

# Low-Cost Aerial Imaging for Small Holder Farmers

Vasuki Narasimha Swamy\*  
UC Berkeley

Akshit Kumar\*  
IIT Madras

Rohit Patil\*  
PESIT

Aditya Jain\*  
IIT Delhi

Zerina Kapetanovic\*  
University of Washington

Rahul Sharma\*  
CMU

Deepak Vasisht\*  
MIT

Manohar Swaminathan  
Microsoft

Ranveer Chandra  
Microsoft

Anirudh Badam  
Microsoft

Gireeja Ranade\*  
University of California Berkeley

Sudipta Sinha  
Microsoft

Akshay Uttama Nambi S N  
Microsoft

## CCS CONCEPTS

• **Human-centered computing** → *Ubiquitous and mobile computing systems and tools*;

## KEYWORDS

Aerial Imagery, Mobile Systems for Sustainability, Innovative Mobile Sensing

### ACM Reference Format:

Vasuki Narasimha Swamy\*, Akshit Kumar\*, Rohit Patil\*, Aditya Jain\*, Zerina Kapetanovic\*, Rahul Sharma\*, Deepak Vasisht\*, Manohar Swaminathan, Ranveer Chandra, Anirudh Badam, Gireeja Ranade\*, Sudipta Sinha, and Akshay Uttama Nambi S N. 2019. Low-Cost Aerial Imaging for Small Holder Farmers. In *ACM SIGCAS Conference on Computing and Sustainable Societies (COMPASS) (COMPASS '19)*, July 3–5, 2019, Accra, Ghana. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3314344.3332485>

**Abstract** – Recent work in networked systems has shown that using aerial imagery for farm monitoring can enable precision agriculture by lowering the cost and reducing the overhead of large scale sensor deployment. However, acquiring aerial imagery requires a drone, which has high capital and operational costs, often beyond the reach of farmers in the developing world. In this paper, we present TYE (Tethered eYE), an inexpensive platform for aerial imagery. It consists of a tethered helium balloon with a custom mount that can hold a smartphone (or a camera) with a battery pack. The balloon can be carried using a tether by a person or a vehicle. We incorporate various techniques to increase the operational time of the system, and to provide actionable insights even with unstable imagery. We develop path-planning algorithms

\*The work was done when the authors were at Microsoft.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

COMPASS '19, July 3–5, 2019, Accra, Ghana

© 2019 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.


ACM ISBN 978-1-4503-6714-1/19/07...\$15.00

<https://doi.org/10.1145/3314344.3332485>

and use that to develop an interactive mobile phone application that provides the user instant feedback to guide users to efficiently traverse large areas of land. We use computer vision algorithms to stitch orthomosaics by effectively countering wind-induced motion of the camera. We have used TYE for aerial imaging of agricultural land for over a year, and envision it as a low-cost aerial imaging platform for similar applications.

## 1 INTRODUCTION

Unmanned Aerial Vehicles (UAVs) and drones have gained tremendous traction as a tool for enabling digital agriculture [13, 25, 31, 32]. In contrast to satellites, which typically captures only about 10 clean images from seeding to harvest [3], a drone can be flown on-demand, and can get very high resolution imagery. They can also be mounted with multi-spectral or hyper-spectral cameras to obtain detailed imagery of the farms [14]. A recent research in networked systems [32] takes the aerial imagery a step further, by combining aerial imagery with ground sensor data, to create detailed precision maps for a farm, such as for soil temperature, pH, and soil moisture, which are then used to enable low-cost precision agriculture. While such systems work well in the developed countries, the startup cost of buying a commercial quadrotor (around \$ 1000) is prohibitive for farmers in developing countries. Additionally, the lack of technical knowledge among the farmers in developing countries means that they often need specialized *technicians* to operate them. This implies



**Figure 1: Mobile TYE: A user images a farm using Mobile TYE**

that for large parts of the world, digital agriculture techniques remain out of bounds.

In this paper, we address this problem by designing an alternative aerial imagery platform, that has low capital investment, low operational costs and high resolution. We present our system, TYE (for Tethered eYE). TYE uses lighter-than-air gas (such as helium) filled balloons to provide lift to an existing imaging infrastructure (such as a smartphone). The balloon is held using a tether at heights from 50 ft to 150 ft depending on the required resolution. The design of TYE has been motivated by the strong penetration of smartphones in the developing world [23], combined with enhanced imaging capabilities of such smartphones.

TYE operates in two modes. First, it has a mobile mode, wherein it operates like a UAV and the tether is held by a person who walks around the field. In this mode, the imagery can be used to generate precision maps. As our aerial imagery platform design does not require power to stay afloat, it enables an additional *static mode*, where the balloon is tethered to a stationary point. In this mode, the balloon can provide the farmer with long term surveillance imagery of the farm. While the design of TYE is simple, the challenge lies in delivering high-fidelity, interpretable aerial imagery. To enable long-duration, low cost, large-area aerial imagery, we must address three fundamental aspects:

**(i) Wind Induced Variability:** Unlike UAVs, low winds cause significant changes in the position (lateral motion and rotation) of the camera that is capturing aerial images using balloons. These unpredictable changes in camera viewpoint due to wind makes imagery difficult to view and analyze. Our solution to this problem is two-fold: (a) we design a custom mount for the payload that can reduce wind-induced variability, and (b) we leverage gyroscopes (present in most smartphones) to eliminate frames with extreme motion while ensuring sufficient coverage. We then use techniques from computer vision to correct for the remaining camera motion in software; thus producing consistent views for the user as if the camera was stable, despite arbitrary camera motions.

**(ii) Spatio-temporal Coverage:** Spatial mapping of a farm requires TYE to be maneuverable like drones. Thus, mobile-TYE requires a person to move the balloon. In our experiments, we observe that users fail to ensure complete coverage of an area if they rely purely on their intuition. This is due to the mismatch between a person's 2D trajectory and the balloon's 3D trajectory. Furthermore, some areas might be inaccessible to the user due to obstructions. For example, in farms, plots are interspersed with trees, bunds, ponds, and electrical poles etc. To address these concerns, we build a smartphone application that runs on a smartphone carried by the user (this is separate from the smartphone mounted on the balloon which is used to image the areas of interest). The application interactively displays the area that has been covered so far and the path to be followed. The application incorporates a novel path-planning algorithm that ensures area coverage, with minimum human motion, in spite of wind-induced balloon motion.

**(iii) Power Management:** For long term imagery (static-TYE), capturing frequent images and sending them to any ground device, such as a PC over LTE, drains power very quickly and limits the uptime to a few hours [26]. One obvious way to extend the battery life is to identify frames where certain pixels have different values

than past frames, and transmit only those frames which might correspond to some form of event or motion. However, this doesn't work for TYE, since wind-induced motion causes most frames to be different from past frames. An alternative is to use complex object detection and motion tracking vision techniques to identify frames in which moving objects were observed. However, these techniques require high compute power, and are either not feasible on a smartphone or take too much time to offset the power savings. In contrast to these approaches, we leverage an important insight to design a low complexity frame selection technique. As the balloon drifts, all pixels in the frame captured by the camera drift in a direction opposite to the balloon. However, if something changes in the environment, the pixels corresponding to the change must have a different drift than the majority of the pixels corresponding to static objects. By identifying these outliers, i.e. pixels that have a drift different from the average drift of the frame, we can identify frames that have an interesting event and only transmit them, thereby reducing the power required for data transmission.

We have implemented TYE as a software-hardware system, using helium balloons and evaluated them in farms across US and India. The main contributions of this paper are:

- We design two modes of TYE: mobile and static which are useful for on-demand imaging of large areas and long-duration imaging for a single area respectively.
- We design new software-hardware mechanisms to reduce wind-induced variability in captured imagery and present a consistent view to the farmer in spite of frequent camera motions.
- We build a novel path-planning algorithm that captures the area being imaged and intelligently mitigates the effects of wind – effectively reducing the time needed to image a given area.
- We design novel low-complexity anomaly-detection algorithms and custom hardware to enhance the uptime of the system.

While a detailed evaluations of TYE are in sections 6 and 8, we list the important results here:

- The hardware modifications reduce the leakage through the surface of the balloon by  $\approx 90\%$ .
- The area imaged in mobile mode using TYE's app and path planning algorithm is 88.2% of the area of interest as opposed to 78.5%, when the user is not using the app. The users walking have to walk 25.2% more steps without the app to guide them.
- The intelligent-frame selection algorithm reduces the average frame transmissions by 94% and the overall battery consumption by 67-88% depending on the image resolution.
- The vision pipeline successfully reduces rolling shutter effects and constructs stable imagery. This imagery can then be combined with sparse sensor measurements to create accurate precision maps such as in [32].

## 2 MOTIVATION & RELATED WORK

The world's food production needs to increase by up to 70% by 2050 to feed the growing population of the world [16] even though the amount of arable land is limited, and the water levels are receding. This challenge is even more severe if we consider nourishing the world, and not just feeding it. One promising approach to address the above challenge is precision agriculture, which refers to the ability to do site specific applications of various farm inputs. Instead

	Height	Duration	Frequency	Area	Resolution	Capital cost
TYE	flexible 20ft-200ft	flexible several days	on-demand	flexible	high, 1 cm per pixel	< \$100
UAV	flexible 20ft-400ft	very limited	on-demand	limited (by battery)	high, 1 cm per pixel	> \$1000
Satellite	fixed	years	fixed	very large	low, > 30 cm per pixel	very high
Towers	fixed	years	always-on	fixed and small	varies with height	high

**Table 1: Comparison of aerial imaging solutions: Unlike existing solutions, TYE delivers long-term, high-resolution aerial imagery at low capital cost.**

of uniformly applying water throughout the farm, a farmer can apply water only where it is needed. Similarly, the farmer can apply pesticide, fertilizer, etc. only where it is needed on the farm. Precision agriculture as a technique has been shown to improve yield, reduce cost, and is also better for the environment.

Precision agriculture requires the creation of a Precision Map, such as a soil moisture map of the farm 6 inches below the ground, or a pest-infestation map. However, creating accurate precision maps is resource intensive. It requires several densely deployed sensors in the field sending a lot of data to the cloud. This would cost thousands of dollars in setup fees and hundreds of dollars in annual subscriptions, per sensor [32]. These make precision agriculture challenging as the labor and operating costs are very high. This has started to change with the advent of personal UAVs. UAVs can be used to create precision maps of the field either in isolation [14, 25, 31] or in combination with ground sensors [32]. This significantly reduces the cost of deploying precision agriculture techniques which has led to increasing adoption of UAVs by farmers. A poll by Farmer’s Journal Pulse [2] showed that around two-third of the farmers surveyed either owned a drone or planned to get one within the year. Among the ones who had drones, 63% operated the drones on their own. However, the high cost of UAVs combined with the limited education levels in the developing world has kept farmers in the developing world behind the technology curve.

Our goal with TYE is to build a low-cost aerial imaging system, that farmers can use to easily obtain an aerial image of the farm at low cost. This image can then be used for scouting, interesting event detection, or for building precision maps. One such useful Precision Map could be for soil moisture, which the farmer can use to decide how much water to apply when and where in the farm, thereby saving irrigation costs, and also being more productive by ensuring that no plant is stressed.

**Prior Work:** Prior studies [5, 6, 17, 22, 27] have used balloon-photography and kite-photography for archaeological surveys, geo-scientific mapping, vegetation analysis, etc. However, these systems were designed for clicking one-time aerial images and do not support long-term large area imagery. In contrast, we propose mechanisms to enable long-term automated visual imagery collection using tethered balloons. Other popular techniques used for aerial imaging are (a) satellites, (b) cameras attached to UAVs, and (c) camera towers. We compare several usability and performance aspects of these systems in Table 1. As noted in the table, satellite imagery has poor image resolution (46 cm per pixel at best), is infrequent (order of days), and requires high capital investments. This makes the acquisition of continuous, long-term, high resolution imagery for applications, such as flood monitoring, surveillance, etc. very expensive for small to medium farmers.

Similarly, UAVs present multiple challenges. First, UAVs consume a large amount of power to stay afloat, resulting in very short battery life (tens of minutes for most commercial UAVs) [9, 32]

which makes UAVs infeasible for applications that require long-term, continuous monitoring such as flood monitoring. Second, UAVs require high capital investment. Commercial UAVs that can last for 30 minutes cost over \$1000 [10]. Third, UAVs involve high operational complexity such as the need for certified operators and stable power infrastructure to recharge batteries which drive up the cost of deployment. These factors make UAV-based aerial imaging infeasible for several developing countries. Thus, in contrast to existing approaches, TYE presents a low-operational complexity, low-cost alternative for long-term and large area aerial imagery.

### 3 TYE DESIGN: OBJECTIVES & OVERVIEW

TYE aims to achieve the following:

- **Aerial Mapping of Large Areas:** To enable a human user to map large areas using the tethered balloon and obtain usable images, we have to address the following challenges: (i) the platform needs to be simple so that a semi-skilled user can assemble and do the imaging, (ii) the user should be able to cover the intended area efficiently despite wind-induced balloon movements as well as obstacles and variability in the terrain, and (iii) the system should generate the orthomosaic of the terrain despite an unstable imaging platform. These challenges are addressed by our design in Section 4.
- **Long-term, Continuous Imaging:** The goal of the static-mode is to provide long-term aerial imaging over a fixed area for a few days. Enabling the system to run for days entails two primary components. First, how does one make the balloon provide enough lift (in spite of leakage) to hold the system aloft for long durations? Second, how does one manage the limited battery that can be carried as payload to maximize the uptime of the imaging device (camera/smartphone)? In section 5, we describe how we solve these challenges using a combination of techniques from computer vision and mobile systems.
- **Output Interpretability:** TYE’s output (from mobile and static modes) should be accessible and interpretable by the farmers. To this end, we design an end-to-end communication pipeline (section 5.1), that allows users to access the captured imagery locally on the edge computer and on the cloud in near real-time. Our design ensures that the system is available locally even when outages happen on the edge-cloud link. Furthermore, wind-induced motions make the captured imagery hard to interpret. We develop vision algorithms to (a) raise alerts when something new is observed (Section 5.3), and (b) stabilize the imagery to a single viewpoint in software (Section 4.3).

### 4 MOBILE-TYE

In precision agriculture, aerial imagery is needed at different periods in a plant’s life so that in combination with the ground sensor data, valuable decision support can be provided to farmers [32]. The average farm size is very small (in the order of a couple of acres) and these farms are clustered in irregular terrains.

Given the low individual incomes of farmers and the lack of access to capital, such precision agriculture efforts can either be implemented by farmer cooperatives, or by agriculture companies, or by local government agencies and not by individual farmers. Hence, mobile TYE needs to be designed as an efficient, low-cost, on demand, aerial image capturing system that can be operated by a semi-skilled person with minimal training and infrastructure requirements. Furthermore, the capital cost and operational cost of the system should be as low as possible. We envision this to be primarily used as an extremely low-cost alternative for UAVs to construct orthomosaics of areas. We first describe the design of the platform in section 4.1. We then describe the key challenges that we solve to deliver the mobile-TYE solution: ensuring good coverage (section 4.2) and extracting stable imagery from wind-induced unstable imagery (section 4.3).




Figure 2: Simple system model

#### 4.1 Mobile-TYE platform design

We use a smartphone on the balloon or imaging and collecting other sensory data. In addition to that, the user who carries the balloon to image the desired area walks with another phone called *guidance phone*. This dual phone system plays a critical role in overcoming the efficient area coverage challenge mentioned earlier. The GPS readings of these two phones along with the known tether length of the balloon enable accurate determination of the balloon phone’s vertical height from the ground level. This information is used by the smartphone application to convey to the user the areas that have been imaged and then guides the user in an efficient manner in spite of wind and obstacles on the ground.

#### 4.2 Path Planning Algorithm

During our initial trials, we observed that the trial participants were not covering the intended area correctly. They were unable to map the balloon’s position and movement to the areas being imaged while moving around at the same time. To help a user effectively cover the area, we develop the path-planning algorithm and implemented it on an android application called “Guidance App” that provides guidance and feedback to the walker in real-time. Note that this is an application that makes suggestions to the farmer and the farmer is free to choose their path. The path suggested is in no way binding. However, the real-time feedback about the amount of area imaged has significant improvement in coverage as we will demonstrate in later sections.

The path-planning algorithm has two main goals: (a) minimize the time taken to cover an area, and (b) acquire sufficient data to successfully create a panoramic view of the area of interest. Traditional aerial imaging systems take steps to ensure that the camera is stable during the flight. However, as wind makes it impossible to fully control the camera in the case of mobile-TYE, we account for it by

estimating the area being mapped continuously. Instead of a rigid path that the user follows, our app provides walkers with real-time updates about the areas being imaged and they can dynamically navigate their path based on the situation. Additionally, we develop a novel process that can deliver complete image coverage of the target area in spite of significant wind-induced camera instability and the movement of the walker. In this section, we describe the algorithm and in section 8.1, we describe the user study and its findings.

The algorithm runs on two devices: the balloon phone and the guidance phone as shown in Figure 2. The user inputs the area that is to be imaged on the guidance phone. The user is then presented with the optimal path which images the entire area (assuming a stable camera) such that the time taken is minimal. The user then follows the given path as closely as possible while dynamically making small adjustments based on the feedback provided from the app. Again, the path is not binding and the user can make adjustments based on the real-time feedback as well as current wind conditions.

##### Functions of the balloon phone

- Take continuous imagery of the area of interest – either using burst photo mode or using a low frame rate video.
- Transmit the GPS coordinate of the camera to the guidance phone whenever there is sufficient stability in the imagery being captures (based on gyroscope measurements).

##### Functions of the guidance phone

- Obtain an estimate of the height of the camera at the beginning of the motion. Obtain the area of interest as well as camera resolution from the user.
- Use the estimate of the height, camera resolution and the area of interest to present a path which minimizes the time taken by a user to cover the area. Figure 3 shows how the camera resolution and height can be used to calculate the field of view of the camera. The estimate is further refined to account for camera rotations (shown by the red area in Fig. 3). This is further explained in section 4.2.1.
- Use the GPS coordinates transmitted by the balloon phone to estimate the area covered and display it on the Guidance App.

4.2.1 *Area captured by an image.* Consider a camera with a vertical field-of-view (FOV) of  $\theta$  and horizontal FOV of  $\phi$  which are along  $b$  and  $l$  respectively. Let the balloon be tethered to a stationary point and flying at a height  $h$  as shown in figure 3. The horizontal length covered  $l$  and the vertical length covered  $b$  is then given by




Figure 3: Lower-bound of area imaged by a single image


$$l = g(h) = 2 * h * \tan\left(\frac{\theta}{2}\right), \quad b = f(h) = 2 * h * \tan\left(\frac{\phi}{2}\right). \quad (1)$$

The platform for mounting the camera is designed in such a way that the camera faces the ground with high probability. However,

there is rotatory motion about the axis normal to the ground (in the plane parallel to the ground) due to which, it is difficult to estimate what area is getting imaged. Therefore, we lower bound the area imaged in the following way. We rotate the rectangle pivoted at the centroid to account for the various orientations that the camera could possibly be in. If we take the intersection of all these rotated rectangles, we get the inner circle of the rectangle (as shown in figure 3) with radius

$$r = \frac{1}{2} \min(b, l) = \frac{1}{2} \min(f(h), g(h)). \quad (2)$$

As the radius of the circle is a function of the height of the balloon and the FOV of the camera, the area imaged by the camera can be lower-bounded by the circle of the appropriate radius. As the height of the balloon, varies, the radius varies as well (as shown in figure 4). As the person moves, the area imaged is shown to the user and they can leverage the visualization to make appropriate decisions. The challenge arising from varying balloon height and rotatory motion is taken care of in the post-processing which we discuss in detail in section 4.3.



**Figure 4:** Area imaged varies with height of the balloon as shown by circles of varying width

**4.2.2 OptimalGuided Algorithm.** The algorithm *OptimalGuided* outputs the optimal path to image an area assuming that there are no random effects associated with wind. *OptimalGuided* algorithm ignores the effect of wind on the balloon path and outputs a deterministic path to be followed which minimizes the time taken to image an area. The algorithm can be described as follows:

- (1) Calculate the convex hull of the area to be imaged.
- (2) Calculate a path in the direction of the shortest ‘width’ of the convex hull taking into account the height of the balloon.

We do not provide a detailed proof of optimality for this algorithm – only give a sketch. Let us assume that if we walk a straight line of length  $l$ , then the area imaged is of size  $l \times w$  where  $w$  is the width of each image. Then, the problem of walking the least distance can be formulated as follows:

**Problem:** Cover the convex polygon with ribbons of width  $w$  such that we minimize the length of the ribbon plus the number of ribbon stripes used to cover the area.

**Solution:** Laying ribbons out in any direction can potentially incur some wastage on the edges. If we ignore those, the area covered by any assembly is the same. Thus, the length of ribbon used which is equal to area divided by  $w$  is also the same. The different ways to cover the area then only differs by the number of stripes. The number of stripes is minimized if we lay them down along the smallest ‘width’ of the polygon. The smallest ‘width’ of the polygon is defined as the smallest edge of all the rectangles which cover the given polygon. One may be able to do another pass, imaging only the most important areas or the areas not initially imaged. The optimal path can be computed through solving a Traveling Salesman Problem. However, in our trials we found that one pass

combined with the user’s dynamical adjustments based on feedback from the Guidance app was sufficient to cover an area.

### 4.3 Gyroscope-based Frame Selection

As mentioned earlier, the balloon’s height varies during the walk. To make sense of the imagery obtained, processing is required. A naive solution to this problem is to use the GPS coordinates of the camera to re-align the images. For example, if the GPS detects that the camera has moved down by 3m, the image should be moved accordingly to accommodate for this translation. However, this strategy has two pitfalls. First, GPS does not determine the camera’s rotation and the images are still prone to rotational transformations. Second, GPS accuracy is limited to a few meters at best, which can lead to more errors in image alignment.


While working with video imagery, an area is captured by several contiguous frames. Thus using all the frames are redundant and to avoid that, vision algorithms skip frames – say use every 20<sup>th</sup> or 30<sup>th</sup> frame for post-processing. However, as mentioned earlier, rotatory motion causes distortions. Specifically, when the cameras have a rolling shutter (like several mobile phone cameras), rotatory motion causes warping which makes the later stages of processing hard. Several works have looked at correcting the rolling shutter artifact [15, 19]. However, an important sensor that can help automated vision algorithms select good frames are gyroscopes. We equip our cameras with gyroscopes (or use smartphones with inertial measurement unit) which capture the amount of rotation experienced by the camera. We use the sensor reading to filter out the bad frames. However, putting a hard threshold on the gyroscope value is ineffective in selecting good frames. If the threshold is too low, then enough frames are not selected and it might prevent enough overlap amongst frames that is crucial for orthomosaic construction. To combat that, we propose the following algorithm. **Frame Selection:** Each video frame has a cost associated with it which is a function  $f(\cdot)$  of the gyroscope reading during the frame  $g_i$  (say the maximum of the absolute value of the three axis gyroscope reading, or the sum of squares of the reading, etc.). Let the ideal number of frames to skip be  $k$  and the allowable range to ensure good enough overlap be  $(\underline{k}, \bar{k})$ . We construct a graph where the nodes are the frames associated with the video and the directed edge from node  $i$  to  $j$  exist only if  $\underline{k} \leq j - i \leq \bar{k}$ . The edges have a weight associated with them given by:

$$e(i \rightarrow j) = \begin{cases} (j - i - k)^2 + \lambda f(g_i) & \text{if } \underline{k} \leq j - i \leq \bar{k} \\ \infty & \text{otherwise} \end{cases} \quad (3)$$

where  $\lambda$  is the weight associated with the node cost. To find the best frames, we use Dijkstra’s algorithm to pick a set of frames which has minimal distortion from rotational motion and enough temporal coverage to ensure good overlap (also explored in [11, 20]). These frames can then be fed into the image misalignment correction pipeline which is described in section 5.4. Once the images have been aligned, they can be stitched together into an orthomosaic.

## 5 STATIC-TYE

The static mode is suitable for applications where the area of interest is constant and the changes are to be tracked frequently for a long period of time such as flood monitoring and tracking various health indicators of the farm. In this mode, human intervention is



**Figure 5: Static TYE: Platform pipeline in the stationary mode of operation**

minimal. In this section, we first describe the design of the pipeline (section 5.1), then we describe the payload capacity of the system and how that dictates its lifetime (section 5.2) and finally describe the novel low-power event detection algorithm that runs locally to detect interesting events in real-time (section 5.3).

### 5.1 Static-TYE pipeline

The system comprises of the following (refer to Fig. 5): (a) helium filled balloon with a suspended smartphone, and (b) a gateway node (computer) which processes the data. The balloon system is tethered to a single stationary point on the ground for an extended period of time (days to weeks).

**Balloon System:** A single latex or mylar balloon filled with helium that is inflated to the required size to mount the payload, and sealed to prevent leakage. The payload consists of a custom lightweight mount, a smartphone to capture images, and a LoRA based device to be an interface between the phone and the gateway.

**Gateway:** The gateway node serves two important purposes. First, it serves as a node with intense computational capabilities for processing the captured images. Second, the gateway uploads data to the cloud where applications use this data to provide new insights. We do local processing to conserve bandwidth as well as to be robust to any outages on the edge-cloud link. This is akin to the design of Farmbeats [32].

### 5.2 Payload Capacity and Uptime

Static-TYE is intended for long-term imaging. We study and model how the payload capacity of balloons change over time. We then use this model to make design choices and to choose the size of the balloon depending on the application. The payload capacity of a balloon is dependent on several factors including: pressure inside the balloon, temperature of the gas inside as well as outside and the volume of the balloon. The payload capacity is derived as a combination of basic physics principles of ideal gas law and Archimedes' principle. We calculate the number of moles of any gas inside a balloon of volume  $V$  at temperate  $T$  and standard atmospheric pressure (101325 pascals), which is a reasonable pressure for the gas inside the balloon as the balloon surface dynamics requires it to be around the same as atmospheric pressure. Thus, from ideal gas law ( $PV = nRT$ ) we get, that the number of moles in a balloon with volume  $V$  and temperature  $T$  (in kelvin) is

$$n = \frac{V}{T} \times \frac{101325}{8.3144598} = 12186.6 \frac{V}{T} \text{ moles} \quad (4)$$

We know from Archimedes' principle that the buoyant force experienced by a body immersed in a denser substance is equal to the weight displaced. Thus we get that the weight displaced by a

'weightless' balloon is

$$\begin{aligned} \text{weight displaced} &= n \times (M(\text{air}) - M(\text{He})) \text{ kg} \\ &= n \times (28.97 - 4.00) \text{ g} = 24.97n \text{ g} \end{aligned} \quad (5)$$

as the average molar mass of air,  $M(\text{air})$  is 28.97g/mole and the molar mass of helium,  $M(\text{He})$  is 4g/mole. Putting Eqs. (4) and (5) together we get,

$$\text{Payload} = 304.3 \frac{V}{T} \text{ kg}. \quad (6)$$

However, the actual capacity of a balloon is reduced due to the weight of the balloon. Considering the weight of the balloon,  $m_b$  kg, the actual payload capacity is,

$$\text{Payload} = 304.3 \frac{V}{T} - m_b \text{ kg}. \quad (7)$$


From our experiments described in section 7.1 we conclude that our modified helium balloon system lose about half their payload capacity in four days. Thus, to have an uptime of a week, one would need a balloon whose initial payload capacity is 4x the actual payload.

### 5.3 Detection of changes in the environment

Capturing frequent aerial images and sending them to a device on the ground, say a PC over LTE, drains a lot of power which limits the uptime of the system. We can extend the battery life (and consequently the uptime) by sending only those frames where any changes have been detected. However, detecting changes is non-trivial. On the one hand, trivially sending frames that are significantly different from previous frames does not work as wind-induced motion causes nearby frames to be significantly different. On the other hand, complex object detection will drain a lot of power which goes against our main goal of power saving. To address these challenges: saving power by sending only the frames that are relevant to the actual changes in the environment and being able to detect these frames despite wind-induced variation, we employ an efficient and low-complexity image processing algorithm. We leverage a key insight to build the algorithm: **when the balloon drifts in one direction, all stationary points drift in the opposite direction. However, the mobile points drift in directions different from the stationary points' drift or the opposite direction.**

Our algorithm is based on the idea of homography from computer vision, which identifies and matches a set of keypoints (points of interest) between consecutive images for image registration and computation of camera motion between two images [30]. In static-TYE, the scene in the image is classified as constant when the keypoints match across images (inliers) or dynamic when there are keypoints that do not match across images (outliers). Depending on the number of outliers, we identify an image to have undergone changes as these outliers are mainly due to the introduction of a new object or movement in the frames. This technique is generic and can be applied to identify changes in the environment in any scenario as the motion of objects of interest will be different from rest of the image content. The algorithm is described as follows:

- Identify the keypoints in consecutive frames using speeded up robust features (SURF) [7]
- Match SURF keypoints between any two consecutive frames



**Figure 6: Detecting Motion:** The figure on the left shows all detected keypoints and the figure on the right shows that all keypoints on moving car are detected as outliers.

- Use random sample consensus (RANSAC) [12] to identify the outliers as they correspond to the objects of interest
- Keep track of the outliers across frames to observe the complete trajectory of the objects of interest such as the motion of a car

Fig. 6 shows the sequence of images in a farm along with inliers and outliers keypoints. We see that when a new object enters the image, the number of outliers increases significantly indicating a change in the image. Upon identification of a change in the environment, the image is subsequently transferred to the gateway node for further processing. This pipeline enables us to save on bandwidth, computation power, and energy on the payload device. In Section. 8.3 we evaluate this algorithm on aerial imagery collected by static TYE.

#### 5.4 Image Misalignment Challenge

Tethered balloons are susceptible to translations and rotations in the air due to the impact of wind. To quantify this impact, we tethered a balloon carrying a GPS equipped camera to a 40m long tether and measured the lateral motion for 10 minutes. The balloon motion covered 20m laterally in maximum point to point distance.

The motion caused by wind makes the imagery collected by TYE hard to interpret. To make sense of the images, the user is forced to constantly re-calibrate their mental mapping between the image plane and the physical world. This makes the user-interface highly cumbersome and non-intuitive. Furthermore, it is difficult to use this data in machine learning algorithms as is for automated processing of data. To construct stable, interpretable images using the data collected, we built a vision pipeline. Specifically, the pipeline follows these steps to re-align images across time:

- **Feature Extraction:** extract the visual features from each image using the DAISY [28, 29] feature extractor
- **Feature Matching:** match the features extracted from one image with the features from another image and reject the features that don't match across images
- **Homography Computation:** compute a homography ( $3 \times 3$  matrix) that maps features in one image to features in another
- **Homography Application:** apply the computed homography computed to the current image and ensure that this image is aligned to the previous image.

The images constructed using this pipeline can then be fed into further processing to extract insights described in Farmbeats [32].

## 6 IMPLEMENTATION

In this section, we describe the details of the implementation including custom hardware, balloon modification, and static and mobile-TYE operation.

### 6.1 Mounting Hardware

Figure 7a shows the hardware used to mount the smartphone onto the balloon. The mount is designed to have the following features: (a) low weight, (b) low cost, (c) simple fabrication using locally available material, (d) provide enough cushion to phone in case of a balloon burst, and (e) provide ventilation to the phone to prevent it from heating.

### 6.2 Balloon Material and Permeability

In section 5.2 we derived the payload capacity of a balloon of volume  $V$  at temperature  $T$ . However, we assumed that the balloon membrane was impermeable – i.e., no molecule can go into the balloon or come out of it. In practice, all materials have gaps – the size of the gaps depend on how closely the atoms are packed which means that some molecules can pass through the balloon.


Helium is the smallest molecule so it can leak through surfaces more rapidly than other gases. The leakage leads to reduction in the amount of helium in the balloon – ultimately the balloon loses the extra buoyancy. To make the system last long, it is essential to choose the balloon material intelligently. While latex balloons are easily available, a simple latex balloon has high leakage as its polymer structure facilitates movement of gas molecules [18, 21]. The rate of leakage depends on several factors which we discuss in Section 8. To use easily available latex balloons and still enable long term imagery, we chemically treat the latex balloon with Hi-Float. Hi-Float coating is easy to obtain and apply. It significantly reduces leakage from simple latex balloons. We evaluate the longevity of Hi-Float coated latex balloons in Section 8.

### 6.3 Operation

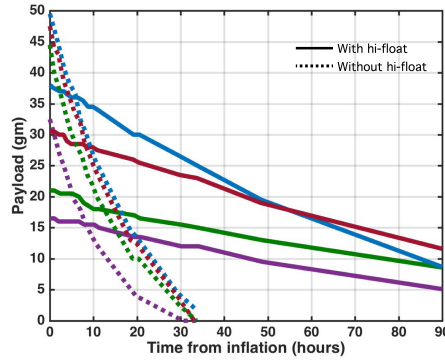
The balloon is inflated to the required diameter and the camera is mounted onto it. This process takes around 15 minutes. In the static mode, the balloon is tied to a stationary point on the ground. In the mobile mode, the guidance phone and the balloon phone are connected through a local Wi-Fi network with the guidance phone being the client and the balloon phone being the server. The guidance phone controls the imaging parameters and functions as a dashboard for the user to track the health of the balloon phone (battery, storage, temperature etc.). The area to be mapped is set using the map-overlay shown in the app and the balloon tether is unreeled to the desired height. The guidance app then shows the path to be followed on the phone along with green circles to indicate the area imaged as the user moves as shown in Fig 7a. The user can walk with the phone and the balloon in one hand using a simple device as shown in Figure 1.

## 7 MICROBENCHMARKS


We present microbenchmarks for evaluation of TYE below.



(a) Implementation: (a) Smartphone Mount, (b) Guidance App used in mobile mode



(b) Air Leakage Prevention: Using Hi-Float coating increases mean uptime of the system by about 90%.



(c) Freebody analysis

## 7.1 Air Leakage

The goal of this experiment was to measure the effect of Hi-Float treatment on the air leakage through the surface of the balloon. We measure the effect by comparing the rate of change of payload capacity of balloons treated with Hi-Float with balloons without Hi-Float over a four day period. Since human operators manually tie the balloons, there can be significant variation between different operators. To capture this variation, we conducted several experiments. These experiments were conducted in controlled environments with small balloons.

Figure 7b plots the payload capacity of the Helium balloons over a 4 day period. As shown in the figure, Hi-float coated balloons last much longer and have a lower degradation of payload capacity over time, i.e., hi-float coating reduces air leakage. Hi-Float based treatment of balloons reduces the leakage through the surface of the balloon by  $\approx 90\%$ . This observation holds across balloons of different initial payload capacities. This validates the design decision to use Hi-Float coating to treat latex balloons.

This experiment also shows that the capacity of a Hi-Float balloon typically decreases to about half over a period of 4 days. This means, for a balloon to last one week, we must over-provision the payload capacity by 4 times. This observation led us to use balloons which are 6 ft in diameter and have a capacity of over 3 kg ( $> 4$  times our payload size).

## 7.2 Path Planning Simulation

To evaluate the efficacy of TYE's path planning algorithm, we perform two different sets of experiments: a real-world user study (detailed in section 8.1) that evaluates TYE's path planning application empirically and a simulation to study the variation in coverage achieved by TYE's algorithm in varying wind conditions. The goal of the simulation is to check if the algorithm can ensure coverage across a broad spectrum of wind conditions. For the simulation, we assumed that the user followed the path suggested by TYE's guidance app (without dynamic adjustments).

**Setup:** Consider the frame of reference centered at the person – the X-Y plane is the ground and the Z-axis is perpendicular to the ground. The origin also serves as the tether point. Let the instantaneous speed of the wind be  $v$  and the angle it makes with the X-axis be  $\phi$  (we restrict the wind to be parallel to the X-Y plane).

Consider a balloon of radius  $r$ . Then the forces acting on the balloon are shown in Figure 7c where  $F_w$  is the force due to the wind,  $B$  is the net lift (buoyancy) acting on the balloon,  $m_b g$  is the weight of the balloon (including any payload), and  $T$  is the tension in the string. The balloon makes an angle  $\theta$  with the Z-axis to counter the wind force. The wind force  $F_w$  is given by

$$\begin{aligned} F_w &= P \cdot A \cdot C_d = \frac{1}{2} \rho v^2 \cdot A \cdot C_d \text{ (substituting } P = \frac{1}{2} \rho v^2) \\ &= 0.613 \cdot A C_d v^2 \text{ (substituting } \rho = 1.216) \\ &= 0.905 r^2 v^2 \text{ (substituting } C_d = 0.47, A = \pi r^2) \end{aligned} \quad (8)$$

where  $C_d$  is the drag coefficient[1] of the wind on a sphere.

Balancing the forces on the balloon along the horizontal and vertical directions, we get:

$$T \cos(\theta) = B - m_b g, \quad T \sin(\theta) = F_w \quad (9)$$

$$\theta = \tan^{-1} \left( \frac{F_w}{B - m_b g} \right) \quad (10)$$

If the length of the tether be  $l$ , the position of the center of the balloon is given by,  
 $(x, y, z) = (l \sin(\theta) \cos(\phi), l \sin(\theta) \sin(\phi), l \cos(\theta))$ .

The above equations show how the speed of the wind as well as the buoyancy effects the displacement of the balloon from its intended position: stronger winds cause large displacements.

**Simulation Results:** We varied the wind speed and estimated the amount of area imaged by TYE under those wind conditions. The results are tabulated in Table 2. The wind speeds considered correspond to a variety of conditions ranging from mild wind (Beaufort number 1) to wind causing dust to rise (Beaufort number 4) [4]. We see that for mild-to-moderate wind conditions ( $< 5$  m/s), the path planning algorithm ensures almost 100% coverage.

## 8 IN-THE-WILD EVALUATION

We conducted field trials for the TYE platform for agricultural monitoring in two farms in the US and India. We performed trials using latex balloons with diameter varying from 2ft (corresponding to payload capacity of 120g) to 6ft (corresponding to payload capacity of 3.2kg). We over-provision the balloon capacity to account for leakage. All balloons were lined with Hi-Float, a liquid solution that forms a coating inside balloons to increase float time. We filled the balloons with industrial grade helium. In all our experiments, we



Speed in m/s (mi/h)	Min Coverage	Max Coverage	Coverage $\leq$ 95%	Coverage $\leq$ 99%
1 (2.24)	100%	100%	0%	0%
2 (4.47)	99.5%	100%	0%	0%
3 (6.71)	98.5%	100%	0%	0.1%
4 (8.95)	96.16%	100%	0%	11%
5 (11.18)	92.10%	100%	10%	60%
6 (13.42)	84.70%	99.85%	49%	98%

**Table 2: Simulation Results: TYE’s path planning algorithm ensures high coverage across diverse wind conditions**

used locally available smartphones or a GoPro camera with a single board computer like Raspberry Pi. We add a 10000mAh battery back to provide additional power. Our payloads weigh from 300g to 500g across our experiments. Our trials lasted for about 4 days on average and collected imagery data for farms about 5 acre in size. Our findings are summarized below<sup>1</sup>:

- The area imaged by using the guidance app is around 88.2% of the area of interest (as opposed to 78.55% without the app) as demonstrated by the user study.
- The hardware modifications reduce the leakage through the surface of the balloon by  $\approx$  90%, thus increasing the system uptime.
- The intelligent-frame selection algorithm reduces the frame transmissions by 94% and the overall battery consumption by 67-88%, depending on the image resolution.
- Our vision pipeline mitigates the rolling shutter effects to produce a stable image.

The rest of the section will describe the experimental evaluation.

### 8.1 Path Planning Algorithm

We developed the path planning algorithm anticipating that variations due to wind as well as the error in human judgment about the area being mapped would result in a bad coverage. To validate our hypothesis and demonstrate the benefits of TYE’s design, we conducted a user study with 5 different users. We asked the users to image an area without TYE’s Guidance App. We described the area to be covered and specified the height at which they can fly their mobile-TYE system. Most of the users were unable to image the area completely and were unable to correct for the variations due to wind. We also used the recorded videos to try to make an orthomosaic but we were unsuccessful as there were holes in the area being imaged<sup>2</sup>.

The users were then presented with the same mobile-TYE system along with the Guidance App and asked to image the same area. This time most of the users were able to image almost 90% of the area and the footage quality was good enough for orthomosaic reconstruction. Averaged across 5 independent runs, the users took 1741 steps to cover 78.55 % area without the guidance App and with the help of guidance App were able to cover 88.20 % of area while taking only 1390 steps.

A sample of the area imaged by a user with and without the Guidance App is shown in Figure 8c and 8b respectively. When asked to cover the area in the blue rectangle, most users strayed off the marked area, while leaving gaps in the intended coverage area. All the users reported difficulties in managing the two tasks

<sup>1</sup>You can find some of the imagery that we collected for evaluating the platform here: <http://bit.ly/2nIMc26>

<sup>2</sup>Orthomosaic re-construction algorithms rely on overlap between images to stitch them together.

simultaneously: keeping track of the balloon position (the user has to crane their neck to spot the balloon), while keeping their eyes on the ground to walk around obstacles and to ensure that they were not trampling on plants. Conversely, the use of the app allowed them to keep their gaze on the ground while glancing at the app to follow the path and be able to make minor path adjustments to ensure coverage.

User	Without App		With App	
	# Steps	% Area	# Steps	% Area
A1	1898	88.59	1607	94.99
A2	1792	68.53	1171	88.72
Mean	1741	78.55	1390	88.2

**Table 3: Mobile TYE User Study: Presents results from a user study with 5 users. Using TYE’s guidance app, users cover more area with fewer steps**

We list results for two representative users in Table 3. User A1 was a conscientious user trying his best to maximize coverage. Yet, A1 could cover just 88.6% of the area without the guidance app. User A2, on the other hand, was less attentive and could not keep track of what area has been covered mentally. In both these cases, the guidance app allowed users to improve area coverage while traveling less number of steps.

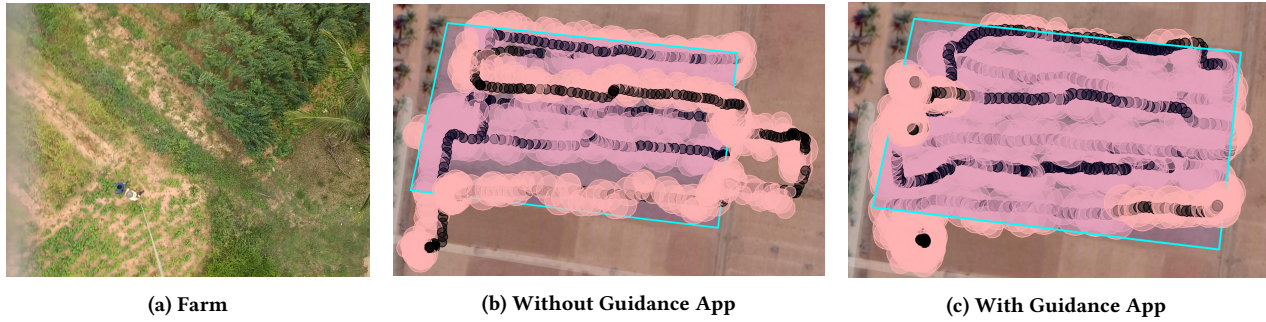
### 8.2 Vision Challenges

As mentioned in Sec. 4.3, gyroscope based frame selection is important when dealing with a lot of rotatory motion. Figure 9a depicts how the gyroscope based rejection picks frame with minimal gyroscope value when the weight  $\lambda$  associated with the gyroscope cost is varied. The higher the value of  $\lambda$ , the higher it prefers to pick very low gyroscope valued frames. The stems represent the frames that are picked, the parameters used are  $k = 30$ ,  $\underline{k} = 20$ ,  $\bar{k} = 40$ .

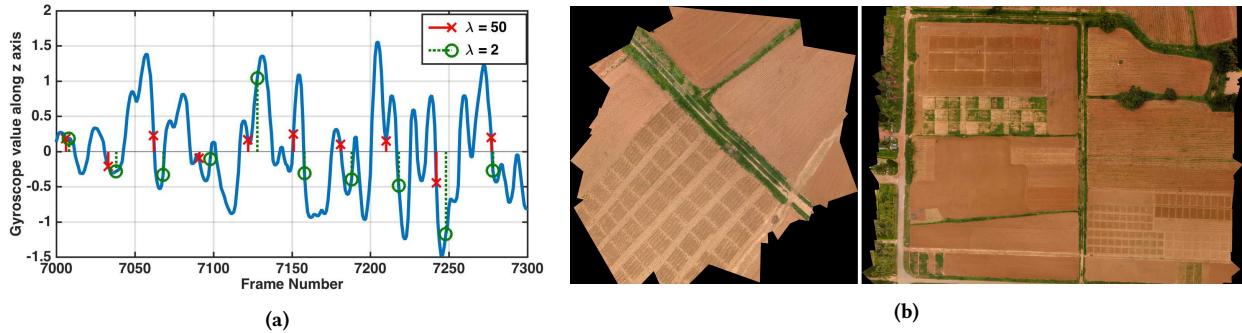
Figure. 9b shows the resulting stitches with and without the gyroscope based rejection of the same area using the same footage. As shown in the figure, the orthomosaic without gyro based frame selection is incomplete, skewed and lacks a stable frame of reference. With gyroscope based frame selection, a stable, correct and complete stitch is generated.

### 8.3 Detection of changes in the environment

In section 5.3 we described our algorithm to detect changes in the environment in static-TYE. We implemented the homography based event detection on an Android smartphone. The homography based event detection was written in C++ using OpenCV [8]. This native code is interfaced with the Android app using Android NDK [24]. The application captures an image and computes the number of outlier keypoints with respect to the previous image. When the ratio of number of outliers to number of inliers crosses a threshold, the image is classified as having gone through sufficient changes



**Figure 8: Mobile TYE User Study:** (a) A representative snapshot of the farm with different types of crops varying in heights and plant state, (b,c) As the user walks without the guidance app, he tends to get lost and walk extra distance without completely covering the intended area (blue rectangle). The guidance app extends the coverage and reduces the distance.



**Figure 9: Resolving Unstable Imagery:** (a) Effect of weighing gyroscope cost function in picking of the frames, (b,) Without TYE’s gyroscope based frame selection, the generated orthomosaic (left) is incomplete, skewed and hard to interpret. When TYE’s frame selection is used, the orthomosaic covers the designated area and produces a consistent view (right).

and transferred to the gateway node. We compared the proposed change detection algorithm with the naive approach – that transfers all images to the gateway node for processing.

Figure 10a shows the percentage of frames that need to be transmitted to the gateway for different values of the threshold parameter (ratio of number of outliers to number of inliers). As the threshold increases, less frames are classified as having undergone significant changes and thus, less frames are transmitted to the gateway.


To benchmark the energy consumption of our approach, we use a Monsoon power monitor to measure the energy consumed by the smartphone in (a) computing if the frame has undergone significant changes, and (b) transmitting a frame using LTE. We use these measurements to compare the energy consumption of TYE’s event detection algorithm (which transmits selected frames) with the naive baseline to transmit all frames. We plot the ratio of the power consumed by TYE’s algorithm to the power consumed by the naive baseline in figure 10b. As seen in the figure, as the threshold increases, TYE’s algorithm transmits fewer frames and consumes less energy. We empirically observe that a threshold of 0.75 is conservative enough to select all frames with any motion. Even with this conservative threshold, for high resolution images, TYE’s algorithm can save 67% power compared to the baseline. For lower resolution imagery, this number goes up to 88%. Of course, when the system is power-limited, the power savings can be increased by increasing the value of the threshold and just selecting the images with significant changes.

## 9 CONCLUDING REMARKS

We present TYE, a platform aimed at small holder farmers to enable them to reap the benefits of precision agriculture by acquiring high resolution aerial imaging and insights over very large areas or over extended periods of time while keeping costs as low as possible. TYE consists of an instrumented aerial camera mounted on a tethered helium balloon. We have developed novel hardware-software innovations to increase uptime and algorithms to ensure coverage of areas in spite of rapid wind-induced motion. We have also built algorithms to alleviate the distortions in imagery caused by the wind-induced motion.

The overarching design goal has been to minimize both the capital costs and the operational costs of the system: the design of the balloon mount and techniques to keep the balloon afloat longer; the use of commodity smartphones keeps maintenance and replacement cost of the imaging platform low; the algorithms to extract stable orthomosaic from a very unstable video feed has kept the mounting simple and lightweight (instead of using powered gyro stabilizers) reducing the capital cost of the mount and the operational cost low since a smaller balloon with lesser helium can keep the payload up; the efficient event detection on the phone reduces the battery use otherwise needed for video transfer keeping the balloon up longer; the simplicity of the balloon platform allows for a semi-skilled person to operate the system reducing the operational cost. The guidance app further reduces the operator time by minimizing the time needed to map a given area.

Our initial deployments in India and US for agricultural applications have shown promising results. We believe TYE will prove to



**Figure 10: Static TYE Evaluation: (a) As the threshold for change detection increases, we select less frames to transmit, (b) We plot the ratio of power consumed by Static TYE phone in event-detection mode to the power consumed if all frames are transmitted over LTE. Even when high-res images are transmitted, our approach saves more than two-third the power at low thresholds of 0.75.**

be the aerial imagery platform of choice for the developing world for applications in precision farming, crowd monitoring, forest and environment monitoring and others.

## REFERENCES

- [1] 2017. Drag coefficient. [https://en.wikipedia.org/wiki/Drag\\_coefficient](https://en.wikipedia.org/wiki/Drag_coefficient). (2017).
- [2] 2017. How Many Farmers Are Really Using Drones - And Who's Flying? <https://dronelife.com/2017/04/17/many-farmers-really-using-drones-whos-flying/>. (2017).
- [3] 2017. Understanding and Evaluating Satellite Remote Sensing Technology in Agriculture. (2017). [http://www.geosys.com/wp-content/uploads/2017/04/Whitepaper\\_SatelliteRemoteSensingTechnology\\_L.pdf](http://www.geosys.com/wp-content/uploads/2017/04/Whitepaper_SatelliteRemoteSensingTechnology_L.pdf)
- [4] 2018. Beaufort Scale. [https://en.wikipedia.org/wiki/Beaufort\\_scale](https://en.wikipedia.org/wiki/Beaufort_scale). (2018).
- [5] James S Aber. 2004. Lighter-than-air platforms for small-format aerial photography. *Transactions of the Kansas Academy of Science* 107, 1 (2004), 39–44.
- [6] James S Aber, Susan W Aber, and Firooza Pavri. 2002. Unmanned small format aerial photography from kites acquiring large-scale, high-resolution, multiview-angle imagery. *International Archives of Photogrammetry Remote Sensing and Spatial Information Sciences* 34, 1 (2002), 1–6.
- [7] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. 2006. Surf: Speeded up robust features. *Computer vision—ECCV 2006* (2006), 404–417.
- [8] Gary Bradski. 2000. The OpenCV Library. *Dr. Dobbs's Journal: Software Tools for the Professional Programmer* 25, 11 (2000), 120–123.
- [9] DJI. 2018. DJI Battery Specifications. (2018). <http://store.dji.com/product/phantom-3-intelligent-flight-battery>.
- [10] DJI. 2018. DJI Phantom Series. (2018). <http://store.dji.com/shop/phantom-series>.
- [11] Pedro F Felzenszwalb and Ramin Zabih. 2011. Dynamic programming and graph algorithms in computer vision. *IEEE transactions on pattern analysis and machine intelligence* 33, 4 (2011), 721–740.
- [12] Martin A Fischler and Robert C Bolles. 1987. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. In *Readings in computer vision*. Elsevier, 726–740.
- [13] Patricia K. Freeman and Robert S. Freeland. 2015. Agricultural UAVs in the U.S.: potential, policy, and hype. *Remote Sensing Applications: Society and Environment* (2015).
- [14] C. M. Gevaert, J. Suomalainen, J. Tang, and L. Kooistra. 2015. Generation of Spectral-Temporal Response Surfaces by Combining Multispectral Satellite and Hyperspectral UAV Imagery for Precision Agriculture Applications. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* (2015).
- [15] Matthias Grundmann, Vivek Kwatra, Daniel Castro, and Irfan Essa. 2012. Calibration-free rolling shutter removal. In *Computational Photography (ICCP), 2012 IEEE International Conference on*. IEEE, 1–8.
- [16] Mitchell C. Hunter, Richard G. Smith, Meagan E. Schipanski, Lesley W. Atwood, and David A. Mortensen. 2017. Agriculture in 2050: Recalibrating Targets for Sustainable Intensification. *BioScience* 67, 4 (2017), 386–391.
- [17] WHITTLES, JH. 1970. Tethered balloon for archaeological photos. *Photogrammetric Engineering* 36, 2 (1970), 181.
- [18] AI Kasner and EA Meinecke. 1996. Porosity in rubber, a review. *Rubber chemistry and technology* 69, 3 (1996), 424–443.
- [19] Chia-Kai Liang, Li-Wen Chang, and Homer H Chen. 2008. Analysis and compensation of rolling shutter effect. *IEEE Transactions on Image Processing* 17, 8 (2008), 1323–1330.
- [20] Tiecheng Liu and John R. Kender. 2007. Computational Approaches to Temporal Sampling of Video Sequences. *ACM Trans. Multimedia Comput. Commun. Appl.* 3, 2, Article 7 (May 2007).
- [21] I Pinnau, JG Wijmans, I Blume, T Kuroda, and KV Peinemann. 1988. Gas permeation through composite membranes. *Journal of membrane science* 37, 1 (1988), 81–88.
- [22] DG Pitt and GR Glover. 1993. Large-scale 35-mm aerial photographs for assessment of vegetation-management research plots in eastern Canada. *Canadian Journal of Forest Research* 23, 10 (1993), 2159–2169.
- [23] JACOB POUHTER. 2016. Smartphone Ownership and Internet Usage Continues to Climb in Emerging Economies. (2016). <http://www.pewglobal.org/2016/02/22/smartphone-ownership-and-internet-usage-continues-to-climb-in-emerging-economies/>
- [24] Sylvain Ratabouil. 2015. *Android NDK: beginner's guide*. Packt Publishing Ltd.
- [25] Catur Aries Rokhmana. 2015. The Potential of UAV-based Remote Sensing for Supporting Precision Agriculture in Indonesia. *Procedia Environmental Sciences* 24 (2015). The 1st International Symposium on LAPAN-IPB Satellite (LISAT) for Food Security and Environmental Monitoring.
- [26] Jeff Sharkey. 2009. Coding for Life: Battery Life, That is. (2009). [https://dl.google.com/io/2009/pres/W\\_0300\\_CodingforLife-BatteryLifeThatIs.pdf](https://dl.google.com/io/2009/pres/W_0300_CodingforLife-BatteryLifeThatIs.pdf)
- [27] Mike J. Smith, Jim Chandler, and James Rose. 2009. High spatial resolution data acquisition for the geosciences: kite aerial photography. *Earth Surface Processes and Landforms* 34, 1 (2009), 155–161. <https://doi.org/10.1002/esp.1702>
- [28] E. Tola, V. Lepetit, and P. Fua. 2010. DAISY: An Efficient Dense Descriptor Applied to Wide Baseline Stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, 5 (May 2010), 815–830.
- [29] E. Tola, V. Lepetit, and P. Fua. 2008. A Fast Local Descriptor for Dense Matching. In *Proceedings of Computer Vision and Pattern Recognition*. Alaska, USA.
- [30] Philip HS Torr and Andrew Zisserman. 2000. MLESAC: A new robust estimator with application to estimating image geometry. *Computer Vision and Image Understanding* 78, 1 (2000), 138–156.
- [31] P. Tripicchio, M. Satler, G. Dabisias, E. Ruffaldi, and C. A. Avizzano. 2015. Towards Smart Farming and Sustainable Agriculture with Drones. In *International Conference on Intelligent Environments*.
- [32] Deepak Vasisht, Zerina Kapetanovic, Jongho Won, Xinxin Jin, Ranveer Chandra, Sudipta Sinha, Ashish Kapoor, Madhusudhan Gumbalapura Sudarshan, and Sean Stratman. 2017. FarmBeats: An IoT Platform for Data-Driven Agriculture. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*. USENIX Association, Boston, MA. <https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/vasisht>