

Powerline Tracking with Event Cameras

Alexander Dietsche, Giovanni Cioffi, Javier Hidalgo-Carrió, Davide Scaramuzza

Abstract—Autonomous inspection of powerlines with quadrotors is challenging. Flights require persistent perception to keep a close look at the lines. We propose a method that uses event cameras to robustly track powerlines. Event cameras are inherently robust to motion blur, have low latency, and high dynamic range. Such properties are advantageous for autonomous inspection of powerlines with drones, where fast motions and challenging illumination conditions are ordinary. Our method identifies lines in the stream of events by detecting planes in the spatio-temporal signal, and tracks them through time. The implementation runs onboard and is capable of detecting multiple distinct lines in real time with rates of up to 320 thousand events per second. The performance is evaluated in real-world flights along a powerline. The tracker is able to persistently track the powerlines, with a mean lifetime of the line 10× longer than existing approaches.

SUPPLEMENTARY MATERIAL

A video showing an application of our powerline tracker is available at <https://youtu.be/KnBJqed5qDI>. The code can be found here https://github.com/uzh-rpg/line_tracking_with_event_cameras.

I. INTRODUCTION

Quadrotors are a fast-to-deploy and a cost-effective solution for powerline inspection [1]. To perform such operation, teams of specialized human operators typically use ropes or scaffoldings to access the powerline infrastructure and manned helicopters for long range inspection. This task requires months of training, is time consuming, and entails high risk for humans. Quadrotors can cut costs by 50% [1], reducing the in-person operation, and improving the data collection. Also, they can drastically reduce the inspection time and increase the frequency and quality of inspection.

Current solutions are transferring from research to industry by partially helping humans in the most challenging scenarios¹ [1], [2]. However, they require a direct visual line of sight and that the quadrotor follows a predefined path to inspect a dedicated zone. The energy sector, in general, and the powerline sector, in particular, raise the fundamental need for closed-loop solutions with highly coupled perception and action [3].

The authors are with the Robotics and Perception Group, Department of Informatics, University of Zurich, and Department of Neuroinformatics, University of Zurich and ETH Zurich, Switzerland (<http://rpg.ifi.uzh.ch>). This work was supported by the National Centre of Competence in Research (NCCR) Robotics through the Swiss National Science Foundation (SNSF), and by the European Union's Horizon 2020 Research and Innovation Programme under grant agreement No. 871479 (AERIAL-CORE).

¹<https://aerial-core.eu>




Fig. 1. Quadrotor with the onboard computer and the event camera (top figure). Bottom figures show the real test scenario during the outdoor experiments. Bottom left: color image with a power line mast. Bottom right: two lines tracked and marked in red with their unique ID number.

Robust and persistent line tracking is beneficial for perception-aware planning [4], [5]. The perception-aware planner uses the information on the position of the line to compute a path that guides the drone to its target location while keeping the desired powerline in the camera field of view. This allows recording images of the powerlines that can be analyzed offline to evaluate the health of the powerline infrastructure.

Onboard perception for quadrotors is currently limited by the payload, where a significant portion of it is used for high-resolution sensors, such as LiDAR and high definition cameras. These sensors are mainly required for data acquisition, high level planning, and offline post-processing, increasing latency when in use by the low level controller. A quadrotor needs low-latency sensory information to perform closed-loop navigation and guarantee safety and quality during inspection.

Event cameras have gained tremendous attention in the last decade in robotics, as well as in computational photography and computer vision [6]. Their inherent robustness against motion blur, low latency, and high dynamic range [7], [8] is attractive for robotic applications. These properties are beneficial for powerline inspection tasks in case the drone needs to perform agile manoeuvres, for example to avoid powerline masts. However, due to their unconventional output, event cameras also present new challenges, and tracking of objects, such as lines, is a difficult task due to the sparsity

Method category	Correctness	Instance	Persistence
Hough transf. [9], [10], [11]	+	-	+
Non parametric [12], [13]	+	-	-
Spatio-temporal [14], [15]	+	+	-
Ours	+	+	+

TABLE I

EXISTING METHODS AND THEIR CATEGORIES. OUR APPROACH BELONGS TO THE SPATIO-TEMPORAL CATEGORY ADDING PERSISTENCE WITH HIBERNATION.

and frequent changes in appearance in the stream of the events.

The difficulty of line tracking using events led previous works [9], [10], [11], [12], [13], [14], [15] to explore different techniques but pursuing a common goal: robustness. We identify the robustness of tracking a line as the combination of three desired characteristics. (1) *correctness* (i.e., the capability of solely tracking lines), (2) *instance* (i.e., the capability of tracking lines with a unique identifier, and (3) *persistence* (i.e., how long the line can be tracked). These characteristics are summarized in Table I.

In this work, we present a method that captures the aforementioned characteristics to robustly track lines for powerline inspection.

The main contributions of this work are:

- a *spatio-temporal* event-based line tracker optimized for powerline inspection. Our algorithm extends [15] with novel processing blocks.
- the use of *hibernation* to improve the persistence of the line tracker by a factor of 10 compared to [15].
- a real-time implementation capable to run onboard a lightweight resource-constrained quadrotor platform.
- experiments on a real world powerline dataset showing superior performance than a custom version of [15].
- the release of the code fully open source.

II. RELATED WORK

The recent literature on event-based line trackers can be divided into three categories:

Hough transform based methods: These methods use polar coordinate to parameterize the lines [16] and search for the lines that contain the largest number of events in resulting parameter space. A Hough transform based line tracker was used for low-latency and high-bandwidth control of a one dimensional dualcopter in [9]. The tracker ran on a sliding window of events, as proposed in [17], and was used in combination with a Kalman filter to estimate the dualcopter state. The work in [10] proposes an application of a Hough transform based tracker for agile aerial robot maneuvering. Consecutive Hough transform based detectors were used in [11] to detect and track poles in a railway localization system.

Although Hough transform based methods have been used for events, they do not fully leverage the asynchronous nature of event cameras. Performing the projection from the pixel to the parameter space and searching such parameter space on an event-by-event basis is computationally costly and

requires fine adjustment of the peak selection and binning of the parameters. Using a sliding window of events helps in reducing the computational cost but introduces latency.

Non parametric based methods: These methods includes a line detector which calculates the spatial derivative of every pixel and clusters the pixels into support regions based on their orientation. The support regions are then used to calculate the line segments [12]. This approach minimizes the need of parameter tuning. A non parametric method was adopted by the event-based line segment detector ELiSeD [13]. ELiSeD attributes an orientation to an incoming event and then clusters together those events that have similar orientation. This method has the tendency to produce small line segments that are short-lived. In our task, the lines of interest have large length and are mostly long-lived.

Spatio-temporal based methods: These methods make the assumption that lines travel approximately at a fixed speed in small time intervals. When such assumption holds, lines describe a plane in the spatio-temporal event space. In [14], the normal vectors in the spatio-temporal space were computed for each incoming event using a local neighborhood. Events with similar normal vectors were clustered together to form lines. The detected lines and inertial measurements were used to estimate the camera poses in an event-based visual odometry algorithm. A spatio-temporal based line tracker algorithm was also proposed in [15]. The algorithm was able to track slowly translating lines for long time. Moreover, each detected line had an internal state and could be uniquely identified. Since these characteristics are desirable for powerline tracking applications, we built our line tracker upon [15].

III. METHODOLOGY

Our approach, based on [15], fits planes in short time periods through the spatio-temporal space of the events. A line is detected if the plane fitting satisfies some requirements (see Sec. III-D). We rely on the assumption that the velocity of the lines is approximately constant in the short time period. Such assumption holds for the lines whose motion mainly consists of translation.

The line tracking algorithm has four main steps. Events are filtered based on a refractory-time filter and a neighborhood filter. Each event that passes the filter becomes a candidate for being added to an existing line. If the line addition fails, the event is a candidate for being added to a cluster. If there are not clusters available for such event, it is left unassigned. New clusters can be generated from the unassigned events using the newest event as the starting point. An overview of the proposed algorithm can be seen in Fig. 2. In parallel to the processing of new incoming events, the existing lines and clusters are periodically checked in a separate thread. These checks are necessary to periodically remove old events, lines, clusters, and update the line and cluster estimates.

A. Filtering

The initial event filtering consists of two independent steps. In the first step, all event firings at a given location

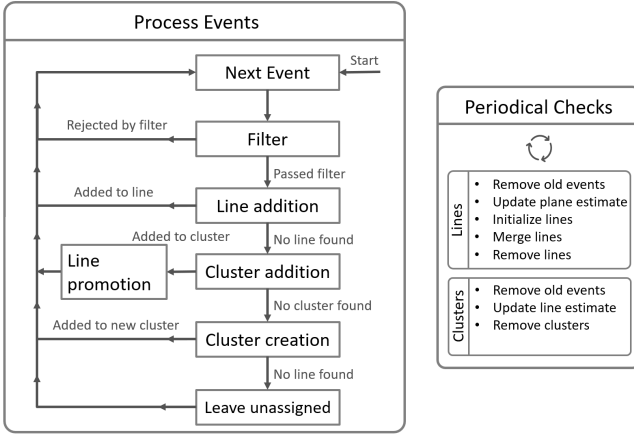


Fig. 2. Overview of the proposed approach, which extends [15].




Fig. 3. Distances a and b are calculated to check if the incoming event (the yellow square) can be added to the line.

are suppressed for a refractory period after an event has been fired at that position. The refractory period is either 8 ms, if an event of the same polarity is fired, or 1 ms, if an event of different polarity is fired. In this way, the filter suppresses spurious event firings of the same polarity. In the second step, the filter checks the 5×5 pixel by 70 ms volume surrounding the incoming event. If there are at least three other events registered in that volume, the event passes the filter. Both filters are implemented using the Surface of Active Events (SAE) [18]. The SAE is comparable with a grayscale image, where instead of intensity values, the newest timestamps are stored. In the refractory time-based filter, we get the newest timestamp from the SAE and compare it with the timestamp of the incoming event. In the neighborhood filter, we get the timestamp of the neighboring events and count the number of timestamps not older than 70 ms.

B. Line Addition

Events that pass the initial filtering process can be added to the existing lines. We evaluate if an incoming event belongs to a line by computing its distances to the line, a , and its distance to the center of the line, b , as depicted in Fig. 3. The event is added to the line if the distance a is below 1.8 pixel, and the distance b is below the line length. If multiple lines fulfill this condition, the event is ignored.

C. Cluster Addition

The cluster addition step tries to add the incoming event, which has not been assigned to a line, to an existing cluster. This step is similar to the line addition step and depends on the same two distances. The distance to the midpoint and the distance to the inferred line must be below a certain threshold for the event to be added to the cluster. If multiple clusters fulfill these conditions, we check if they can be merged. This check consists of calculating the relative angle of the respective inferred lines. If such angle lies below 15 deg, the clusters are merged. If the event is successfully added to a cluster and that particular cluster has collected more than 35 events, the cluster can be promoted to a line.

D. Line Promotion

A cluster is promoted to a line if the events belonging to that cluster lie on a plane. The covariance matrix of the event coordinates is computed, and if its smallest eigenvalue is below a certain threshold, the cluster is promoted to a line. To calculate the covariance matrix, the coordinates are centered to have zero mean. Let $\mathbf{X} \in \mathbf{R}^{N \times 3}$ be the matrix that contains the centered coordinates of the N events that belong to the cluster. We then perform the eigen-decomposition of the covariance matrix $\mathbf{C} = \frac{1}{N-1} \mathbf{X}^T \mathbf{X} = \mathbf{Q} \mathbf{\Lambda} \mathbf{Q}^T$. Since the covariance matrix is symmetric and positive semi-definite, $\mathbf{\Lambda}$ is guaranteed to contain positive real eigenvalues. If the smallest eigenvalue is smaller than 1.2 pixel the cluster will be promoted to a line. The fitted plane is located at the center of gravity of the events, with its normal \mathbf{n} being the eigenvector corresponding to the smallest eigenvalue. The detected line is then obtained by calculating the intersection of the fitted plane with the xy spatial plane at current time t . We are interested in the line direction \mathbf{d} , the midpoint \mathbf{p} , and the line length l . The direction \mathbf{d} is the intersection of the xy plane at t and the fitted plane. The xy plane has normal $\mathbf{e}_t = [0 \ 0 \ 1]$ and the fitted plane has normal $\mathbf{n} = [n_x \ n_y \ n_t]$. The line direction is then calculated as:

$$\mathbf{d} = \mathbf{n} \times \mathbf{e}_t = [n_2 \ -n_1 \ 0]^T. \quad (1)$$

To calculate the midpoint \mathbf{p} , we project the center of gravity \mathbf{g} of the events onto the line. The vector \mathbf{s} which points from \mathbf{g} to \mathbf{p} lies within the fitted plane and is thus perpendicular to the normal vector \mathbf{n} . The vector \mathbf{s} is also perpendicular to \mathbf{d} , since the projection of a point onto a line is always orthogonal to the line direction, and it is computed as:

$$\mathbf{s} = \mathbf{n} \times \mathbf{d} = [n_x n_t \ n_y n_t \ -n_x^2 - n_y^2]^T. \quad (2)$$

The point \mathbf{p} is then found as:

$$\mathbf{p} = \mathbf{g} + \frac{t - \bar{t}}{-n_x^2 - n_y^2} \mathbf{s}, \quad (3)$$

where $\frac{t - \bar{t}}{-n_x^2 - n_y^2}$ is a scaling factor.

The calculation of the line length is based on the assumption that the events are uniformly distributed along the line. Using the formulation of the standard deviation of the Uniform distribution, Γ , $\sigma = \sqrt{\mathbb{E}[\Gamma^2] - (\mathbb{E}[\Gamma])^2}$ and the fact that the




Fig. 4. The detected line, in red, is the intersection of the fitted plane and the xy plane at the current time. The point p denotes the midpoint of the line, which is the projection of the center of gravity of the events, g , along the direction s onto the line. The vector n is the plane normal, and the vector d is the direction of the line.




Fig. 5. Directed chain growth depends on the relative position of the last and current chain element. If no event is found in the search space (green) the extended search space (blue) will be checked.

data is centered (i.e., $\mathbb{E}[\Gamma] = 0$), the line length is $l = \sqrt{12}\sigma$. We compute σ by projecting the largest and second largest eigenvectors scaled by their eigenvalues onto the line. Fig. 4 depicts the quantities involved in the line promotion step.

E. Cluster Creation

Events that are added neither to existing lines nor to clusters are left unassigned and are stored in a SAE. We use the event position as a starting point to find a chain of connected events in the SAE. There are two separate SAEs, one for each polarity. The first chain element is the incoming event. The second chain element is the youngest event in the 3×3 neighborhood that is not older than a certain threshold. From there on, the chain grows in a directed fashion. Depending on the relative position of the last and current element, the next element is searched in a predefined search space. If no element is found, this space is extended. The predefined search spaces are depicted in Fig. 5, where both the initial search pattern (green) and the extended search pattern (blue) are shown. The chain growth is stopped once no more events are found or if the chain reaches a predefined length.

F. Initialization

We propose to add an initialization step in order to be more conservative about the line detection. The proposed tracker does not make any distinction between lines that have many evenly distributed events and lines that have just few events that are far apart one another. In view of the powerline inspection task, we are more interested in the former ones. To




Fig. 6. Line states and their transition conditions.

that end, newly created lines start out in the *initializing* state and collect events during an initialization period. We check the connected length at the end of the initialization period. The connected length is calculated by projecting the events along the fitted plane onto the line. The projected events are then binned into bins of size 2 pixel. The connected length is the length of the longest non-empty bin chain with gaps of max one empty bin. If the connected length is above 70 pixel, the line receives a unique line ID and goes into the *active* state. Otherwise, the line is discarded.

G. Hibernation

Hibernation is introduced to make the line tracker robust against lines that change direction, which corresponds to change in appearance in the stream of events. In the case of powerline tracking onboard a quadrotor, the change of direction of the lines is due to the flying path of the robot. At the moment of a change of direction, just few events are fired and the last estimate of the event plane is not perfectly perpendicular to the xy plane since events only up to a certain point in the past are considered. Consequently, for a short period of time, the estimated line keeps moving in the original direction. Once new events are fired, the estimated line has moved too far in the opposite direction and is lost. A new line instead will be detected. Hibernation alleviates this problem by calculating the midpoint of the line as the orthogonal projection of the center of gravity of the events onto the current xy plane instead of the projection onto the fitted plane. Thus, a hibernated line stand still until it enters the *active* state again, once new events are added.

Another issue that the tracker would face without hibernation is that in the case of slow direction changes no events are fired for an extended period. The lines would lose many events since old events are removed periodically and only very few ones are added. Once a line contains just few events, it will be deleted. Hibernation helps in this case since old events of hibernated lines are not removed.

The condition for the line state to change between the *active* and *hibernated* state depends on the event density. The event density is the number of events younger than 25 ms divided by the length of the line. If the event density is below $0.08 \left[\frac{\# \text{ events}}{\text{pixel}^2 \cdot \text{ms}} \right]$, the line state changes from *active* to *hibernated*.

H. Periodical Checks

The tracker performs periodical checks on the current clusters and lines. Events belonging to either clusters or

lines that are older than 50 ms are removed. Clusters whose the newest event is older than 40 ms are also deleted. For the remaining clusters, their inferred lines are updated by incorporating the newly added events. The periodical checks on the lines depend on the line state. For *initializing* lines, we check if their initialization period has ended, in that case we compute their connected length. If the connected length exceeds 70 pixel, the line will change from state *initializing* to *active*.

The periodical checks for *active* lines consist of updating the fitted plane estimate and checking the line length. If the line is too short, it is discarded. Finally, *hibernated* lines are checked for the time they have been hibernated. If such time exceeds 1 second, the line is deleted. We also check for the age of the newest event, which has to be below 40 ms. Finally, we check if two lines can be merged. The merge conditions are fulfilled when the angle difference of the inferred lines is below 8.0 deg and the midpoint distances are below 3.5 pixel. If that is the case, the line with the lower ID will receive the events of the other line, which is deleted.

I. Limitations

The tracker performs at its best in scenes where the lines of interest move in a mostly translating fashion. Fast rotating lines lead to detection problems due to the underlying assumption that the lines form planes in the spatio-temporal space. In texture-rich environments where there are a lot of events fired but no obvious lines are present, the tracker has a tendency to detect spurious lines. Under these circumstances, there are enough events for lines to pass initialization step and to continue existing. With a steady flow of incoming events, the tracker see no reason to remove the lines. Hibernation exacerbates this issue by preventing the tracker from quickly deleting these lines. However, this is not an usual case during powerline inspection tasks. Our tracker may not be able to track the line when only few events are fired. This occurs in the case when the drone flies perfectly parallel to the powerline.

IV. EXPERIMENTS

A. Powerline dataset

We evaluated the proposed line tracker on a dataset recorded onboard a manually piloted quadrotor tracking powerlines. The quadrotor was equipped with an iniVation Dynamic and Active Vision Sensor (DAVIS) 346 event camera. The camera resolution was 346×260 pixel. A Matrix Vision MLC200wG Bluefox camera was also mounted to record higher quality standard frames than those provided by the DAVIS camera. The frames of the Bluefox camera were only used for visualization and evaluation purposes. The main parameters used for the tracker are in the Appendix (see Sec. VI).

1) *Evaluation*: Since, to the best of our knowledge, it does not exist an event-based dataset with ground truth for powerlines, we used the frames captured by the Bluefox camera to qualitatively evaluate the accuracy of the proposed line tracker. This was done by overlaying the detected lines




Fig. 7. Detected lines on the powerline dataset recorded onboard a quadrotor. Red lines are in the *active* state, yellow lines are in the *hibernated* state, and grey lines are in the *initializing* state.




Fig. 8. Detected lines on the powerline dataset without hibernation. The images are the same as in Fig. 7.

on the grayscale standard images, see Fig. 7. The tracker was able to successfully extract lines from the event stream and keep track of them. The detected lines are shown at different times during the flight in Fig. 7. The tracker shows good persistence, i.e., once a line is detected, it is tracked for a long period of time. This result can be visualized in the accompanying video by the mostly constant line IDs throughout the sequence. The few cases where the lines are lost are due to very few event firings which occur when the drone flies almost parallel to the line. This problem could be addressed by combining the event-based with a standard frame-based line tracker, using an approach similarly to the one proposed in [19]. The tracker persistence is quantified by the mean lifetime of a line. For the flight in Fig. 7, the mean lifetime was 31.12 s.




Fig. 9. Swift change of direction with hibernation (top row) and without hibernation (bottom row). The line moves to the left direction (left column), the line stands still at the exact moment the direction changes (middle column), and the line moves to the right (right column).

2) *Benefits of Hibernation:* To quantify the impact of hibernation on the tracking performance, we compared the mean line lifetime with and without hibernation on the real powerline dataset. The tracker without hibernation corresponds to our implementation of [15], which ignores event polarity and includes the initialization block. With hibernation, the mean lifetime of a line is 31.12 s. Without hibernation, the lines exist on average for 2.97 s. The detected lines without hibernation are depicted in Fig. 8. The higher line IDs in [15] show that the proposed hibernation improves the persistence of the tracker. As it is usual during powerline inspection, the dataset contains a lot of slow direction changes. Thus, there are some prolonged periods where just few events are emitted. *Hibernated* lines help dealing with those eventless periods.

To illustrate the advantage of hibernation in the case of swift direction changes, the tracker was also evaluated with and without hibernation on a dataset with more aggressive direction changes. In this dataset, a rope is suspended from the ceiling whilst the camera yaws from left to right and back in an oscillatory fashion. In Fig. 9, the bottom row shows the issue of swift direction changes. With hibernation, in the top row, the line goes into the *hibernated* state at the exact moment when the camera changes its direction and thus will have its midpoint projected orthogonally onto the xy plane. Instead, the line without hibernation continues to project its midpoint along the fitted event plane which here keeps moving in the left direction. The line is detached from the actual line and is deleted after a short period of time.

B. Hardware Test

In this set of experiments, we show that our tracker can run onboard a resource quadrotor platform in a closed-loop controller and analyze the computational cost of the different parts of the proposed algorithm. The platform used for these experiments was a quadrotor equipped with a Nvidia Jetson TX2. To record the events, the DAVIS 346 camera was used. The quadrotor setup is depicted in Fig. 1 (top picture). The state estimation was provided by an external motion capture system. The line tracker ran onboard together with a model predictive controller (MPC) [3]. The desired yaw rate was




Fig. 10. Rotating square followed by the quadrotor. The drone yaws counterclockwise from left to right.

Processing step	Processing time [ns]
Filtering	276.6
Line addition	1,369.8
Cluster addition	744.0
Cluster creation	585.7

TABLE II

MEAN PROCESSING TIME PER EVENT FOR EACH STEP OF THE PROPOSED LINE TRACKER ALGORITHM.

computed by a PD controller. The error fed into the controller was the pixel difference between the center of the image and the x coordinate of the midpoint of the reference line. In this way, we wanted to keep the reference line in the center of the image. The reference line was selected as the longest vertical line in the scene. The drone was placed in the middle of a wooden square suspended from the ceiling by a rope. The square was then wound up and released. The quadrotor tried to follow one of the sides of the rotating square. In Fig. 10, the top row shows the captured images with the detected lines overlaid. The green line is the reference line that the drone tried to follow. The drone successfully followed the square for multiple rotations as shown in the accompanying video.

1) *Computational Cost:* This section evaluates the computational cost of the tracker running onboard the quadrotor platform. We aimed to get an estimate of the bandwidth of the incoming events that the tracker can handle. The mean processing times of the different processing steps of the proposed line tracker are in Table II.

Most time is spent trying to add an event to an existing




Fig. 11. Processing time of the line addition and cluster addition steps as a function of the number of lines and clusters, respectively. The red line connects the means.

line, followed by the time needed to add an event to a cluster. The time needed for the filtering step and the cluster-creation step is independent of the number of lines and clusters. The line addition step and the cluster addition step are linearly dependent on the number of lines and clusters, respectively, as shown in Fig. 11. The large variance of the data can be explained by the mutex locks. The lines and clusters are objects that are used in the periodical update, which runs on a different thread. Whenever a periodical check occurs, the line addition step and the cluster addition step are blocked. To calculate the bandwidth that the tracker can handle, an assumption of the expected number of lines and clusters has to be made. Since our tracker is primarily designed for powerline inspection tasks, the dataset of the real powerline was used to estimate the expected number of lines and clusters, which is 3 lines and 4 clusters. The mean processing time for the line addition step is $1.376 \mu\text{s}$, and for the cluster addition step is $0.859 \mu\text{s}$. Thus, the overall mean processing time for environments similar to the real powerline data set is $3.097 \mu\text{s}$. This means the tracker can handle approximately 320,000 events per second. During the hardware tests in Section IV-B the incoming events stream briefly exceeded 400,000 events per second with 4 lines and 8 clusters. Thus, for a short period of time, the tracker was not able to keep up with all the incoming events.

V. CONCLUSION

This paper presented a powerline tracker for quadrotors using event cameras. We proposed to add a new line state, hibernation, to cope with direction changes of the lines. We showed that hibernation increases the tracker persistence, up to $10\times$, in terms of the lifetime of the detected lines with respect to the state of the art. In addition, we showed that our algorithm is able to run onboard resource constrained platforms as quadrotors, and tested it in a closed-loop controller experiments where the quadrotor needs to keep a reference line in the center of the image. As future work, we plan to include the proposed line tracker in a perception-aware planner.

ACKNOWLEDGMENT

We would like to thank Christian Pfeiffer and Juan Simon Jaime for their help with piloting the quadrotor to record the powerline dataset. We would also like to thank Philipp Foehn for his help with the MPC controller.

REFERENCES

- [1] "Drone industry insights - drones in the energy industry: The energy drone operator benchmark," <http://droneii.com>.
- [2] "Danmarks dronestrategi: Danish strategy for drone technology development and use," <https://ufm.dk/publikationer/2016/filer/dronestrategi-2016.pdf>.
- [3] D. Falanga, P. Foehn, P. Lu, and D. Scaramuzza, "Pampc: Perception-aware model predictive control for quadrotors," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*. IEEE, 2018.
- [4] G. Costante, C. Forster, J. Delmerico, P. Valigi, and D. Scaramuzza, "Perception-aware planning," *arXiv*, 2016, retrieved on July 29th, 2021. [Online]. Available: <https://arxiv.org/abs/1605.04151>
- [5] Z. Zhang and D. Scaramuzza, "Perception-aware receding horizon navigation for mavs," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.

- [6] G. Gallego, T. Delbruck, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. Davison, J. Conradt, K. Daniilidis, and D. Scaramuzza, "Event-based vision: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2020.
- [7] S. Sun, G. Cioffi, C. de Visser, and D. Scaramuzza, "Autonomous quadrotor flight despite rotor failure with onboard vision sensors: Frames vs. events," in *IEEE Robot. Autom. Lett.*, 2021.
- [8] D. Falanga, K. Kleber, and D. Scaramuzza, "Dynamic obstacle avoidance for quadrotors with event cameras," *Science Robotics*, 2020.
- [9] R. S. Dimitrova, M. Gehrig, D. Brescianini, and D. Scaramuzza, "Towards low-latency high-bandwidth control of quadrotors using event cameras," in *IEEE Int. Conf. Robot. Autom. (ICRA)*. IEEE, 2020, pp. 4294–4300.
- [10] A. G. Eguluz, J. Rodriguez-Gómez, J. Martinez-de Dios, and A. Ollero, "Asynchronous event-based line tracking for time-to-contact maneuvers in uas," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*. IEEE, 2020.
- [11] F. Tschopp, C. von Einem, A. Cramariuc, D. Hug, A. W. Palmer, R. Siegwart, M. Chli, and J. Nieto, "Hough²map—iterative event-based hough transform for high-speed railway mapping," *IEEE Robot. Autom. Lett.*, 2021.
- [12] R. G. Von Gioi, J. Jakubowicz, J.-M. Morel, and G. Randall, "Lsd: a line segment detector," *Image Processing On Line*, vol. 2, pp. 35–55, 2012.
- [13] C. Brändli, J. Strubel, S. Keller, D. Scaramuzza, and T. Delbruck, "Elised—an event-based line segment detector," in *Int. Conf. on Event-based Control, Comm., and Signal Proc. (EBCCSP)*. IEEE, 2016, pp. 1–7.
- [14] C. L. Gentil, F. Tschopp, I. Alzugaray, T. Vidal-Calleja, R. Siegwart, and J. Nieto, "Idol: A framework for imu-dvs odometry using lines," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*. IEEE, 2020.
- [15] L. Everding and J. Conradt, "Low-latency line tracking using event-based dynamic vision sensors," *Frontiers in neurorobotics*, vol. 12, p. 4, 2018.
- [16] R. O. Duda and P. E. Hart, "Use of the hough transformation to detect lines and curves in pictures," *Communications of the ACM*, vol. 15, no. 1, pp. 11–15, 1972.
- [17] E. Mueggler, B. Huber, and D. Scaramuzza, "Event-based, 6-dof pose tracking for high-speed maneuvers," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*. IEEE, 2014, pp. 2761–2768.
- [18] E. Mueggler, C. Bartolozzi, and D. Scaramuzza, "Fast event-based corner detection," in *British Mach. Vis. Conf. (BMVC)*, 2017.
- [19] A. R. Vidal, H. Rebecq, T. Horstschaefer, and D. Scaramuzza, "Ultimate slam? combining events, images, and imu for robust visual slam in hdr and high-speed scenarios," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 994–1001, 2018.

VI. APPENDIX

A. Main tracker parameters

Cluster parameter	Value	Description
creation number events	7	min number of events to create a cluster from a chain.
addition threshold	1.3 pixel	max distance to inferred line for an event to be added
merge angle	15 deg	max angle difference between clusters to be merged
cleanup event age	50 ms	max age of events before they are removed
deletion no events	40 ms	if there have not been any event for longer than this time, the cluster is deleted

Line Parameter	Value	Description
promotion threshold	1.2 pixel	value of the smallest eigenvalue for line promotion
promotion number of events	35	minimum number of events in cluster for line promotion
initialization length	70 pixel	minimum length of the line to pass initialization
initialization period	90 ms	initialization period
addition threshold	1.8 pixel	max distance to the inferred line for an event to be added
merge angle	8 deg	max angle difference between lines to be merged
merge distance	3.5 pixel	max distance of the line mid point to the other line to be merged
hibernation density	$0.08 \frac{\# \text{ events}}{\text{pixel}^2 \cdot \text{ms}}$	min line density before being hibernated
cleanup event age	50 ms	max age of the events before they are removed
deletion no events	200 ms	if there have not been any events for longer than this, the line is deleted