



Using event cameras for imaging through atmospheric turbulence

Nicolas Boehrer, Robert Nieuwenhuizen, Judith Dijk

► To cite this version:

Nicolas Boehrer, Robert Nieuwenhuizen, Judith Dijk. Using event cameras for imaging through atmospheric turbulence. COAT-2019 - workshop (Communications and Observations through Atmospheric Turbulence: characterization and mitigation), ONERA, Dec 2019, Châtillon, France. 10.34693/COAT2019-S4-005 . hal-03146226

HAL Id: hal-03146226

<https://hal.science/hal-03146226>

Submitted on 18 Feb 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Using event cameras for imaging through atmospheric turbulence

Nicolas Boehrer, Robert Nieuwenhuizen, Judith Dijk
TNO, PO Box 96864, 2509 JG The Hague, The Netherlands

ABSTRACT

Long range horizontal path imaging through atmospheric turbulence is hampered by spatiotemporally randomly varying shifting and blurring of scene points in recorded imagery. Typical mitigation strategies employ software algorithms that combine optical flow to estimate and reduce tip/tilt aberrations with lucky patch identification and data fusion to filter higher order aberrations. In practice, accurate and fast optical flow estimation in turbulence faces is highly challenging. Here we investigate if these challenges can be overcome by using a neuromorphic camera. These sensors measure logarithmic changes in scene radiance in each pixel instead of the radiance itself. The changes are thresholded to produce a stream of events, which enables an ultrafast time resolution on the order of microseconds and enables efficient optical flow estimation at very high frame rates. Here we report on our initial experiments, where we have used a neuromorphic camera to image through turbulence in a controlled indoor setting. We analyze how the sensor responds to the turbulence induced apparent scene motion and propose an algorithm for computing high resolution images of the scene from the event stream in combination with intensity frames. Our initial results do not show a statistical relation between event counts and lucky patches. The image reconstruction based on the event stream did show promising improvements in output image sharpness and stability compared to the raw image stream, with many opportunities for further improvement.

Keywords: Neuromorphic sensors, event cameras, anisoplanatic turbulence, image reconstruction

* E-mail: judith.dijk@tno.nl

1. INTRODUCTION

Light traveling through the atmosphere encounters turbulent regions which modify the optical path length [1]. As the light propagates, the effect of turbulent regions accumulate leading to a random phase distortion of the wave front which causes time-varying blurs and shifts in the image recorded by a camera. Atmospheric turbulence therefore limits the effective resolution of optical imaging in many long range observation applications like surveillance or astronomy.

The effect of turbulence can be mitigated by either using hardware capable of measuring and correcting for the wave front distortion while recording (adaptive optics) or by using software: image processing techniques [2]. Most image processing techniques rely on a combination of motion compensation, sharp region (often called lucky region) identification, multi-frame data fusion and deconvolution [3, 4, 5, 6, 7, 8, 9, 10, 11, 12]. Those algorithms are applied on classical cameras in which a total frame of pixels is recorded. However, the frame rate of classical cameras is typical too low to capture all the dynamics of the turbulence induced changes in the images. As a result, local regions in single frames combine instances where they are temporarily sharp with instances where they are less sharp, and tip/tilt aberrations are averaged to further blur these regions.

Neuromorphic cameras do not record an entire frame but instead output an asynchronous stream of changes, so-called events. Contrary to classical frame recordings that record a lot of redundant data, i.e. image regions that stay constant, this new recording technique allows to record only the local changes between frames so that the bandwidth and the recording resources are best used to record the local dynamics of the scene. With its low latency and high temporal sampling, the neuromorphic camera is therefore expected to be well suited to record the temporal variations of the atmospheric turbulence otherwise unseen by conventional camera. This additional information may improve the quality of the restored turbulence free image.

Here we investigate how the event stream may be used to for this turbulence mitigation. In section 2 and 3 we present the event camera used and describe our experimental setup. In section 4 we investigate if the events themselves can be related to properties of the local shifts and blurs in the scene caused by turbulence. Subsequently in section 5 we derive an algorithm to reconstruct high resolution frames by combining the event stream with conventional intensity frames. In

section 6 we show the results of this initial algorithm. To our knowledge, there is currently no published approach for turbulence correction using a neuromorphic sensor in the image plane of a camera.

2. EVENT CAMERA

Unlike conventional cameras that record entire frames synchronously, event camera only record logarithmic intensity changes. The event camera encodes the changes as an asynchronous series of spikes called events. Each event $e_k = [t_k, u_k, p_k]$ carries the time stamp t_k , the pixel location $u_k = [x_k, y_k]$ and the polarity (on or off) of the log intensity change. As shown in Figure 1, an event is produced when the difference between the memorized log intensity and the current log intensity exceeds a certain threshold (controlled by the user). When the threshold is passed, the time stamp, the pixel position and the polarity is emitted and a the current log intensity value is memorized for the monitoring.

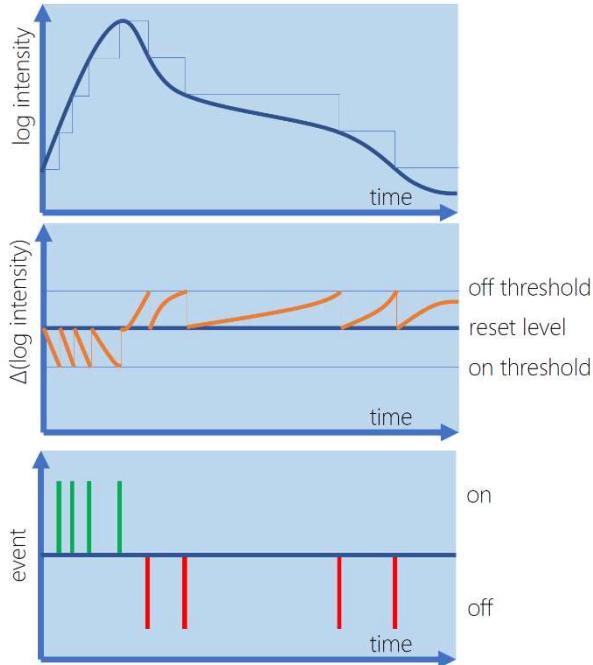


Figure 1: DVS Pixel operation

With a temporal sampling that adapts natively to the actual dynamic of the scene, the event camera offers multiple advantages: high temporal resolution (a few microseconds) with a low latency, a very large dynamic range (120dB) and a low power and bandwidth consumption.

Event cameras belong to the family of gradient cameras, which react to intensity changes and thus a need scene with enough contrast (mostly edges) and movement from the camera or the scene. Even though neuromorphic imaging is relatively recent in the field, it gained a lot of attention and event cameras are already available as commercial product. Many applications and research directions have been investigated. See [16] for a recent survey.

Among all research directions that have been investigated, the ones with strong camera motion or object motion are the most promising. Atmospheric turbulence produces time varying intensity changes in the form of apparent scene motion. Continuous motion recording therefore implies the potential offered by an event camera to outperform conventional frame approaches.

3. EXPERIMENTAL SETUP

For the research reported here we used the DAVIS346 *dynamic vision sensor* from iniVation [17]. **Error! Reference source not found.** shows the specification of this camera. The DAVIS346 has the particular benefit of recording both events and intensity frames.

Table 1: Specifications of the DAVIS346 dynamic vision sensor of iniVation.

Feature	Value
Resolution	346×260
Time resolution event stream	$1 \mu\text{s}$
Frames	Gray scale
Dimensions	$40 \times 60 \times 25 \text{ mm}^3 (\text{H} \times \text{W} \times \text{D})$
Weight	100 g (without lens)

In order to collect test material, we recorded video sequences through man-made, indoor turbulence. We produced turbulent air flows with a hot plate placed close to the camera and placed a flat test chart with enough contrasted content at a distance of about 4m. We used the camera with the Avenir lens 16-160mm f/2.0 at 100 and 160mm. The lens was focused on a siemens star placed on the target before recording. After a few attempts with different heaters, we decided to use a hot plate without fan, as a fan reduces the turbulence. The resulting flow pattern seems like a reasonable approximation of operationally relevant turbulence. The magnitude of shifts and blurs as well as length scales in the images appeared consistent with previously recorded data in a NATO-SET trial in Quebec [18] (NATO Quebec).

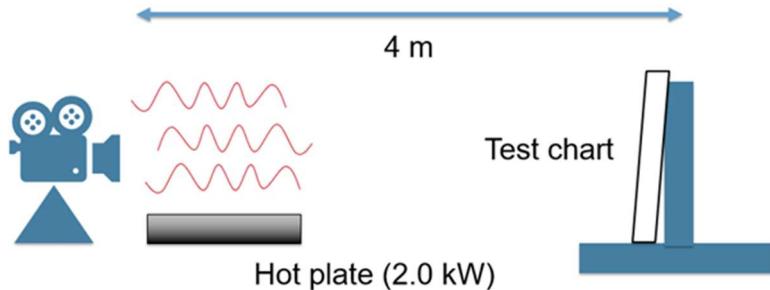


Figure 2 Illustration of the experimental setup.

4. EVENT STATISTICS

In section 1 we proposed that event cameras may be beneficial for capturing the dynamics of the turbulence induced changes in images in two respects: identifying lucky patches and measuring the turbulence induced shifts in the images. In this section we will explore these hypotheses by analysing the spatiotemporal statistics of the event counts.

We first investigated how the event stream relates to the instantaneous sharpness (i.e. luckiness) of the images. To assess this we focused on the variations of a small region of interest (ROI) over 300 frames. The effect of turbulence is a combination of a local spatial shift and a local blur. Our hypothesis is that sharp zones that change only due to shifts produce higher event counts than zones that change due to a variation in blur. The sharp zones can be considered as lucky image patches that once stabilized could deliver sub-pixel sampling to the corrected image estimate. However, when comparing the event counts to the local contrast in an ROI, which is mostly determined by the local sharpness, no evidence of a correlation between these quantities is found as shown in Figure 3.

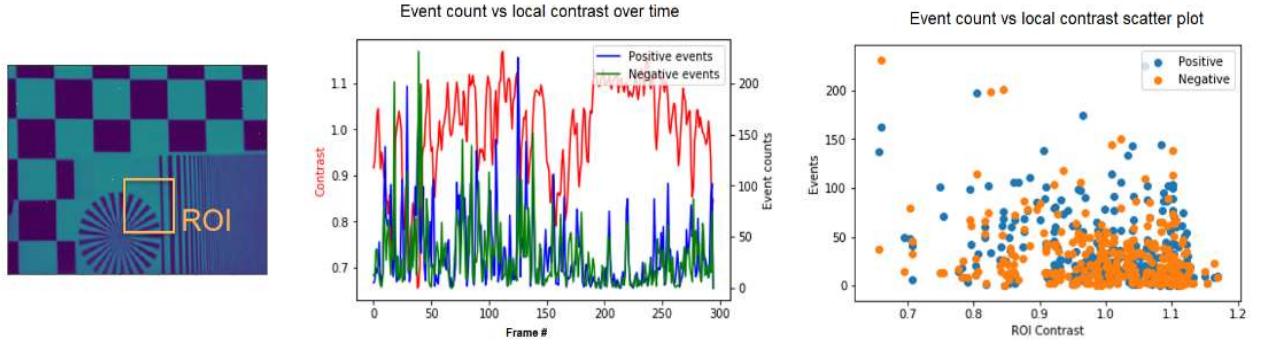


Figure 3 Statistics of event counts in relation to local image sharpness.

To further analyse the relation between sharpness and event counts we characterized the influence of the lens aperture on the event stream. We measured the event count for different f-numbers and noticed that the event count increased with the f-number. Imaging with a smaller aperture visibly resulted in a video disturbed more by local motion and less by local blur variations. This also resulted in lower event rates as shown in Figure 4.

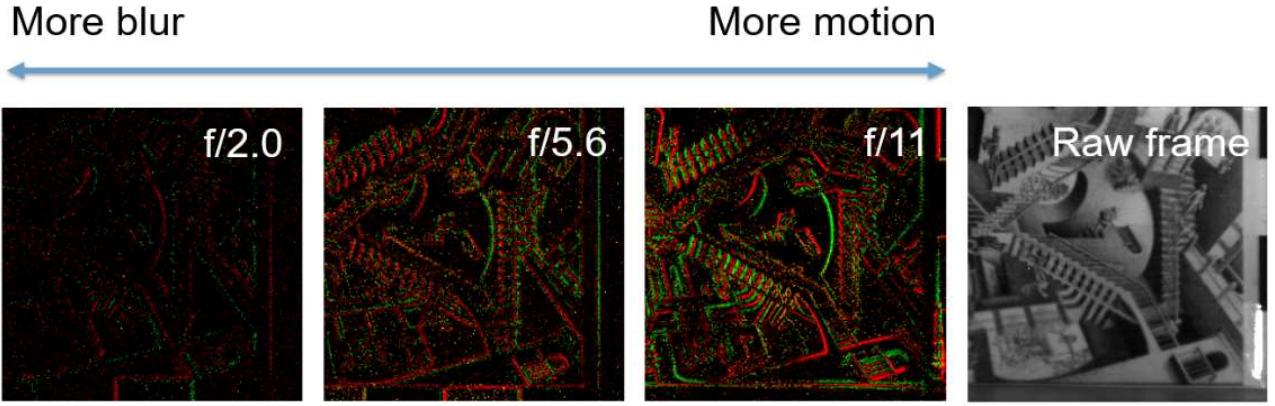


Figure 4 Visualization of local event counts between two intensity frames for different aperture settings. Positive events are shown in green, negative events in red.

Next we considered if the event camera allows for more accurate measurement of local image shifts due to turbulence. The event stream has the potential to contain high frequent information about local intensity changes. To evaluate if the event stream contains more information than intensity frames about the local shifts induced by turbulence, we measured the frequency at which we could locally estimate the motion. Linearly moving edges typically produce events which correspond to 3d points that reside on a plane in x-y-t space. The slope of the plane corresponds to the motion velocity. Therefore local motion estimation with the event stream consists of fitting a plane on the recent points around the current event location. Figure 5 shows how often a weak plane fit can be done relative to the rate of intensity frames, assuming that we need at least 3 points in a 5x5 neighborhood.

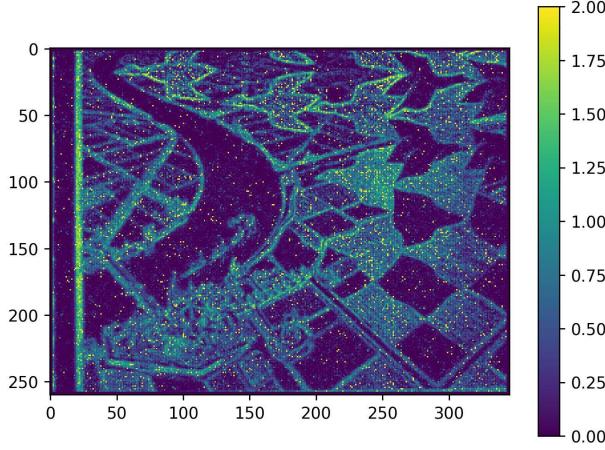


Figure 5: Average number of possible fit per frame using a 5×5 neighborhood

Some regions show rates up to 2 motion estimation (plane fit) per frame, but rates generally stay below the 1 estimation per frame rate. Even though the event stream may contain high frequent changes, the turbulence phenomenon does not seem to generate those high frequent variations and the intensity images already carry most of the local shift information.

5. IMAGE RECONSTRUCTION ALGORITHM

The preceding section did not directly show a relation between event counts and the sharpness of image patches. However, for image reconstruction it is highly desirable to identify image regions (in space and time) with relatively high sharpness and build the final corrected output out of the subset of “lucky” regions. Therefore we implemented an image reconstruction algorithm in which we more directly measure if regions are “lucky” by measuring their instantaneous sharpness. Figure 6 gives an overview of our proposed algorithm.

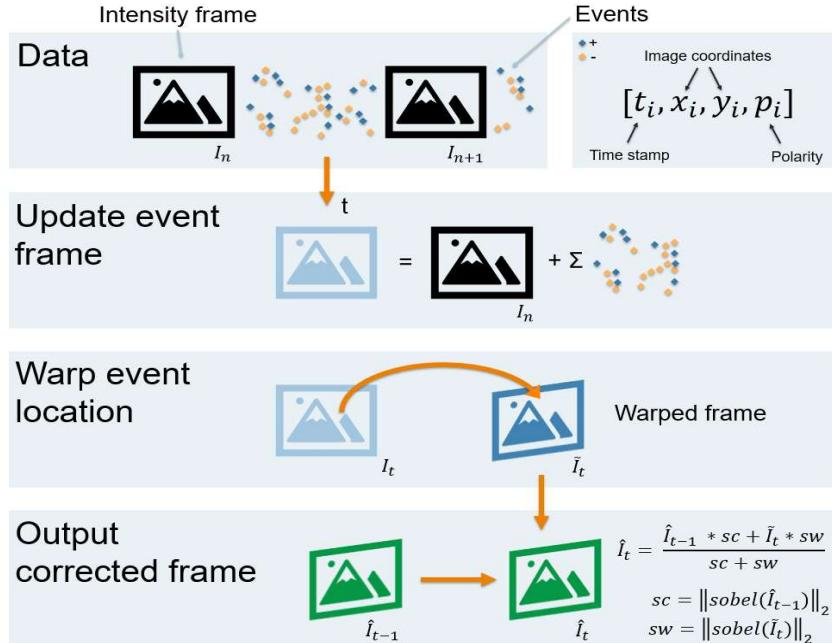


Figure 6 Schematic overview of the image reconstruction algorithm.

We derive our algorithm by modeling the event camera as a high speed camera that records only changes. The algorithm first extracts a virtual “high speed” frame at a time t by reconstructing the intensity from the previous intensity frame and the sequence of events up to the time tag t .

Given a set of events $\{k\}$ from $e_k = [t_k, u_k, p_k]$ that occurred at time $t_k \in [t_n, t]$ with $t < t_{n+1}$, location $u_k = [x_k, y_k]$ with a polarity $p_k \in [+1, -1]$ and the corresponding constant relative brightness change C (here identical for positive and negative polarity), one can construct the difference image relative to the last intensity frame:

$$\delta I(x) = C \sum_k p_k \delta(x - u_k)$$

As illustrated in Figure 6, the difference image is used to iteratively update the raw log intensity image and create a virtual frame I_t at any point between intensity images I_n :

$$I_t = I_n + \delta I$$

The contribution of each ON or OFF event is calibrated with the first intensity frames so that updating I_n with the contributions of all events between frame I_n and I_{n+1} best matches the intensity measured at frame I_{n+1} .

Once the virtual raw frame is available, the algorithm removes the local shifts induced by the turbulence. Before processing the sequence, a reference image is first computed by averaging all intensity frames. The local shifts between each intensity frame and the reference image are estimated using Farnebäck dense optical flow [19].

The virtual raw frame is warped to the reference space based on the optical flow interpolated between two intensity frame times.

$$\tilde{I}_t = I_t(W(u, t))$$

Here the continuous warp is the linear interpolation between the two optical flows from two subsequent frames:

$$W(u, t) = W(u, t_n) * \alpha + W(u, t_{n+1}) * (1 - \alpha)$$

$$\alpha = \frac{t - t_n}{t_{n+1} - t_n}$$

With $W(u, t_i), \{i \in \mathbb{N} \mid 0 < i < N\}$, N the number of intensity frames, the warp computed from the optical flow between each intensity frame and the average reference frame.

In the last step, the algorithm updates the final turbulence corrected image based on a weighting scheme that emphasizes lucky regions. We use the Sobel response as a sharpness (and therefore luckiness) metric.

$$\hat{I}_t = \frac{\hat{I}_{t-1} H(\hat{I}_{t-1}) + \tilde{I}_t H(\tilde{I}_t)}{H(\hat{I}_{t-1}) + H(\tilde{I}_t)}$$

With

$$H(I) = \sqrt{H_x(I)^2 + H_y(I)^2}$$

$$H_x(I) = \begin{bmatrix} +1 & 0 & -1 \\ +2 & 0 & -2 \\ +1 & 0 & -1 \end{bmatrix} * I$$

$$H_y(I) = \begin{bmatrix} +1 & +2 & +1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix} * I$$

6. IMAGE RECONSTRUCTION RESULTS

We are interested in applying the lucky imaging technique described in the previous section onto the event stream to:

- Take into account the local turbulence motion generating the events. We want to verify if we are able to capture high frequency motion that would make the edge position estimate converge towards the true edge position.
- Benefit from the temporal resolution of the camera. The virtual frame may contain more lucky regions than the ones visible in the intensity frames

We ran the pipeline on multiple targets recorded with the test setup described above. We compared our results to the state of the art turbulence mitigation method based on Hardie's super-resolution algorithm [20, 6].

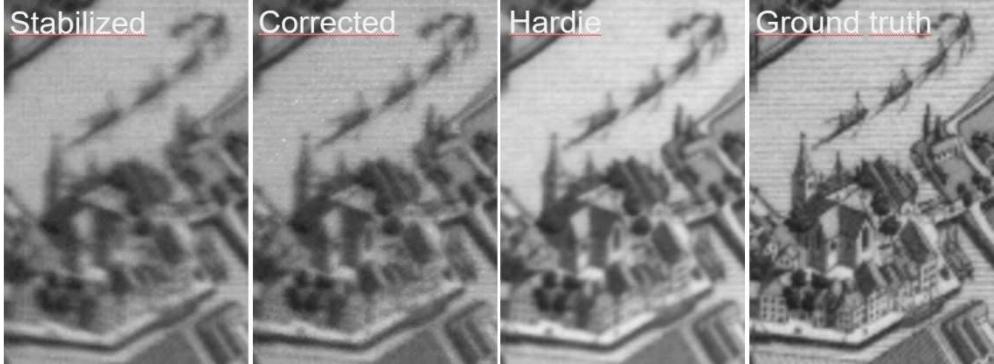


Figure 7 Comparison of image quality of the event based image reconstruction (corrected frames) with the Hardie super-resolution based reconstruction and ground truth.

When compared to a simple stabilization, the turbulence correction using events allows to recover some fine details (see Figure 7, windows, wake on the river) but the gain of event based lucky imaging isn't significant when compared to more classical intensity frame based approach e.g. the Hardie approach.

We also looked at the temporal stability of our estimation pipeline. When looking at videos, we observe that a temporally varying correction quality could be more disturbing than the actual turbulence effects it tries to correct. Therefore we inspected the temporal constancy of the output sequences of our pipeline. This confirmed the stability of the iterative correction: compared to the output of a simple geometric stabilization obtained by only warping raw frames, the turbulence correction reduces the effect of local shift and blur induced by the turbulence, leading to a more stable output. Figure 7 shows an example of a region of interest in one of our result videos, which shows that the output frames of our pipeline (the corrected frames) indeed exhibit a more constant image quality than the stabilized raw frames.

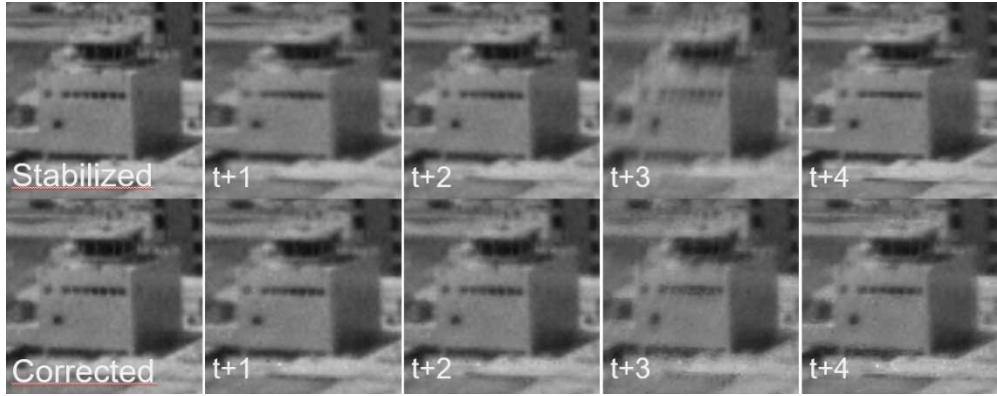


Figure 8 Comparison of the stability over multiple frames between locally stabilized raw frames and the event based image reconstruction (corrected frames).

7. CONCLUSIONS & DISCUSSION

In this paper we report on experiments with a neuromorphic camera for turbulence mitigation. First we investigated the properties of the event data stream in relation to the sharpness of the reference intensity frame. Our analysis did not show a correlation between event counts and the sharpness variations of the intensity frames. Nevertheless, we were able to construct a lucky imaging pipeline that makes use of events and therefore benefits from the high temporal resolution of the camera to better record the continuous variations of the atmospheric turbulence. We applied thus lucky imaging pipeline on video's with "artificial" turbulence. The improvements are relatively modest when compared to intensity frame-based correction e.g. Hardie. There may be different reasons for that:

- The turbulence in our experimental setup is generated with a hot plate close to the sensor. This may not be a good approximation to typical atmospheric turbulence found in practical scenarios, which may exhibit faster temporal dynamics due to crosswinds for example. For such high frequency variations the neuromorphic camera would have an advantage over a classic camera.
- The current turbulence motion compensation relies on a dense optical flow computed between the intensity frames and the reference frame [19]. The motion to compensate the event position is then interpolated between the frames. This linear interpolation may be inaccurate.
- The sharpness measure used for updating the corrected images relies on a Sobel response. The Sobel operator is designed to measure continuous gradients rather than the spatially discrete event features.

In future research we will record experimental data with natural turbulence in an outdoor imaging setup to further verify our current conclusions. On the image reconstruction side we will explore different possible improvements: spatial localization of the event updates, 2) sharpness measures and 3) weighting schemes for the event-based and intensity frame-based updates in the reconstruction algorithm.

ACKNOWLEDGEMENTS

The work was funded by the Office of Naval Research Global.

REFERENCES

- [1] M. C. Roggemann, B. M. Welsh and B. R. Hunt, *Imaging through turbulence*, CRC press, 2018.
- [2] A. W. M. Eekeren, K. Schutte, J. Dijk, P. B. W. Schwering, M. Iersel and N. J. Doelman, "Turbulence compensation: an overview," in *Infrared Imaging Systems: Design, Analysis, Modeling, and Testing XXIII*, 2012.

- [3] M. A. Vorontsov and G. W. Carhart, "Anisoplanatic imaging through turbulent media: image recovery by local information fusion from a set of short-exposure images," *JOSA A*, vol. 18, pp. 1312-1324, 2001.
- [4] B. Fishbain, L. P. Yaroslavsky and I. A. Ideses, "Real-time stabilization of long range observation system turbulent video," *Journal of Real-Time Image Processing*, vol. 2, pp. 11-22, 2007.
- [5] M. Aubailly, M. A. Vorontsov, G. W. Carhart and M. T. Valley, "Automated video enhancement from a stream of atmospherically-distorted images: the lucky-region fusion approach," in *Atmospheric Optics: Models, Measurements, and Target-in-the-Loop Propagation III*, 2009.
- [6] A. v. Eeckeren, M. Kruithof, K. Schutte, J. Dijk, M. v. Iersel and P. Schwering, "Patch-based local turbulence compensation in anisoplanatic conditions," in *Proc. SPIE*, 2012.
- [7] C. S. Huebner, "Turbulence mitigation of short exposure image data using motion detection and background segmentation," in *Infrared Imaging Systems: Design, Analysis, Modeling, and Testing XXIII*, 2012.
- [8] N. Anantrasirichai, A. Achim, N. G. Kingsbury and D. R. Bull, "Atmospheric turbulence mitigation using complex wavelet-based fusion," *IEEE Transactions on Image Processing*, vol. 22, pp. 2398-2408, 2013.
- [9] O. Oreifej, X. Li and M. Shah, "Simultaneous video stabilization and moving object detection in turbulence," *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, pp. 450-462, 2013.
- [10] X. Zhu and P. Milanfar, "Removing atmospheric turbulence via space-invariant deconvolution," *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, pp. 157-170, 2013.
- [11] N. Anantrasirichai, A. Achim and D. Bull, "Atmospheric turbulence mitigation for sequences with moving objects using recursive image fusion," in *2018 25th IEEE International Conference on Image Processing (ICIP)*, 2018.
- [12] R. Nieuwenhuizen, J. Dijk and K. Schutte, "Dynamic turbulence mitigation for long-range imaging in the presence of large moving objects," *EURASIP journal on image and video processing*, vol. 2019, p. 2, 2019.
- [13] J. Gilles, T. Dagobert and C. De Franchis, "Atmospheric turbulence restoration by diffeomorphic image registration and blind deconvolution," in *International Conference on Advanced Concepts for Intelligent Vision Systems*, 2008.
- [14] E. Chen, O. Haik and Y. Yitzhaky, "Detecting and tracking moving objects in long-distance imaging through turbulent medium," *Applied optics*, vol. 53, pp. 1181-1190, 2014.
- [15] K. K. Halder, M. Tahtali and S. G. Anavatti, "Geometric correction of atmospheric turbulence-degraded video containing moving objects," *Optics express*, vol. 23, pp. 5091-5101, 2015.
- [16] G. Gallego, T. Delbrück, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. Davison, J. Conradt, K. Daniilidis and D. Scaramuzza, "Event-based Vision: A Survey," 2019.
- [17] Inivation, "Inivation launches next generation event based dynamic vision sensor," Inivation AG, 15 05 2018. [Online]. Available: <https://inivation.com/inivation-launches-next-generation-event-based-dynamic-vision-sensor/>. [Accessed 24 02 2020].
- [18] M.-T. Velluet, C. Bell, J.-F. Daigle, J. Dijk, S. Gladysz, A. Kanaev, A. Lambert, D. Lemaster, G. Potvin and M. Vorontsov, "Data collection and preliminary results on turbulence characterisation and mitigation techniques," in *Electro-Optical and Infrared Systems: Technology and Applications XVI*, 2019.
- [19] G. Farnebäck, "Two-Frame Motion Estimation Based on Polynomial Expansion," *Image Analysis*, pp. 363-370, 2003.
- [20] R. C. Hardie, K. J. Barnard, J. G. Bognar, E. E. Armstrong and E. A. Watson, "High-resolution image reconstruction from a sequence of rotated and translated frames and its application to an infrared imaging system," *Optical Engineering*, vol. 37, 1998.