

Change	None	Trajectory					Brightness				Average
Validation Dataset	Orig.	1	2	3	4	5	6	7	8	9	All
No Adaptation	44.32	41.01	34.63	40.00	25.48	34.89	22.12	31.27	37.12	35.36	33.54
Mummadi *et al.* [10]	-	41.03	44.17	45.01	41.01	40.43	25.07	33.97	43.33	43.28	39.70
URIE [17]	-	34.24	34.71	35.11	30.76	33.13	12.50	21.88	33.83	32.59	29.86
SENTRY [13]	-	40.63	43.79	44.62	39.62	39.00	21.49	32.78	42.74	42.61	38.59
Tent [20]	-	39.77	43.60	44.23	40.20	39.36	25.13	33.33	42.50	42.39	38.95
Ev-TTA	-	**42.51**	**45.18**	**45.29**	**41.37**	**40.97**	**25.68**	**35.25**	**44.15**	**43.88**	**40.48**

Table E.8. Online evaluation results of time surface [8] on N-ImageNet [6] and its variants.

Change	None	Trajectory					Brightness				Average
Validation Dataset	Orig.	1	2	3	4	5	6	7	8	9	All
No Adaptation	47.90	44.33	33.50	40.17	23.72	37.19	21.57	30.31	36.63	35.18	33.62
Mummadi *et al.* [10]	-	47.26	47.35	47.47	44.29	45.64	25.56	36.63	46.60	46.20	43.00
URIE [17]	-	42.87	43.76	44.90	40.50	40.82	22.05	33.55	43.37	41.70	39.28
SENTRY [13]	-	47.66	47.32	47.45	42.55	45.25	22.04	34.66	46.12	45.84	42.10
Tent [20]	-	44.65	45.94	45.78	42.10	43.91	27.12	35.11	44.96	44.55	41.57
Ev-TTA	-	**49.58**	**47.67**	**48.36**	**45.59**	**46.72**	**30.07**	**39.30**	**48.24**	**47.76**	**44.81**

Table E.9. Offline evaluation results of sorted time surface [1] on N-ImageNet [6] and its variants.

Change	None	Trajectory					Brightness				Average
Validation Dataset	Orig.	1	2	3	4	5	6	7	8	9	All
No Adaptation	47.90	44.33	33.50	40.17	23.72	37.19	21.57	30.31	36.63	35.18	33.62
Mummadi *et al.* [10]	-	44.49	45.03	45.15	41.68	42.84	26.07	35.10	44.15	44.08	40.95
URIE [17]	-	40.13	40.60	41.29	37.30	39.24	20.80	30.45	39.54	40.41	36.64
SENTRY [13]	-	43.98	44.10	44.79	39.97	42.01	21.71	32.38	43.03	43.56	39.50
Tent [20]	-	43.40	44.24	44.30	40.70	42.18	25.64	34.08	43.16	43.13	40.09
Ev-TTA	-	**46.02**	**45.29**	**45.91**	**42.53**	**43.90**	**26.70**	**36.17**	**45.00**	**45.22**	**41.86**

Table E.10. Online evaluation results of sorted time surface [1] on N-ImageNet [6] and its variants.

Change	None	Trajectory					Brightness				Average
Validation Dataset	Orig.	1	2	3	4	5	6	7	8	9	All
No Adaptation	48.43	45.17	36.58	42.28	26.57	38.70	24.39	32.76	38.99	37.37	35.89
Mummadi *et al*. [10]	-	48.02	47.41	47.98	44.37	46.39	28.52	38.60	47.21	46.80	43.92
URIE [17]	-	43.78	43.31	44.03	41.04	40.71	16.80	28.61	43.52	41.33	38.13
SENTRY [13]	-	48.33	47.70	48.38	43.71	46.28	25.25	37.51	47.53	46.73	43.49
Tent [20]	-	46.32	46.17	46.64	42.74	44.56	28.20	36.93	45.59	45.32	42.50
Ev-TTA	-	**48.53**	**47.75**	**48.38**	**45.35**	**47.26**	**31.02**	**39.07**	**48.19**	**47.66**	**44.80**

Table E.11. Offline evaluation results of DiST [6] on N-ImageNet [6] and its variants.

Change	None	Trajectory					Brightness				Average
Validation Dataset	Orig.	1	2	3	4	5	6	7	8	9	All
No Adaptation	48.43	45.17	36.58	42.28	26.57	38.70	24.39	32.76	38.99	37.37	35.89
Mummadi *et al*. [10]	-	45.85	45.73	46.25	42.39	44.13	27.82	36.48	45.19	44.85	42.08
URIE [17]	-	40.88	41.04	42.03	37.68	40.17	20.14	30.18	41.44	40.14	37.08
SENTRY [13]	-	45.47	45.58	46.12	41.55	43.52	24.57	35.03	45.03	44.69	41.28
Tent [20]	-	43.27	44.10	44.37	40.61	41.78	25.52	34.10	43.39	43.23	40.04
Ev-TTA	-	**46.32**	**46.05**	**46.57**	**43.23**	**44.58**	**28.05**	**36.98**	**46.03**	**45.64**	**42.61**

Table E.12. Online evaluation results of DiST [6] on N-ImageNet [6] and its variants.

References

[1] I. Alzugaray and M. Chli. Ace: An efficient asynchronous corner tracker for event cameras. In *2018 International Conference on 3D Vision (3DV)*, pages 653–661, 2018.

[2] Jonathan Binas, Daniel Neil, Shih-Chii Liu, and Tobi Delbruck. Ddd17: End-to-end davis driving dataset, 2017.

[3] G. Cohen, S. Afshar, G. Orchard, J. Tapson, R. Benosman, and A. van Schaik. Spatial and temporal downsampling in event-based visual classification. *IEEE Transactions on Neural Networks and Learning Systems*, 29(10):5030–5044, 2018.

[4] Daniel Gehrig, Mathias Gehrig, Javier Hidalgo-Carrió, and Davide Scaramuzza. Video to events: Recycling video datasets for event cameras. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, June 2020.

[5] Yuhuang Hu, Jonathan Binas, Daniel Neil, Shih-Chii Liu, and Tobi Delbruck. Ddd20 end-to-end event camera driving dataset: Fusing frames and events with deep learning for improved steering prediction, 2020.

[6] Junho Kim, Jaehyeok Bae, Gangin Park, Dongsu Zhang, and Young Min Kim. N-imagenet: Towards robust, fine-grained object recognition with event cameras. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2146–2156, October 2021.

[7] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

[8] Xavier Lagorce, Garrick Orchard, Francesco Galluppi, Bert Shi, and Ryad Benosman. Hots: A hierarchy of event-based time-surfaces for pattern recognition. *IEEE transactions on pattern analysis and machine intelligence*, 39, 07 2016.

[9] Ana I. Maqueda, Antonio Loquercio, G. Gallego, N. García, and D. Scaramuzza. Event-based vision meets deep learning on steering prediction for self-driving cars. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5419–5427, 2018.

[10] Chaithanya Kumar Mummadi, Robin Hutmacher, Kilian Rambach, Evgeny Levinkov, Thomas Brox, and Jan Hendrik Metzen. Test-time adaptation to distribution shift by confidence maximization and input transformation, 2021.

[11] P. K. J. Park, B. H. Cho, J. M. Park, K. Lee, H. Y. Kim, H. A. Kang, H. G. Lee, J. Woo, Y. Roh, W. J. Lee, C. Shin, Q. Wang, and H. Ryu. Performance improvement of deep learning based gesture recognition using spatiotemporal demosaicing technique. In *2016 IEEE International Conference on Image Processing (ICIP)*, pages 1624–1628, 2016.

[12] Etienne Perot, Pierre de Tournemire, Davide Nitti, Jonathan Masci, and Amos Sironi. Learning to detect objects with a 1 megapixel event camera. *arXiv preprint arXiv:2009.13436*, 2020.

[13] Viraj Prabhu, Shivam Khare, Deeksha Kartik, and Judy Hoffman. Sentry: Selective entropy optimization via committee consistency for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8558–8567, October 2021.

[14] Henri Rebecq, Daniel Gehrig, and Davide Scaramuzza. ESIM: an open event camera simulator. *Conf. on Robotics Learning (CoRL)*, Oct. 2018.

[15] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.

[16] B. Son, Y. Suh, S. Kim, H. Jung, J. Kim, C. Shin, K. Park, K. Lee, J. Park, J. Woo, Y. Roh, H. Lee, Y. Wang, I. Ovsiannikov, and H. Ryu. 4.1 a 640×480 dynamic vision sensor with a 9μm pixel and 300meps address-event representation. In *2017 IEEE International Solid-State Circuits Conference (ISSCC)*, pages 66–67, 2017.

[17] Taeyoung Son, Juwon Kang, Namyup Kim, Sunghyun Cho, and Suha Kwak. Urie: Universal image enhancement for visual recognition in the wild. In *ECCV*, 2020.

[18] Stepan Tulyakov, Daniel Gehrig, Stamatios Georgoulis, Julius Erbach, Mathias Gehrig, Yuanyou Li, and Davide Scaramuzza. TimeLens: Event-based video frame interpolation. *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.

[19] Dennis D. Wackerly, William Mendenhall III, and Richard L. Scheaffer. *Mathematical Statistics with Applications*. Duxbury Advanced Series, sixth edition edition, 2002.

[20] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *International Conference on Learning Representations*, 2021.