Reconstructing the flight kinematics of swarming and mating in wild mosquitoes

Sachit Butail[1], Nicholas Manoukis[4], Moussa Diallo[3],
José M. Ribeiro[2], Tovi Lehmann[2] and Derek A. Paley[1]

***Abstract*—We describe a novel tracking system for reconstructing three-dimensional tracks of individual mosquitoes in wild swarms and present the results of validating the system by filming swarms and mating events of the malaria mosquito *Anopheles gambiae* in Mali. The tracking system is designed to address noisy, low frame-rate (25 frames per second) video streams from a stereo camera system. Because flying *An. gambiae* move at 1–4 m/s, they appear as faded streaks in the images or sometimes do not appear at all. We provide an adaptive algorithm to search for missing streaks and a likelihood function that utilizes streak endpoints to extract velocity information. A modified multi-hypothesis tracker probabilistically addresses occlusions and a particle filter estimates the trajectories. The output of the tracking algorithm is a set of track segments with an average length of 0.6–1 seconds. The segments are verified and combined under human supervision to create individual tracks up to the duration of the video (90 s). We evaluate tracking performance using an established metric for multi-target tracking and validate the accuracy using independent stereo measurements of a single swarm. Three-dimensional reconstructions of *An. gambiae* swarming and mating events are presented.**

I. Introduction

Quantitative observations of the flight patterns of wild mosquitoes are critical to expanding our understanding of swarming and mating behavior [1], [2], [3], [4], [5], [6]. Female *Anopheles gambiae* find male swarms in order to mate [7], [5]. A single mating event results in all of the fertilized eggs that a female mosquito lays in her lifetime [8], [9]. Although the basis of mate selection has generated much interest [7], [10], [11], [8], [6], generation of three-dimensional trajectory data of mosquitoes in wild swarms has not been previously accomplished. As in earlier work on midges [12], such trajectory data can provide valuable insight into the dynamical aspects of collective behavior [13], [14]. Past studies on swarming insects [2], [3], [5], [15] focused on two-dimensional trajectories or three-dimensional positions. Recent advancements in high-resolution filming, computer vision, and estimation techniques have increased the degree of automation in data collection, and have made available large datasets for subsequent analysis, such as those developed for starlings [16] and fruit flies [17]. Similar analyses of malarial mosquitoes may inform the first steps towards strategies of vector control [7], [8].

Multi-target tracking systems have been developed for other animals. In two dimensions, ants have been tracked using a joint-state particle filter and interaction models [26]. In three-dimensions, up to a hundred bats have been tracked with three cameras using a Kalman filter in conjunction with a multi-dimensional assignment strategy [18], [19]. Fruit flies have been tracked in an acrylic box by setting up the problem of data association across views and time in the form of a global optimization problem that is solved at every step [20]. Real-time tracking systems for flies were developed using an extended Kalman filter in [21], [22]. Each of these tracking systems implements a nonlinear filtering or optimization method with specialized likelihood functions, data association strategies, and/or experimental design. However, the targets are large with a dark center or appear in an arena constructed to minimize noise, unlike wild mosquitoes.

Filming wild mosquitoes poses special challenges such as low natural lighting and a cluttered dynamic background. At least two cameras are needed to reconstruct the three-dimensional position of individual mosquitoes. A multi-camera setup with a large baseline reconstructs positions accurately, but may be difficult to implement in the field. In a multi-target scenario, one must also address the data-association problem, which entails assigning image blobs to targets across multiple views and time steps. With the typical number of mosquitoes in the swarms we studied[1], the data-association problem is non-trivial. The challenge lies in tracking small, highly maneuverable targets that appear as dots or faded streaks in noisy images with frequent occlusions.

This paper describes an automated multi-target tracking system that reconstructs the three-dimensional flight kinematics of individual mosquitoes in wild swarms. We collect data using two cameras operating synchronously at 25 frames per second. (The frame rate is limited by the ambient light.) The cameras and a laptop are powered by an uninterrupted power supply (UPS) for up to thirty minutes. The mosquitoes appear as dark streaks or dots on a light background. At high speeds, the mosquito streaks fade, making them hard to detect and even harder to track. Because the swarms are dense, occlusions are frequent and often appear in both camera frames. We tested the system by filming swarms and mating events of *An. gambiae* in a rural village in Mali in August 2010. Fig. 1 shows a pair of magnified and enhanced sample frames from this field

[1]Department of Aerospace Engineering, University of Maryland, College Park, MD 20742, USA. Send correspondence to dpaley@umd.edu.

[2]Laboratory of Malaria and Vector Research, National Institute of Allergy and Infectious Diseases, Bethesda, MD 20892, USA

[3]Malaria Research and Training Center, Faculté de Médecine, de Pharmacie et d'Odontostomatologie, Université de Bamako, Bamako, Mali

[4] US Pacific Basin Agricultural Research Center, Agricultural Research Service, U.S. Department of Agriculture, Hilo, HI 96720, USA

[1]Typical swarms in the field site where we filmed ranged between 30–100 mosquitoes, however other sites are known to have swarms of up to 1000 mosquitoes.

experiment.

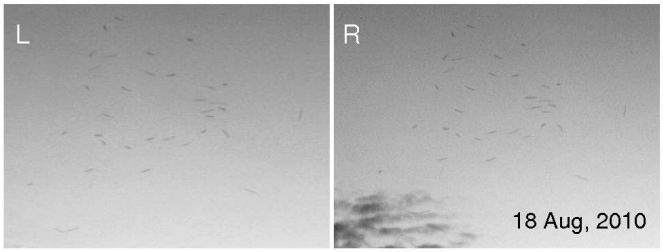

Fig. 1. The pair of images above are magnified and enhanced versions of raw footage obtained from the authors' field work in Mali.

In order to track each insect in a wild swarm, we implemented a probabilistic multi-target tracking framework. Specifically, the contributions of this paper are as follows: (1) we provide a measurement likelihood function that utilizes the properties of image streaks such as midpoint and endpoint locations to extract insect position and velocity; (2) we provide methods to improve data association in noisy images by adaptively seeking missing measurements and splitting occluded blobs into individual measurements; and (3) we present validated tracking results in the form of three-dimensional trajectories of wild mosquito swarming and mating events filmed in Mali in August 2010. Although we describe the experimental method and tracking algorithm for mosquito swarms, the techniques presented in this paper may be beneficial for generating trajectory data for other insect swarms in the field or laboratory.

The tracking system is implemented in MATLAB and consists of two parts: an automated component that outputs track segments called tracklets and a human-supervised component that is used to verify and combine the tracklets into full-length tracks. Tracklets produced by the automated component typically range between 15–25 frames (0.6–1 s) long and can be used to extract position and velocity data for 80% of the swarm. The human-supervised component uses a particle filter to combine tracklets into individual mosquito tracks. It takes approximately 20 minutes to generate a 10-second track (250 frames). When validated using data filmed in Mali in August 2010, the tracking system produced 30–40 second trajectories of individual mosquitoes in swarms of 6–25 mosquitoes. We have reconstructed six swarms and six mating events from these data. We evaluate the performance of the automated component of the tracking system using an established metric based on position error and the number of targets tracked (cardinality); tracking accuracy was also evaluated using two independent rigs to simultaneously track the same swarm.

The paper is organized as follows: Section II provides a background on multi-target tracking, data association, and tracking performance evaluation. Section III presents the novel components of the tracking system, including the likelihood function and occlusion reasoning. Section IV describes the data collection, validation methods, and performance evaluation; it also includes representative kinematic data for wild mosquito swarming and mating events. Section V summarizes the paper and our ongoing analyses of the kinematic data.

II. Probabilistic tracking and data association

Our aim in designing the mosquito tracking system was to combine nearly indistinguishable measurements available from stereo images recorded at discrete times into trajectories that represent real mosquitoes (targets). We represent target i at time step k by the state vector $\boldsymbol{X}_i[k] \in \mathbb{R}^6$, which contains the target's instantaneous three-dimensional position and velocity. In a Bayesian framework, a tracking algorithm recursively iterates through two steps: the update step and the predict step. The update step uses a measurement model to revise the estimate based on new observations. The predict step integrates a motion model to obtain the target state at the time of the next measurement. The measurement $\boldsymbol{Z}_i[k] \in \mathbb{R}^6$ in our case consists of the two-dimensional positions of the midpoint and two endpoints of an elongated blob in an image that corresponds to the motion-blurred silhouette of a flying mosquito in each of the two images. Assuming motion model $\boldsymbol{F}$ and measurement model $\boldsymbol{H}$, the state of target i satisfies

$$\begin{aligned} \boldsymbol{X}_i[k] &= \boldsymbol{F}(\boldsymbol{X}_i[k-1], \boldsymbol{w}) \\ \boldsymbol{Z}_i[k] &= \boldsymbol{H}(\boldsymbol{X}_i[k], \boldsymbol{n}), \end{aligned} \tag{1}$$

where $\boldsymbol{w}$ and $\boldsymbol{n}$ denote disturbance and noise values, respectively.

Because of noise, disturbances, and approximations in $\boldsymbol{F}$ and $\boldsymbol{H}$, the state estimate is a random quantity represented in the form of a probability density function (pdf). We recursively construct the filtering pdf $p(\boldsymbol{X}[k]|\boldsymbol{Z}^k)$ of the joint state $\boldsymbol{X}[k]$ at timestep k given the set $\boldsymbol{Z}^k$ of all measurements up to k using the conditional probabilities $p(\boldsymbol{Z}_i[k]\big|\boldsymbol{X}_i[k])$ and $p(\boldsymbol{X}_i[k]\big|\boldsymbol{Z}^{k-1})$. The probability $p(\boldsymbol{Z}_i[k]\big|\boldsymbol{X}_i[k])$, known as a likelihood function, is the conditional probability of measurement $\boldsymbol{Z}_i[k]$ given the state $\boldsymbol{X}_i[k]$. The filtering pdf can be obtained with minimum mean-square error provided the models $\boldsymbol{F}$ and $\boldsymbol{H}$ are linear and the noise values $\boldsymbol{w}$ and $\boldsymbol{n}$ are Gaussian. Otherwise, it is possible to use suboptimal methods such as an extended Kalman filter [23], which predicts and updates the estimate using a first-order linearization of (1), or a particle filter [24], [25], which represents the target state as a point-mass distribution. A particle filter is attractive in that it relaxes most restrictions on the target and measurement models and the disturbance and measurement noise, but a particle filter is computationally burdensome for a large, joint, state space. We address the computation-size problem by making the assumption that the targets do not interact at short time-scales (<40 ms), which allows us to use a separate filter for each target. Particle-filtering methods also allow us to encode extra information such as the velocity of the mosquito using the streak endpoints. (See Supplementary document for a description of particle filtering.)

A multi-target tracking system must associate measurements and targets. A target-based method associates each target to a measurement [23], whereas a measurement-based method associates each measurement to a target [27]. A measurement-based method can inherently handle a variable number of targets, which may appear and disappear from the field of view. The reliability of the association depends on the proximity of the actual measurement to the predicted measurement, which is

produced from the target estimate using the motion model and the measurement model; measurement proximity is determined using the position likelihood function.

Our choice of a data-association strategy is based on speed, variability, and density of targets in the image. A nearest-neighbor association is target-based and associates the predicted measurement to the nearest measurement. It works well in low-target densities with high frame rates [28], [21], but results in duplicate tracks and incorrect associations at high target densities. A global nearest-neighbor (GNN) association avoids duplicate assignments by minimizing a global assignment [29]. GNN has been successful in tracking dense aggregations [17], [30] in which the number of targets are fixed and move in two dimensions (so that target overlap is rare), however, the possibility of a variable number of targets and frequent occlusions make it difficult to use GNN without additional heuristics. Short-duration occlusions can be addressed using motion coherence [31], [18], [20], whereas long-duration occlusions can be overcome by methods that minimize a global cost function over all measurements in a sliding window [19]. However, it is not clear how an offline global optimization method might address a low probability of detection which is common in our datasets. Instead, we selected a measurement-based method called the multiple-hypothesis tracker (MHT), which looks into future assignment probabilities before making a decision on the current assignment [27]. Within the MHT we use a motion model at each step to search for missing measurements.

A hypothesis in MHT is a combination of measurement-target assignments that satisfy the following two rules [27]: (i) a target is not associated to more than one measurement and (ii) a target is only associated to a measurement that lies within its gating volume. The gating volume or validation region is generated from the difference $\nu = \boldsymbol{u} - f(\boldsymbol{r})$ between the position measurement $\boldsymbol{u} \in \mathbb{R}^2$ and the predicted position measurement $f(\boldsymbol{r}) \in \mathbb{R}^2$. Let S be the covariance of ν, which is also called the innovation. If a measurement that is normally distributed about the true value lies within the validation region, the weighted norm $\|\nu\|^2 = \nu^T S^{-1} \nu$ satisfies $\|\nu\| < t_{gate}$. (The quantity $\|\nu\|$ is also called the Mahalanobis distance [32].) For example, a threshold value $t_{gate} = 16$ defines a region around a predicted measurement with 99.97% probability of containing the actual measurement [23]. A measurement may be assigned to an existing target, a new target, or a false alarm. As time progresses each hypothesis gives rise to successive hypotheses resulting in an exponential growth in time. Hypothesis reduction strategies include applying a threshold on track probability, choosing a few best hypothesis [32], and clustering the targets [27]. (See Supplementary document for a description of MHT.)

Measures of effectiveness to evaluate a multi-target tracking algorithm include the following [33]: track initiation delay (timeliness), position and velocity errors (accuracy), fragmentation and identity swaps (continuity), and number of targets tracked (cardinality). Performance evaluation also includes visual verification, running the algorithm on simulated data [18], comparison with manually generated ground truth, and reconstructing the same event from independent camera systems [31]. We evaluate the performance of our tracking algorithm using manually generated ground-truth as well as filming using two independent camera systems. For comparing tracking results with ground truth we use the Optimal SubPattern Assignment (OSPA) metric [34], which is a well-established metric for evaluating multi-target tracking algorithms [33] that allows comparison of sets with differing cardinality. (See Supplementary document for details on computing OSPA.)

III. The mosquito tracking system

The mosquito tracking system takes a sequence of stereo image pairs as input and produces three-dimensional tracks as output. Fig. 2 depicts a block diagram of the tracking system, which was created in MATLAB and includes an automated tracking algorithm and a human-supervised component. Image streaks are modeled as straight lines; we extract the midpoint and endpoints as measurements. We find missing measurements using a gating volume generated around predicted measurements. Measurement pairs, i.e., one from each camera, that satisfy the epipolar constraint [38] are selected for data association. We again use a gating volume to assign measurements and targets to independent sets called clusters. Instead of generating definite tracks, hypotheses connecting measurements to targets are propagated to the next step using a particle filter. Based on the probability of each hypothesis at the current time step, the number of hypotheses at a previous time step are reduced to a single assignment. A particle filter verifies and combines tracklets under human supervision and the combined tracks are passed through a Kalman smoother. The tracking algorithm is summarized in Table I.

TABLE I

Mosquito tracking algorithm

Input: Sequence of synced images from a stereo-camera setup, camera calibration matrices, parameters in Table II

Output: Estimated three-dimensional mosquito trajectories

For each time step k:

1: *Extract measurements:* Model each blob as a straight line and find the midpoint and endpoints.

2: *Find missing measurements, if any:* Ensure that each hypothesized target has at least one measurement within the gating volume; if not, lower the intensity threshold. If a measurement is found append it to the existing set of measurements.

3: *Validate:* Use the epipolar constraint (3) to generate valid measurement pairs, one from each camera view.

4: *Cluster:* Use gating volume of each target within a cluster to add measurements to that cluster. A cluster is the smallest set of measurements and targets that exist independently; combine/divide existing clusters as needed.

5: *Compute hypotheses:* Generate hypotheses for each cluster, and compute probabilities.

↪ *Resolve occlusions:* If an occlusion is detected split the image blob into individual streaks as described in §III-C and recompute the hypotheses.

6: *Hypothesis reduction:* Based on the most probable hypothesis at k and scanback range N_s, reduce the number of hypotheses at $k - N_s$ to a single assignment.

7: *Initialize & update:* Initialize tentative targets from unassociated measurement pairs; resample target states based on hypotheses using the three-dimensional estimate and velocity likelihood function (7). Each new target forms a new cluster.

8: *Predict:* Use the constant velocity motion model with random (Gaussian) disturbance to propagate hypotheses to timestep $k + 1$.

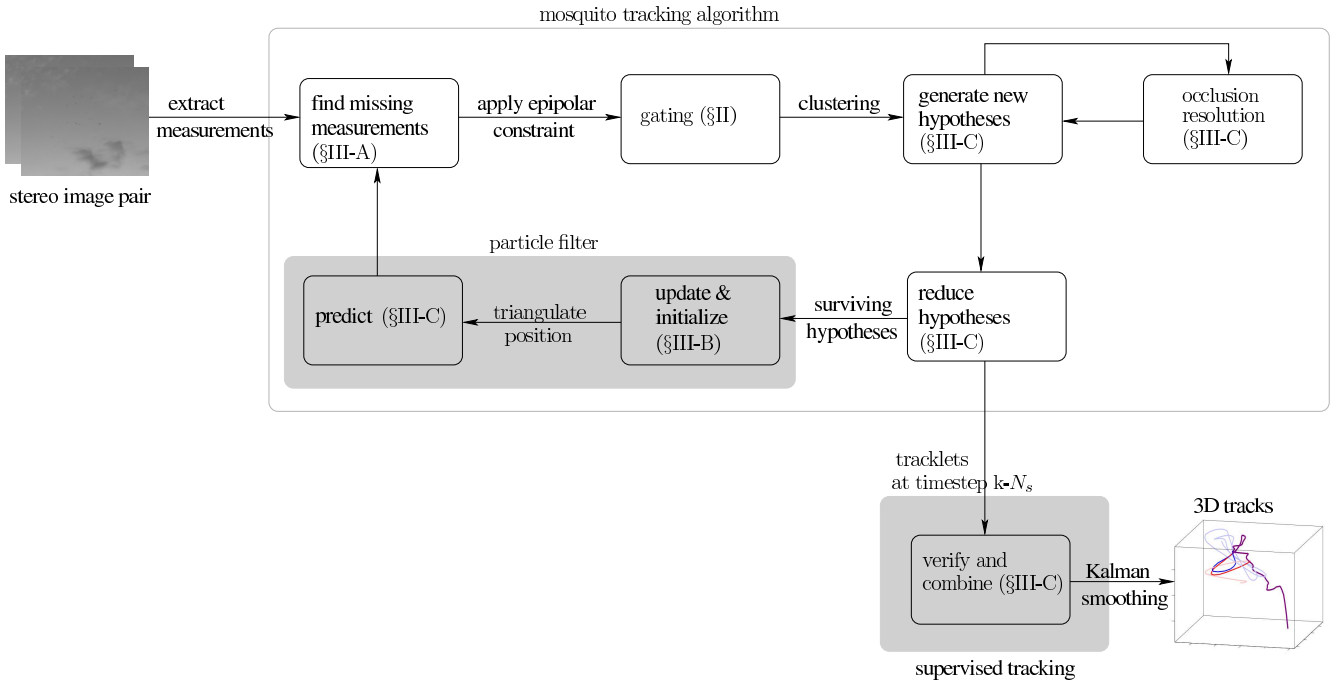

Fig. 2. Block-diagram of the mosquito tracking system.

The remainder of this section describes the novel aspects of the tracking system that we designed to improve its accuracy and level of automation. First, we describe an image-processing technique to find missing measurements during image segmentation. We then describe the measurement model that is used to extract velocity information from streaks. Lastly, we present the data-association method, including the strategy to detect and address mosquito occlusions.

A. Extracting measurements

During observation of mosquito swarms, which typically appear silhouetted in front of swaying trees under a cloudy sky, it may not be possible to use a static (mosquito-free) background to segment the mosquitoes out of the image stream. Instead we create a dynamic background by choosing the highest intensity point within a sliding window [35]. Let $B_{u,v}$ be the background image value at the pixel position (u, v) and $t_{win} = 2d+1$ be the width of the sliding window centered at time step k. The background value at time k is

$$B_{u,v}[k] = \max_{i \in [k-d, k+d]} B_{u,v}[i]. \tag{2}$$

The foreground F is obtained by subtracting the background B from the current image I and applying an intensity threshold t_{int}, i.e., $F_{u,v}[k] = \max(I_{u,v}[k] - B_{u,v}[k], t_{int})$. We automatically select the value of t_{int} by running the background subtraction algorithm recursively on different segments of the image sequence until the number of blobs detected are within an acceptable range of the expected number of mosquitoes. We extract blobs using the *regionprops* routine in MATLAB, which performs connected-component labeling to extract features such as centroid, area, and bounding ellipse. We remove large insects and birds from the foreground by applying a threshold on the blob area. (See Table II for the values of the threshold parameters.)

Due to the duration of the camera exposure ($\delta t_e = 25$ ms)[2], fast mosquitoes (1–4 m/s) appear as elongated image blobs or streaks. Depending on the mosquito speed, the streaks may fail to appear in the foreground for a given value of t_{int}. Existing strategies for low signal-to-noise environments include the track-before-detect approach [36], which permits raw sensor data as the input. The success for track-before-detect relies on the low target density and relatively straight movement of targets in the measurement space [37]. However, using raw sensor data is not a viable option for mosquito tracking, because it generates more false targets than observed in a single noisy image. Instead, we search for the missing streak in a new foreground generated using a threshold $t'_{int} = 0.75 t_{int}$. The search is performed within the gating volume of the predicted measurement. If a missing measurement is found, it is added to the list of existing measurements.

A measurement $\boldsymbol{Z}^c = [\boldsymbol{e}^c_-, \boldsymbol{u}^c, \boldsymbol{e}^c_+]^T$ from camera c contains the image locations of a streak's start $\boldsymbol{e}^c_-$, midpoint $\boldsymbol{u}^c$, and end $\boldsymbol{e}^c_+$. These values are extracted from a blob by modeling it as a straight line along the major axis of the bounding ellipse. The streak therefore represents a perspective projection of the mosquito trajectory for the duration of exposure δt_e. Let $\tilde{\boldsymbol{u}}^c = [(\boldsymbol{u}^c)^T, 1]^T$ be the homogeneous representation of $\boldsymbol{u}^c$. A pair of measurements with midpoints $\boldsymbol{u}^1$ and $\boldsymbol{u}^2$, one from each camera, must satisfy the epipolar constraint [38]:

$$\left|(\tilde{\boldsymbol{u}}^2)^T F \tilde{\boldsymbol{u}}^1\right| < t_e, \tag{3}$$

where $F \in \mathbb{R}^{3\times 3}$ is the fundamental matrix for the stereo

[2]The duration of exposure (25 ms) is less than the time between frames (40 ms). The remaining time (15 ms) is for image processing.

camera calibration and $t_e \ll 1$ is a value that depends on calibration accuracy. Measurement pairs from a true target satisfy the above constraint; clutter or mismatched measurement pairs should not. We use the midpoint and endpoint locations to define a likelihood functions for position and velocity.

B. Position and velocity likelihood functions

A constant-velocity model suffices to describe the mosquito motion during the exposure, $\delta t_e = 25$ ms. (The streaks are well approximated by straight lines on the image plane.) Let $\boldsymbol{r} \in \mathbb{R}^3$ be the three-dimensional location of the midpoint of a streak. The start and end of the streak are located at $\boldsymbol{r}_- = \boldsymbol{r} - \dot{\boldsymbol{r}}\frac{\delta t_e}{2}$ and $\boldsymbol{r}_+ = \boldsymbol{r} + \dot{\boldsymbol{r}}\frac{\delta t_e}{2}$, respectively. The corresponding point on the image plane is given by the perspective projection model [38],

$$f^c(\boldsymbol{r}) = \Big(\frac{w_1}{w_3}, \frac{w_2}{w_3}\Big), \tag{4}$$

where $\boldsymbol{w} = P\boldsymbol{r} \in \mathbb{R}^3$, and P is the camera projection matrix. Let $\mathbb{N}(\boldsymbol{u}; f(\boldsymbol{r}), \Sigma)$ denote a normal density function evaluated at $\boldsymbol{u}$ with mean $f(\boldsymbol{r})$ and covariance matrix $\Sigma \in \mathbb{R}^{2\times 2}$. Assuming that the measurement is normally distributed about the true value, the likelihood of midpoint $\boldsymbol{u}^c$ given $\boldsymbol{r}$ is

$$P^c_{mp}(\boldsymbol{u}^c \big| \boldsymbol{r}) = \mathbb{N}(\boldsymbol{u}^c; f^c(\boldsymbol{r}), \Sigma_{mp}), \tag{5}$$

We set the diagonal entries of Σ_{mp} equal to the length of the major and minor axes of the streak's bounding ellipse in the streak frame; the off-diagonal entries are zero.

As with the midpoint likelihood function, we assume the endpoint likelihood function is based on a normal density function. However, due to uncertainty in the labeling of the start and end of the streak, the endpoint likelihood function is bimodal. The directional ambiguity is described by a sum of conditional probabilities on the order of endpoints. Let Σ_{ep} be the covariance of the endpoint position in pixels (computed empirically). The endpoint likelihood function is

$$\begin{aligned} P^c_{ep}(\boldsymbol{e}^c_-, \boldsymbol{e}^c_+ \big| \boldsymbol{r}, \dot{\boldsymbol{r}}) =& \mathbb{N}(\boldsymbol{e}^c_-; f^c(\boldsymbol{r}_-), \Sigma_{ep})\mathbb{N}(\boldsymbol{e}^c_+; f^c(\boldsymbol{r}_+), \Sigma_{ep}) + \\ & \mathbb{N}(\boldsymbol{e}^c_-; f^c(\boldsymbol{r}_+), \Sigma_{ep})\mathbb{N}(\boldsymbol{e}^c_+; f^c(\boldsymbol{r}_-), \Sigma_{ep}), \end{aligned} \tag{6}$$

where $\boldsymbol{r}_\pm = \boldsymbol{r} \pm \dot{\boldsymbol{r}}(\delta t_e/2)$. The combined effect of using a pair of points in the endpoint likelihood function (6) is to reduce the set of velocity values along the camera axis, which is otherwise unobservable.

The combined position and velocity likelihood function is

$$\begin{aligned} P(\boldsymbol{Z}|\boldsymbol{X}) &= \prod_{c=1,2} P(\boldsymbol{e}^c_-, \boldsymbol{u}^c, \boldsymbol{e}^c_+ | \boldsymbol{r}, \dot{\boldsymbol{r}}) \\ &= \prod_{c=1,2} P^c_{mp}(\boldsymbol{u}^c \big| \boldsymbol{r}) P^c_{ep}(\boldsymbol{e}^c_-, \boldsymbol{e}^c_+ \big| \boldsymbol{r}, \dot{\boldsymbol{r}}). \end{aligned} \tag{7}$$

Fig. 3 shows the combined position and velocity likelihood function. The likelihood function (7) is used to weight the particles in the resample step of the particle filter. We update the position estimates using triangulation [39], thereby effectively marginalizing out the position from the combined position and velocity filtering pdf.

A velocity likelihood function improves the reliability of data association by placing predicted measurements closer to the actual measurements. We compared the absolute velocity estimation error between a standalone position likelihood function and the combined position and velocity likelihood function (7). To create ground-truth data we isolated a single mosquito track in both camera frame for 8 seconds. We then interpolate the position values to every 1/800th of a second. These values were then used to create an artificial mosquito streak during the time of exposure δt_e from a 1 cm sphere. We then project the streak on a pair of white synthetic left and right camera images with resolution 1392 $\times$ 1024 pixels. To achieve a faded-streak effect, we reduce the intensity value of a pixel by 30 every time it is visited on the screen during the time of exposure. Mosquito motion normal to the image plane results in darker, shorter streaks, where as motion parallel to the image plane results in lighter, longer streaks (Fig. 4a). We tracked this dataset using multiple Monte-Carlo runs of a particle filter. The combined position and velocity likelihood function performed better than the standalone position likelihood function, with an average improvement in mean absolute velocity error of 27% (see Fig. 4b).

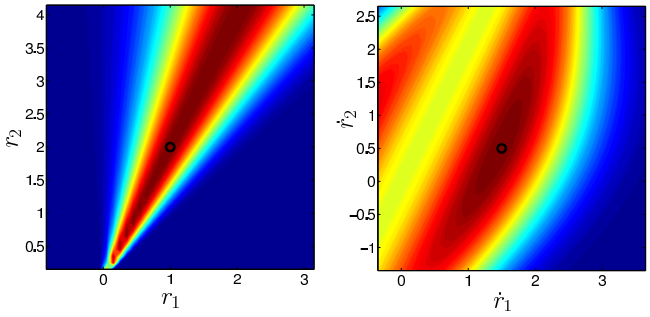

Fig. 3. Top-down view of the (a) position and (b) velocity likelihood functions on a plane orthogonal to the image and parallel to the camera axis. The camera is located at $(r_1, r_2) = (0, 0)$; the black circle is the true value.

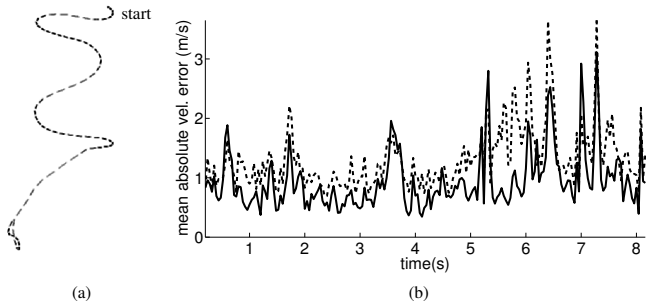

Fig. 4. (a) Two dimensional projection of a 1 cm sphere tracing the interpolated path of a mosquito. (b) Mean of absolute velocity estimate error over multiple Monte-Carlo runs using a standalone position likelihood function (dash-dot) compared to a combined position and velocity likelihood function (solid).

C. Data association and occlusion resolution

Prior to weighting a target distribution with a likelihood function, we must first address the data-association problem. The mosquito data-association problem is challenging due to the variable number of targets. To mitigate the uncertainty

in association (for example, did the paths of two mosquitoes cross each other, or was it a close encounter?), we use a deferred-logic method called the multiple hypothesis tracker (MHT) [27]. Each assignment of measurements to targets is set aside as a hypothesis and acted upon in a future time-step when we are more certain. The certainty is computed using the probability of a hypothesis that depends on the innovation $\nu^c = \boldsymbol{u}^c - f^c(\boldsymbol{r})$ of each measurement-target assignment in the hypothesis, the probability of detection of actual targets, and the covariance of the predicted measurement S.

We reduce the number of hypotheses by clustering and prune them by selecting a few best hypotheses based on their probability at each step. Clustering is performed by dividing the measurement and hypothesized targets at each step into independent sets. At each time step, measurements are associated to each cluster based on the combined gating volume of all targets within the cluster. Measurements that do not belong to any cluster form their own clusters. Two clusters that consist of the same measurement are combined to form a single cluster. Similarly, we split clusters that consist of targets only assigned to a single measurement. Hypotheses are computed for each cluster independently. Hypotheses within a large cluster (more than 10 measurements) are limited to a single localized global nearest neighbor assignment [29]. Using a single scanback [27] at each step, we choose the hypothesis with the highest probability to reduce to one the number of hypotheses at the previous step. Child hypotheses resulting from a pruned parent hypotheses are also removed.

New targets are automatically initialized from unassigned measurements and confirmed if they are tracked for more than three frames. New target distributions are sampled from a normal distribution with a low standard deviation in position (5 mm) about the triangulated point, and a large standard deviation in velocity (500 mm/s) about zero. The combined likelihood function resamples the distribution to equally favor particles getting projected on either side of the streak in the next timestep.

Occlusions are not directly addressed as part of any data-association strategy, because existing strategies assume that motion coherence will automatically associate the right tracks in a future timestep. In our case, occlusions undermine the velocity estimate, making future associations less reliable. An occlusion is detected if (a) two measurement pairs within a hypothesis consist of the same measurement from a single camera, or (b) multiple hypotheses assign the same measurement to two or more targets. We interpret an occlusion as a combination of individual streaks, which are then used to extract velocity information as described in §III-A.

In order to cluster the pixels in an occlusion blob we use the information about the number of mosquitoes hypothesized in the occlusion as well as their position and velocity estimates to model the blob as a mixture of Gaussians. An expectation-minimization algorithm [40] uses position estimates for initial means and velocity estimates for covariance in each dimension to hard-cluster the pixels into individual streaks. This set of individual streaks is used as an initial guess to soft-cluster[3] the pixels into more accurate overlapping streaks. Using the shortest distance of a pixel from the line that passes through the split streak, we allow multiple assignments of each pixel to individual streaks. Fig. 5 shows four instances of splitting an occluded blob into two individual mosquito streaks.

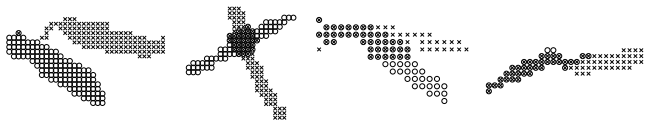

Fig. 5. Four examples of occlusion resolution by soft-clustering the occluded blobs into two individual mosquito streaks. Each streak is denoted by a marker type.

IV. Data collection & Tracking results

To validate the mosquito-tracking system and film mosquito swarms in the field, we used a pair of phase-locked Hitachi KP-F120CL cameras in a stereo configuration. Each camera captured 10-bit images at 25 frames per second and 1392 $\times$ 1040 pixel resolution. Fig. 6 shows a schematic of the data collection system. The video streams were recorded using a 2.8 GHz quad core laptop, an Imperx FrameLink Express frame grabber (Imperx Inc, Boca Raton, FL USA), and Streampix 5 software (Norpix Inc, Quebec, Canada). Each camera was calibrated onsite using a checkerboard and the MATLAB Calibration Toolbox [41]. Reprojection error, which is a measure of calibration accuracy, was in sub-pixels for each camera. Relative camera orientation and position was determined by extrinsic calibration by taking multiple pictures of a stationary checkerboard with both cameras. During filming, the camera height, azimuth, and elevation were recorded to create a ground-fixed reference frame. We used a Kestrel 4500 portable weather station (Nielsen-Kellerman, Boothwyn, PA USA) to sample other environmental factors such as wind velocity and humidity at 0.1 Hz.

Filming was done in the village of Donéguébougou, Mali in Western Africa. Donéguébougou is 29 km north of Bamako and has been the site of previous research on *An. gambiae* mosquitoes [5], [10]. Swarms formed approximately 20 minutes after sunset, initially with only one or two males then increasing in numbers, and lasted for 20 minutes. Most couples were seen 5-10 minutes after the swarm was first observed. Couples formed only for a few minutes during this period, then were no longer observed, though the males continued to swarm for many minutes after the last couple had formed. We filmed swarms of *An. gambiae* that formed over bare ground or markers.

Female mosquitoes are difficult to detect and track because they fly faster than the average male (see tracking results), and appear as a faint streak much of the time. However, a mosquito couple is distinguishable to the human eye due to its distinct flying pattern and darker appearance against the sky. Upon spotting a couple we noted the frame number displayed on

[3]Soft clustering allows a single pixel to be assigned to more than one cluster, whereas hard-clustering assigns each pixel to exactly one cluster.

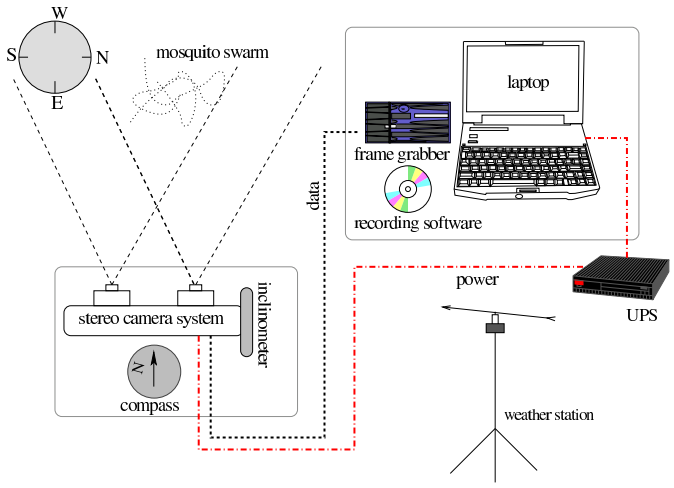

Fig. 6. The data collection system consists of the stereo camera rig and a laptop powered by an uninterrupted power supply (UPS). A capture signal is sent from the frame grabber to record frames in sync. A Kestrel weather station records environmental factors such as temperature and wind velocity at 0.1 Hz.

the laptop screen. The couples were located after filming by manually reviewing the video footage at the designated frame. Out of the two mating mosquitoes, the female mosquito was identified as the mosquito that entered the swarm last. We tracked the pair, first as a couple and then individually, by playing the sequence backwards. Parameter values used for data collection and tracking are described below and in Table II. The validation and evaluation of tracking performance follow.

TABLE II
PARAMETER VALUES USED FOR DATA COLLECTION AND TRACKING

Parameter	Value	Description
b	20 cm	Stereo camera configuration baseline
t_{win}	7 frames	Sliding window for segmentation
Σ_{ep}	diag{4,4} pixels^2	Covariance of endpoint error
σ_w	100 m^2/s^4	Covariance of disturbance
δt_e	25 ms	Duration of camera exposure
t_{gate}	16	Threshold for gating volume
t_e	.5	Threshold on epipolar constraint
t_{area}	(20, 150)	Minimum and maximum blob areas
N_s	1 frames	Scanback for MHT
N_p	200	Number of samples in particle filter

A. Parameters used for data collection and tracking

The camera baseline b, i.e., the distance between cameras, affects the disparity Δu in pixel positions of an object in a stereo camera setup [38]. A large disparity reduces uncertainty along the camera axis, which in turn improves accuracy as well as the ability to resolve occlusions. For a stereo-camera configuration with focal length f and no vertical offset between centers, the baseline and disparity are related according to $\Delta u = (bf)/z$ [38], where z is the distance along the camera axis of the target from the stereo setup. The overlap between camera views is $(I_w - \Delta u)/I_w$, where I_w is the image width resolution in pixels. A large overlap is desirable for maximum coverage. Since the majority of swarms were filmed with $1.5 \leq z \leq 2.5$ m, we selected a baseline of 20 cm to achieve 80–90% overlap and 3–5 pixel difference between two mosquitoes that are 3 cm apart (approx. 2 body lengths) along the camera axis.

In addition to the intensity threshold t_{int} described in §III-A, foreground segmentation requires setting the sliding window t_{win} and a threshold on area of the blobs t_{area}. We selected $t_{win} = 7$ frames centered on the current frame, although swarms filmed at short ranges required a sliding window in the range of 3–5. We computed the area-threshold limits $20 \leq t_{area} \leq 150$ from several different swarms to achieve the best rejection of noise as well as large insects.

The covariance $\Sigma_{ep} = \mathrm{diag}\{4, 4\}$ pixels^2 for endpoints was computed by manually selecting the endpoints of streaks in a random sampling of frames and comparing with the calculated value. The disturbance $\boldsymbol{w}$ for the constant-velocity model was sampled from $\mathbb{N}(0, 100\ \mathrm{m}^2/\mathrm{s}^4)$, whose covariance was found by fitting a normal distribution to the acceleration values of manually generated tracks.

B. Validation and performance of tracking system

We tested the position accuracy of our tracking system using a calibration checkerboard with squares of known dimensions by manually clicking pairs of points whose separation distance was in the range 3–40 cm. This method yielded an error of 5 $\pm$ 5 mm for 50 pairs. Average position error by tracking an artificial mosquito projected on the stereo images as in §III-B over multiple runs was 5 $\pm$ 4 mm. We also reconstructed tracks from a single swarming event on Aug. 29 using two independent stereo camera rigs. We created a common reference frame by measuring the height, azimuth and elevation of the cameras. The videos were time-synced using a laser pointer flashed at the end of the sequence. The mean distance between independent tracks of the same mosquito (200 data points) was 4.4 $\pm$ 1.3 cm, although up to 3 cm error can be attributed to the inter-frame time difference between the camera systems (caused due to delay within a single frame that was used to match the laser flash). A mosquito flying at an average speed of 1.5 m/s will cover 3 cm in 1/50th of a second.

Fig. 7 shows the results of using the OSPA metric (see Supplementary document, equation (4)) to compare tracks from the multi-target tracking system to the manually generated ground-truth. We tested two swarms with 10 and 20 mosquitoes, respectively. The order parameter and the cut-off parameter for computing OSPA values were set 2 and 50 mm respectively. Decomposing OSPA into position and cardinality errors shows that the average root mean square (RMS) position errors in the 10- and 20-mosquito swarms were 2.17 $\pm$ 0.58 cm and 2.3 $\pm$ 0.46 cm, respectively. Correspondingly, average absolute position errors for the 10- and 20-mosquito swarms were 1.74 $\pm$ 0.56 cm and 2.03 $\pm$ 0.47 cm, respectively. A low cardinality error was often accompanied by relatively high position error during periods when the swarm was dense, because of occlusions and false tracks. As would be expected for a stereo setup, position error was highest (44%) in the range measurement (along the camera optical

axis) as compared to either of the other two dimensions. OSPA was larger for the 20-mosquito swarm, mainly due to cardinality errors. The position error is likely a consequence of image noise, which resulted in partially segmented streaks. (We mitigate this problem by filtering trajectory data using a Kalman smoother.) Average reprojection error on the images was less than 2 pixels.

The labeling error, which captures track continuity and identity swaps, was computed separately. An identity swap results in a labeling error of 2 before or after the swap in the sequence. Track fragmentation results in a labeling error of 1 after the disconnect occurs. We randomly selected 100 instances of 25 continuous frames in a swarm of 10 mosquitoes. The average labeling error (most of which was due to track fragmentation) was 2.1 ± 1.4 tracks. A simple average of track lengths across six swarms ranged between 15–25 frames corresponding to 0.6–1 s. Track fragmentation occurs due to early terminations, which can be caused by the following:

- Partially segmented streaks due to noise, cloudy background, and clutter. Partially segmented streaks in one frame often violate the epipolar constraint. Decreasing the intensity threshold to get full streaks adds noise to the measurements. (A possible solution that we are exploring in ongoing work is to reconstruct the streak using velocity estimates.)
- Occlusion between a tracked and untracked target. Occlusions between a tracked (known) targets and an as yet uninitialized target are not detected. The success rate of surviving such an occlusion depends on the motion of the tracked target after the occlusion. A maneuver or successive occlusion may terminate the track.

C. Tracking results

This section presents a subset of the three-dimensional trajectory data generated using the mosquito tracking system. We filmed twenty-one swarms and thirteen mating events between August 17, 2010 and September 3, 2010. Out of the twenty-one swarms, eighteen formed over bare ground and three formed over natural markers. (A natural marker is an area of high contrast with the rest of the ground such as a patch of grass.) *An. gambiae* can be divided into two incipient species namely the M and S molecular forms [43]. In [10] a strong association between the swarming marker type and molecular form has been found. The M form was found to swarm over natural markers, whereas the S form swarms over bare ground [10]. We collected a few mosquitoes from each swarm and performed a polymerase chain reaction (PCR) test to determine the molecular form. All sequences presented in this paper were of type S. Each day two teams of 3–5 people with identical camera rigs selected separate swarming sites for filming. The swarming sites were usually within a few hundred meters of each other. Swarming sites were surveyed the day before to record average swarm size and location. Filming locations spread throughout the village (see Fig. 8) were chosen based on swarm size (less than about 100 mosquitoes for tractability in tracking) and the presence of few trees or houses in the background (i.e., in the direction of the setting sun). Once filming began, 60–90 s stereo video sequences were recorded as 10-bit synchronized tiff images on separate solid state drives. (The drives were backed up daily on to two separate disks.) A filming session typically produced 5–8 video sequences before it became too dark to film.

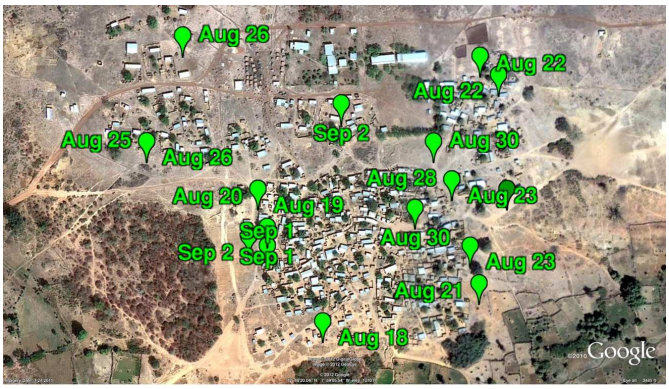

Fig. 8. GPS measurements of filming locations in Mali, Africa ©2012 Google, ©2012 DigitalGlobe.

To create representative trajectory dataset, we selected six video sequences that contain a mating event. We call these the mating sequences. We refer to the mating mosquitoes as the female and the focal male. We also selected six other video sequences with no female present, called the male-only sequences, to produce full-length trajectories of swarming behavior. Trajectory data presented here are from swarms filmed on August 20, 21, 25, 26, 28, and 29 and September 1. Male-only sequences last between 20–35 seconds, whereas mating sequences start a few seconds prior to the detection of female within the field of view and end when the couple flies out of field of view (0–5 s).

Fig. 9 shows the position and velocity of a randomly selected male *An. gambiae* in the Aug. 29 male-only sequence, which was a swarm that formed over bare ground (S molecular form). The swarm consisted of 20 mosquitoes at the beginning of the sequence and dropped to 19 after 10 seconds. The mosquito movement is characterized by quasi-periodic motion in each of the three spatial dimensions. The instantaneous mean position of the mosquitoes in the swarm, i.e., the swarm centroid, is also shown. The origin of the inertial frame is located at ground level under the camera rig; the inertial frame is oriented along east-west, north-south, and vertical directions. The 3σ bounds for position and velocity of all of mosquitoes in this swarm are shown in gray. Swarm size (twice the 3σ bounds) averaged 1.17 m in the horizontal plane and 0.56 m in the vertical. The average swarm size across all planes ranged between 0.52–1.86 m. The average height of the swarm was 1.89 m. The average velocity along each dimension is close to zero with a highest standard deviation in the east-west direction (0.514 m/s) followed by the north-south (0.332 m/s) and the vertical (0.281 m/s).

Fig. 10 shows the ratio between horizontal and vertical speed for each swarm. The Aug. 28 sequence was filmed on a day with relatively high wind (approx. 0.8 m/s) as compared to other sequences. The mosquito movement for that swarm

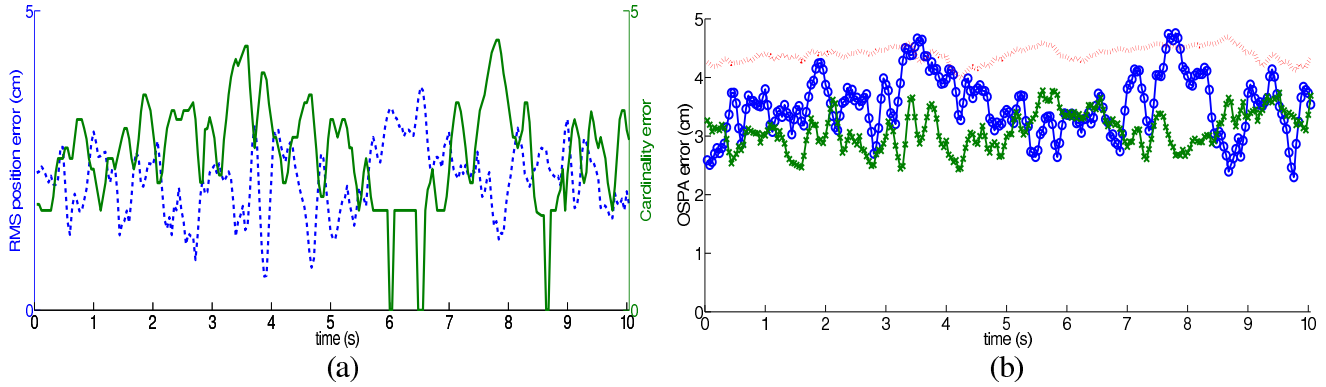

Fig. 7. (a) Position (dashed blue) and cardinality error (solid green) for a swarm of 10 mosquitoes. (b) OSPA error for different methods and swarm sizes: nearest neighbor [42] (dotted) for a swarm of 20 mosquitoes and single-scan MHT for two swarms of sizes 10 (crosses) and 20 (circles), respectively.

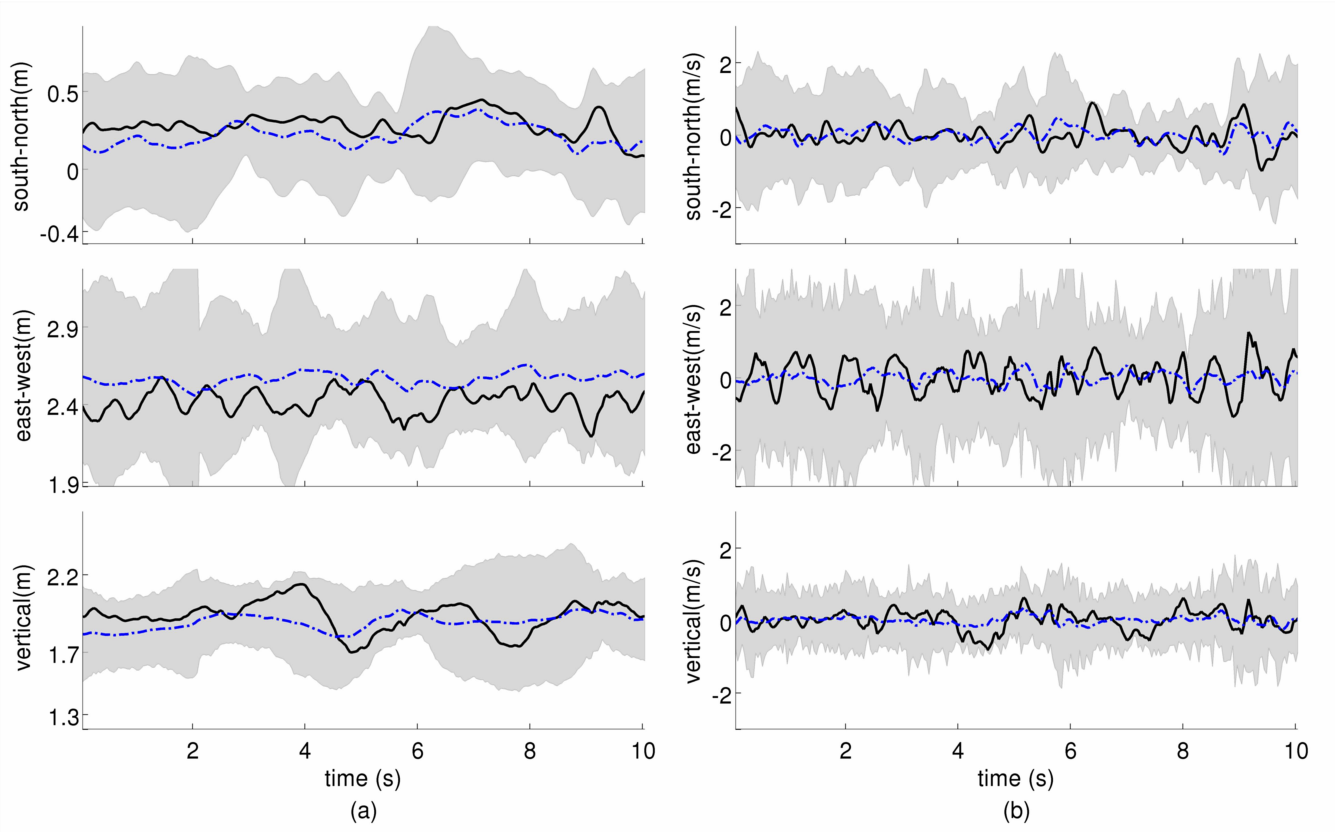

Fig. 9. Instantaneous three-dimensional (a) position and (b) velocity tracks of a male mosquito (black solid) in the Aug. 29 sequence (S molecular form). Mean (dotted blue) and 3 σ bounds (gray) for all mosquitoes also shown. Also see Supplementary video 1.

was characterized by a rolling motion in the direction of the wind and relatively higher vertical velocities. In five out of the six swarms that we used to generate male-only sequences, we witnessed mating events at a later time. The horizontal and vertical speeds of female mosquito that formed couples are also plotted in Fig. 10. Non-parameteric Kruskal-Wallis tests on each dataset show that the average male and female speeds in the same sequence are significantly different for each sequence. The maximum p-value among all mating sequences was 0.0003. (In contrast the maximum p-value for male speeds during the same mating sequence was 0.051.)

Fig. 11 shows the position and velocity of a female mosquito that formed a successful couple in the Aug. 29 sequence. The mating sequence was filmed about a minute after the male-only sequence on the same date. The female appeared in the field of view 5 seconds prior to couping. The movement of the female crosses the 3σ boundaries of the swarm in the north-south dimension. The average speed of the female was higher than the male mosquito until just before the couple forms, when the focal male speeds up. The vertical movement shows that the female stayed predominantly at heights corresponding to the lower half of the swarm. A three-dimensional reconstruction

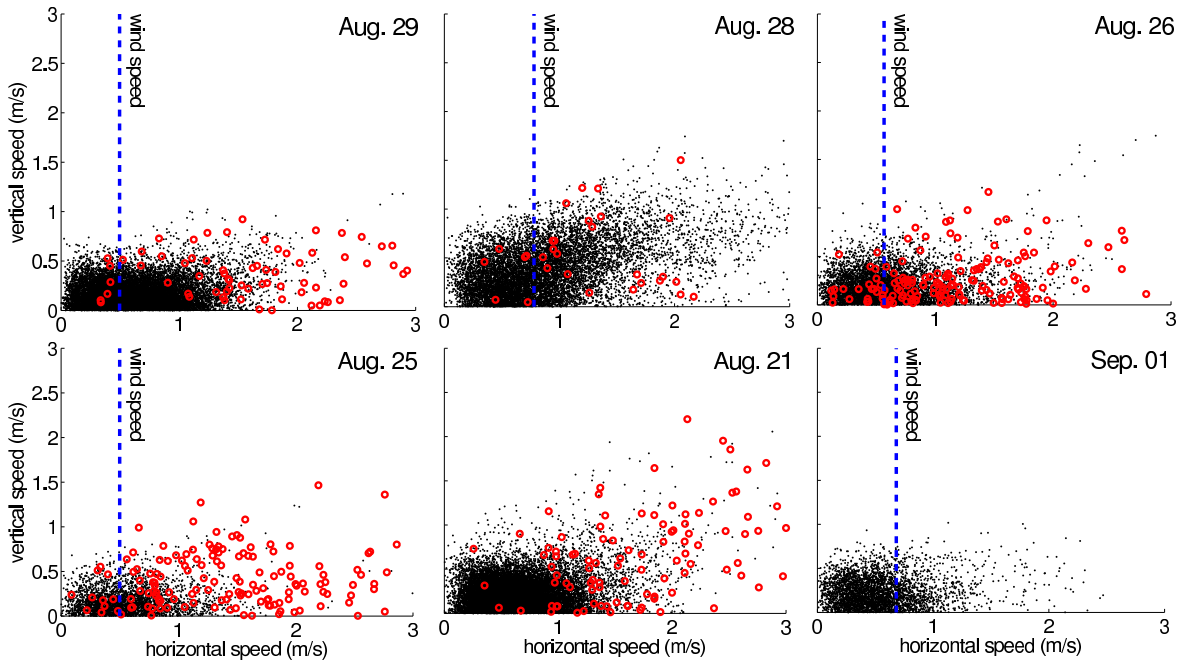

Fig. 10. Instantaneous horizontal versus vertical speed of male mosquitoes in six sequences (black dots). Female speeds (red circles) and average wind speed (blue dashed lines) are shown, if available.

of the mating mosquitoes in six mating sequences is shown in Fig. 12. Across all mating sequences, the female mosquito covered an average 59% more distance than the focal male during the same time interval.

Fig. 13 shows the separation distance and speeds from six mating sequences. The amount of time we observed the females in the swarm before forming a couple was up to 5 seconds. In each mating sequence that lasted longer than 0.5 seconds, the number of close encounters (moments when the separation distance between the mating mosquitoes dropped below 3 body lengths, or 4 cm) with the successful male mosquito was in the range 3–6.

V. Conclusions and Ongoing Work

We describe a tracking system to reconstruct the three-dimensional trajectories of wild mosquito swarms. We address noisy images by adaptively seeking missing measurements and exploit streak orientation and length to extract velocity information. A probabilistic data association method that uses multiple hypothesis (MHT) is modified to address occlusions. We evaluate the system using an established multi-target tracking metric and validate using independent measurements of the same swarm. Tracking results are presented in the form of three-dimensional trajectories of swarming and mating mosquitoes. To date, the data produced from the tracking system described in this paper are an order of magnitude larger (97 trajectories and 55,000 position points) than the last published result [5] on reconstruction of wild mosquito swarms, and the first to contain three-dimensional trajectories rather than three-dimensional positions. In ongoing work, we are investigating these trajectories to characterize swarming and mating behavior.

As part of ongoing work on the tracking system we are working to include the streak intensity in the image as part of the likelihood function. This will help predict the appearance of a mosquito on the image plane as a function of its velocity, thereby allowing the possibility of streak retrieval. Such an approach would for example reduce track terminations and create longer tracklets. Another aspect of the tracking system that we are investigating is the automatic detection of mating events. In order to avoid sifting through video streams to locate mating events, the distinct flying pattern and appearance of the mating couple can be used for automatic detection and backwards tracking of the female. With enough mating events a higher-order motion model (that depends on more than one previous time-step) will automatically predict and detect mating events.

VI. Acknowledgement

This work is built on the prior experience on filming mosquito swarms at the University of Bamako in Mali. Specifically, we thank Alpha Yaro, Adama Dao, and Sekou Traoré for their support. We gratefully acknowledge the travel support from Robert Gwadz of the Laboratory of Malaria and Vector Research at the National Institute of Allergy and Infectious Diseases. Thanks to Diana Huestis for performing the PCR tests. Finally, we would like to thank the residents of Donéguébogou for allowing us to film.

References

[1] J. Charlwood and M. Jones, "Mating in the mosquito, *Anopheles gambiae* sl," *Physiological Entomology*, vol. 5, no. 4, pp. 315–320, 1980.

[2] G. Gibson, "Swarming behaviour of the mosquito *Culex pipiens quinquefasciatus*: a quantitative analysis," *Physiological Entomology*, vol. 10, no. 3, pp. 283–296, Sep. 1985.

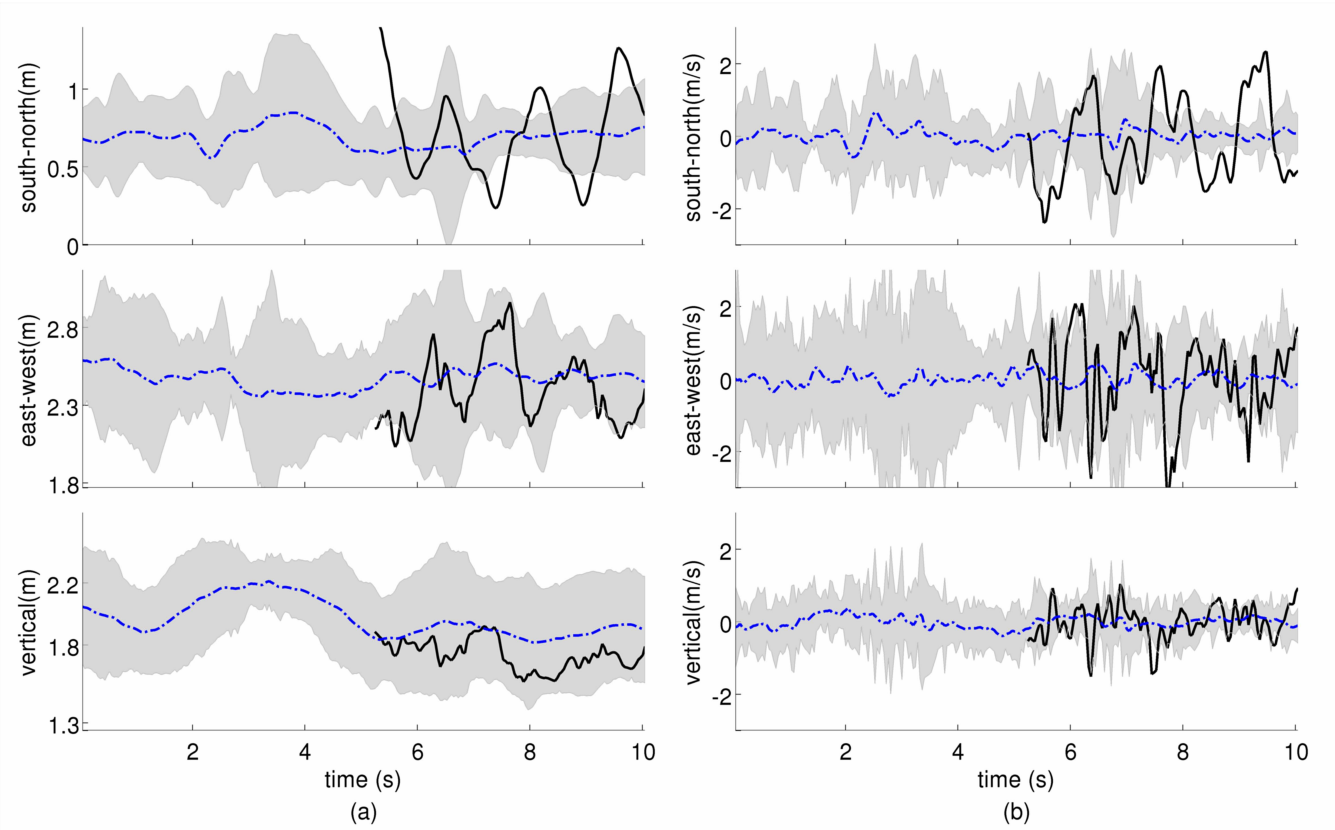

Fig. 11. Instantaneous three-dimensional position (a) and velocity (b) of a female mosquito that coupled in the Aug. 29 sequence, filmed 70 seconds after the male position and velocity shown in Fig. 9. Mean (dotted blue) and 3σ bounds (gray) of all males are also shown.

[3] T. Ikawa, H. Okabe, T. Mori, K. Urabe, and T. Ikeshoji, "A method for reconstructing three-dimensional positions of swarming mosquitoes," *J. of Insect Behavior*, vol. 7, no. 2, pp. 237–248, 1994.

[4] J. Charlwood, J. Pinto, C. Sousa, H. Madsen, C. Ferreira, and V. do Rosario, "The swarming and mating behaviour of *Anopheles gambiae* ss (Diptera: Culicidae) from Sao Tome Island," *J. of Vector Ecology*, vol. 27, no. 2, pp. 178–183, 2002.

[5] N. C. Manoukis, A. Diabate, A. Abdoulaye, M. Diallo, A. Dao, A. S. Yaro, J. M. C. Ribeiro, and T. Lehmann, "Structure and Dynamics of Male Swarms of *Anopheles gambiae*," *J. Medical Entomology*, vol. 46, no. 2, pp. 227–235, 2009.

[6] A. Diabate, A. S. Yaro, A. Dao, M. Diallo, D. L. Huestis, and T. Lehmann, "Spatial distribution and male mating success of *Anopheles gambiae* swarms," *BMC Evolutionary Biology*, vol. 11, no. 1, p. 184, Jun. 2011.

[7] J. Charlwood and J. Pinto, "Male size does not affect mating success (of *Anopheles gambiae* in Sao Tome)," *Medical and Veterinary Entomology*, vol. 16, no. 1, pp. 109–111, 2002.

[8] J. Thailayil, K. Magnusson, H. Godfray, A. Crisanti, and F. Catteruccia, "Spermless males elicit large-scale female responses to mating in the malaria mosquito *Anopheles gambiae*," *Proceedings of the National Academy of Sciences*, vol. 108, no. 33, pp. 13 677–13 681, 2011.

[9] J. Bryan, "Results of consecutive matings of female *Anopheles gambiae* species B with fertile and sterile males," *Nature*, vol. 218, no. 489, 1968.

[10] A. Diabaté, A. Dao, A. S. Yaro, A. Adamou, R. Gonzalez, N. C. Manoukis, S. F. Traore, R. W. Gwadz, and T. Lehmann, "Spatial swarm segregation and reproductive isolation between the molecular forms of *Anopheles gambiae*," *Proceedings of the Royal Society B: Biological Sciences*, pp. 4215–42 222, Sep. 2009. [Online]. Available: http://rspb.royalsocietypublishing.org/content/276/1676/4215.short

[11] C. Pennetier, B. Warren, K. R. Dabiré, I. J. Russell, and G. Gibson, ""Singing on the wing" as a mechanism for species recognition in the malarial mosquito *Anopheles gambiae*." *Curr Biol*, vol. 20, no. 2, pp. 131–136, Jan. 2010.

[12] A. Okubo, "Dynamical aspects of animal grouping: swarms, schools, flocks, and herds," *Advances in Biophysics*, vol. 22, no. 0, pp. 1–94, 1986.

[13] N. Shimoyama, K. Sugawara, T. Mizuguchi, Y. Hayakawa, and M. Sano, "Collective motion in a system of motile elements," *Physical Review Letters*, vol. 76, no. 20, pp. 3870–3873, 1996.

[14] J. K. Parrish and W. M. Hammer, *Animal groups in three dimensions*. Cambridge University Press, 1997.

[15] A. Okubo and H. C. Chiang, "An analysis of the kinematics of swarming of *Anarete Pritchardi Kim* (Diptera: Cecidomyiidae)," *Researches on Population Ecology*, vol. 16, no. 1, pp. 1–42, 1974.

[16] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, and Others, "Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study," *Proc. Nat. Acad. Sciences*, vol. 105, no. 4, pp. 1232–1237, 2008.

[17] K. Branson, A. A. Robie, J. Bender, P. Perona, and M. H. Dickinson, "High-throughput ethomics in large groups of Drosophila," *Nature Methods*, vol. 6, pp. 451–457, 2009.

[18] Z. Wu, N. I. Hristov, T. L. Hedrick, T. H. Kunz, and M. Betke, "Tracking a large number of objects from multiple views," in *Int. Conf. on Computer Vision*. Citeseer, 2009, pp. 1546–1553.

[19] Z. Wu, N. Hristov, T. Kunz, and M. Betke, "Tracking-reconstruction or reconstruction-tracking? Comparison of two multiple hypothesis tracking approaches to interpret 3D object motion from several camera views," in *Workshop on Motion and Video Computing*. IEEE, 2010, pp. 1–8.

[20] H. S. Wu, Q. Zhao, D. Zou, and Y. Q. Chen, "Automated 3D trajectory measuring of large numbers of moving particles," *Optics Express*, vol. 19, no. 8, pp. 7646—-7663, Apr. 2011.

[21] A. D. Straw, K. Branson, T. R. Neumann, and M. H. Dickinson, "Multi-camera real-time three-dimensional tracking of multiple flying animals," *J. Royal Society Interface*, vol. 8, no. 56, pp. 395–409, 2010.

[22] D. Grover, J. Tower, and S. Tavaré, "O fly, where art thou?" *J. of the Royal Society Interface*, vol. 5, no. 27, p. 1181, 2008.

[23] Y. Bar-Shalom, *Tracking and data association*. San Diego, CA, USA: Academic Press Professional, Inc., 1987.

[24] N. J. Gordon, D. J. Salmond, and A. F. M. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *IEEE Proc. on Radar and Signal Processing*, vol. 140, no. 2, pp. 107–113, Apr. 1993.

[25] N. D. Freitas and N. Gordon, *Sequential Monte Carlo Methods in Practice*. Birkhäuser, 2001.

[26] Z. Khan, T. Balch, and F. Dellaert, "MCMC-based particle filtering for tracking a variable number of interacting targets," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 27, no. 11, pp. 1805–1918, 2005.

[27] D. Reid, "An algorithm for tracking multiple targets," *IEEE Trans. on Automatic Control*, vol. 24, no. 6, pp. 843–854, 1979.

[28] A. Veeraraghavan, M. Srinivasan, R. Chellappa, E. Baird, and R. Lamont, "Motion based correspondence for 3D tracking of multiple dim objects," in *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP*, 2006.

[29] H. W. Kuhn, "The Hungarian method for the assignment problem," *Naval Research Logistic Quarterly*, vol. 2, pp. 83–97, 1955.

[30] D. Swain, I. Couzin, and N. E. Leonard, "Real-time feedback-controlled robotic rish for behavioral experiments with fish schools," *Proceedings of the IEEE*, vol. 100, no. 1, pp. 150–163, 2011.

[31] M. Betke, D. E. Hirsh, A. Bagchi, N. I. Hristov, N. C. Makris, and T. H. Kunz, "Tracking large variable numbers of objects in clutter," in *IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2007, pp. 1–8.

[32] I. J. Cox and S. L. Hingorani, "An efficient implementation of Reid's multiple hypothesis tracking algorithm and its evaluation for the purpose of visual tracking," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 18, no. 2, pp. 138–150, 1996.

[33] B. Ristic, B. Vo, D. Clark, and B. Vo, "A metric for performance evaluation of multi-target tracking algorithms," *IEEE Trans. on Signal Processing*, vol. 59, no. 7, pp. 3452–3457, 2011.

[34] D. Schuhmacher, B. Vo, and B. Vo, "A consistent metric for performance evaluation of multi-object filters," *Signal Processing, IEEE Transactions on*, vol. 56, no. 8, pp. 3447–3457, 2008.

[35] A. Cavagna, I. Giardina, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic, "The STARFLAG handbook on collective animal behaviour: 1. Empirical methods," *Animal Behaviour*, vol. 76, no. 1, pp. 217–236, 2008.

[36] D. J. Salmond and H. Birch, "A particle filter for track-before-detect," *In Proc. of American Control Conference*, vol. 5, pp. 3755–3760 vol.5, 2001.

[37] S. Blackman, "Multiple hypothesis tracking for multiple target tracking," *IEEE Aerospace and Electronic Systems Magazine*, vol. 19, no. 1, pp. 5–18, 2004.

[38] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2004.

[39] R. Hartley and P. Sturm, "Triangulation," *Computer vision and image understanding*, vol. 68, no. 2, pp. 146–157, 1997.

[40] A. Dempster, N. Laird, and D. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. of the Royal Statistical Society. Series B (Methodological)*, pp. 1–38, 1977.

[41] J.-Y. Bouguet. Camera Calibration Toolbox for Matlab. [Online]. Available: http://www.vision.caltech.edu/bouguetj/calib_doc/index.html

[42] S. Butail, N. Manoukis, M. Diallo, A. S. Yaro, A. Dao, S. F. Traore, J. M. C. Ribeiro, T. Lehmann, and D. A. Paley, "3D tracking of mating events in wild swarms of the malaria mosquito *Anopheles gambiae*," in *IEEE Conf. Engineering in Medicine and Biology Society*, no. 2, Boston, USA, 2011, pp. 720–723.

[43] G. Favia, A. della Torre, M. Bagayoko, A. Lanfrancotti, N. Sagnon, Y. T. Touré, and M. Coluzzi, "Molecular identification of sympatric chromosomal forms of *Anopheles gambiae* and further evidence of their reproductive isolation." *Insect molecular biology*, vol. 6, no. 4, pp. 377–383, Nov. 1997.

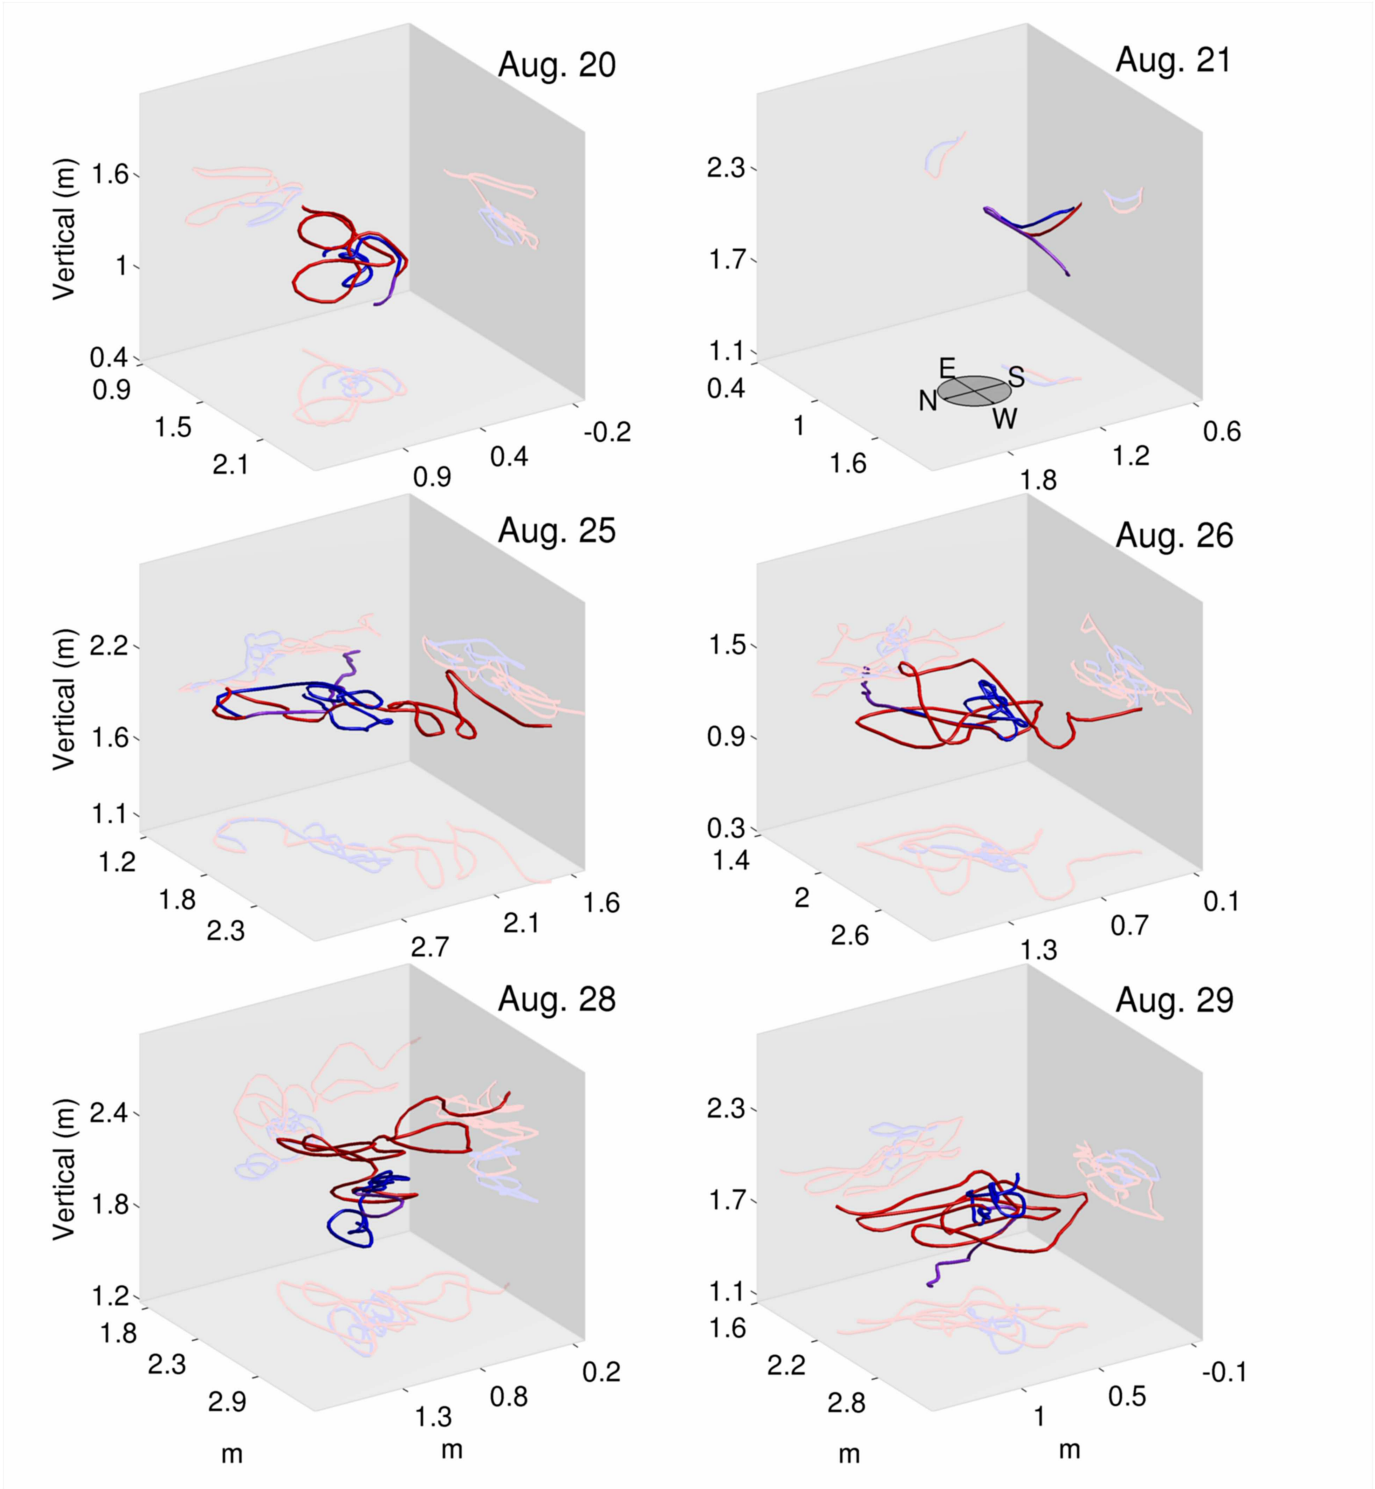

Fig. 12. Three-dimensional reconstruction of *An. gambiae* mating events in the wild. The female mosquito track (red) and male mosquito track (blue) are shown. The couple is shown in purple. Pre-coupling tracks are projected onto two-dimensional planes on each side. (The Aug. 29 mating event is shown magnified in Supplementary document. Also see Supplementary video 2.)

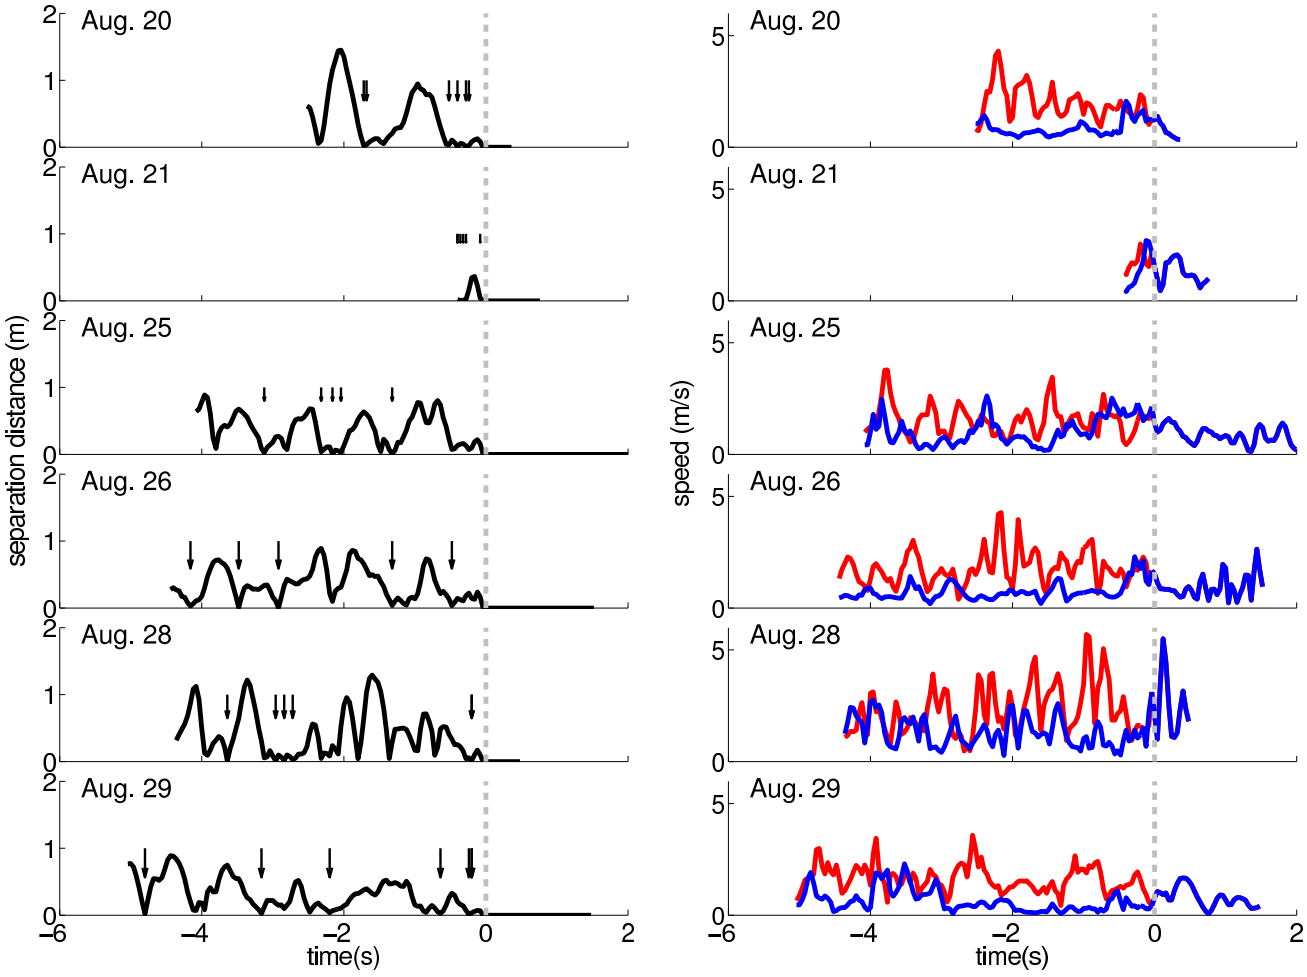

Fig. 13. Relative distance (a) and speeds (b) of mating male and female *An. gambiae* mosquitoes (a) in six mating sequences. Time 0 s occurs when the separation distance first drops below 2 cm. Arrows depict close encounters (separation distance less than 4 cm).