ABSTRACT

Title of Thesis: ANALYZING HONEY BEE FLIGHT WITH EVENT-BASED VISION

Team HiveMIND
Gemstone Honors Program
2024

Thesis Directed By: Dr. Timothy K. Horiuchi
Associate Professor, Department of Electrical and Computer Engineering

An estimate of bee hive activity allows beekeepers and researchers to better understand trends in a colony's health. This work presents a system utilizing an event-based vision sensor (e.g., Dynamic Vision Sensor, or DVS) to track flying bees in real-time with the intent of accurately monitoring the flow of bees in and out of an *Apis mellifera* colony. Neuromorphic event-based vision sensors like the DVS are well-suited to the detection of small, fast-moving bees with minimal latency due to the asynchronous pixels. Rather than processing and

transferring full images, these pixels detect changes in brightness independently, only sending updates where movement occurs, dramatically reducing the computational load.

Using this spatio-temporal input, event-based algorithms are able to track fast-moving bees in real-time to determine the position of the bee relative to the hive entrance, and by defining a boundary, count the number of bees leaving and returning. Due to the sensor's temporal resolution, the flapping bee wing can be observed in flight and its wingbeat frequency can be estimated during tracking in real-time. To evaluate the proposed event-based tracking system, a side-by-side comparison with a frame-based camera at an active colony was performed. Real-time tracking of trends in bee activity should provide early warning signs of problems such as robbing, swarming, absconding, etc. Detailed analysis of wingbeat frequency may eventually provide a real-time detection system for invading insects.

ANALYZING HONEY BEE FLIGHT WITH EVENT-BASED VISION

by

Team HiveMIND

Ayman Fatima, Kalonji Harrington, Riya Kukadia, Matthew Lynch, Zain Majumder, Rohan Mathur, Daniel Park, Richard Strucko, Elijah Taeckens, Stefan Traska, Matthew Tremba

Thesis submitted in partial fulfillment of the requirements of the

Gemstone Honors Program, University of Maryland

2024

Advisory Committee:

Dr. Timothy K. Horiuchi, Mentor

Dr. Pamela Abshire

Dr. Cornelia Fermüller

Dr. Giacomo Indiveri

Karen Rennich

Dr. Dennis vanEngelsdorp

Acknowledgements

We would like to thank Dr. Lovell, Dr. Lansverk, Dr. Hill, Dr. Skendall, and everyone at the Gemstone for their incredible guidance and support throughout the research process. We would also like to thank all of our discussants for taking the time to read our work and supporting our team by attending the thesis conference. Additionally, we would like to thank Dr. Dennis vanEngelsdorp, Karen Rennich, Eric Malcolm, Mark Dykes, and everyone at the UMD Bee Lab for providing guidance and allowing us to study their bee hives, without which we could not have accomplished any of our research. We are also grateful for our team librarian Jordan Sly for assisting us with resources to find scholarly sources and helping us set up Zotero to organize our references. Lastly, we are most grateful for the support and guidance that we have received from our mentor Dr. Timothy Horiuchi, who helped us focus our research goals, gave valuable feedback to improve our work, and went above and beyond the requirements of Gemstone mentors to ensure that we had all the resources we needed to succeed.

Additionally, we would like to acknowledge our sources of funding. We are grateful to have received funding from a variety of sources, including the Gemstone department, our LaunchUMD fundraising donors, the Do Good Institute through the Do Good Showcase, and UMD Libraries through the Gemstone Libraries Award.

Table of Contents

Chapter 1. Introduction

Apis mellifera (L.), the "western honey bee" or "European honey bee," plays a critical role in our food supply. By some estimates, honey bees are responsible for 80% of pollination in agricultural settings, which accounts for $6.4 billion worth of crop yield in the United States alone [1], [2]. Consequently, many people around the world depend on bees for employment and livelihoods, especially in developing and rural areas [3]. In some areas of the world, the decline in honey bee populations have motivated an increase in hand pollination as an alternative [4].

The importance of honey bee pollination also extends to natural ecosystems. Among studied plant-pollinator interactions, *A. mellifera* provides 13% of visits to all flowering plants worldwide — the most of any single species of pollinator [5]. For 4.5% of plant taxa honey bees were the only recorded pollinators, which emphasizes their role in maintaining biodiversity [5], [6].

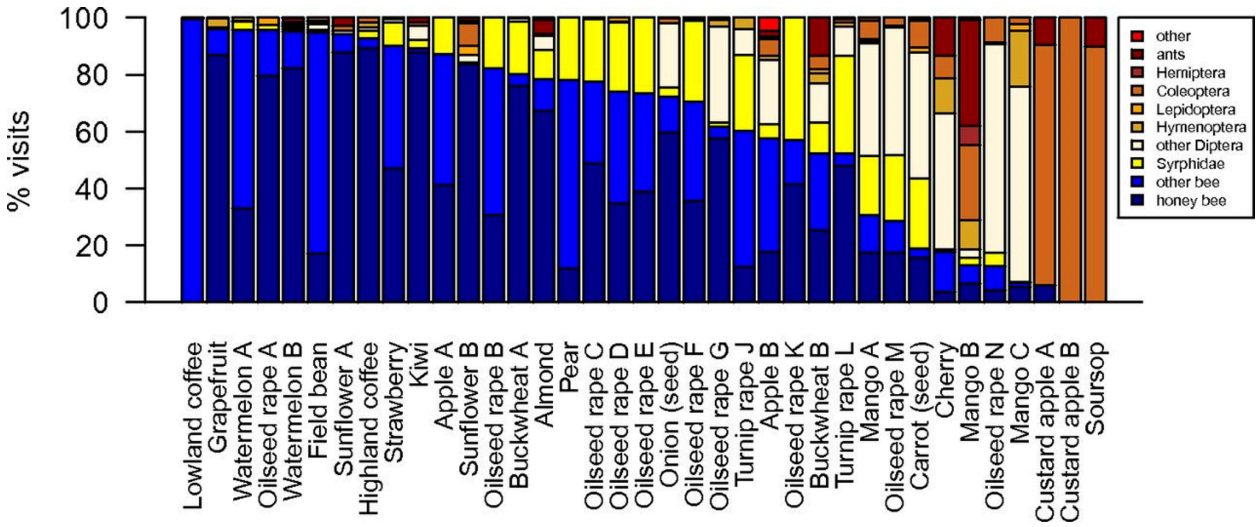

Fig. 1. Insect contribution to flowering plant pollination. Honey bees (dark blue) and other bees (blue) make up the majority of pollination for most crops studied.

Unfortunately, bees face a number of threats. In 2020, beekeepers in the U.S. lost 43.7% of managed hives, the second highest loss rate since researchers started keeping records in 2010 [7]. There are a variety of factors contributing to this decline, including parasites, pathogens, pesticides, habitat loss due to climate change, and more [8], [9], [10].

Given that bees play a vital role as pollinators and are under increasing threat, a number of techniques for monitoring and protecting their health have been developed. Metrics like hive weight, humidity, temperature, and CO_2 emissions can all be used to check for abnormal patterns in a colony's activity [11], [12], [13]. Audio sensors have also been used to monitor hives [14], [15]. In individual bee behavioral studies, it is common to tag individual bees (e.g., visual marker, radio transmitter, or radar reflector) and then use the tag to recapture them at a later date/location, or to track them in real-time. However, these methods are limited because they do not scale well and do not address the population of the colony as a whole.

A technology better suited for the task of monitoring bees is an event-based vision sensor. Event-based vision sensors differ from traditional video cameras in that they do not use image frames. Instead, each individual pixel is equipped with two *neurons* (inspired by the biological retina), which fire in response to changes in light intensity. One of the neurons generates "events" (a signal analogous to a biological neuron action potential, aka "spike") at a rate proportional to the rate of increase in light level while the other generates events at a rate proportional to the rate of decrease in light level. The resulting output from the array of independently-functioning pixels is a stream of events.

By designing the sensor chip readout system to only communicate events from pixels where objects are moving (i.e., where light intensity is changing), small moving objects can be

tracked with high temporal precision and low-latency. This pixel-focused method of recording allows for the detection of distinct events that are only microseconds apart [16]. Therefore, event-based vision sensors are well-suited to the task of bee monitoring, which involves recording small changes at high speed. These relatively sparse events can be processed so quickly that algorithms can operate on the sensor output to measure individual bee metrics in real time.

We utilize the advantages of event-based vision to develop algorithms which count entrances/exits from the hive and detect the wingbeat frequency of individual bees in flight.

Chapter 2. Literature Review

2.1. Bees

Bees are insects most well known for their role as pollinators [17], [18], [19], [20]. They are close relatives of wasps and ants [21]. Together they form order Hymenoptera, section Aculeata, a taxonomic group most famously characterized by its stingers [21].

Although all insects play important roles throughout the biosphere, bees' role as widespread pollinators make them unique. Bees rely on flowers as their principal food source and have evolved several adaptations to efficiently interact with them. Nectar, a bee's primary source of carbohydrates, is gathered with a proboscis, a long tube through which fluids may be drawn in [21]. Pollen, a bee's primary source of protein, is collected on external hairs, placed in pollen baskets (corbiculae) on the back legs, and later cleaned off and packed into cells [21]. In turn, plants have co-evolved to take advantage of bee foraging by placing their sex organs near the nectar source so that visiting bees brush against them, thereby transferring pollen and facilitating reproduction [21].

This unique relationship makes bees the world's foremost pollinating insect. The vast majority of the world's tropical forests are pollinated by bees, and even in temperate climates smaller plants like shrubs and wildflowers depend on bees for reproduction [21]. Of particular interest to humans are agricultural crops such as fruits, vegetables, nuts, fiber, and forage. Thirty-nine of the top 57 monocultures in agriculture require pollination by bees [22], and in the United States agricultural bee pollination provides economic gains of up to $5.7 billion per year [23].

2.1.1. General Biology

As members of class *Insecta,* bees are characterized by a segmented exoskeleton with three pairs of legs, three main body regions, and one pair of antennae [21], [24]. More specifically, bees may be differentiated from other insects by their narrow waist, double wings, and, in females, ovipositors modified into stingers [24].

2.1.2. Flight

Insect flight mechanics have been extensively studied [17]-[21]. Insects oscillate their wings on a slightly tilted vertical axis, rotating the surface so the leading edge faces downward on the downstroke and upwards on the upstroke [30]. This motion, which can occur up to 300 times per second [28] and averages between 200 and 225 times per second in honey bees [31], induces a variety of aerodynamic effects such as leading-edge vortices and the clap-and-fling effect which are theorized to make efficient flight possible [32]. These parameters are dynamically altered to provide control of direction and speed. Honey bees in particular are known to resist strong wind gusts and hover by altering stroke amplitude and wingbeat frequency [29].

2.1.3. Threats to Honey Bees

Honey bees face many threats, including parasites, pathogens, pesticides, habitat loss due to climate change, lack of floral resources, and more [8], [9], [10]. While these threats are well-understood individually, their interactions are not. For example, colony collapse disorder (CCD) is a phenomenon in which worker bees abandon a colony despite appearing to be in good health otherwise, leaving plenty of brood and a minimal number of dead workers. Researchers

believe that CCD is caused by a combination of many individual stressors, but the subject remains an area of active research [33].

One such individual threat is the *Varroa* mite, an external parasite which feeds on the bees throughout its life cycle [9]. Female mites lay eggs inside the sealed brood cells, which hatch and begin to feed on immature bees in either the larval or pupal stage (collectively referred to as the brood). Once matured to adulthood, the mites attach to adult bees where they continue to feed [34]. While damaging by themselves, *Varroa* mites also act as the vector for many diseases, including acute bee paralysis virus (ABPV), deformed wing virus (DWV), and more. In addition to transmitting these diseases, the mites also weaken their hosts' immune systems, exacerbating the dangers from both the infections and other threats like pesticides [9].

Viral and bacterial pathogens also pose considerable threats. Deformed wing virus (DWV) can result in stunted wings that render the infected bee unviable and causes them to die shortly after hatching. Acute bee paralysis virus (ABPV) can also damage the bee's ability to fly, and is extremely infectious [9].

Non-biological threats to bees include pesticides, weather, and habitat loss. Bees, especially those in managed colonies that are transported for pollination events, come into contact with pesticides while foraging from treated plants. Neonicotinoids in particular can impair reproduction, reduce mobility and cognition, and increase susceptibility to infection [8]. Bees are also sensitive to both temperature and precipitation, which can affect their ability to forage and store enough food to survive the winter. Habitat changes can also affect colonies quite significantly through urbanization, pollution, agriculture (especially monocropping), all of which can disrupt colonies' well-being [8].

2.1.4. Colony Behaviors

Honey bees live in colonies, highly-structured social organizations composed of a single queen bee and many thousands of her offspring that together perform the tasks necessary for survival. Since bees cannot succeed individually for extended periods, the colony functions as a "superorganism" which acts collectively to ensure the continued survival of its individual members through reproduction, brood rearing, hive construction, hive defense, and foraging [35].

At the biological level, bees are divided into three distinct castes: queens, workers, and drones. Queen bees are the principal reproductive female and are responsible for laying the colony's eggs and producing pheromones which maintain social cohesion [36]. Worker bees are sterile females that handle the colony's labor needs. A worker's specific role is determined by her age: cell cleaning between 1 and 3 days old, nursing between 4-12 days old, comb-building, nectar processing, guarding, and other in-hive tasks between 12-21 days old, and foraging until death around 30 days old [37], [38], [39]. Drone bees are males whose sole job is to pass on the colony's genes by mating with queens from other colonies, facilitating the reproduction of the entire species [36].

To fulfill these roles effectively in a large colony, individual bees rely on decentralized, local-information systems to guide their behavior. In the context of a bee colony, decentralization means that information and decision-making are dispersed among all individuals, rather than being concentrated in a single "leader" [40]. As a consequence, no one individual is thought to have a complete picture of the colony's needs or the ability to create orders based on global understanding. Therefore, effective collaboration requires bees to rely

on the information available to them as individuals, here termed “local information” [41]. While these general principles govern all colony activities, the specific mechanisms vary depending on the application.

These decentralized processes govern the colony’s defense. Intruders such as insects and vertebrates are first detected by guard bees, a group of roughly 75 worker bees which monitor the entrance of the hive for external organisms [42], [43]. For larger threats like vertebrates, guard bees enter the hive and disperse alarm pheromones from their stinger by beating their wings. This recruits a large number of bees, who fly out of the hive to harass or sting the intruder. Upon stinging, the bee releases yet more alarm pheromones which attract further defenders to the target [42]. By using precise tactics, teamwork, and multiple levels of escalation, the colony is able to deal with a variety of threats while minimizing worker losses. Finally, foraging is coordinated through division of labor and inter-bee communication.

2.1.5. Africanized Subspecies

In an effort to increase honey production in the Americas, Brazilian researchers imported *A. m. scutellata L.* African honey bees, which then escaped and created a hybrid with the European honey bee, known as the Africanized honey bee [44]. This hybrid has since spread to most of South and Central America, as well as parts of North America, due to the dominance of the African genotype.

These bees inherited a variety of traits allowing them to be more resilient and have better honey production in tropical climates, including smaller hives and different foraging tactics [45]. However, some adaptations make for much more difficult management, such as the increased defensiveness and a tendency to abscond, or abandon a nest due to lack of floral

resources or because of external disturbances [45], [46]. After considerable adjustment, including improved hives and equipment, Brazilian beekeepers were able to adapt to the behavior of these bees, allowing them to take advantage of their superior performance in tropical climates [47]. However, the traits that make Africanized bees more suitable to tropical environments also make them less desirable in temperate ones [45].

2.2. Beehive Monitoring Techniques

Researchers use a myriad of tools to study the movement and life cycle of different bee species as analyzing the behavioral patterns of bees is critical to improving conservation efforts. One of the most important tasks in maintaining the health of a colony is ensuring that the queen bee is healthy. If the queen is missing, out of sperm, deceased, or diseased, then the beekeeper must replace her as quickly and efficiently as possible — a costly process that can reduce hive productivity [48]. Routine inspection is therefore essential, but can become a stressor if done too frequently and can have a detrimental effect on the colony's wellbeing as a result [49].

2.2.1. Traditional Methods

Depending on their goals, prior studies have employed different methods of studying bees. Colony health may be monitored by weight, temperature, sound analysis, and CO_2 levels; foraging movements may be monitored by manual mark-and-recapture methods, radio tracking, and tagging; *Varroa* mite infestations using alcohol or sugar shakes, brood sampling, or sticky paper analysis; and identification of Africanized honey bees with biochemical and molecular

analyses as well as morphometric techniques [50]. Each method has its own benefits and drawbacks.

Weight can be used to detect swarming and determine when the honey is ready to be harvested. This requires the ability to weigh the entire hive (on the order of tens of kilograms) with precision greater than 100 grams to be useful year-round. Additionally, automatic systems often require hive modifications to accommodate the scale mechanism [12].

Temperature, measured with either thermometers in the hive or thermal imaging cameras outside it, can predict large-scale swarming and absconding behaviors. However, internal thermometers are somewhat invasive, and thermal imaging may have low contrast under certain weather conditions [51].

In an attempt to create a non-invasive method of determining the health of a colony, researchers from Gdańsk University of Technology and the University of Spain, and Università Politecnica delle Marche analyzed sounds emanating from a hive to gauge the colony's wellbeing. In addition to recording the sounds with a specialized microphone, their study also incorporated the use of temperature and humidity sensors to help with further classification of the data. Theoretically their system may be able to detect various illnesses, such as the presence of *Varroa destructor*, but due to the lack of data about those events (none of their bees were infected or faced critical situations) they were unable to rigorously test their system although the system is currently being improved with further research. Their study did, however, show that they can pattern-match sound samples to differentiate between typical behavior and anomalies such as a queenless colony or swarming [52].

CO_2 sensors can indicate infection by *Varroa* mites or other parasites by acting as a proxy for population. Fewer bees in the hive means less respiration and therefore lower CO_2 concentrations. Thus by monitoring CO_2 concentrations over time, beekeepers can identify anomalies that accompany parasitic infections [53].

While studying wild bee colonies, researchers have applied different techniques to track the long-range movement patterns of individual bees. In 2019, researchers from the University of California, Davis published an in-depth review of several methods used to study the movement of Apidae *Bombus* spp. (bumble bees) and while wild bumble bees differ from domesticated European honey bees, the benefits and flaws of these tools are still relevant [54].

One of the most commonly used methods is traditional mark and recapture. This manual method consists of capturing bees near the hive, marking them with paint or other identifiable tags, and releasing them with the goal of later identifying individuals when recaptured at different locations in the field [54]. While it benefits from being the simplest and least invasive method to study the spatial range and flight patterns of bees, mark and recapture fails to provide any further information other than displacement from the hive and also depends on the researchers capturing specific bees to gather statistical data [54].

A solution to the search effort demanded by traditional mark and recapture is genetic mark-recapture, otherwise known as sibship assignment. This method forgoes having to recapture individual bees by utilizing similarities in the genetic makeup of workers from the same colony. Due to their haplodiploid breeding (male drones only have one set of genes while the queen and workers have two sets), sibling worker bees share roughly 75% of their genes [54]. Sibship assignment allows researchers to determine where bees of a given hive travel and

is the only method that can estimate queen dispersal in bee species like the bumble bee where the previous queen, workers, and drones die in the winter [54]. While improving on traditional mark and recapture, sibship assignment makes it difficult to determine the location of wild bee hives and can be prone to significant error and bias. Errors in assignment can lead to researchers mixing data from different hives and all forms of mark and recapture introduce bias as bees tend to visit the same areas throughout a given season [54].

A more effective way to determine foraging characteristics such as the maximum range of a colony is with homing studies. This method of study focuses on the bee's ability to spatially locate themselves and successfully return to their hive. By displacing bees a known distance from their hive, researchers have studied how different environmental factors (landscape context, pesticide exposure, nutrition, etc) impact the navigational abilities of bees. While useful for some applications, homing studies are affected by researcher bias, as determining the exact reason why a bee might have succeeded/failed in returning to its hive is difficult [54].

Another way researchers track individual bees is with radio telemetry and harmonic radar. Both methods rely on small tags that are attached to a bee and can provide information about optimizations to flight paths, foraging distance, lifelong movement patterns, and more. The key difference between radio telemetry and harmonic radar is that radio telemetry is a form of active tagging in which the tag utilizes a battery to record and transmit data whereas harmonic radar is a form of passive tagging by adding a radar reflective object. Both methods also introduce the drawback of physically changing the bee's movement as the weight of the tag affects the speed of the bee and causes it to rest for extended periods. Harmonic radar is also

limited in the environments where it can be used, as a constant line of sight is needed to be able to track the bee's movement with a radar dish [54].

Radio frequency identification (RFID) and optical tags are other forms of tagging that are more specialized for studying bees close to the hive. Figure 2 shows how researchers from the Jiangxi Agricultural University used RFID to track the activity level of honey bees before inclement weather [55].

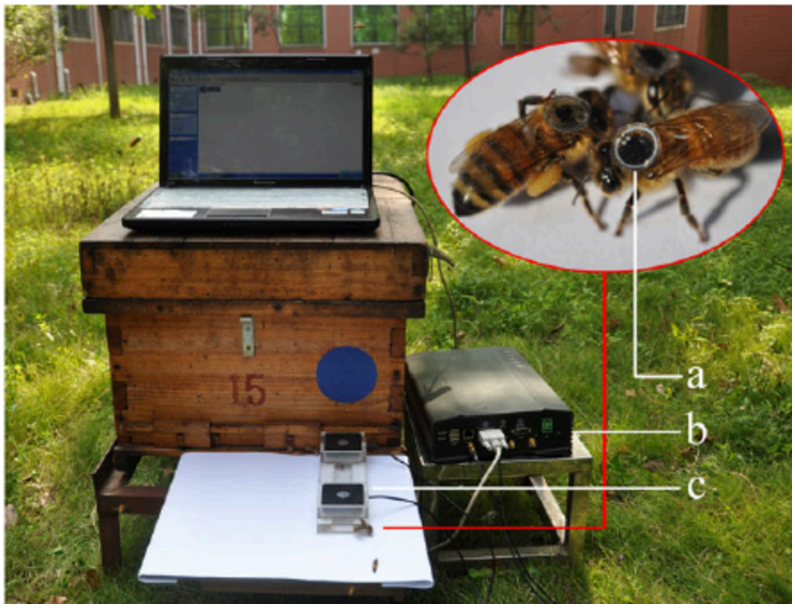

Fig. 2: Researchers used RFID and optical tags on bees to track hive activity. a: tag; b: recorder; c: reader [55]

While unfeasible for long range detection, RFID and optical tags can be used to measure activity within the hive using a scanner that can detect when bees with specific tags pass by. These forms of tagging are better at measuring time intervals and can provide information about foraging times and other routines that bees follow but neither RFID or optical tags provide specific information about bee activity [54].

With all of these methods of studying bee movement, most of them have noticeable drawbacks that prevent them from being widely applicable. This in particular makes it difficult

to observe the behavior of bees and introduces barriers to understanding how to combat threats to colonies and aid in general bee management.

Due to laws governing the movement of different bee species across state lines and concerns about protection of native bees, it is important for beekeepers to differentiate between Africanized and European honey bees. Early methods of identification involved expensive biochemical and molecular analyses that required time and access to a laboratory [50]. Over time, faster and cheaper morphometric techniques have become more prevalent, resulting in the "Fast Africanized Bee Identification System" (FABIS) for official identification in the USA [50]. Even these more recent methods require trained specialists for manual identification making automatic techniques much more appealing [50].

A lesser-explored method for identifying Africanized honey bees is through differences in wingbeat frequency. Studies have shown that African and Africanized bees have higher wingbeat frequencies than European honey bees [56]. By collecting acoustic data, the difference in wingbeat frequency can be detected [57]. However, for the best results, background noise needs to be isolated. This requires live bee specimens to be captured and contained near a microphone [56], a much more challenging and disruptive endeavor, especially considering the defensive behavior of Africanized bees.

2.2.2. Conventional Machine Learning

There have been many attempts in the past to use machine learning to track or monitor bees. For example, one study focused on recognizing pollen-bearing bees from videos of the bees as they enter the hives. The goal of the study was to obtain data about the bees' foraging behavior and task specialization [58]. They did this by detecting the presence of pollen, as the

presence indicates success in the pollen foraging tasks. To obtain the data for study, the researchers used a conventional camera connected to a video capture system which monitored a ramp which all bees must pass through to enter or exit the colony. They then analyzed the data using a variety of machine learning classifiers, and the approach using convolutional neural networks was the most accurate, reaching an accuracy of 91.16%. However, the classification was not done in real time. The research team first obtained the dataset, and manually annotated it to guarantee a proper dataset. This demonstrates that using conventional cameras and convolutional neural networks were quite capable of determining whether bees are carrying pollen, as shown in Figure 3, though not in real time [58].

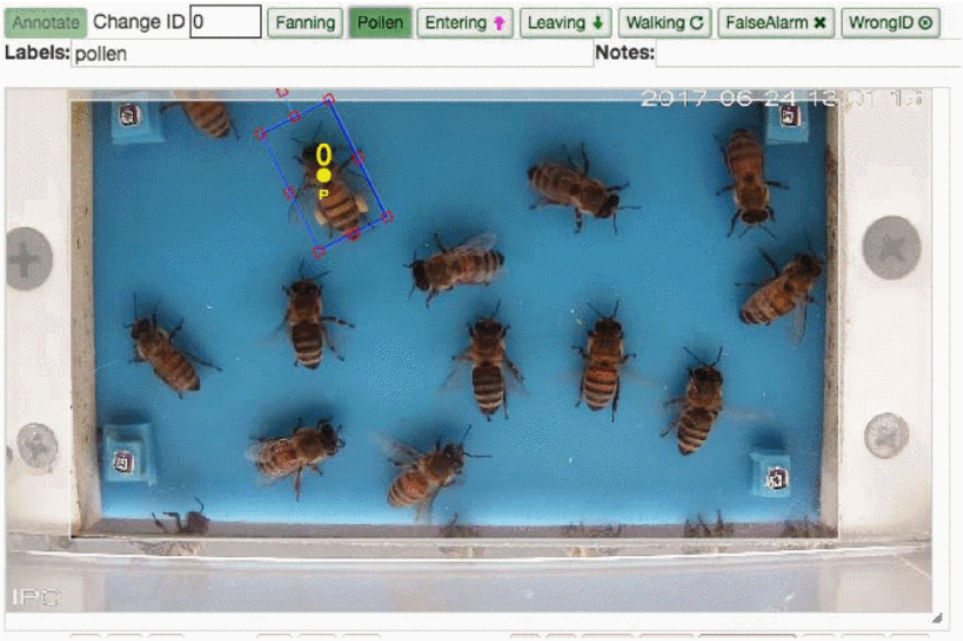

Fig. 3: A highlighted bee with pollen (pollen baskets on legs) among other bees [58]

Another study focused on tracking individual honey bees among wildflower clusters with computer vision to monitor pollinators [59]. They mentioned how challenging it is to track unmarked insect pollinators in uncontrolled outdoor environments. Despite the difficulties associated with such an experimental environment, the research group was able to develop a

deep-learning based detection algorithm that tracked honey bees at a detection rate of 86.6%, meaning that the measured bee position matched human observations in 86.6% of frames. The algorithm the researchers used tracked a single insect from when it first enters the frame until it exits, and continuously reapplies itself as more insects enter. However, this required over 3 hours of processing to complete tracking in a 6 minute video segment. With the results, the research group commented on how computationally inexpensive detection methods can increase the rate of detection and save computational costs. This study shows how the tracking of bees can be performed well with conventional cameras [59].

Other researchers have also attempted to design a computer vision algorithm to track and count bees. While accurately able to detect and track bees, the relatively slow framerate of 60 FPS was a limitation [60]. There was no report for computation time for this method of tracking.

The FABIS technique for identifying Africanized honey bees has also benefited from advances in image recognition. Researchers have been successful in using image recognition to improve upon the efficiency of using morphological tests. This technique was able to achieve a 98.05% accuracy using the FABIS standards [50]. The primary drawback of this method is that it requires wings samples to be collected and analyzed under a microscope, making it only useful for post-processing when a beekeeper has reason to suspect a colony is Africanized.

2.2.3. Dynamic Vision Sensor (DVS)

While neuromorphic technologies are rapidly being applied to many different aspects of computing, the development of biologically-inspired sensors is rapidly developing [61]. Neuromorphic sensors are designed to mimic the behavior of biological sensory organs, notably

the method of encoding signals and the signal processing used. The most popular area of research in neuromorphic sensing is vision. These sensors attempt to replicate the function of biological eyes, and are referred to as dynamic vision sensors (DVS) [62].

Neuromorphic, event-based vision sensors differ from conventional cameras in several ways. Conventionally, cameras represent video as a sequence of still image frames. Frame-based video generally contains large amounts of redundant information, since the light intensity at every pixel is reported in each frame repetitively, even if their value remains unchanged from the previous image. By contrast, DVS cameras use the sensor communication bandwidth to only transmit information about changes in light intensity where it occurs in the image [16]. This means that if a scene in a video is static, the DVS will not send new data, saving enormous amounts of data. Additionally, the pixels in a DVS behave asynchronously, meaning that they are activated independently of one another. This contrasts with frame-based cameras, in which every pixel records data simultaneously at regular intervals.

The asynchronous behavior of the pixels in a DVS grants several advantages over traditional cameras in certain applications. DVS cameras have extremely high temporal resolution and low latency [63], allowing them to detect distinct events on very small timescales, often separated by just a few microseconds. Additionally, DVS cameras have high dynamic range [64], meaning that they can adapt to different light levels, and perform well in both bright and dark environments. Finally, DVS cameras can run on relatively low amounts of power, often using less than 100 mW, far less than a standard frame based camera [65].

This technology does not exist without drawbacks, however. The changes in luminosity created by a moving background insert unnecessary data that must be accounted for in object

detection, just like the redundant data a traditional camera collects. Researchers have developed successful algorithms capable of reducing this noise, but their implementation is nontrivial and removes some of the benefits event-based cameras offer [66]. Low latency is also sacrificed if a large number of simultaneous events occur, causing a buffer to accumulate that results in large timestamp errors [67]. Since these are fixed timestamps from the data transmission, this not only affects real-time applications, but also post-processing. The novelty of event-based cameras also presents some challenges, as most systems are not designed to take advantage of the camera's capabilities. Thus, new conceptions of how data is represented in space-time and new algorithms for interpreting that data must be examined and implemented to solve problems for which traditional cameras already have platforms built [62].

Despite these disadvantages, DVS cameras have been applied to several problems that traditional cameras struggle with. For example, researchers have used DVS cameras to track the movement of satellites above Earth's atmosphere during both the night and the day [68]. Satellites are notoriously difficult to track due to their small size and rapid movement, and can normally only be tracked at night. DVS cameras are able to adapt to surrounding light levels to see satellites during the day. Additionally, the DVS is sensitive to small changes in luminosity, so it was able to easily track the satellite despite its rapid movement. DVS cameras have also been used to enhance optical flow [66]. Optical flow is the task of modeling the directions and velocities of objects in a visual scene, and is usually very computationally intensive. DVS cameras focus on objects that are changing or moving, so they can estimate optical flow in a scene much more quickly and with far less computational power. DVS cameras have also been

used to model a wide variety of other problems that require high speed monitoring of moving targets, such as controlling a robotic goalie arm or tracking the speeds of cars in traffic [62].

2.2.4. DVS Applicability

As previously described, DVS sensors are made up of individual pixels that send on and off signals during changes in light levels as it senses motion in the environment. The main advantage of this feature is that there is no opportunity for motion blur, allowing for accurate representations on environmental changes. Specifically, a number of American researchers have found that the DVS would be useful in understanding ultra-efficient flying organisms like flies, birds, and bees [69]. This advantage also allows for fast paced asynchronous changes to be analyzed efficiently, such as the proposed use of hybrid tracking algorithms that can immediately recognize when a bee has left the camera frame and when it has reentered. Such tracking would help deepen our understanding of pollination in a noisy environment, and would be challenging without a fast paced computer vision device.

Event-based vision has often been used for tasks involving rapid tracking in varied lighting conditions, such as tracking ping-pong balls or even tracking satellites in daylight [63] [68]. These abilities will be immensely helpful at interpreting fast paced changes, such as bees flying rapidly against a changing background or detecting the wing beats of bees, making the DVS more impactful in monitoring bees than traditional CV technology.

In comparison to the 60 FPS limitation of conventional methods, the DVS is capable of detecting changes in light intensity with microsecond latency [16], implying that the DVS could be of benefit in improving current computer vision systems for monitoring bees, especially their fast paced behaviors.

In fact, event-based vision has already been used in the task of tracking flying insects. Janick Cardinale used an early model of the Dynamic Vision Sensor to track a fruit fly in a similar manner to what this paper proposes [70]. His paper utilized a cluster tracking algorithm, which tracks groups of nearby spikes by shifting the center towards more recent events. The trackers are then put through a Kalman filter to provide a smoother and more accurate estimate of the tracker's movement. This method was promising, but was limited by the 128 x 128 pixel spatial resolution of the original sensor and tuning of the Kalman filter to prevent divergence. It also did not handle the case of occlusions, instead combining trackers and later creating a new one if two flies occluded from the view of the sensor.

2.3. Wingbeat Frequency and Species Detection

The wingbeat frequencies of insects can give valuable information about species and behavior. Differences in wingbeat frequencies between different species of bees have been documented [56], as well as differences between different species and across different environments [71]. The ability to automatically detect differences between species in real-time could be very valuable for protecting beehives from invaders such as wasps and hornets. If an invader could be accurately and quickly identified, a targeting system could be used to kill an invading species before it can cause damage to the colony.

Wingbeat frequencies are also a method of communication for bees. A study of the wingbeats of individual bees within hives found that bees with different roles tended to have different wingbeat frequencies, and that they often used a variety of wingbeat patterns to communicate information [72]. This contrasts with a previous belief that bees only move their wings to cool the hive, remove CO_2, or dry honey. Expanding experiments to include precise

measurement of flying bees allows researchers to dive deeper into how bees may communicate while in flight.

2.3.1. Previous Wingbeat Detection Methods

There are several existing methods that are used to measure wingbeat frequency in honey bees and other flying insects. Some of the oldest methods for frequency estimation involve acoustic recordings of the sound of an insect's wings [73], [74]. This is an approach that works well in lab settings, but can run into problems when recording insects in their natural setting due to the rapid attenuation of the sound produced by insect wings when the distance is any more than a few centimeters [75]. Additionally, frequency detection by sound can make it difficult to isolate the wingbeat frequency of an individual bee, since the signals will overlap when multiple bees fly close to each other. More recent frequency detection attempts have focused on optical sensors or lasers [75], [71], [76], high-speed cameras [72], [77], or optical backscattering [78].

One study used a scanning laser source in front of a phototransistor array to extract the wingbeat signal of the bee [75]. As the bee passes between the laser and the phototransistors, its wings partially occlude the laser, generating a signal that can be extracted and analyzed to extract the wingbeat. A similar approach was used in [71] to measure changes in the wingbeat frequencies of several different bee species over a span of several months. These optical sensors have the advantage of being accurate and low cost, but they do have some notable disadvantages.

One major problem is that the flight pattern of the bee must be constrained to pass between the laser and the phototransistor, which are generally placed within a few tens of

centimeters of each other. The authors of [71] got around this problem by placing the sensors at the front of the bee hives, where bees are somewhat constrained to fly anyway. However, placing the sensors in an area of frequent flight presents another issue, which is that multiple bees may attempt to pass through the sensor simultaneously. When this happens, the recorded signal will contain elements from each bee's wingbeats, making it hard to distinguish the presence of individuals. Additionally, research has shown that laser scanning can lead to damage of insect wings, especially at high intensities, which are often used to improve signal to noise ratio [79].

Other attempts have tried to track wingbeat frequency using high speed cameras. One study used computer vision to isolate a bumblebee from its background and used optical flow fields to track wing displacement [77]. The frequency was then obtained by applying a short-time Fourier transform to the wing displacement signal. This innovative approach addresses some of the issues of other optical methods, since it does not necessarily require that bees be constrained in flight and can distinguish individual bees. However, the system faces many of the same limitations as other computer vision systems discussed in the previous section, given the high computational demand of trying to apply computer vision at high speeds.

2.3.2. Event-Based Vision for Frequency Detection

The DVS has several advantages that make it ideally suited to measure the wingbeat frequency of bees and address some of the issues with other methods. First, as mentioned in the tracking section, the DVS can track fast-moving flying objects in real time [63], removing any need to constrain the flight of the bees. This also allows the system to track the wingbeats of

multiple bees simultaneously, which can be difficult with optical or acoustic methods. Additionally, the sensitivity to changes in light and high temporal resolution offered by the DVS makes it a natural choice for monitoring periodic systems.

Event-based vision has previously been used to monitor the frequency of high speed systems. In one study, a DVS was used to track the frequency of the blades of a quadcopter in real-time as the quadcopter performed various aerial maneuvers [80], demonstrating the versatility of the DVS for real-time frequency tracking of flying objects.

Event-based vision has also been applied to wingbeat detection in insects. Cardinale's paper tracked and analyzed fruit fly wings using an early version of the DVS [70]. To analyze the frequency, Cardinale detects and tracks the wings themselves, then extracts frequency from changes in the wing's orientation. This orientation was then tracked over time, with frequency and amplitude estimated using both a weighted running average and an extended Kalman filter. However, these analyses were seemingly only applied to a single fruit fly, thus avoiding the problem of occlusions and allowing the sensor to focus on a single target. This meant that despite the low spatial resolution of the 128 x 128 pixel model, the wing edges could still be reliably detected and tracked, which is not possible when tracking a large open area.

Chapter 3. Methods

3.1. Event Based Sensors

There are various commercial vendors for event-based vision sensors that offer sensors with different resolutions and development packages. These include the Metavision sensors (developed by Prophesee and Sony) and Inivation's Dynamic Vision Sensor (DVS).

Pilot studies with a DVS128 sensor (Inivation) allowed for familiarity with the event-based processing and development of the initial software, however, it quickly became clear that the low resolution of the sensor (128 pixels by 128 pixels) was insufficient for tracking bees at longer ranges. With the limited resolution, accurately tracking bees over a larger region of interest was difficult.

Fig. 4. The image above depicts early proof of concept tests with the DVS 128 at hives owned by the University of Maryland's beekeeping club. The low-resolution sensor had to be very close to the hive to detect bees.

For this reason, later research applied the results from the DVS128 to the higher-resolution DVXplorer. The higher resolution (640 x 480) allows the sensor to capture more detail and detect smaller bees, but introduces higher latency when a large number of pixels are activated, a drawback that must be compensated for in other aspects of the design. To reduce the risk of a large number of the sensor's pixels from activating at once, the sensor is placed about 45 cm from the entrance to the hive.

3.2. Tracking Algorithms

Object tracking has been extensively studied within the field of computer vision. Typically, the process of tracking an object begins by isolating the desired object from the background. However, when using event-based sensors to track moving objects, the target is already isolated, since the static background will not cause the DVS to generate events. As a result, the process of tracking an object can be simplified to a simple, iterative algorithm.

The algorithm is designed to be primarily event-based, allowing it to save on power and computation costs when the frequency of events goes down. To allow the system to operate in real-time, however, a few computationally expensive operations have to be performed at a fixed rate, since performing these computations after every event causes significant delays.

To track the bees in flight, the system initiates objects referred to as trackers. Trackers are square objects that maintain a position and velocity that update dynamically whenever a new event occurs inside or around the tracker.

The algorithm for tracking bees has 3 main components: tracker initialization, tracker updates, and tracker removal. The tracker initialization component focuses on detecting when a

new bee enters the field of view. Moving bees generate quick sequences of events, so the tracker initialization algorithm searches for regions of the screen in which several spikes occur within a short period of time. Tracker updates ensure the trackers follow the bee's movements, even during brief periods where it may not generate spikes. Tracker removal ensures that once a bee leaves the frame or no longer generates any spikes for an extended period, it is no longer tracked.

3.2.1. Time Surfaces and Initialization

The Dynamic Vision Sensor display consists of an array of 640 by 480 pixels. The tracker initialization algorithm splits this display into square *2r* by *2r* pixel non-overlapping regions. Each region acts as a modified version of a time surface. Time surfaces are a method used for extracting spatiotemporal information from event-based visual data [81]. A standard time surface consists of a region in which the value of the surface at a location **x** is determined by the time since the most recent event at that location. As a result, time surfaces can offer great detail about the history of events on a local scale. However, for the simple task of detecting an object, this level of spatial resolution is not useful. Instead, the algorithm assigns a single value $S_i(t)$ to each 2r pixel by 2r pixel region of the DVS output, which is determined by:

$$S_i(t+1) = e^{-t/\tau} S_i(t) + \nu \sum_{x=-r}^{r} \sum_{y=-r}^{r} e_{x,y}(t) \tag{1}$$

where $e_{(x,y)}(t)$ is a binary value that represents whether the pixel at location (x,y) produced an event at time *t*, ν is a scaling constant, and τ is a time constant defining the update period.

The nature of the time surfaces are such that the value $S_i(t)$ will be large if there is a large amount of activity in the i-th region within a relatively short period of time determined by the time constant τ. As a result, the value can be used to detect sudden increases in activity levels. When $S_i(t)$ surpasses a threshold value, a new bee tracking object is initiated, assuming that no existing trackers are currently overlapping with the region of interest. A maximum of 20 trackers were allowed to exist simultaneously to prevent excessive latency, meaning that a new tracker would not be initiated if there are already 20 existing ones.

The bee tracker initiation process is affected by the choice of constants. Without loss of generality, the threshold for tracker initiation can be given a value of 1, since the behavior of the time surface can be scaled using the constant ν. A similar approach is taken in [82]. Choosing a larger value of ν will increase the impact of one event, which effectively decreases the amount of events required to initiate a tracking object. The initiation process is also affected by the time constant τ. Increasing the value of τ increases the amount of time that an individual event has a noticeable effect on the time surface. In other words, an increase in τ increases the amount of time that can occur between events while still initiating a tracker.

Increasing the value of τ decreases the decay rate of the time surface, allowing an individual event to have a noticeable effect for longer. This has the effect of increasing the amount of time that can occur between events while still increasing the pixel's value and moving towards the threshold to initialize a tracker.

The values of the constant were chosen experimentally to match recordings of bees in flight. Sample data of bees flying in front of a hive was collected to optimize these parameters. Based on the apparent speed and size of the average flying bee in the sensor, the parameters

were adjusted so a reasonable tracking box could be formed around the spike clusters. The equilibrium for trade offs such as maximum tracker size were reached by finding the point where the algorithm was still fairly consistent at tracking but allowed for the most flexibility in abnormal cases such as a bee flying closer to the sensor. ν was set to 0.15 after observing the typical number of events generated by a bee in a given region, while τ was chosen based on the amount of time it typically takes for a bee to move from one time surface to the next.

It should be noted that this method is not unique to tracking bees, and will conceivably track any object of similar size and speed. Depending on parameter values, it could even track objects with different apparent sizes and speeds. Additionally, it is unable to track bees that crawl along the surface of the beehive, since the algorithm is tuned to track objects that are moving faster and generate more events. Future research will work on other methods of differentiating bees from other potential trackable flying objects, such as measuring the wingbeat frequency of the bees.

3.2.2. Tracking

After a tracker is initiated, its position is updated constantly using an event-based update method. If an event occurs inside the tracking region, the position is updated according to the following equation:

$$\vec{x} = \alpha\vec{x_e} + (1-\alpha)\vec{x} + \vec{v}(t - t_{prev}) \tag{2}$$

where **x** is the position of the tracker, $\mathbf{x_e}$ is the position of the new event, **v** is the velocity of the tracker, t- t_{prev} is the time since the previous event, and scaling constant $\alpha \in (0,1]$. A similar method was used to track the movement of ping-pong balls in [63]. Each tracker maintains a

velocity value, which is updated at a fixed time rate Δt and given by:

$$\vec{v}(t) = (\vec{x}(t) - \vec{x}(t-1))/\Delta t \tag{3}$$

The velocity is critical to ensure smooth movement and minimize the chance of disruption when two bees occlude. When position updates are solely event based, the tracker can struggle to keep up with fast moving bees unless α is close to 1. Of course, raising the value of α presents its own issues, as the tracker movement will then become very sensitive to noise and can easily be disrupted. Monitoring velocity and including it in the position update equation allows the trackers to develop momentum, such that they can move very quickly if needed but are unlikely to suddenly change directions or respond to noise. Nevertheless, the choice of the constant α is still relevant, as too small of a value will prevent the tracking object from being able to respond to changes in direction in bee flight.

Additionally, it is very common for two or more bees to occlude from the perspective of the camera, passing through the same point in the field of view simultaneously. When this happens, the events generated by each bee will momentarily overlap before separating again. It is not possible to know with full certainty which bee is which after an occlusion, but a reasonable assumption is that the bees will continue traveling in the same direction as they were before the occlusion. Velocity is again useful here, as each tracker will have a tendency to travel in the same direction as it was before the occlusion, allowing it to stay with the original bee.

The tracking algorithm is also capable of responding to changes in the apparent size of a bee. As bees fly towards or away from the lens, their apparent size will scale. The default

starting value for a tracker is *r* = 20 pixels. This value was chosen after examining recordings of the bees and manually experimenting to determine the minimum value of *r* such that the majority of events corresponding to an average-sized bee are contained within the tracker. The value of r is then updated using an event-based equilibrium model. The radius shrinks at an unchanging exponential rate given by:

$$r(t) = r(t_{prev})e^{\frac{t-t_{prev}}{\tau_s}} \tag{4}$$

and grows whenever an event occurs in a border region outside the tracker boundaries between *r* and 4/3 *r*. The maximum value of *r* is fixed at 40 pixels to prevent uncontrolled tracker growth, and the growth rate slows as the tracker size approaches the maximum.

These two equations oppose each other such that the tracker will eventually reach an equilibrium where growth due to events outside the tracker is offset by the constant shrinkage, keeping the radius roughly constant.

This allows the algorithm to adapt to apparent changes in the size of a bee, but only to a certain point. If a bee flies away from the sensor, it will eventually not generate enough spikes for any tracking to take place. And as a bee approaches the sensor, it grows in apparent size, eventually taking up the entire field of view. In such an instance, the entire time surface would be contained in a single tracker, which would be kept active by every spike that occurs, including noise within the sensor's output. Thus a maximum tracker size is enforced to prevent the creation of large, self-sustaining trackers.

3.2.3. Tracking Box Removal

If a bee is no longer in the field of view, the tracking object associated with it should be removed. When a bee leaves the field of view, the tracking object follows it to the edge of the field of view and then stops, since no new events are being generated. To prevent the tracking object from persisting indefinitely, trackers must maintain a minimum event rate for a tracking object to exist. If the event rate drops below this minimum, the tracker is removed. This not only addresses bees leaving the field of view, but also removes the tracker if a bee lands on the hive and stops moving, since the bee will no longer generate events and as a result cannot be accurately tracked. If a bee lands and later takes flight, a new tracking object will be instantiated to track it once it resumes flying. Likewise, if a bee exits the frame of view and then returns, it will be treated as a new bee entering the field of view.

As with all choices of constants in this algorithm, the choice of minimum event rates involves tradeoffs. If the rate is too low, then tracking objects may continue to persist after a bee is gone due simply to events generated by background noise. However, if the threshold is too high, a tracking object might be removed while a bee is still in flight and in view. The appropriate value was determined experimentally.

3.3. Benchmarking

To analyze the algorithm and the limitations of the DVS under controlled conditions, an artificial test with a bee-like projectile was developed. The experiments were used to obtain data under ideal conditions, such as good lighting, high contrast, and proper focus. In particular, the distance from the sensor was matched to the hive scenario as well as the situation of

occlusions by multiple “bees”. These benchmark experiments enabled researchers to more precisely assess the performance capabilities of both the Dynamic Vision Sensor (DVS) and the tracking algorithm.

A range of techniques were used to simulate the flight patterns of bees. The primary test was tracking straight-line movements. This involved the use of small gel pellets (“Orbeez” ™) launched from a projectile device. To enable some degree of repeatability, a commercially-available launcher (aka “Orbeez gun”) was purchased to propel the small gel projectiles in front of the DVS. The projectiles were soaked in water overnight to produce translucent light blue water-gel balls approximately 7 mm in diameter, as pictured in Figure 5. The gel launcher lent itself to analyzing a variety of metrics such as: the highest speed (in pixels per second) that the algorithm could detect, as well as the fastest movement the DVS could detect. By introducing artificial obstacles to block the field of view (i.e., occlusions), the team could determine how the algorithm would respond to a “bee” temporarily “disappearing” and then “reappearing”. The effect of various conditions such as lighting and focus could be determined by using the gel launcher as a consistent simulator.

The logistics and results of these techniques are discussed in further detail below.

Fig. 5. Measurement of gel-ball, about 7mm [83]

3.3.1. General Test Configuration

For assessing how the algorithm would respond to bees colliding or occluding each other, several tests were performed with two objects meeting head-on. The first test involved attaching plastic toy flies to a thin piece of thread and swinging them in front of the DVS. A white poster board served as the background, consistent with the background used in the beehive recordings, which will be discussed further below. Maneuvering the plastic flies proved challenging as the thread tended to intertwine upon impact, compelling the team to explore the use of the gel launcher. The launcher proved to be a much better solution to testing speed and the limitations of the DVS and the tracking algorithm considering the smaller size and faster speed than manual flying objects.

Arranging collisions with the gel spheres proved impractical due to their smaller size and lower visibility (a translucent light blue) and thus larger objects at further distances were tested (white ping pong balls attached to thread). Ultimately, these tests were not fruitful due to the inability to control the movements of the balls.

Recording collisions was useful to understand how the DVS and algorithm would interpret head-on collisions. This was occasionally observed with the bees at the hive, however, occlusions tend to be more common. Occlusions were tested by placing a fixed object in the field of view and firing the gel launcher behind the object.

Initial testing early in the project occurred outdoors during the brighter morning hours where light levels average 8,800 lux. In these initial tests, the visual background was not uniform and occasionally had movement disturbances (i.e. trees moving in the wind,

pedestrians, etc.). For more controlled testing, testing was moved indoors with uniform intensity backgrounds for high-contrast. The light levels were kept constant throughout all the tests at approximately 200 lux. The distance between the DVS and the poster board was kept between 18 - 22 inches, consistent with the distances observed at the beehives.

3.3.2. Observations and Discussion

To vary the speed of the gel projectiles, they were launched at varying distances from the DVS in a series of outdoor tests. From around 2 feet, the algorithm was able to clearly estimate a speed of approximately 7000-10000 pixels per second. At a distance of around 4 inches, the algorithm was not able to record enough data points to make a confident estimation of speed in pixels per second. However, the DVS was capable of capturing the projectiles at that speed. At a distance of 6.5 feet, both the DVS and the algorithm were able to capture the gel projectiles. At this distance, the algorithm estimated the speed of the projectiles to range from 5000 to 7500 pixels per second, and the gel balls appeared relatively smaller in the visualization. Because of the widened field of view, the poster board could not be used to create a perfectly controlled background. In this experiment, 45 gel balls were launched at various rates of fire. As the direction of the gels was not entirely consistent, the algorithm was modified to determine if a cluster crosses a vertical boundary line in the center of the field of view, with a 10 pixel buffer zone on either side. The algorithm counted 46 total crossings and 44 net crossings. The one negative crossing was caused by environmental noise since the background could not be fully controlled. These experiments were conducted outdoors with light levels around 8000 lux.

Additional experiments were conducted indoors at a distance of 2 feet from the DVS. Indoors with fluorescent lighting the light levels were much lower at 200 lux. Even though the

gel projectiles were traveling at approximately the same speed, the DVS itself was not able to capture the projectiles indoors. Although not fully investigated, it is assumed that the lower indoor light levels produced significantly fewer events at the higher speeds of the projectiles, even though events are plentifully generated with slower moving stimuli (e.g., waved hand) and the DVS pixels adapt to different light levels.

3.4. Recordings of Bees at the Beehive

While the target simulators can model many aspects of bee flight that are important for tracking, there are some aspects that cannot be represented easily. Erratic movements, changes in direction, and large numbers of simultaneous targets are among the features of a live recording that are challenging to produce in a target simulator. Once the basic algorithms and hardware were adequately debugged, data was collected live from the active beehives maintained by the University of Maryland Bee Laboratory (https://www.umdbeelab.com/).

3.4.1. Camera Apparatus

To obtain long data recordings from real beehives, the DVS must be prepared for outdoor operation for an extended period of time, along with a traditional CMOS camera (i.e., USB webcam) for visual image verification. An apparatus was developed to hold both during long periods of data collection. This apparatus needs to be able to withstand reasonable climate conditions, including light wind and rain.

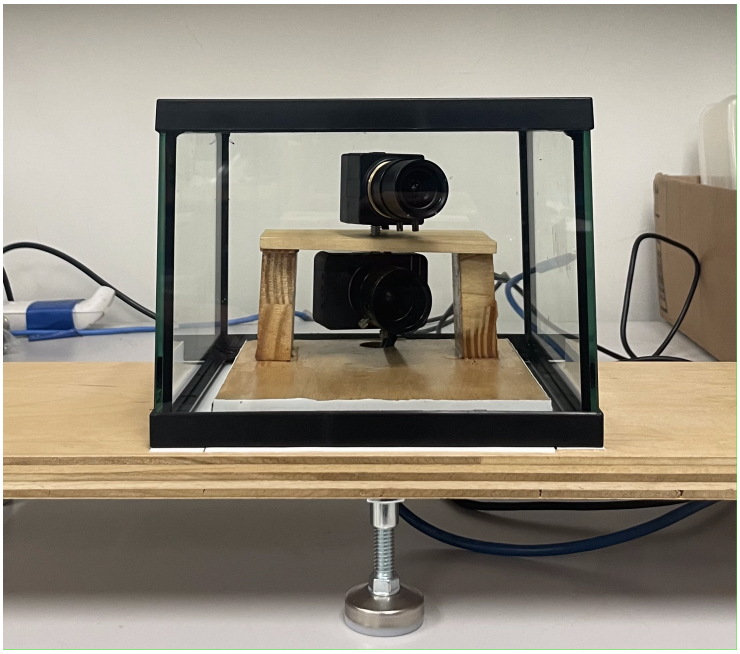

Fig. 6. The image above is the camera apparatus implemented. The webcam sits on the top of the bracket and the DVXplorer sits below the bracket. The cameras are protected by the aquarium and sealed with silicone tape.

The apparatus also needs to include a mechanism for quick and easy alignment and leveling as well as repeatable focus and aperture setting procedures. As the DVS and traditional camera are operated independently with different but overlapping fields of view, it is important that the apparatus allow alignment of these fields of view as easily as possible, particularly along the rotational axes to each other and to the beehive itself.

The total footprint of the apparatus is roughly based on a 1’x2’ sheet of ¾” plywood. To enable tripod-like leveling, three threaded legs were added, two at the rear corners and one at the middle of the front edge. Tightening or loosening the screws will change its level alignment. With the assistance of two bubble levels, the apparatus can stay level for consistency between measurements and alignment with the hive.

To protect the DVS and camera from weather conditions, several features were added. A waterproof seal with transparent material was provided with a small pet aquarium and silicone tape. The camera lenses were able to catch light with little to no obstruction through the aquarium while still being protected from the elements. The flexibility of the silicone tape allows for the aquarium to be pressed into it to form a proper waterproof seal. The entire wooden surface of the apparatus was also twice coated with a waterproof sealant to prevent warping and other forms of water damage.

For further protection, a roughly 9”x9” sheet of ¾” plywood was centered beneath the aquarium, providing a lift of ¾” from any leaking water to further seal off the DVS and camera. For accessibility of the camera cabling, a 1” diameter hole was drilled through the base to snake the DVS and camera wires out of.

Finally, on top of the raised plywood, a brace was constructed of scrap wood. The brace is centered roughly 6” back from the front of the raised plywood base. The legs of the brace are made of ½”x2” wood cut into two 1 ½” strips that were glued to the board so as to suspend the main brace 2” off of the base. The main brace is a ¼”x2” plank of wood glued to both legs, with a ¼” diameter hole threaded into the center.

The threaded hole receives a threaded rod capable of receiving the camera from above and the DVS from below via the tripod screw holes on both devices. The choice to use a single threaded rod was for alignment. A single shared rod locks X-axis and Z-axis alignment, thus ensuring that the only alignment instability is locked to the Y-axis. To limit this instability, rubber washers were inserted between the brace and the devices. This added friction to the assembly, keeping the devices aligned along the Y-axis once they were aligned by hand.

Despite alignment and leveling, the field of view of the DVS and camera are not exactly the same, and the boundaries the algorithm uses cannot be perfectly projected onto the camera's view. To account for this, the DVS boundary is defined such that it appears in the camera's field of view. The location of the bounding box around the hive entrance in the camera's view was also recorded for verification.

Fig. 7. A demonstration of the setup of the DVS in front of the hive. The white poster board was used to provide a uniform, high-contrast background. The DVS is enclosed in a glass box to allow long-term recordings without concerns about weather damage. The box was tested to ensure that glare and condensation would not have an impact on its performance.

3.4.2. The Visual Recording Environment

Event-based sensors require spatial contrast in the image to detect light intensity changes due to movement. Thus, the background behind the objects of focus can have a significant impact on the sensor's ability to detect and track these objects. The performance of the algorithm was evaluated both with a controlled background provided by a white poster board, as shown in Figure 7, and an uncontrolled background of a hill next to a building. The controlled background offers more contrast, but requires adding an additional component to the hive. It was also observed to alter the bees' flight patterns, with many of them hovering around the poster board. The uncontrolled background minimizes the changes to the environment, but also risks creating regions where a bee does not generate spikes because it does not contrast with the background.

During a recording session, the apparatus is placed on cinder blocks to the side of the hive entrance, pointed towards the bees entering and exiting the hive. The camera is positioned so that it points away from other hives, removing the effects of out-of-focus bees in the background. The device is placed between six inches and two feet away from the hive, after which the exact distance is recorded. A white poster board is present across from the apparatus to reduce background noise in some recordings. General weather conditions and light levels are recorded for each session, and both the event-based sensor and camera are manually focused before each recording. The light levels varied based on time of day and amount of shade, with the lowest recorded values being 4.07 klux and the highest being 46.3 klux.

For shorter recordings, the system is supervised until a satisfactory duration is recorded, at which point the devices are stopped. Such recordings last between twenty minutes and an hour, and the data from these recordings is saved onto a connected laptop.

For longer recordings, the sensor is attached to a Raspberry Pi 4™ and a large battery power bank (Jackery™ portable power station) instead of a laptop. The apparatus was set up in the morning near the hive at around 7 AM. Neither event data nor webcam data is recorded for these sessions to save on storage and computation time. Instead, the total and net counts are stored until the evening when the system is recovered. The system is not left in place overnight, since bees are not very active and the vast majority of detections would likely be noise.

The proposed method for tracking bee populations is not a perfect strategy. It is possible for a bee to enter or exit the hive without being detected by the sensor. For example, an exiting bee could go around the hive by moving away from the sensor and going backward. For the bees at the University of Maryland Bee Lab, this movement was unusual, as exiting bees were more likely to fly straight out from the entrance. On the other hand, a bee did occasionally enter the hive by crawling around from the back of the hive, which meant that it would be nearly impossible for it to be visible on the DVS or the camera.

Chapter 4. Counting

4.1. Counting Algorithms

After the bees have been tracked, the movement of bees into and out of the hive needs to be recorded. Originally, this was done by considering a vertical line near the hive entrance; a bee that crosses the boundary is considered to have either entered or exited the hive based on its direction of travel across the line. This approach is likely to miss bees that enter or exit the view of the DVS from above or below, as these bees would not be detected moving across the line. To account for this, the algorithm instead considers a box that surrounds the hive entrance. This 2D box represents a projection of a 3D boundary extending out from the vision sensor (see Figure 8 below). The boundaries are such that the box is square and approximately 50 cm wide at a z-distance corresponding to the middle of the hive entrance. Due to the nature of projection, the boundary widens as the distance from the vision sensor increases. Additionally, a buffer zone on either side of the boundary to prevent hysteresis if a bee hovers very close to the boundary (see Figure 9 below).

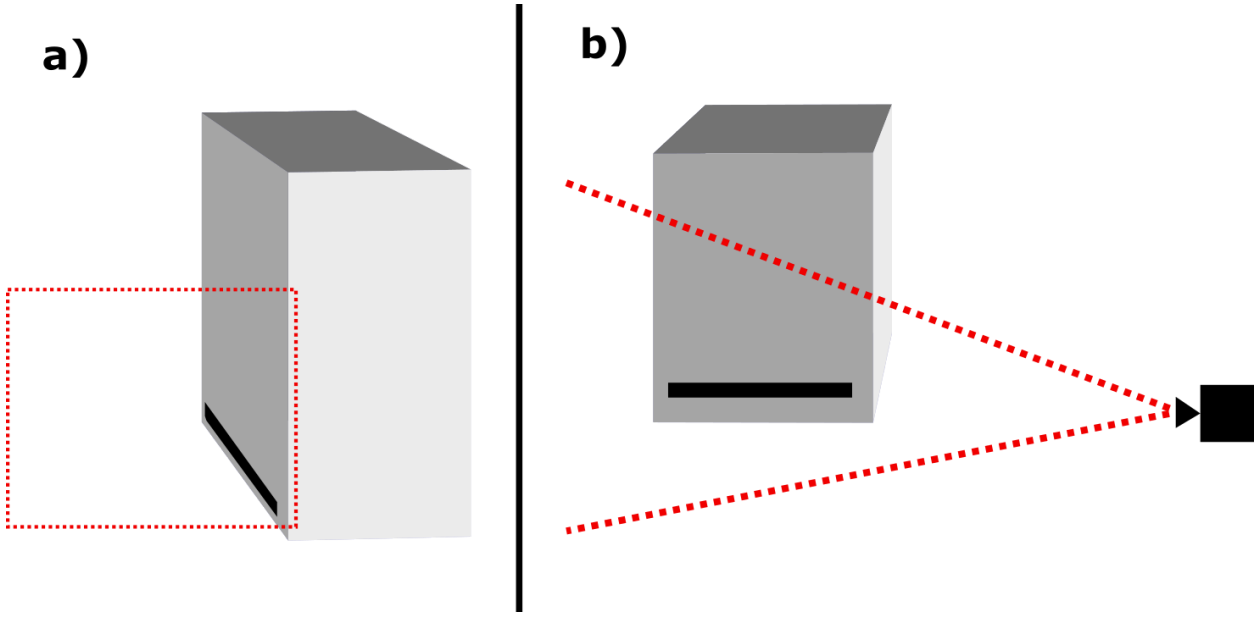

Fig. 8. A graphical depiction of the projection of the bounding region onto the DVS field of view. Fig. 8a shows the bee hive as viewed from the perspective of the DVS, where the square box shown is monitored to count bees entering and leaving. Fig. 8b shows the hive setup viewed head-on. The diagonal red lines represent the region enclosed by the bounding box at different depths from the DVS. As a result, the bounding box encloses a smaller physical area the closer one gets to the DVS.

When each bee is detected, it is first determined whether it is inside or outside the box. This value remains constant unless the bee is detected on the other side of both buffer boundaries, at which point it is considered to have crossed the boundary. The value indicating bee position in relation to the box is updated accordingly. The algorithm can then maintain a running count of the number of bees that enter and exit the hive entrance area.

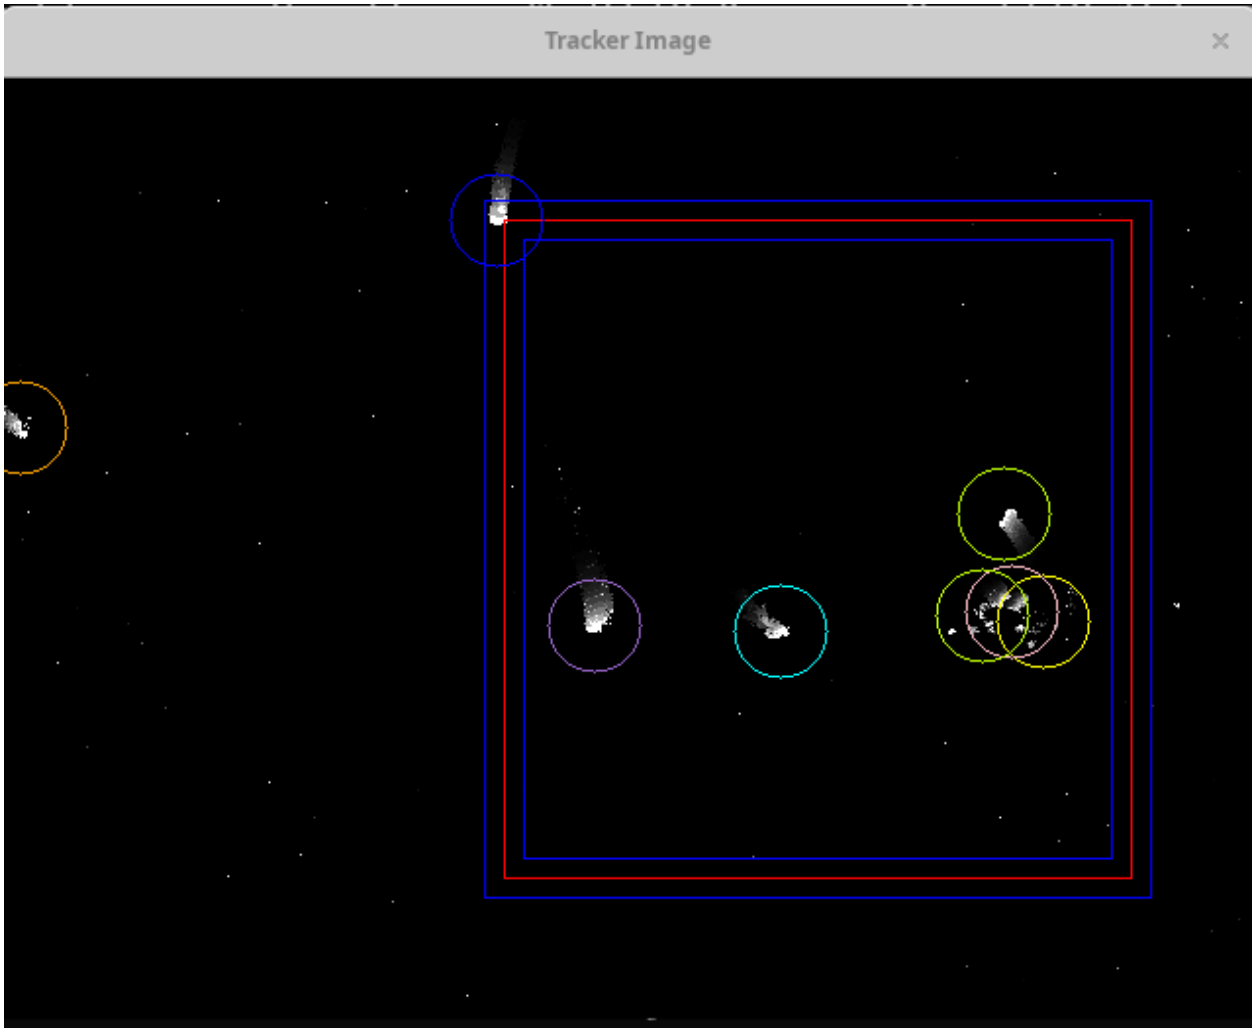

Fig. 9. A screenshot of the hive setup as shown by the DVS visualizer. The colored circles show regions of activity where the tracking algorithm has identified a bee. The red line represents the bounding box that a bee must cross to enter or leave the hive. The two blue boxes inside and outside the red box show the hysteresis threshold. A bee must cross both blue lines to be counted as having crossed the boundary.

This method does have some limitations. A bee flying away from the sensor may not be counted as entering or leaving the hive. A bee leaving the hive and flying away from the sensor would shrink until it disappears without crossing any boundary. A bee flying towards the sensor still must cross the threshold at some point to leave the frame. However, if it is sufficiently fast and close, it is possible the bee could leave the frame before the update phase is executed and

a crossing is detected. Both behaviors are very unlikely and occur very infrequently in live recordings.

4.2. Challenges to Accurate Bee Counts

There exist certain conditions under which the event-based recording will not produce the correct bee count.

One clear source of interference is the movement of objects or background elements in the recording. At a minimum, the system must be oriented to avoid viewing moving pedestrians, vehicles, or animals. Wind on background objects like tree branches, bushes, or tall grass will produce events in the recordings, triggering the tracking algorithm to generate erratic false positive counts of bees crossing the threshold. Additionally, the sensor detects moving objects based on detecting changes in pixel intensity which only occurs when the object has a light intensity different from its background. If the natural background has regions with light intensities comparable to the bees, the moving bee will not produce events in that part of the image. While the DVS is exquisitely sensitive to tiny changes in intensity, and natural backgrounds generally produce good event data, the uncertainty created by moving backgrounds and variable contrast was problematic in the initial studies. To mitigate this problem, a fixed white poster board that covered the background of each recording was used. In addition to reducing the triggering of unrelated events, the board also makes manual verification easier, as bees are easier to track with the poster board when manually counting bee crossings on traditional camera recordings. Sometimes, the poster board was not used. The primary reason for this was the noticeable effect that the nearby board had on the flight patterns of the bees. After placing the poster board near the hive, there would be significantly

more flight activity, and bees would hover around the poster board instead of flying into and out of the hive as normal. This behavior did not seem to stop even after the board was placed near the hive for a long time period. The non-board recordings provide an opportunity to see how the recording procedure works under normal bee flight conditions.

Another factor which could affect the recording quality is the weather. The event-based sensor generally works better under brighter light due to the higher contrasts that can exist in those conditions. Rain can produce vertical trails of events in the image that can be tracked, potentially creating false counts. For shorter recordings, the light level was recorded manually, but light level varied throughout the day during the longer recordings and were not recorded for the duration of these recording periods, mostly due to the difficulty of regularly measuring light levels throughout the day. Other weather factors were also left unrecorded.

Factors regarding how the DVS apparatus is set, like recording distance and sensor focus, can affect the performance of the algorithm. It was not easy to maintain a constant focus across recordings, as the DVS does not provide functionality for measuring the exact focus settings. During recordings, focus was manually adjusted to provide the clearest view of objects at the approximate distance of where the bees would be flying. Bees that appear larger in the field of view generate more events in close proximity, and are therefore more likely for the counting algorithm to detect. The algorithm also requires the bee to be in the field of view for enough time to be detected before and after it crosses the boundary of the counting box. These conflicting factors imply that the DVS should have a broad view of the hive without being too far away from it. After balancing these factors, the apparatus was consistently placed approximately two feet from the hive entrance.

4.3. Validation

To determine the accuracy of the tracking algorithms, a manual counting process conducted by a human is required. It is necessary for the people counting to carefully analyze frame-by-frame on a traditional camera recording of bees entering and leaving the hive, as it is one of the only ways to verify the algorithm's accuracy. The DVS cannot be used to verify this data since its resolution isn't as high as a traditional camera, and it can only see moving objects and not stationary ones, which would make it difficult to see when they enter a specific region, such as around the beehive.

A 5.0 megapixel webcam operating at 30 frames per second simultaneously recorded the bees alongside the DVS to enable manual validation of the tracking results. The frame rate of the traditional camera was usually high enough to manually track the bees for verification, and the resolution was high enough to see the body of the bee at all times in the recording while it was in the field of view. However, at times there would be confusion with tracking overlapping, clustered bees, or bees that blend in with objects in the background, such as trees or buildings. Additionally, the frame rate is low enough for bees to appear to move significant distances when comparing adjacent frames. The motion of bees could still be tracked from one frame to the next, but the exact motion of bees between frames was not detectable via the webcam.

Fig. 10. These two images are adjacent frames from the webcam recording for box-based validation. The bees' positions vary widely in these two frames, as some bees cross paths, while others end up overlapping.

In the apparatus, a traditional camera is set up directly above the DVS to ensure the two cameras' fields of view are as similar as possible. After recording the bees, a clip that is a few minutes long in the recording session is taken from both the DVS and traditional camera. Then, the bounding area, also known as the threshold, is added to the recording from the traditional camera through video editing, while also adjusting for the difference in field of view. The threshold is a boundary with a 10-pixel buffer that the algorithm uses to determine whether a bee has entered or left the hive. A crossing is counted once the bee has fully crossed this threshold. This threshold, after being edited into the video from the traditional camera, can be used to track bees entering and leaving the hive manually during the counting process. The person verifying the counting of bees goes frame by frame, tracking each bee that crosses the threshold in the clip. They manually enter each crossing, with separate tallies for bees entering or leaving from the left, right, top and bottom threshold.

After manual verification, the tallies are added up to compare to the tracking algorithm. Originally, a single-stage verification process was used. The algorithm would run through the recording from the DVS, and track both the number of bees detected and the net number of

bees crossing the threshold in and out of the hive. Finally, the human counting and algorithm counting would be compared, and the accuracy of the program would be computed. We eventually changed to using a two-stage verification process which makes it easier to determine exactly which bees are not being detected. Instead of only looking at the webcam algorithm, validators would also observe the corresponding DVS recording. After ensuring that the two recordings have been properly synced up, differences could be noticed between the webcam and the DVS. There are multiple ways an error can occur: a bee may not be detected by the DVS, or a detected bee may not be properly tracked by the algorithm. The validator can keep track of the issues causing discrepancies between the two recordings.

As shown in Figure 10, this threshold can take the appearance of a box that surrounds the hive entrance. However, alternative threshold shapes exist, such as a single vertical line that a bee crosses when it leaves the hive straight from the entrance. We found that there were sometimes issues with determining the exact alignment of the box from the points of view of both the DVS and webcam. Thus, to test the accuracy of the algorithm, we instead used a vertical threshold near the entrance of the hive, as seen in Figure 11. This figure also demonstrates the differences between the fields of view of the webcam and the DVS.

Fig. 11. A webcam recording of bees with a vertical boundary added for verification purposes. The red highlighted areas on the left, right, and top sides of the frame signify regions which are not in the field of view of the DVS. While performing the two-stage verification process, validators use the region which is not highlighted to determine where the DVS would report each bee's location.

4.4. Results

To determine the accuracy of the DVS bee counting algorithm, we compared a 90-second sample of "bee crossing" events as counted by the DVS against the manual count that we performed from analyzing the webcam footage, as mentioned above. A bee crossing event is when a bee flies across the vertical boundary in the middle of the screen. We controlled the background using a white poster board, increasing the visual contrast for a higher confidence in the accuracy of both the DVS and manual counts. The results of this comparison are shown in Figure 12. Due to the frame rate of the webcam, crossing times cannot be

determined at the microsecond precision the DVS reports, so the counting events were binned into 1-second intervals.

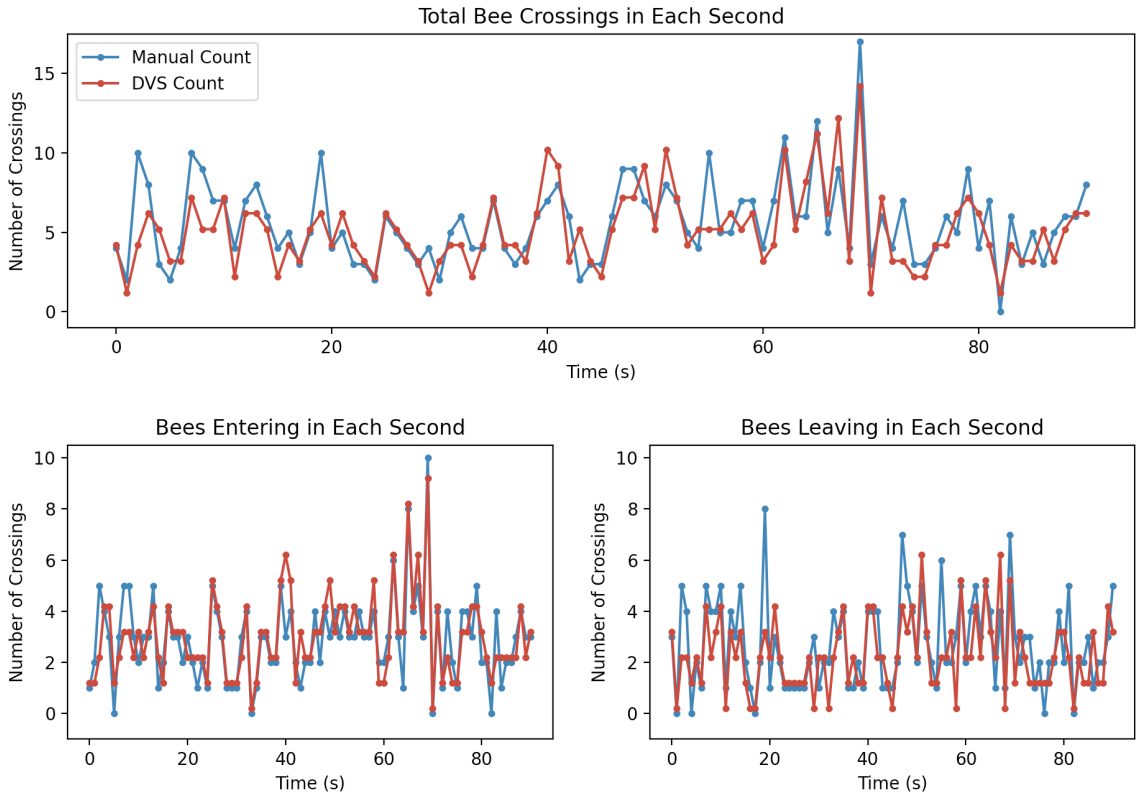

Fig. 12. A comparison of the bees counted by the manual webcam verification process (blue) and the DVS algorithm (red). Bee crossing events are binned in 1 second intervals for ease of comparison. The red lines showing the count performed by the DVS are offset by 0.2 from their true values so as to not to obscure the manual count when the two values line up exactly. The top graph shows the total bee crossing events counted by each method (i.e. any time a bee crosses the boundary), while the bottom two show the counts of bees entering and exiting the hive, respectively. These graphs indicate that the DVS counting algorithm generally matches up with the manually verified counts, but sometimes misses bees leaving the hive.

To quantify the algorithm's performance, we measured the true-positive, false-positive, and false-negative rates for bees entering and exiting the hive. A true-positive occurs whenever the DVS and manual counting both count a bee entering or exiting the hive within the same second. A false-positive occurs if the DVS counts a bee crossing the boundary that the manual verification does not count, while a false-negative occurs if the DVS fails to count a bee crossing

that was counted during manual verification. Over the 90 seconds there were 233 true-positive bee entering achieved an F-score of F-score of 0.89. For bees leaving had an F-score of 0.80, due to a low recall of 0.71. The full results are shown in table 1.

October 7th Results			
	Entering	Leaving	Total
True Positive	233	177	410
False Positive	25	17	42
False Negative	33	72	105
Precision	0.9031	0.9124	0.9071
Recall	0.8759	0.7108	0.7961
F-Score	0.8893	0.7991	0.8480

Performance of counting algorithm on data with controlled background. Entering denotes bees entering the hive, leaving is for bees leaving, and total represents the total number of crossings during the time periods. Each datum is the number of crossings in a binned 1-second period.

Chapter 5. Wingbeat Frequency

5.1. Periodicity in Event-Based Data

The ability to accurately determine the wingbeat frequency of a bee using neuromorphic vision is contingent on the assumption that the event-based data will contain some type of periodicity that corresponds to the activity of a bee's wings. We hypothesized that the rapid wing movement would generate periodic sequences of events. To determine the nature of this periodic signal, we analyzed the event history of multiple tracked bees. We initially looked for patterns such as offset periodic signals in events found in different regions of the bee, which we assumed would occur based on the fact that the wings of a bee have been shown to move in a circular pattern [25]. Our results indicated, however, that the strongest periodic signal could be found in the total sum of events across the entire bee within a fixed time window of 1 ms. Figure 13 shows this signal for an example bee.

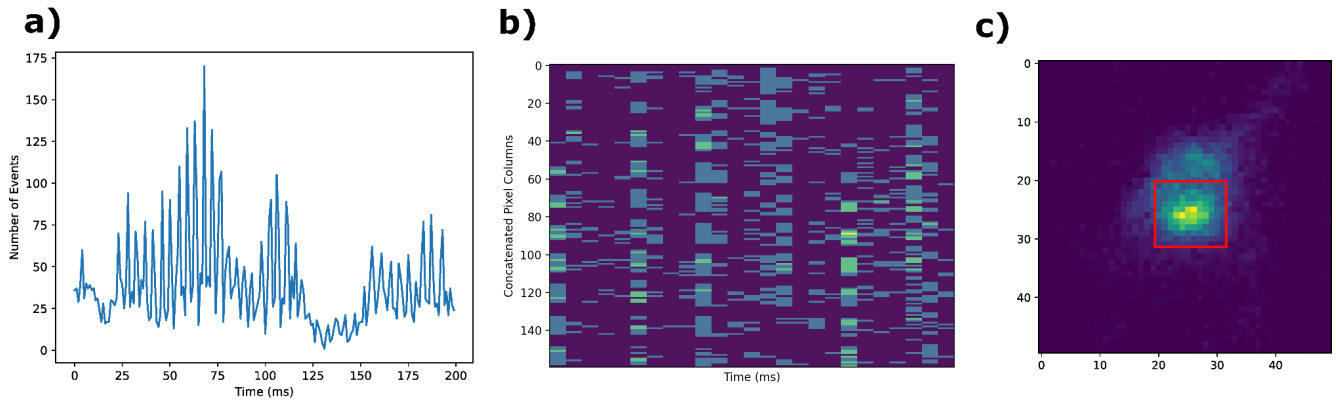

Fig. 13. a) The total amount of events occurring within the entire 40 by 40 pixel region surrounding a bee, binned in 1 ms intervals. This signal visually shows strong periodicity, oscillating regularly between samples with high and low event counts.
b) A heat map showing the presence of events in each pixel within the center of the ROI over the time period from 50-75 ms after the initiation of the bee tracker. Each column is made by concatenating ten rows of pixels from the center of the ROI. This again shows the presence of periodicity, this time indicated by the visual appearance of

prominent vertical lines in the image. These vertical lines represent moments during which many pixels throughout the center of the ROI emitted events, followed immediately by periods during which few pixels initiated events. c) An image showing the relative firing frequency of each pixel in the ROI. The red box indicates the region that was selected and unwrapped to form columns in image b). We selected only the region with highest activity to most clearly show the periodicity effect.

5.2. Frequency Extraction Methods

We tested several different methods for extracting the frequency of the wingbeats from the sequence of events. The advantages and disadvantages of each method are outlined below.

5.2.1. Fourier Analysis

Given the strong periodic nature of the signal generated by summing events across the bee, we extract frequency by determining the discrete Fourier transform of this signal. DVS events occur asynchronously, so for an arbitrarily small time window the sum of all events is likely to be close to zero. To extract a signal with clear periodicity, a bin size of 1 ms was used; that is, the signal S(t) at time *t* represents the number of events that have occurred in the region of interest within the 1 ms interval surrounding *t*. Wingbeat frequencies of honey bees are known to be in the range of 150-250 Hz [28], so we are confident that a 1000 Hz sampling frequency is high enough to avoid any Nyquist sampling errors.

An important parameter for the discrete Fourier transform is the number of samples to collect before performing the transform. There is a clear tradeoff between accuracy and latency; more samples leads to a higher frequency sampling resolution and decreases the chance of generating a very noisy Fourier transform, but it also takes more time to collect these samples. From our observations, an individual bee was typically only in the field of view for 500-1000 ms, which puts an upper bound on the number of samples that can be collected. Our results determined that noisiness of the DFT stayed roughly constant as long as the sample length

stayed at or above 250 samples, allowing for four frequency readings per second (see Validation for more details). To determine the frequency of the wingbeat, we simply choose the frequency corresponding to the DFT coefficient with the highest magnitude.

5.2.2. Delay Timing

If the events in a region are generated by a repeating pattern, spikes should exhibit a similar pattern. Measuring the time difference between repetitions in the pattern yields an estimate of the frequency at which it occurs. Patterns are identified based on “transition” events, where the previous event at the same location had a different polarity than the current event [80]. A running average of the time delays between transition events yields an estimate of the periodicity in the region where these events occur.

Previous implementations of this method have used it to identify the location of a known, clear signal such as an LED [80]. For detecting the frequency of bees, the problem is reformulated to estimate the frequency of a noisy signal in a known region. The bee also does not remain in the same location and orientation within the tracking box at all times, further increasing the complexity of the problem. Multiple strategies were attempted to combat the particular difficulties of finding the wingbeat frequency of a tracked bee.

As a bee moves and changes direction, it is unlikely that the same regions of the tracking box will be activated, so individual pixels may not generate enough events for the frequency to be detected. However, since preliminary analysis showed the total number of spikes fluctuating, transitions could occur over the tracking box overall. Thus, the first method was to implement the transition-based delay timing treating the entire tracking box as a single pixel, denoted the

Cluster-wide Transition method. The running average of the time difference between transitions was recorded and used to later report the frequency.

The Cluster-wide Transition method has the potential to be susceptible to noise and the effects of the trailing edge of the bee. To avoid this, we also test a Average Grid Transition method, where the tracker can be broken into subregions, each of which acts as a pixel timing its transitions. Each region is a 3x3 box of individual pixels, where a difference in polarity between any of the spikes in the subregion is considered a transition for the region. A 7x7 grid of these regions are stored in each tracker, and when the frequency is reported, the tracker averages the record from each region that encountered enough transitions.

The final method, the Time Surface Delay method, was designed to take advantage of the burst of spikes that occur within a tracker as the number of spikes fluctuates. It uses a time surface with a fast decay rate to determine when a number of spikes have occurred in rapid succession. When the surface value crosses a threshold, the time delay is recorded just like the other methods. A minimum amount of time between successive activations must also be enforced, since simultaneous events may both cause value to surpass the threshold, despite being a part of the same cycle. The threshold is currently adjusted manually, but could be automatically determined using the peaks of warm up cycles.

5.2.3. Forced Harmonic Oscillators

The sequence of events generated by a bee can be used to drive a harmonic oscillator. Harmonic oscillators are periodic systems that can be modeled by differential equations that determine a resonant frequency. When the harmonic oscillator receives inputs at regular intervals close to the resonant frequency, the amplitude of the system grows. If it receives

inputs at intervals that do not match the resonant frequency, the amplitude of the model will not grow as quickly, but may still grow if it is close to the resonant frequency.

As a result, an individual harmonic oscillator is not an accurate frequency measurement. If an outside frequency is close enough to the resonant frequency of the oscillator it can occasionally hit a peak that would surpass a threshold. However, by modeling several oscillators simultaneously with different resonant frequencies, the frequency of the wingbeat can be better approximated. The oscillator with the closest resonant frequency will grow the fastest, achieving the highest peak of the series. By selecting the oscillator with the largest amplitude, the most accurate measure of the frequency of events can be obtained.

The main advantage of the forced oscillator method is that it does not require a fixed sampling rate or number of samples to determine the frequency; it is entirely event based. However, a major disadvantage is that it has low frequency resolution, requiring an individual oscillator model for each desired frequency. Tracking too many frequencies can be very computationally expensive, so the system must be limited to monitoring a small range of frequencies. Based on known average ranges of the wingbeat frequencies of bees, we chose 12, 5 Hz bins between 190 and 250 Hz.

5.3. Validation - Confidence Metric

Our recordings showed a large variation in the strength of the periodic signal associated with the wingbeat frequency of a bee. This variation could be caused by factors such as the orientation of the bee with respect to the camera, the length of time that the bee is in the field of view, the speed of the bee, and the effect of occlusion due to another bee. Fig. 14 shows

examples of a strong and weak periodic signal, along with their corresponding Fourier transforms.

One of the main applications of this method is the potential to monitor beehives for potential invaders, such as wasps or hornets which have a very different wingbeat frequency. For this to be effective, there must be a high degree of confidence that the measured wingbeat frequency is accurate, so that the system cannot misclassify a honey bee as an invader. This requires a method of differentiating strong periodic signals from weak periodic signals.

To accomplish this, we introduce a confidence metric, C, defined as follows:

$$C = \sum_{f=0}^{F_{max}} \frac{|X[f/f_s]|}{|X[f_{peak}/f_s]|} * |f - f_{peak}| \quad [6]$$

where |X[k]| represents the magnitude of the kth Fourier coefficient, f_s is the sampling frequency, and f_{peak} is the frequency corresponding to the Fourier coefficient with the largest magnitude. This metric is effectively a measure of kurtosis; the value will be smallest when the Fourier transform contains one global of much larger magnitude than all other frequency values, and largest when there are multiple local maxima of similar magnitude that are spaced far apart in frequency. Figure 14 shows some examples of Fourier transforms and their corresponding confidence metric.

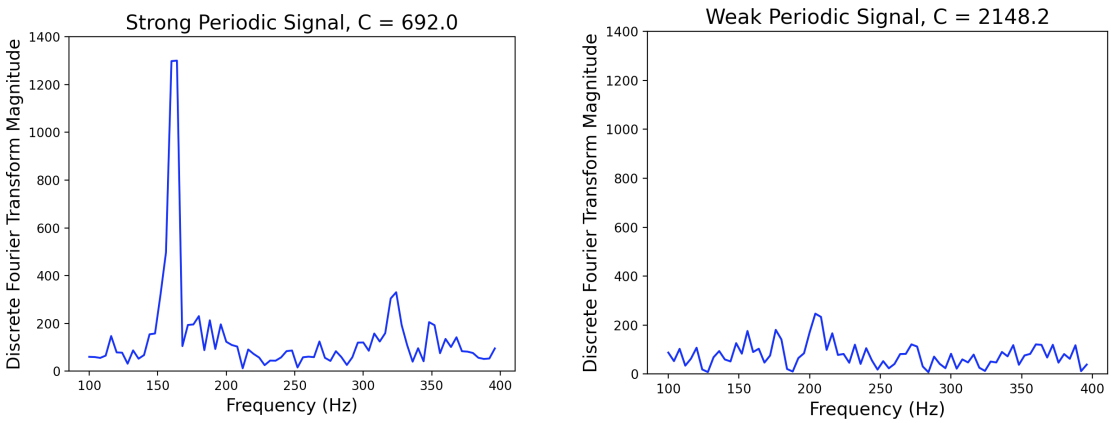

Fig. 14. A comparison of frequency spectra extracted from the event sequences of two different bees. The spectrum on the left has a relatively low C-value (C = 692.0), and shows a well-defined peak at a frequency of 170 Hz. The spectrum on the right has a much higher C-value (C=2148.2) and is visually very noisy, with no obvious frequency peak. This demonstrates that spectra with lower C-values tend to have obvious peaks in the frequency domain, indicative of a strong periodic signal. This provides a way to automatically distinguish between clean and noisy signals to ensure that only those signals for which a clear frequency peak can be established are used when tracking wingbeat frequencies.

In general, our results indicate that it is not possible to accurately determine the wingbeat frequency if the confidence metric value is greater than 2000. Approximately 54% of the bees we sampled had a confidence value <2000, meaning the wingbeat could accurately be determined. Of course, this choice of cutoff is arbitrary, and anyone who uses this frequency detection algorithm may choose a more stringent cutoff value depending on their use.

5.4. Results

Each method was first evaluated on a controlled LED experiment to evaluate baseline accuracy. As shown in Figure 15, the Fourier Transform method was the most effective method at detecting a clear LED signal. However, it does have drawbacks. Its report speed is much slower, often only reporting 2 or 3 times before a bee leaves the frame due the amount of data and computation time needed for the analysis. The Average Grid Transition and Forced

Oscillator methods offer a higher reporting speed, at the cost of accuracy. For the LED data file, the Fourier method reported 77 times, while the Average Grid Transition and Forced Oscillator methods reported 3909 and 2934 times respectively.

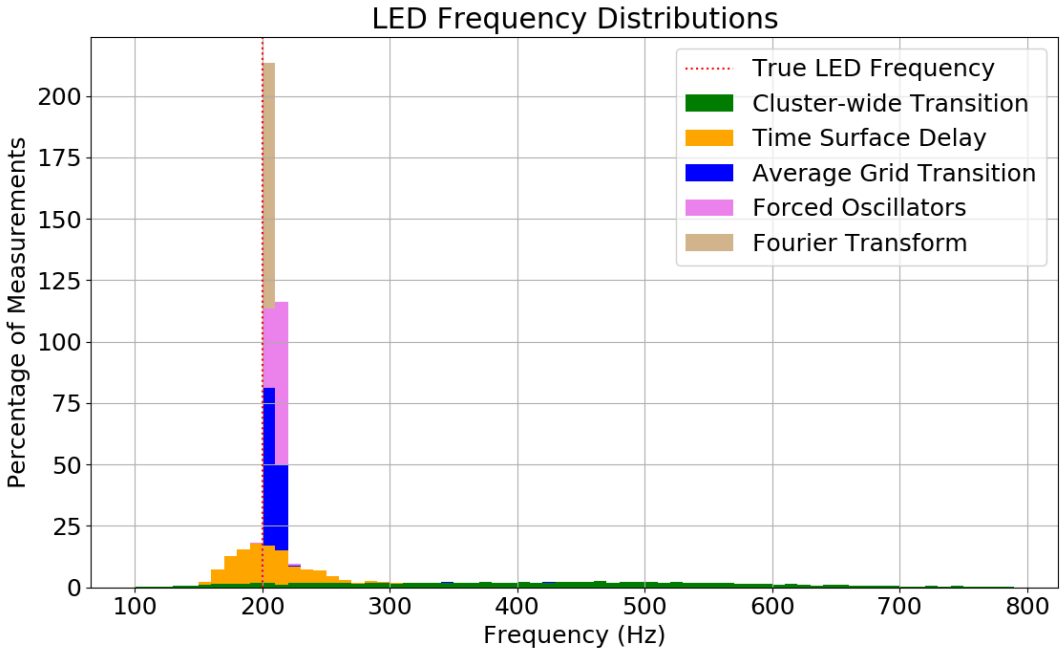

Fig. 15: Stacked distribution of reported frequencies on the LED flickering at 200 Hz for different frequency detection methods. The Cluster-wide Transition and Time Surface Delay methods were not able to reliably detect the LED's blinking frequency. The Average Grid Transition and Force Oscillator methods got a better approximation, but were overall unreliable with fluctuating readings. The Fourier Transform was the most accurate and reliable of the methods reaching within 10 Hz of the actual frequency for 100% of its reports.

Of the five methods examined, the Fourier-transform frequency extraction method provided the most reliable estimates of frequency of the LED. However, the wingbeat frequency of a bee is not as stable and ground truth is not known. Although we do not have a direct way to measure the accuracy of each method, since there are no ground truth labels, we can use indirect measures of reliability. For example, the Fourier transform method showed greater consistency across consecutive measurements (i.e. it was less likely for the same bee to

suddenly and abruptly change its wingbeat frequency). Additionally, the Fourier transform method generally produced clean signals as measured by the confidence metric. The distribution of confidence metrics is shown below in Figure 16.

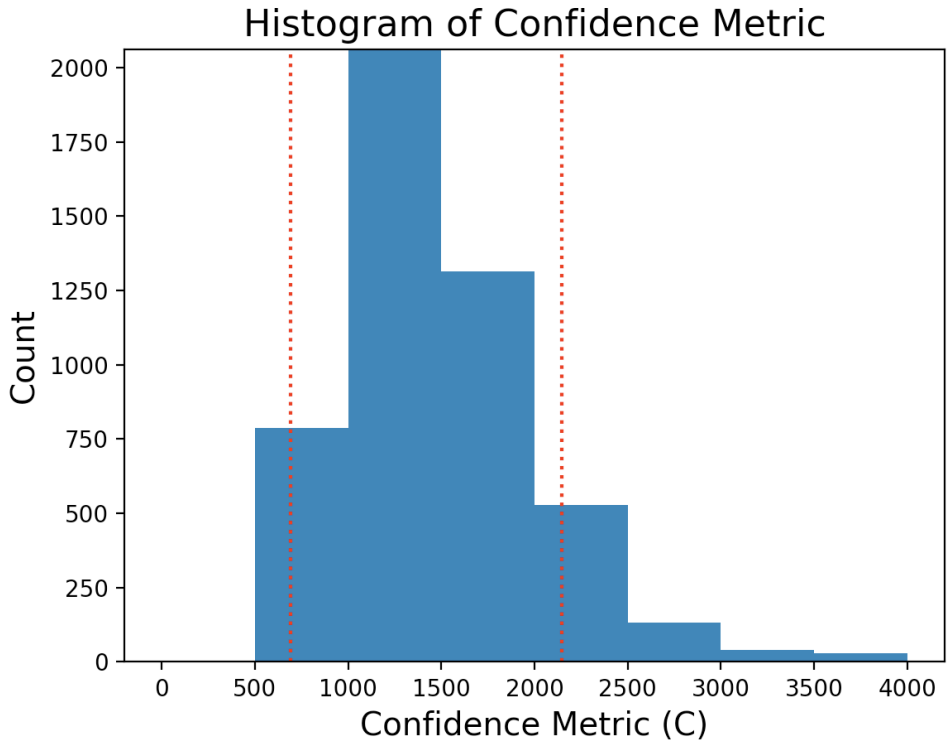

Fig. 16. This histogram shows the confidence metric for the Fourier analysis of 5,000 bees wingbeats. The red lines correspond to the clean and noisy samples shown previously in Fig. 14. These results demonstrate that the majority of the bees sampled (83%) are within the C<2000 threshold for reliable wingbeat measurement.

Across a sample of 5,000 bees, we found the above distribution for the confidence metric using the Fourier-transform method. In this sample, 83% of the bees were below the chosen cutoff threshold of 2000, meaning that it is reasonably likely that the frequency measurements for those 83% matched the true wingbeat frequency.

Chapter 6. Discussion

6.1. Counting

Bee colonies are best able to care for themselves when they have an adequate population for the size of their hive. As such, colony population is an important metric for beekeepers and entomologists to track the overall health of a bee colony. Changes in population as measured by the algorithm and the DVXplorer help to track overall population changes when given a baseline, such as a purchased colony with a known number of bees or an initial weight measure. Additionally, changes in bee population from entering and leaving the hive can provide helpful metadata. A large excess of bees leaving over those that enter can be signs of swarming, a bee behavior whereby the colony splits into two new colonies. An excess can also be a sign of high forager deaths, meaning the colony will be missing important resources. An excess of bees entering over those leaving can be signs of invasion, whereby wasps, bees, or other insects invade the hive in search of resources. All of these things are damaging for the health of the colony, and knowledge of them is important to protect beekeepers and help research entomologists.

In this work, we demonstrated an event-based algorithm to track changes in colony population. The counting algorithm showed the potential to work in the field, with manual verification showing that the system performs well under ideal conditions, but still has some issues to be addressed. The relatively low error rate for bees entering the hive shows that the system is capable of counting bees in real-time, and the fact that there were approximately equal numbers of false positives and false negatives demonstrates that the algorithm does not

have a bias towards overcounting or undercounting bees entering the hive. The high false-negative rate for bees leaving the hive, however, presents a problem since it creates an imbalance when tracking the net change in population. This is most likely due to the fact that activity near the hive entrance limits the number of tracking objects that can be initiated in that location due to crowding; by the time a bee has gotten far enough away from the hive entrance to generate its own tracker, it may be too close to the boundary to be counted.

Additionally, attempts to verify the performance of the algorithm may further be hindered by problems with the manual verification process itself. It is possible that human error occurred during the counting stage. The traditional camera frame rate does not have the temporal resolution of the DVXplorer camera. Sometimes during the counting process, many bees were active in the current frame; In the next frame, all the bees would have moved significantly (occasionally switching positions) due to their high speeds. This made counting the bees difficult in certain circumstances, and may have affected the human counting totals. Also, overlaps of bees occurred during several frames, and so people verifying, as well as the algorithm, may have gotten confused in tracking.

Another limitation is the field of view difference in the traditional camera and the DVXplorer, as the manual counting on the traditional camera may see more or less of the beehive or its surroundings depending on its sensor and lens. This can affect the manual verification process due to certain bees being included or excluded that may or may not have been counted by the algorithm. Being able to spot bees in the recording when there is a complex background is another barrier. For example, there may be shadows of bees or dark leaves in the background which can be difficult to distinguish from a bee. The team solved this

problem with the use of a poster board as the background, which made it much easier to manually count the bees.

The heart of this prototype system is the DVXplorer (inivation.com) and its connected real-time computing system. Although most of the studies presented here were performed on laptop computers, the Raspberry Pi 4™ -based system (raspberrypi.com) was demonstrated to be adequate for real-time counting of bees if extensive event data collection and video data is not recorded and stored. Other elements of the system: batteries, waterproof enclosure, structural elements (wooden platform, cinder blocks, cables, etc.) raise the cost of the prototype system, but the DVXplorer dominates the overall cost. The overall cost of the system is approximately $5400.

The low-cost, USB web camera used for the validation studies and documentation and would not necessarily be part of future systems. Although the cost of this prototype system is large due to the DVXplorer (estimated to be $4600), recent, wider-scale adoption of event-based sensors is anticipated to dramatically lower the cost of such a system in the relatively near future.

6.2. Wingbeat Frequency

Wingbeat frequency is of great interest for both entomological research and technical applications. Identifying wingbeat frequency can provide insight into insect flight mechanisms [84] and reactions to environmental conditions [71]. For bees in particular, wingbeat frequency plays an important role in communication (for example, as a component of the waggle dance) [72] and the pollen-gathering via behavior of buzzing [85], [86].

In addition to their intrinsic biological interest, these relationships have been exploited for a variety of practical applications. Wingbeat frequency has been used to differentiate between species for the purposes of intelligent trapping [87], monitoring [88], and general differentiation between species [89]. Many of these applications rely not only on accurate measurements, but also on speed in order for quick intervention to take place. The Fourier Transform technique of wingbeat frequency offers very accurate measurements with a marker for confidence, but may need to be supplemented with a faster technique if immediate intervention is crucial.

To fully assess the accuracy and applicability of these methods, a ground truth measurement must be established. Future work would include testing the event-based frequency detection methods on a single bee while using another method to determine ground truth measurements. Ideally, these measurements could also be expanded to include multiple bees , particularly to include cases where occlusions occur and a tracker will briefly be viewing another bee. Additionally, ground truth data would be useful for determining whether the signal we are measuring truly represents the wingbeat frequency. Since our system only looks for general periodicity, it is possible that it is picking up on half or double the true frequency. Ground truth would allow for better assessment of these possibilities.

Currently, there are two primary techniques for determining wingbeat frequency: optical and audio. Audio systems record insects using microphones and analyze the resulting sound files using typical signal processing techniques [85], [86], [89]. Though cheap and (relatively) easy to implement, these systems have several limitations. They can only identify the frequency of one individual at time and typically restrict motion through the use of a tube or other small

opening. For audio systems in particular, the microphone must be placed close to the bee (in some cases less than 1 cm) to minimize environmental noise. Our system's ability to automatically identify the wingbeat frequency of multiple individuals in unrestricted flight therefore opens the door to many new avenues of research and applications.

Optical systems consist of a light source such as a laser or infrared LED and photoresistor which records the fluctuation in light caused when a bee crosses the beam [88], [90]. Less commonly, optical systems may employ high-speed video with manual analysis [72] or doppler radar [91]. Event-based vision joins this category of wingbeat detection offering low latency and high temporal precision for flying bees in a relatively open area.

From the entomological perspective, monitoring multiple bees can extend previously impractical measurements to a greater portion of the hive. Manual verification of wingbeat frequency by high-speed video is time-consuming, imposing an additional constraint on data gathering that often limits sample sizes. Our system alleviates that constraint by quickly and automatically extracting the frequency, allowing for greater sample sizes and a more thorough understanding of the colony. Since wingbeat frequencies are strongly characteristic of species and quickly computable, they could also be used as a lightweight verification for other systems. For example, future work may include enhancing our own counting algorithm by excluding and terminating trackers which do not contain the strongly periodic signal characteristic of a bee. Finally, one of the primary advantages of our system is that it allows bees to be monitored with relatively few restrictions to their flight, which more closely mimics the natural conditions researchers typically hope to replicate.

In terms of applications, wingbeat frequency shows potential as an indicator of colony health. Both parasites and exposure to pesticides have been shown to lower wingbeat frequencies in species such as bumblebees and mosquitoes, and machine learning systems can use wingbeat frequency to accurately identify infection/exposure [92], [93]. Our wingbeat detection could therefore be used by beekeepers to monitor the hive for threats in real time. Additionally, the ability to differentiate between species could be used to detect intruders and deploy a deterrent such as a laser before they enter the hive. Wingbeat frequency detection is particularly useful in this regard because different intruders require different responses; the periodic signal of a wasp might trigger a laser, while the more aperiodic signal of a mouse might trigger a physical barrier. Further experimentation as to whether the event-based methods proposed are accurate enough to differentiate various species of insect is necessary to assess how applicable the DVS is to these scenarios.

References

[1] N. Carreck and I. Williams, “The economic value of bees in the UK,” *Bee World*, vol. 79, no. 3, pp. 115–123, Jan. 1998, doi: 10.1080/0005772X.1998.11099393.

[2] S. A. M. Khalifa *et al.*, “Overview of Bee Pollination and Its Economic Value for Crop Production,” *Insects*, vol. 12, no. 8, p. 688, Jul. 2021, doi: 10.3390/insects12080688.

[3] Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services, IPBES, “Assessment Report on Pollinators, Pollination and Food Production,” Zenodo, Dec. 2016. doi: 10.5281/ZENODO.3402856.

[4] A. Wurz, I. Grass, and T. Tscharntke, “Hand pollination of global crops – A systematic review,” *Basic Appl. Ecol.*, vol. 56, pp. 299–321, Nov. 2021, doi: 10.1016/j.baae.2021.08.008.

[5] K.-L. J. Hung, J. M. Kingston, M. Albrecht, D. A. Holway, and J. R. Kohn, “The worldwide importance of honey bees as pollinators in natural habitats,” *Proc. R. Soc. B Biol. Sci.*, vol. 285, no. 1870, p. 20172140, Jan. 2018, doi: 10.1098/rspb.2017.2140.

[6] G. Papa *et al.*, “The Honey Bee Apis mellifera: An Insect at the Interface between Human and Ecosystem Health,” *Biology*, vol. 11, no. 2, p. 233, Feb. 2022, doi: 10.3390/biology11020233.

[7] N. Steinhauer, M. Wilson, D. Aurell, S. Bruckner, and G. Williams, “United States Honey Bee Colony Losses 2022-2023: Preliminary Results from the Bee Informed Partnership,” Jun. 2023. [Online]. Available: https://beeinformed.org/wp-content/uploads/2023/06/BIP-2022-23-Loss-Abstract.pdf

[8] P. Hristov, R. Shumkova, N. Palova, and B. Neov, “Honey bee colony losses: Why are honey bees disappearing?,” *Sociobiology*, vol. 68, no. 1, p. e5851, Feb. 2021, doi: 10.13102/sociobiology.v68i1.5851.

[9] E. Genersch, “Honey bee pathology: current threats to honey bees and beekeeping,” *Appl. Microbiol. Biotechnol.*, vol. 87, no. 1, pp. 87–97, Jun. 2010, doi: 10.1007/s00253-010-2573-8.

[10] D. Goulson, E. Nicholls, C. Botías, and E. L. Rotheray, “Bee declines driven by combined stress from parasites, pesticides, and lack of flowers,” *Science*, vol. 347, no. 6229, p. 1255957, Mar. 2015, doi: 10.1126/science.1255957.

[11] A. Zaman and A. Dorin, “A framework for better sensor-based beehive health monitoring,” *Comput. Electron. Agric.*, vol. 210, p. 107906, Jul. 2023, doi: 10.1016/j.compag.2023.107906.

[12] S. Cecchi, S. Spinsante, A. Terenzi, and S. Orcioni, “A Smart Sensor-Based Measurement System for Advanced Bee Hive Monitoring,” *Sensors*, vol. 20, no. 9, p. 2726, May 2020, doi: 10.3390/s20092726.

[13] B. Bąk, J. Szkoła, J. Wilk, P. Artiemjew, and J. Wilde, “In-Field Detection of American Foulbrood (AFB) by Electric Nose Using Classical Classification Techniques and Sequential Neural Networks,” *Sensors*, vol. 22, no. 3, p. 1148, Feb. 2022, doi: 10.3390/s22031148.

[14] S. Ferrari, M. Silva, M. Guarino, and D. Berckmans, “Monitoring of swarming sounds in bee hives for early detection of the swarming period,” *Comput. Electron. Agric.*, vol. 64, no. 1, pp. 72–77, Nov. 2008, doi: 10.1016/j.compag.2008.05.010.

[15] M. Z. Sharif, F. Wario, N. Di, R. Xue, and F. Liu, “Soundscape Indices: New Features for Classifying Beehive Audio Samples,” *Sociobiology*, vol. 67, no. 4, pp. 566–571, Dec. 2020, doi: 10.13102/sociobiology.v67i4.5860.

[16] P. Lichtsteiner, C. Posch, and T. Delbruck, “A 128×128 120 dB 15 μs latency asynchronous temporal contrast vision sensor,” *IEEE J Solid-State Circuits*, vol. 43, no. 2, pp. 566–576, 2008.

[17] J. Ollerton, R. Winfree, and S. Tarrant, “How many flowering plants are pollinated by animals?,” *Oikos*, vol. 120, no. 3, pp. 321–326, Mar. 2011, doi: 10.1111/j.1600-0706.2010.18644.x.
[18] R. Rader *et al.*, “Non-bee insects are important contributors to global crop pollination,” *Proc. Natl. Acad. Sci.*, vol. 113, no. 1, pp. 146–151, Jan. 2016, doi: 10.1073/pnas.1517092112.
[19] M. A. Aizen, L. A. Garibaldi, S. A. Cunningham, and A. M. Klein, “How much does agriculture depend on pollinators? Lessons from long-term trends in crop production,” *Ann. Bot.*, vol. 103, no. 9, pp. 1579–1588, Jun. 2009, doi: 10.1093/aob/mcp076.
[20] “Honey Bee Colonies,” United States Department of Agriculture, Aug. 2021. Accessed: Oct. 05, 2021. [Online]. Available: https://usda.library.cornell.edu/concern/publications/rn301137d?locale=en
[21] C. D. Michener, *The Bees of the World*. Baltimore, Md: Johns Hopkins University Press, 2000.
[22] G. R. Williams *et al.*, “Colony Collapse Disorder in Context,” *BioEssays*, vol. 32, no. 10, pp. 845–846, Oct. 2010, doi: 10.1002/bies.201000075.
[23] E. E. Southwick and L. Southwick, “Estimating the Economic Value of Honey Bees (Hymenoptera: Apidae) as Agricultural Pollinators in the United States,” *J. Econ. Entomol.*, vol. 85, no. 3, pp. 621–633, Jun. 1992, doi: 10.1093/jee/85.3.621.
[24] “Insect Biology : A Primer.” Accessed: Feb. 05, 2022. [Online]. Available: https://biocontrol.entomology.cornell.edu/bio.php
[25] D. L. Altshuler, W. B. Dickson, J. T. Vance, S. P. Roberts, and M. H. Dickinson, “Short-amplitude high-frequency wing strokes determine the aerodynamics of honeybee flight,” *Proc. Natl. Acad. Sci.*, vol. 102, no. 50, pp. 18213–18218, Dec. 2005, doi: 10.1073/pnas.0506590102.
[26] S. Łopuch and A. Tofilski, “Direct Visual Observation of Wing Movements during the Honey Bee Waggle Dance,” *J. Insect Behav.*, vol. 30, no. 2, pp. 199–210, Mar. 2017, doi: 10.1007/s10905-017-9610-8.
[27] H. Iwamoto and N. Yagi, “The Molecular Trigger for High-Speed Wing Beats in a Bee,” *Science*, vol. 341, no. 6151, pp. 1243–1246, Sep. 2013, doi: 10.1126/science.1237266.
[28] N. P. Goyal and A. S. Atwal, “Wing Beat Frequencies of Apis Cerana Indica and Apis Mellifera,” *J. Apic. Res.*, vol. 16, no. 1, pp. 47–48, Jan. 1977, doi: 10.1080/00218839.1977.11099859.
[29] J. T. Vance, I. Faruque, and J. S. Humbert, “Kinematic strategies for mitigating gust perturbations in insects,” *Bioinspir. Biomim.*, vol. 8, no. 1, p. 016004, Jan. 2013, doi: 10.1088/1748-3182/8/1/016004.
[30] “Biology and physics of locust flight. I. Basic principles in insect flight. A critical review | Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences.” Accessed: Feb. 03, 2022. [Online]. Available: https://royalsocietypublishing.org/doi/abs/10.1098/rstb.1956.0007?casa_token=d9LX9r-qcfcAAAAA:daqlRpOnSfJbqmndBBCxvAIBbSz9dKV_ShAlGr8aJIP-LXDhQjePZjE7aMin8nhVohlxkbKVwg-Odg
[31] P. C. Witherell and H. H. Laidlaw, “Behavior of the honey bee (Apis mellifera L.) mutant, diminutive-wing,” *Hilgardia*, vol. 45, no. 1, pp. 1–29, Apr. 1977, doi: 10.3733/hilg.v45n01p001.
[32] J. T. Vance, D. L. Altshuler, W. B. Dickson, M. H. Dickinson, and S. P. Roberts, “Hovering Flight in the Honeybee Apis mellifera: Kinematic Mechanisms for Varying Aerodynamic Forces,” *Physiol. Biochem. Zool.*, vol. 87, no. 6, pp. 870–881, Nov. 2014, doi: 10.1086/678955.
[33] D. vanEngelsdorp *et al.*, “Colony Collapse Disorder: A Descriptive Study,” *PLoS ONE*, Aug. 2009, doi: https://doi.org/10.1371/journal.pone.0006481.
[34] P. Rosenkranz, P. Aumeier, and B. Ziegelmann, “Biology and control of Varroa destructor,” *J.*

Invertebr. Pathol., vol. 103, pp. S96–S119, Jan. 2010, doi: 10.1016/j.jip.2009.07.016.
[35] T. D. Seeley, “The Honey Bee Colony as a Superorganism,” *Am. Sci.*, vol. 77, no. 6, pp. 546–553, 1989.
[36] A. N. Mortensen, B. Smith, and J. D. Ellis, “The Social Organization of Honeybees,” *Univ. Fla. Inst. Food Agric. Sci.*, Mar. 2019, [Online]. Available: https://edis.ifas.ufl.edu/publication/IN1102
[37] B. R. Johnson, “Division of labor in honeybees: form, function, and proximate mechanisms,” *Behav. Ecol. Sociobiol.*, vol. 64, no. 3, pp. 305–316, Jan. 2010, doi: 10.1007/s00265-009-0874-7.
[38] B. R. Johnson, “Within-nest temporal polyethism in the honey bee,” *Behav. Ecol. Sociobiol.*, vol. 62, no. 5, pp. 777–784, Mar. 2008, doi: 10.1007/s00265-007-0503-2.
[39] T. D. Seeley, “Adaptive significance of the age polyethism schedule in honeybee colonies,” *Behav. Ecol. Sociobiol.*, vol. 11, no. 4, pp. 287–293, Dec. 1982, doi: 10.1007/BF00299306.
[40] S. C. Pratt, “Decentralized control of drone comb construction in honey bee colonies,” *Behav. Ecol. Sociobiol.*, vol. 42, no. 3, pp. 193–205, Mar. 1998, doi: 10.1007/s002650050431.
[41] B. R. Johnson, “Global information sampling in the honey bee,” *Naturwissenschaften*, vol. 95, no. 6, pp. 523–530, Jun. 2008, doi: 10.1007/s00114-008-0354-3.
[42] M. Nouvian, J. Reinhard, and M. Giurfa, “The defensive response of the honeybee Apis mellifera,” *J. Exp. Biol.*, vol. 219, no. 22, pp. 3505–3517, Nov. 2016, doi: 10.1242/jeb.143016.
[43] M. D. Breed, G. E. Robinson, and R. E. Page, “Division of labor during honey bee colony defense,” *Behav. Ecol. Sociobiol.*, vol. 27, no. 6, pp. 395–401, Dec. 1990, doi: 10.1007/BF00164065.
[44] R. C. Garcia *et al.*, “Honey and propolis production, hygiene and defense behaviors of two generations of Africanized honey bees,” *Sci. Agric.*, vol. 70, no. 2, pp. 74–81, Apr. 2013, doi: 10.1590/S0103-90162013000200003.
[45] M. L. Winston, “The Biology and Management of Africanized Honey Bees,” *Annu. Rev. Entomol.*, vol. 37, no. 1, pp. 173–193, Jan. 1992, doi: 10.1146/annurev.en.37.010192.001133.
[46] E. Guzmán-Novoa and R. E. Page, “Backcrossing Africanized Honey Bee Queens to European Drones Reduces Colony Defensive Behavior,” *Ann. Entomol. Soc. Am.*, vol. 86, no. 3, pp. 352–355, May 1993, doi: 10.1093/aesa/86.3.352.
[47] D. De Jong, “Africanized honey bees in Brazil, forty years of adaptation and success,” *Bee World*, vol. 77, no. 2, pp. 67–70, Jan. 1996, doi: 10.1080/0005772X.1996.11099289.
[48] S. S. Schneider and G. DeGrandi-Hoffman, “Queen replacement in African and European honey bee colonies with and without afterswarms,” *Insectes Sociaux*, vol. 55, no. 1, pp. 79–85, Feb. 2008, doi: 10.1007/s00040-007-0973-2.
[49] A. Terenzi, S. Cecchi, S. Orcioni, and F. Piazza, “Features Extraction Applied to the Analysis of the Sounds Emitted by Honey Bees in a Beehive,” in *2019 11th International Symposium on Image and Signal Processing and Analysis (ISPA)*, Dubrovnik, Croatia: IEEE, Sep. 2019, pp. 03–08. doi: 10.1109/ISPA.2019.8868934.
[50] T. M. Francoy *et al.*, “Identification of Africanized honey bees through wing morphometrics: two fast and efficient procedures,” *Apidologie*, vol. 39, no. 5, pp. 488–494, Sep. 2008, doi: 10.1051/apido:2008028.
[51] A. Zacepine and T. Karasha, “Application of Temperature Measurements For Bee Colony Monitoring: A Review,” presented at the ENGINEERING FOR RURAL DEVELOPMENT, May 2013.
[52] T. Cejrowski, J. Szymański, H. Mora, and D. Gil, “Detection of the Bee Queen Presence Using Sound Analysis,” in *Intelligent Information and Database Systems*, vol. 10752, N. T. Nguyen, D. H. Hoang, T.-P. Hong, H. Pham, and B. Trawiński, Eds., in Lecture Notes in

Computer Science, vol. 10752. , Cham: Springer International Publishing, 2018, pp. 297–306. doi: 10.1007/978-3-319-75420-8_28.
[53] P. De Souza *et al.*, “Low-Cost Electronic Tagging System for Bee Monitoring,” *Sensors*, vol. 18, no. 7, p. 2124, Jul. 2018, doi: 10.3390/s18072124.
[54] J. M. Mola and N. M. Williams, “A review of methods for the study of bumble bee movement,” *Apidologie*, vol. 50, no. 4, pp. 497–514, Aug. 2019, doi: 10.1007/s13592-019-00662-3.
[55] X.-J. He, L.-Q. Tian, X.-B. Wu, and Z.-J. Zeng, “RFID monitoring indicates honeybees work harder before a rainy day: Weather forecast in honeybees,” *Insect Sci.*, vol. 23, no. 1, pp. 157–159, Feb. 2016, doi: 10.1111/1744-7917.12298.
[56] J. F. Harrison, O. R. Taylor Jr., and H. G. Hall, “The Flight Physiology of Reproductives of Africanized, European, and Hybrid Honeybees (*Apis mellifera*),” *Physiol. Biochem. Zool.*, vol. 78, no. 2, pp. 153–162, Mar. 2005, doi: 10.1086/427057.
[57] H. T. Kerr, M. E. Bukanau., and K. H. Valentine, “Method and device for identifying different species of honeybees,” 4,876,721, Oct. 24, 1989 Accessed: Feb. 04, 2022. [Online]. Available: https://patft.uspto.gov/netacgi/nph-Parser?Sect1=PTO1&Sect2=HITOFF&d=PALL&p=1&u=%2Fnetahtml%2FPTO%2Fsrchnum.htm&r=1&f=G&l=50&s1=4876721.PN.&OS=PN/4876721&RS=PN/4876721
[58] I. F. Rodriguez, R. Megret, E. Acuna, J. L. Agosto-Rivera, and T. Giray, “Recognition of Pollen-Bearing Bees from Video Using Convolutional Neural Network,” in *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, Mar. 2018, pp. 314–322. doi: 10.1109/WACV.2018.00041.
[59] M. Ratnayake, A. Dyer, and A. Dorin, “Tracking individual honeybees among wildflower clusters with computer vision-facilitated pollinator monitoring,” *PLoS ONE*, Feb. 2021, [Online]. Available: https://doi.org/10.1371/journal.pone.0239504
[60] B. Magnier, G. Ekszterowicz, J. Laurent, M. Rival, and F. Pfister, “Bee Hive Traffic Monitoring by Tracking Bee Flight Paths:,” in *Proceedings of the 13th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, Funchal, Madeira, Portugal: SCITEPRESS - Science and Technology Publications, 2018, pp. 563–571. doi: 10.5220/0006628205630571.
[61] C. D. Schuman *et al.*, “A Survey of Neuromorphic Computing and Neural Networks in Hardware,” *ArXiv170506963 Cs*, May 2017, Accessed: Oct. 05, 2021. [Online]. Available: http://arxiv.org/abs/1705.06963
[62] G. Gallego *et al.*, “Event-based Vision: A Survey,” *IEEE Trans. Pattern Anal. Mach. Intell.*, pp. 1–1, 2020, doi: 10.1109/TPAMI.2020.3008413.
[63] T. Delbruck and P. Lichtsteiner, “Fast sensory motor control based on event-based hybrid neuromorphic-procedural system,” in *2007 IEEE International Symposium on Circuits and Systems*, May 2007, pp. 845–848. doi: 10.1109/ISCAS.2007.378038.
[64] H. Rebecq, R. Ranftl, V. Koltun, and D. Scaramuzza, “High speed and high dynamic range video with an event camera,” *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019.
[65] A. Amir *et al.*, “A Low Power, Fully Event-Based Gesture Recognition System,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI: IEEE, Jul. 2017, pp. 7388–7397. doi: 10.1109/CVPR.2017.781.
[66] A. Lakshmi, A. Chakraborty, and C. S. Thakur, “Neuromorphic vision: From sensors to event-based algorithms,” *WIREs Data Min. Knowl. Discov.*, vol. 9, no. 4, p. e1310, 2019, doi: 10.1002/widm.1310.
[67] C. Aerne, “Understanding the Performance of Neuromorphic Event-based Vision Sensors.” iniVation, May 2020. Accessed: Oct. 04, 2021. [Online]. Available: https://inivation.com/wp-content/uploads/2020/05/White-Paper-May-2020.pdf
[68] G. Cohen *et al.*, “Event-based Sensing for Space Situational Awareness,” *Proc. Adv. Maui*

Opt. Space Surveill. Technol. Conf., 2017.
[69] A. Mitrokhin, C. Fermüller, C. Parameshwara, and Y. Aloimonos, “Event-Based Moving Object Detection and Tracking,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2018, pp. 1–9. doi: 10.1109/IROS.2018.8593805.
[70] J. Cardinale, “Tracking objects and wing beat analysis methods of a fruit fly with the event-based silicon retina,” University of Zurich, 2006. [Online]. Available: https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=e8a502b485d04fec223bc38efc85bde82650aae8
[71] A. R. S. Parmezan, V. M. A. Souza, I. Žliobaitė, and G. E. A. P. A. Batista, “Changes in the wing-beat frequency of bees and wasps depending on environmental conditions: a study with optical sensors,” *Apidologie*, vol. 52, no. 4, pp. 731–748, Aug. 2021, doi: 10.1007/s13592-021-00860-y.
[72] S. Łopuch and A. Tofilski, “Use of high-speed video recording to detect wing beating produced by honey bees,” *Insectes Sociaux*, vol. 66, no. 2, pp. 235–244, May 2019, doi: 10.1007/s00040-018-00678-5.
[73] M. C. Kahn and W. Jr. Offenhauser, “The Identification of certain West African Mosquitoes by Sound.,” *Am. J. Trop. Medice*, vol. 29, no. 5, pp. 827–836, 1949.
[74] D. R. Raman, R. R. Gerhardt, and J. B. Wilkerson, “Detecting Insect Flight Sounds in the Field: Implications for Acoustical Counting of Mosquitos,” *Transcations Am. Soc. Agric. Biol. Eng.*, vol. 50, no. 4, pp. 1481–1485, 2007, doi: 10.13031/2013.23606.
[75] Y. Chen, A. Why, G. E. A. P. A. Batista, A. Mafra-Neto, and E. Keogh, “Flying Insect Classification with Inexpensive Sensors,” *J. Insect Behav.*, vol. 27, pp. 657–677, 2014, doi: 10.1007/s10905-014-9454-4.
[76] J. van Roy, J. De Baerdemaeker, W. Saeys, and B. De Ketelaere, “Optical identification of bumblebee species: Effect of morphology on wingbeat frequency Author links open overlay panel,” *Comput. Electron. Agric.*, vol. 109, pp. 94–100, Nov. 2014, doi: 10.1016/j.compag.2014.09.014.
[77] J. Santoyo, W. Azarcoya, M. Valencia, A. Torres, and J. Salas, “Frequency analysis of a bumblebee (Bombus impatiens) wingbeat,” *Pattern Anal. Appl.*, vol. 19, pp. 487–493, Jul. 2015, doi: 10.1007/s10044-015-0501-3.
[78] K. M. Fristrup, J. A. Shaw, and M. J. Tauc, “Development of a wing-beat-modulation scanning lidar system for insect studies,” in *Lidar Remote Sensing for Environmental Monitoring 2017*, U. N. Singh, Ed., San Diego, United States: SPIE, Aug. 2017, p. 15. doi: 10.1117/12.2274656.
[79] T. Samson, S. Sepehrirahnama, B. Halkon, J. Lai, and S. Oberst, “LASER INTENSITY-INDUCED DAMAGE EFFECTS ON DYNAMIC CHARACTERISATION OF WINGS OF THE EUROPEAN HONEYBEE (APIS MELLIFERA),” *Proc. Int. Congr. Sound Vib.*, vol. 28, Aug. 2022.
[80] S. Hoseini and B. Linares-Barranco, “Real-Time Temporal Frequency Detection in FPGA Using Event-Based Vision Senso,” *Proc. IEEE Int. Conf. Intell. Comput. Commun. Process.*, pp. 271–278, 2018, doi: 10.1109/ICCP.2018.8516629.
[81] X. Lagorce, G. M. Orchard, F. Galluppi, B. Shi, and R. Benosman, “HOTS: A Hierarchy of Event-Based Time-Surfaces for Pattern Recognition,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 7, pp. 1346–1359, Jul. 2017, doi: 10.1109/TPAMI.2016.2574707.
[82] E. O. Neftci, H. Mostafa, and F. Zenke, “Surrogate Gradient Learning in Spiking Neural Networks,” 2019, doi: 10.48550/ARXIV.1901.09948.
[83] “Caissa Splatter Ball Gun Gel Ball Blasters AEG AKM-47.” [Online]. Available: https://caissatoy.com/products/caissa-gel-blaster-pistol-akm-47?variant=40856309399757
[84] N. S. Ha, Q. T. Truong, N. S. Goo, and H. C. Park, “Relationship between wingbeat frequency and resonant frequency of the wing in insects,” *Bioinspir. Biomim.*, vol. 8, no. 4, p. 046008, Oct. 2013, doi: 10.1088/1748-3182/8/4/046008.

[85] A. Burkart, K. Lunau, and C. Schlindwein, “Comparative bioacoustical studies on flight and buzzing of neotropical bees,” *J. Pollinat. Ecol.*, vol. 6, Jan. 2012, doi: 10.26786/1920-7603(2011)17.
[86] P. Volponi, L. Dapporto, and M. Skowron Volponi, “The BUZZOMETER system: In situ audio recordings of pollinators in flight,” *Methods Ecol. Evol.*, vol. 14, no. 12, pp. 2985–2993, Dec. 2023, doi: 10.1111/2041-210X.14224.
[87] D. F. Silva, V. M. A. D. Souza, G. E. A. P. A. Batista, E. Keogh, and D. P. W. Ellis, “Applying Machine Learning and Audio Analysis Techniques to Insect Recognition in Intelligent Traps,” in *2013 12th International Conference on Machine Learning and Applications*, Miami, FL, USA: IEEE, Dec. 2013, pp. 99–104. doi: 10.1109/ICMLA.2013.24.
[88] K. Rydhmer *et al.*, “Automating insect monitoring using unsupervised near-infrared sensors,” *Sci. Rep.*, vol. 12, no. 1, p. 2603, Feb. 2022, doi: 10.1038/s41598-022-06439-6.
[89] S. Kawakita and K. Ichikawa, “Automated classification of bees and hornet using acoustic analysis of their flight sounds,” *Apidologie*, vol. 50, no. 1, pp. 71–79, Feb. 2019, doi: 10.1007/s13592-018-0619-6.
[90] G. E. A. P. A. Batista, Y. Hao, E. Keogh, and A. Mafra-Neto, “Towards Automatic Classification on Flying Insects Using Inexpensive Sensors,” in *2011 10th International Conference on Machine Learning and Applications and Workshops*, Honolulu, HI, USA: IEEE, Dec. 2011, pp. 364–369. doi: 10.1109/ICMLA.2011.145.
[91] R. Wang, C. Hu, X. Fu, T. Long, and T. Zeng, “Micro-Doppler measurement of insect wing-beat frequencies with W-band coherent radar,” *Sci. Rep.*, vol. 7, no. 1, p. 1396, May 2017, doi: 10.1038/s41598-017-01616-4.
[92] V. Chatzaki, M. Montoro, R. El-Rashid, A. B. Jensen, and A. Lecocq, “A New Approach for Detecting Sublethal Effects of Neonicotinoids on Bumblebees Using Optical Sensor Technology,” *Insects*, vol. 14, no. 8, p. 713, Aug. 2023, doi: 10.3390/insects14080713.
[93] D. Park *et al.*, “An investigation of Dirofilaria immitis infection and its effects on mosquito wingbeat frequencies,” *Vet. Parasitol.*, vol. 283, p. 109112, Jul. 2020, doi: 10.1016/j.vetpar.2020.109112.