PROCEEDINGS OF SPIE

SPIEDigitalLibrary.org/conference-proceedings-of-spie

Lessons learned towards the immediate delivery of massive aerial imagery to farmers and crop consultants

Gooch, Amy, Petruzza, Steve, Gyulassy, Attila, Scorzelli, Giorgio, Pascucci, Valerio, et al.

Amy A. Gooch, Steve Petruzza, Attila Gyulassy, Giorgio Scorzelli, Valerio Pascucci, Lance Rantham, Weston Adcock, Calvin Coopmans, "Lessons learned towards the immediate delivery of massive aerial imagery to farmers and crop consultants," Proc. SPIE 11747, Autonomous Air and Ground Sensing Systems for Agricultural Optimization and Phenotyping VI, 117470A (12 April 2021); doi: 10.1117/12.2587694

SPIE.

Event: SPIE Defense + Commercial Sensing, 2021, Online Only

Lessons learned towards the immediate delivery of massive aerial imagery to farmers and crop consultants

Amy A. Gooch[a], Steve Petruzza[b], Attila Gyulassy[c], Giorgio Scorzelli[c], Valerio Pascucci[c,a], Lance Ramthun[d], Weston Adcock[a,d], and Calvin Coopmans[b]

[a]ViSUS LLC, Salt Lake City, Utah
[b]Utah State University, Logan, Utah
[c] Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, Utah
[d]Crop Solutions LLC, Jonesboro, Arkansas

ABSTRACT

In this paper, we document lessons learned from using ViSOAR Ag Explorer™ in the fields of Arkansas and Utah in the 2018-2020 growing seasons. Our insights come from creating software with fast reading and writing of 2D aerial image mosaics for platform-agnostic collaborative analytics and visualization. We currently enable stitching in the field on a laptop without the need for an internet connection. The full resolution result is then available for instant streaming visualization and analytics via Python scripting.

While our software, ViSOAR Ag Explorer™ removes the time and labor software bottleneck in processing large aerial surveys, enabling a cost-effective process to deliver actionable information to farmers, we learned valuable lessons with regard to the acquisition, storage, viewing, analysis, and planning stages of aerial data surveys.

Additionally, with the ultimate goal of stitching thousands of images in minutes on board a UAV at the time of data capture, we performed preliminary tests for on-board, real-time stitching and analysis on USU AggieAir sUAS using lightweight computational resources. This system is able to create a 2D map while flying and allow interactive exploration of the full resolution data as soon as the platform has landed or has access to a network. This capability further speeds up the assessment process on the field and opens opportunities for new real-time photogrammetry applications.

Flying and imaging over 1500-2000 acres per week provides up-to-date maps that give crop consultants a much broader scope of the field in general as well as providing a better view into planting and field preparation than could be observed from field level. Ultimately, our software and hardware could provide a much better understanding of weed presence and intensity or lack thereof.

Keywords: gigapixel, 2D images, on-demand, visualization and analysis, agriculture, UAV, sUAS, drone

1. INTRODUCTION

In this paper, we document lessons learned from using our software in the fields of Arkansas in the 2018, 2019, and 2020 growing seasons. Our insights come from creating software with fast reading and writing of 2D aerial images for platform agnostic collaborative analytics and visualization. Our software provides in field image stitching at a coarse resolution, progressively solving image mosaicing for thousands of aerial images in minutes. The full resolution result is then available for instant streaming visualization and analytics via Python scripting.

More recently, we prototyped a system to perform on-board, real-time stitching and analysis of sUAS-borne imagery using lightweight computational resources. This system is able to create a 2D map while flying and allow interactive exploration of the full resolution data as soon as the platform has landed or has access to a network. This capability further speeds up the assessment process on the field and opens opportunities for new real-time photogrammetry applications.

Further author information: (Send correspondence to Amy Gooch, ViSUS LLC)
Amy Gooch: E-mail: amy@visus.net

Autonomous Air and Ground Sensing Systems for Agricultural Optimization and Phenotyping VI,
edited by J. Alex Thomasson, Alfonso F. Torres-Rua, Proc. of SPIE Vol. 11747, 117470A
 · doi: 10.1117/12.2587694

1.1 Technical Innovation Foundation

Our agriculture solution is based upon the OpenViSUS™ I/O library,[1] which instead of writing data from imaging scans to traditional formats (like TIFF and RAW), translates aerial data on-the-fly and writes data to IDX file format, a hierarchical multi-resolution format, inherently suitable for visualization and analysis tasks. Consequently, OpenViSUS™ provides both fast write speeds at scanning time as well as fast parallel, serial, and even remote data read access during post-processing. Supported by several different DoE programs (SciDAC, PSAPP, and the Co-Design centers), OpenViSUS™ was initially developed to address the leading edge in terms of speed and scalability.

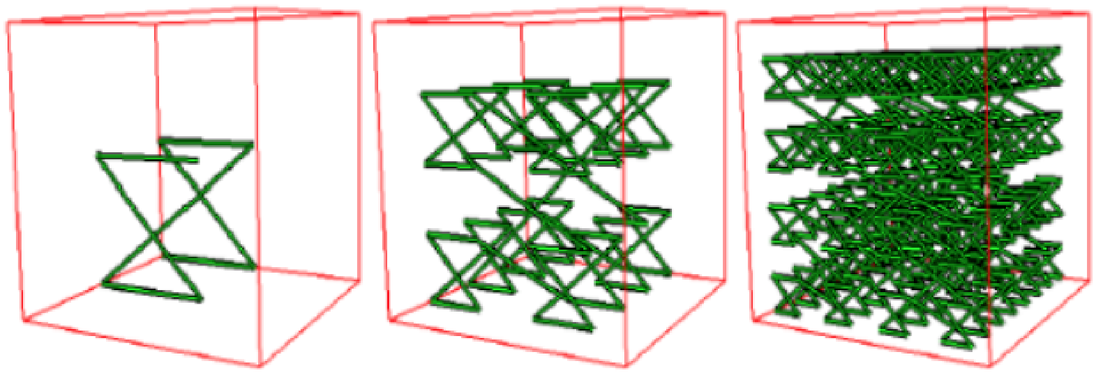

Figure 1. The multi-resolution, cache-oblivious IDX data format written by OpenVisus™ framework deploys a hierarchical variant of the Lebesgue space filling curve (commonly referred to as HZ order) at the data access layer.

Innovation: The multi-resolution, cache-oblivious IDX data format written by OpenVisus™ framework deploys a hierarchical variant of the Lebesgue space filling curve (commonly referred to as HZ order) at the data access layer (Figure 1. This layout instills both spatial and hierarchical locality, meaning that spatially coherent queries in the data domain map to contiguous regions in the memory layout, making optimal use of deep computer memory hierarchies. In addition to OpenViSUS™ being the fastest library to date for writing the results of large-scale image acquisition,[2–7] the files written also convey benefits for later analysis and visualization. In particular there are three key features that make the IDX format very attractive for post-processing analysis and visualization workflows:

1. Queries in IDX files for blocks of data are returned in a coarse-to-fine order, enabling interactive analysis and visualization with progressive refinement, no matter the data size or location: local disk access or remote transmission over a network are handled uniformly and do not require any large data reorganization (technically this termed a "cache oblivious" approach). This also enables a very lightweight server to service remote data requests.

2. IDX HZ-order indexing, implemented similar to Z-order indexing used in classical database approaches with a simple sequence of bit-string manipulations, results in both fast writes and fast reads.

3. The IDX layout avoids data replication, in contrast to other hierarchical approaches, which both reduces size-on-disk and eliminates the performance penalties associated with guaranteeing consistency during dynamic updates. As a consequence of the layout, scanned data in the IDX format can be hierarchically traversed from coarse to fine resolutions and can progressively update any output data structures derived from the data. The ViSUS™ web-server, a light Apache module, allows fast remote access IDX data, enabling interactive visualization and exploration without the need to move large datasets. The data order and progressivity of refinement enables a streaming framework for real-time access to massive remote data, for example, having been demonstrated in the context of remote monitoring of petascale combustion simulation output.

Figure 2. The ViSUS ViSOAR Ag Explorer has demonstrated that actionable information can be obtained directly on the field thanks to its fast stitching and data streaming technology, as shown with this example of a georeferenced mosaic of aerial images registered on Google Maps.

OpenViSUS™ supports an innovative way of storing data, where data can come from heterogeneous and diverse sources and can be stored on different targets. The only assumption we imposed on the data is that (1) data is organized in a continuous stream of samples coarse-to-fine, meaning that the probability of requesting some subset of the data is higher for samples at the beginning of the stream, and lower for the samples at the end of the stream; and (2) it is possible during the traversal to skip chunks of unnecessary data. Furthermore, OpenViSUS™ allows the possibility to have a virtually infinite stream, i.e. it allows the possibility to add as many details as needed in only specific local regions. Consider, for example, the case of a world satellite map (Google Earth, 3 Petabytes) where we could stitch several world cities in real-time and go even at centimeters level-of-details (LOD). These techniques, combined with a fast and cache-friendly data access mechanism, allow the user to interactively explore and edit massive imagery, with the illusion of having a full solution at hand. Figure 2 provides an example of an aerial IDX dataset layered over Google maps.

Additionally OpenViSUS™ provides scalability across a wide range of platforms, from mobile to laptop to office desktop to high performance computing clusters. Also, OpenViSUS™ provides a Python API for data streaming access (e.g., to a streaming data portal) and a set of examples of real visualization and analysis use cases created in collaboration with users of different communities. The OpenViSUS™ Python API makes it easy to create scripts for processing NDVI, thresholding, region finding and counting, as shown in Figure 3.

2. DEVELOPMENT OF THE IDEA

In 2015 we started a journey via the NSF I-Corp program, looking for a problem in need of our solution for fast reading and writing of data. After looking through many domains including the mature market of histopathology, mismatch of technologies for applications in architecture, established solutions in web mobile development, and established players in oil and gas, we started examining in depth the domain of agriculture. While starting with a solution and looking for a problem is not ideal, our technical innovation may never have been applied to the field of agriculture without the process of the NSF I-Corp program. We talked to more than a hundred people in order to understand their pains while internally considering how we could help provide gains. From our interviews, we heard, "It's not about money, it's about optimization" and "Growers are competitive! They are willing to do anything to beat their neighbor." We learned there was a software bottleneck (time and labor) in processing large aerial surveys, and minimizing this bottleneck could enable a cost-effective process to deliver actionable information to farmers. In turn, farmers and farm consultants could optimize the use of fertilizers and more finely control the amount of pesticides and herbicides necessary to increase crop yield. Furthermore, farmers mitigate costs and losses by being able to spot problem areas, minimize the spread of plant diseases, and identify issues such as standing water, irrigation malfunctions, and persistent automated machinery errors in planting

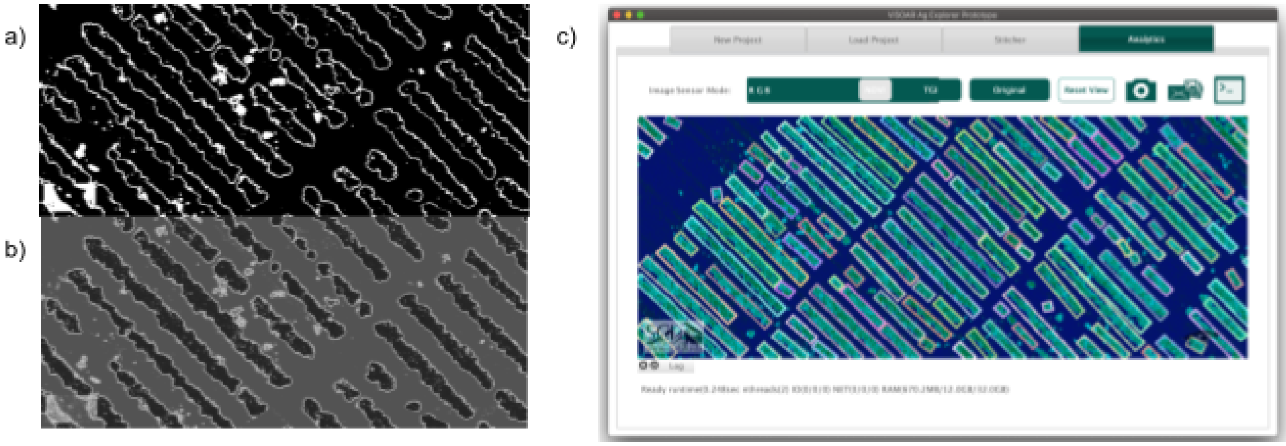

Figure 3. Examples of Python post processing in ViSOAR Ag Explorer™: a) Thresholding, b) Region Finding, c) Contour finding and counting

or cultivation. We believed we could develop new algorithms for on-the-fly orthorectification, stitching, and normalization of aerial image mosaics. The OpenViSUS™ technology that underlie our technical development is designed from the ground up to process massive data with less memory and increased speed relative to other approaches. These performance gains are enabled by a proprietary streaming cache-oblivious generic image representation, one that enables multi-channel giga- and tera-pixel images to be treated as ordinary images.

As we dug into the domain of agriculture, we initially thought the farming equation was simple. You add seed, fertilizer, pesticide, a farm, and out comes crops. We thought the basic equation of farming involved paying for inputs, applying them to the land, and then selling the resulting crop. However, the domain of precision agriculture aims to help optimize both the inputs and outputs in farming, often making prescriptions for where exactly in the field to put how much, allowing farmers to put money into where it is needed most. At the upper end, farmers can save 40% on these inputs by only applying where it is needed, and such information that can prevent crop loss, as well as increased yields around 15%.

In the domain of agriculture, everything ACTUALLY depends on WHERE do you apply HOW MUCH of WHAT and WHEN. In fact, this is so complex that farmers rely on crop scouts and farm advisors, people who have devoted their entire careers to answering these questions. One can almost think of the crop scout as a black box, where we seek to help provide the information of what goes on in that box in order to get a recommendation, and this is where aerial imagery is starting to make an impact. Where are the crops stressed? Where should the farmer apply pesticides, where should the farmer tackle weeds? How much fertilizer should be recommended?

Consider a typical corn quarter-section, about 160 acres (or about 160 football fields), identifying pests or crop stress is like finding a "needle in a haystack." This is where delivering aerial imagery provides REAL VALUE to crop advisors, and improves their recommendations to farmers.

In late 2015, we were awarded the NSF SBIR Phase I: Immediate Delivery of Massive Aerial Imagery to Farmers and Crop Consultants (Award # 1549187) and in 2017 we received follow up Phase II funding (Award #1738448). These grants aided us in building ViSOAR Ag Explorer gigapixel imaging software, which enables farmers and crop consultants to **fully utilize massive aerial surveys seamlessly in management decisions**, a transformative capability leading to reduced inputs (fertilizers, herbicides, pesticides) and increased yields. Data from aerial scans becomes most useful when acquired with high frequency, as trends over time tell a farmer where problems may arise. To be cost-effective, the marginal cost of each successive aerial scan must be negligible. While small unmanned aerial systems (sUAS) promise to deliver low cost through automation, the current state-of-the-art software and algorithms for ingesting the imagery have pain points that potentially obstruct adoption of the entire workflow. The software step of image stitching (i.e., building maps using the images acquired by the drone) is long and labor intensive, expensive computational hardware increases a user's barrier to entry, and artificial limits on image resolution reduce the information that can be derived. Several companies are turning to cloud-based software to address these issues with raw compute power. However, farmers and rural communities remain the most underserved in terms of internet connectivity, the long data transmission times (possibly days) being a roadblock in adoption of the technology.

Our innovation, ViSOAR Ag Explorer, overcomes the obstacles of limited compute resources and limited connectivity to the cloud. New algorithms built on the stream processing image infrastructure developed by OpenViSUS™ enable automated image stitching on any computer configuration, from commodity hardware such as a farmer's laptop PC, to mobile devices, with no loss in image resolution. Furthermore, the software will deliver interactive exploration and analysis of gigapixel images on any device (including mobile) and adaptive analysis that can be performed on-the-fly. The net effect will be to hide the complexity of handling massive image mosaics, saving both labor and time. ViSOAR Ag Explorer can be deployed both as a stand-alone application and as a library with an API for easy integration with end-to-end precision agriculture systems.

3. IN FIELD TRIALS

As part of our NSF ICorp customer search in 2016 we met Lance Ramthun, owner of Crop Solutions LLC and forged a multi-disciplinary collaboration that led to us being able to shadow him through his daily tasks and eventually we were granted access to scanning the fields he manages. Lance's team consists of full time management employees and seasonal employees. Full time management positions are tasked with day-to-day operation procedures and decision-making responsibilities. The scope of their customer base is such that they have to employ part time seasonal employees to scout fields so they are able to keep schedules and deadlines. Quality control is of the utmost importance, wavering in consistency or attention is not acceptable. Our in field trials we sought to understand if ViSUS technology will at the least improve those factors, and the degree to which they could replace the need for seasonal employment in the future.

Our initial guess started with the idea that ViSUS has developed technology for greatly reducing the time and resources needed to stitch data from UAV acquired imagery. ViSOAR Ag Explorer by ViSUS LLC provides a solution for stitching of full resolution aerial imagery in the field without requiring an internet connection, making it an indispensable part of the software stack underlying every agricultural UAS system sold. Agisoft Metashape and Pix4D each take approximately 4 hours to process 100 images on a high-end workstation, or up to 8 hours for imagery per flight; far too time consuming for large-scale deployment. At the same time we were developing ViSOAR Ag Explorer products like Pix4D Fields and Slant Range developed in field stitching software. Although our in field stitching time is similar, our software is differentiated by having the full resolution imagery available upon stitching, not downsampled data. Our goal is to provide in field image stitching at a coarse resolution that progressively solves 600 images in 10 minutes providing a full resolution view of the data, enabling instant analytics. Our current image registration is fast, able to handle thousands of images (from different flights). Current timings are about 10 minutes to stitch 400 images on a regular laptop, but on a fast solid-state drive we can process about 1000 images in 12 minutes.

Multiple seasons were required to flesh out our software pipeline and ultimately understand if ViSUS technology could enable Crop Solutions staff to be able to return quality data quickly and in the field with little to no internet access. Our breakthrough technology could increase the quality control measures demanded within crop consulting operations and set them ahead of the industry. Weather is one variable that cannot be controlled in commercial crop production, damage from storms is inevitable. Crop Solutions provides their growers with damage estimates to prepare them for commodity contract estimates not fulfilled and insurance quote negotiation. Performing this task is very time consuming and detailed, and is also provided as an added value. Damage estimates are a perfect fit for sUAS and ViSUS. One of our long term goals is to provide a pipeline with an sUAS that will enable a full field view in combination with the ViSUS technology the ability to stitch and view data acquired, enabling the delivery of concise damage estimates same day to crop consultant customers, with much less effort and overhead from the company.

Additionally, we believe that Lance's expertise and network can show ViSUS several new areas of exploration or aspects of the industry that ViSOAR Ag Explorer may help with. Technology is on the forefront of agriculture everyday and crop consultants pride themselves on being the first to successfully implement new ways of thinking and performing.

Figure 4. Image of the equipment set up in the back of a crop scout's utility vehicle. The cases hold the drones, repair kits, and extra batteries. The battery charger is plugged in and ready to charge batteries.

3.1 Data Gathering Equipment

Our equipment for one drone pilot, parts of which are shown in the back of the crop scout's truck in Figure 4, consists of:

- Drone with 6 batteries and charging stations; (We have tried DJI Phantom 4, DJI Phantom 4 Pro and Parrot Bluegrass)
- An iPad for flying the drone with Pix4D Fields or Drone Deploy
- Our field laptop is a 2020 Dell XPS 15 7590 with an Intel Core i7- 9750H up to 4.5GHz and GTX 1650 4GB and 32GB RAM and 2TB SSD
- 4TB and 10 TB external storage
- Google Drive Account for keeping all notes and shared data
- Micro SD cards for storage inside drone (multiple 128 GB, 64 GB, 32 GB)

3.2 Data Gathering for the 2019 season

Once we had our ViSOAR Ag Explorer prototype for stitching aerial images built, we attempted to start field testing in 2018, but could not find a drone pilot until the 2019 and 2020 growing seasons in Arkansas as part of the agile development of our software. We found difficulty in hiring a drone pilot to collaborate with our Crop Consultant in Jonesboro, Arkansas because drone operators and pilots are still grouped in select areas around the country. We made a connection with the local university, but our timing did not line up with students' availability in 2018, since the growing season ran February to August. With a lesson learned, we started our hiring more than 8 months out. However, we also learned that the cost of flying a drone is much more than we expected. Professional UAV pilots make about $20/hr and some operations require the pilot to have someone else in the field to watch out for obstacles (crop dusters, flocks of birds, etc.). Still, we were only able to hire a full time drone pilot to work with us in Arkansas from April to July 2019, and then scrambled to hire a replacement drone pilot to fly July to the end of the season.

3.3 The 2020 Growing Season

Learning from our experience in the 2019 growing season, we decided that it would be necessary to find and to train our own drone operator. The drone operator needed to primarily be a crop scout, someone who could handle Arkansas summers, extended hours, and field work, with the addition being willing to get the drone operator license and perform the necessary tasks. Late in 2019 Lance found the crop scout Weston Adcock and we helped him spend a couple of months studying for the certification, working with a local college professor

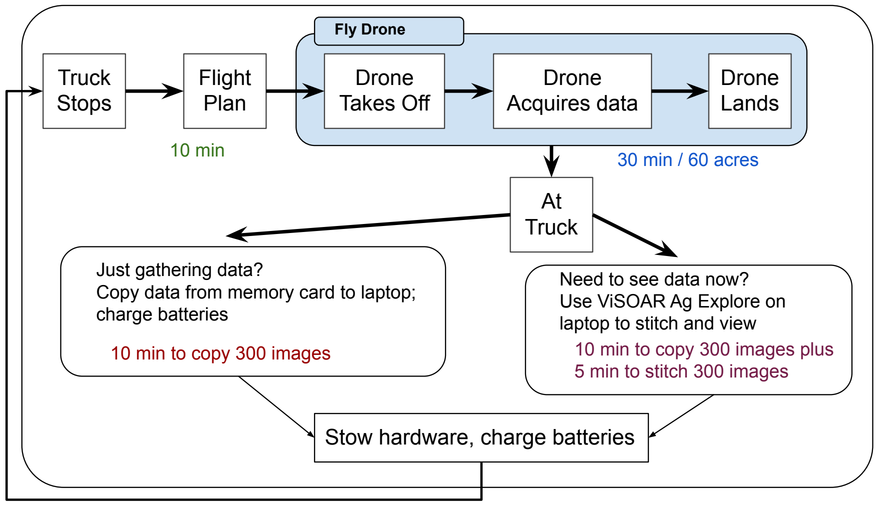

Figure 5. Flow chart of activities for flying the drone in the field and capturing data with significant times for each task noted. All timings are averages and may be affected by weather and hardware.

who taught students how to get their drone pilots license, and he passed the test on the first try. The test will need to be retaken every two years in order to maintain his certification.

Weston spent winter and early spring practicing gathering data under the guidance of Lance Ramthun and developing a daily flying routine, shown in Figure 5. The about 10 minute flight preparation routine consisted of deciding what field to fly considering the weather and field conditions, setting up a field plan using off the shelf software such as Pix4d Fields or Drone Deploy on an iPad, and setting up the drone. When ready to fly, the pilot started up the software flight plan software and watched the streaming view of the drone during flight image acquisition. Some fields require two or three battery changes to acquire the whole field. After the drone landed the battery is pulled from the drone to let cool before recharging, the drone is carefully stowed in protective boxes, the data card is pulled and copied over to an external drive attached to a laptop [5-10 minutes] and a clean data card is used for every flight. If needed, the data was then stitched in the field using ViSOAR Ag Explorer. Our initial solution also included the ability to email screenshots to the crop consultant, but these are static images and not at full resolution. Figure 6 provide information on acreage and dates flown from June 2020 to October 2020.

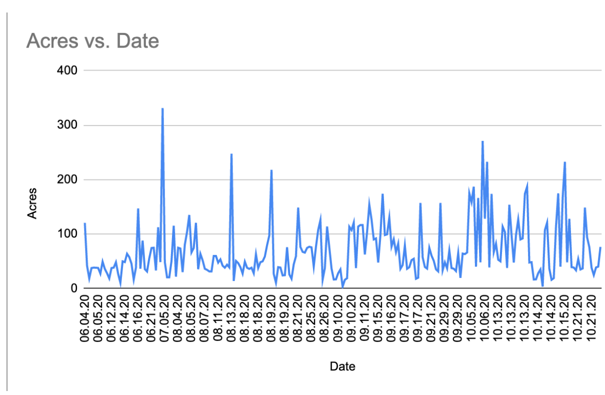

Figure 6. Plot of acres versus date from June 2020 to October 2020. We conducted 226 flights, scanning and processing the data for 16,000 acres. The maximum acres we were able to capture in one day was about 1200 acres and the largest single scan was 330 acres. We also collected data February thru May of 2020, but the data logging was insufficient for gathering statistics about acreage, etc.

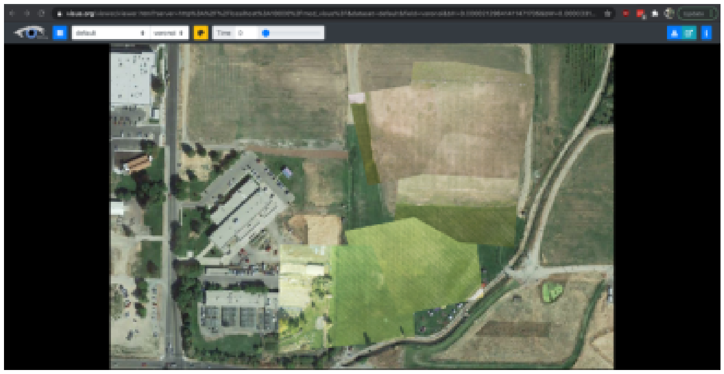

Figure 7. ViSUS Web portal interactive view available 3 minutes after a ViSOAR Ag Explorer On-Board 10 minute test flight

3.4 On-Board Computing

Experiments in the field have also highlighted the fact that data transfer and external (i.e., outside the drone platform) processing accounts for a large amount of down time, leaving the operator waiting for the drone to land and then waiting again for the transfer of the data on to a computer or to the cloud for further processing. This scenario is certainly not ideal for obtaining actionable information in a timely manner, especially in critical scenarios. Advances in edge computing (low-energy, low-weight computational hardware) open an opportunity to minimize the existing time delay from the moment of the drone survey to the spatial information availability for decision making. This motivated us to use our technology to close that gap and bring fast image data registration and interactive streaming visualization on board of an aerial platform. The proposed technology has the potential to enormously benefit the situational awareness capability, integrating maps and information from various sensors on and off-board, including satellite imagery. Interactive analysis can be easily employed to identify targets and features, and compare with other sensors data or historic information.

In the last few months we have developed a prototype software able to perform real-time 2D map generation using a sUAS from the USU AggieAir[8] fleet equipped with a single RGB camera. The acquisition of the aerial data is in raw format from a set of scientific cameras assembled into a payload and stored on-board of a high-density storage unit such as M.2 flash. Performed by AggieAir custom payload software,[9] this capture process is a combination of software and hardware. In our software pipeline, we have used the OpenViSUS™ multi-resolution data format that allows interactive visualization of the resulting real-time map in streaming without need to transfer all of the data at full resolution to the user. Figure 7 shows the interactive visualization on a browser of the images acquired during a test flight on the Utah State University campus over a Google Maps satellite image layer. The test flight had a duration of about 10 minutes and we were able to view the entire 2D map interactively around 3 minutes after the drone landed and connected to a wireless network. The current prototype used in this test flight was equipped with a Raspberry Pi 4 with a Picam v2 installed with the ViSOAR stitching software, enabling the generation of the entire map while flying. Preliminary experiments on a better computing unit (i.e., Intel NUC) show that we could achieve the same results even faster, saving computational power for further analysis, such as multi-sensor features cross-correlation and inclusion. The developed prototype represents a good proxy for more powerful platforms. We are currently developing interfaces for our software to interoperate in a real-time communication middleware such as ROS2 (Robot Operative System).[10] The use of standard frameworks for real-time operations together with a set of simulators and field testing will guarantee a robust, easy to integrate and interoperable solution.

4. LESSONS LEARNED

We have learned a lot through our field trials on the day to day of using our software and understand how it *actually* benefits crop consultants. Since we have been riding with crop consultants and have a first hand experience on their operations, we have gathered an in-depth knowledge of many areas where visualization technology can play a critical role in the crop growing business. One of the main lessons learned is there is not "a one size fits all solution" visualization software answer to the agricultural sector. Another important lesson has been learning about the different "players" in the industry and how visualization could disrupt or advance existing business services. Particularly with extensive agricultural practices where crop consultants recommendations have a significant impact on yield and profit margins.

4.1 Acquisition

Acquiring data is time consuming. A typical flight requires 10 minutes of hardware and software preparation time per field. Capturing a field will often require multiple battery changes, cooling periods, and data movement off the card that resides in the drone. Another big issue for us was getting data saved and off drone due in part to how long it takes. If there was a request for a fix in this area it would be either bigger batteries or a drone with more flight time and an streamline method for moving data off the drone, including a way to consistently name datasets with field names and dates.

During the 2019 season, we experienced a wealth of problems mostly centered around the hardware, data acquisition, weather, battery life, and management of the drone pilot. We bought a DJI Phantom 4 and Parrot Bluegrass drone with a built in Sequoia sensor. We had pretty good luck with the DJI Phantom 4, but we could not get the Parrot drone to reliably record the data. We returned the drone twice under warranty, but never successfully got the drone to record data.

Data acquisition also requires flexibility due to weather delays. Weather is always an issue, especially wind which can change fast and result in altering the altitude of the drone 2-5 feet and reducing the battery life of the drone. In practice we were able to get 12-15 min of flight time on moderately windy days whereas we got about 25 min per battery without wind. For 57 acres, a flight took up to 30 minutes at 100 feet altitude with 45% overlap. The drone cannot be flown when the wind is more than 8m/sec or 18.6m/h and higher wind speeds often resulted in issues like the drone landing far from home position.

Crop consultants we worked with worked 6 days a week from 7am to 5pm. However different fields benefit from different weather conditions. For example, it is ideal to capture wheat fields in full sun with as little wind as possible between 10am-2pm because there are less shadows and the wheat shows up better under full sun. However, overcast skies are better for capturing flooded rice fields in order to remove side effects of the glare of the sun.

During the 2019 and 2020 seasons we also had two drones fall out of the sky, a problem that we believe may have been due to battery issues, although both batteries were sufficiently charged upon retrieval of the broken drones. Additionally flying drones requires quite a bit of duplicity of batteries, we keep an average of six charged batteries per pilot to keep our drones flying for a day in the field. A typical drone battery lasts for 12 to 20 minutes flying at 100 feet altitude above the ground level with ground resolution of 3.7 cm/px with 45% overlap. The charge time for the drone battery is about 2 hours 30 minutes. In the 2020 season, we found flying at 200ft with an image overlap of 70% worked best with our software.

One of the big surprises in our journey is that it is easier to grow your own drone pilots, and it is easier to have a domain expert (in our case a crop scout) who is additionally trained as a drone pilot then it is to take an existing drone pilot and train them in the domain. Our crop consultant collaborator calls this the problem of detail oriented versus task oriented. In agriculture, the task is to gain an understanding of what decision and recommendation need to be made and given to the farmer at the end of the day. In our experience, we found that drone pilots trained in college programs were too worried about details of flying, the drone, the sensors, and focused less on why they were flying and what information they needed at the end of the day. Additionally it is difficult to find someone who can work full time March to August, as this does not conform to any academic or other professional timelines.

Data management becomes a huge problem from just noting the time, duration, weather, and problems with the hardware/software, as well as the need for a well defined filing system for naming the data fields such that the fields are named with the same string over the whole season. Over the 2020 season, we gathered more than 8.6 GB of data. A real solution to this problem will require a single software solution that does flight planning and logging of this information with a seamless, low burden, approach. It has also been difficult to document where the imagery helped because the daily grind and the business of the crop consultant. A lot of the memorable moments or actionable information was used as part of the process but was not ever noted. It is much easier to let software developers know when software fails, than it is to let them know when it worked. The only stories we got on successes came long after the fact. For example as shown in Figure 8, we were told when either the software or the hardware was to blame for discoloration in the imagery or when the files were missing in a stitched image. However we only recently learned about the time that an aerial survey and ViSOAR Ag Explorer software was used to inform the farm owner of irrigation issues (Figure 9).

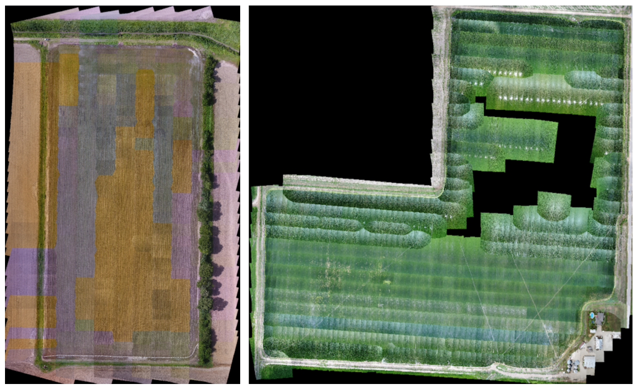

Figure 8. Problems noticed with fast stitching in field: Left) Discoloration when the blue sensor was not working properly. Right) Missing pictures due to UAV image capture issues.

4.2 Storage

Moving data is one of the biggest bottlenecks in acquiring aerial imagery. Data has to be moved off the drone, onto the laptop or cloud for processing, potentially relocated for sharing, and stored for long term backup. Towards next season, we are working on rclone or Google Drive scripts to put data up to cloud. Additionally, we rely heavily on Google Drive, created an account to be used by our drone pilots for everything so we had the data when our drone pilots left.

4.3 Computation/stitching lessons learned

Our original proposal for ViSOAR Ag Explorer software assumed farmers need the information right there in the field at time of capture. In our field studies with one crop consultant in Arkansas, it turns out that this argument is an exception, not the rule. In practice, our crop scout who was flying the drone did not actually stitch data in the field, except on rare occasions, when called upon by the crop consultant to examine major issues like irrigation. Typically our crop scout drone pilot would optimize to use good weather days to just fly and capture as many fields as possible, and then spend a day or two stitching the data when the weather was not ideal.

For example, the drone was used to scan a field looking for dry spots in a multiple inlet irrigation system that was not complete, i.e. the water was not flowing to the end of the furrows of soil filled with water, shown in Figure 9. The area was flown over, the images were stitched, and crop consultants were able to meet with the farmer at the end of the day to go over the data on the field laptop, identifying tire ruts that were blocking the flow in lots of rows. Without aerial imagery, such a task would require walking each row manually, resulting in hours of labor exchanged for one 20 minute flight.

Our field studies have shown us that we need flexibility in our approach for handling data processing and analytics. We need to expand our software to handle batch processing of scanned field data, providing the ability to stitch them offline and ready the data for loading in the desktop application or distributed via our web-based data portal for sharing the data. Additionally, we need to minimize the number of times we copy the data or move the data. Currently our crop scout drone pilot clears the camera cards after each flight, copying the data to the laptop (5-10 min after every field). One limiting factor is whether we have to wait for an internet connection

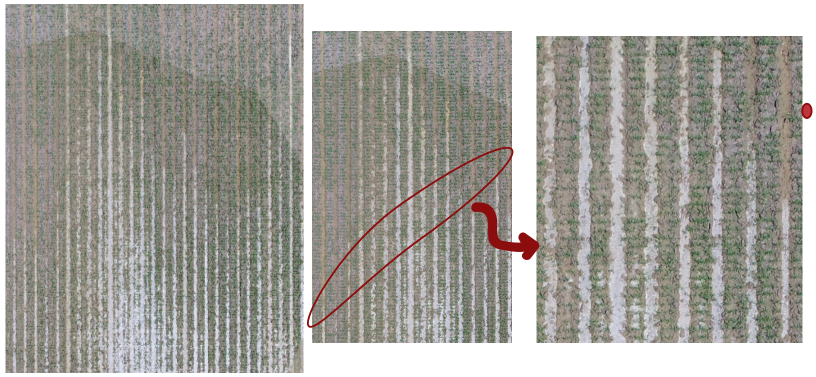

Figure 9. A success story of our aerial acquisition: our drone was used to scan a field looking for dry spots in a multiple inlet irrigation system that was not completing, i.e. the water was not flowing to the end of the furrows of soil filled with water. The area was quickly flown over, the images were stitched but not blended, and crop consultants were able to meet with the farmer at the end of the day to go over the data on the field laptop, identifying tire ruts that were blocking the flow in lots of rows. The tire rut starts at the arrow on the left side of the last image and ends at the dot on the right. Without aerial imagery, such a task would require walking each row manually, resulting in hours of labor exchanged for one 20 minute flight.

at the office, because copying to the cloud in the field may be prohibitive (remote access, cellular not internet). Other options include having enough memory cards to last all day and then move the data later.

Ultimately, we believe that on-board computing is critical for streamlining the process.

4.4 Viewing

Sharing of the newly collected data is one of the biggest bottlenecks in scanning of aerial data. We made a button to email screenshots from the crop scout to the crop consultant but this loses the ability to see the data from the macro to the micro. Our ViSOAR Data Portal software solution provides a client-server web-embedded web widget for interactive multi-layer visualization. Solutions allowing interactive exploration of large imagery exist, but often limited with slow reading or writing of data, especially if high fidelity information is required. Old solutions often use a pyramid of image tiles for multi-resolution visualization, inevitably including data replication, increasing the data storage requirements, and ultimately limiting the visualization to 2D datasets. During several years of research and work with domain scientists we have developed a hierarchical multi-resolution data format, IDX, and the OpenViSUS™ software framework to efficiently access and stream very large 2D and 3D datasets for interactive analysis and visualization. We are working on a solution that uploads the newly in-field stitched aerial data to the cloud when the laptop is parked at the office with an internet connection and emails the crop consultant the ViSOAR Data Portal link for viewing the next day. Our in field stitching puts immediate viewing of data in the hands of the crop scout who passes on information to the crop consultant, and our data portal will allow the crop consultant to see macro and micro details in the field to assist with long term planning.

4.5 Analysis

As a multidisciplinary collaboration, it was difficult for the computer scientists building the software to understand what crop consultants need. It is about more than just "Look at the imagery" in order to find the detail (irrigation, etc.). The highest priorities include being able to perform stand counts and being able to write a variable rate prescription for inputs. These will be critical additions to add to our software over the next season.

4.6 Plan

Adding aerial image capture to the traditional crop consultant model adds more logistics and may be a net neutral in terms of benefit. We believe it adds a positive in terms of the ability to document and inform discoveries, but at an added cost of time, labor, decision making and planning. It creates added value but at the cost of added complications (management of personnel, tasks, hardware, software). The Crop Scout/Drone operator needs to be able to work independently, deciding what fields need to be flown, and to be able to reliably handle the whole pipeline from documentation of the parameters of the flight, the data acquisition and storage, as well as sharing and creation of actionable data.

4.7 General

There is no doubt in our experience that visualization will be a disruptive technology but the key question is where the cost of the services will be accrued. To a large extent it will depend on each type of crop and the size of each operation. In large farms operations, 10,000 acres and larger, growers will incorporate the technology in-house. For growers with less than 10,000 acres that outsource crop management to scout consultants, the adoption of the technology will largely depend on the scout business model. For those scouts with few clients and acres there is very little incentive to switch to a visualization technology model. Our experience tells us that visualization is only attractive when a scout consultant is trying to expand their practice beyond 25,000 acres that are not adjacent to each other. Any scout consulting practice beyond that many acres tends to increase the operational challenges and compromise the quality of the consultant recommendation. Still how visualization could be embedded into the consulting practices largely will be dependent on the type of crops and farms sizes. We foresee our software solution able to tap into this segment of the market where customization of drone technology and scout consultants operation management software could be easily integrated with the ViSUS platform.

Additionally, it is worth considering the future of agriculture. As younger generations of farmers take over land holdings, how can we help make aerial imagery more effective, easier to acquire, and more efficient at providing information. "...young farmers are also very similar to their non-farming peers: tech-savvy, purpose-driven, and entrepreneurial".[11] However, evaluating costs is a big part of the optimization: "Even if they want to try a few autonomous technologies, these new farmers can't afford them." Although many believe that "...autonomous technologies can drastically improve farmers' quality of life," it still remains to be proven that aerial imagery is a major disruptor in agriculture. We thought we would be able to strongly prove this for at least one crop consultant over the growing seasons of 2019 and 2020, but we learned that we are still a ways from making such a definitive observation.

5. CONCLUSION

As we conduct field trials, flying over 1500-2000 acres per week in Arkansas, we seek to provide a much broader scope of the field in general as well as providing a better view into planting and field preparation than could be observed from field level. Later in the season the imagery can provide a much better understanding of weed presence and intensity or lack thereof. Our platform needs to be optimized to adequately handling data movement and needs improvements in providing actionable information. Our software, ViSOAR Ag Explorer™ removes the software bottleneck (time and labor) in processing large aerial surveys, enabling a cost-effective process to deliver actionable information to farmers. However, our explorations of on-board computing with edge processors have shown that such a hardware and software configuration is critical for reducing the time to acquire aerial imagery. Such an easy to use hardware and software configuration has major potential to allow farmers and farm consultants to optimize the use of fertilizers and more finely control the amount of pesticides and herbicides necessary to increase crop yield. Furthermore, farmers mitigate costs and losses by being able to spot problem areas, minimize the spread of plant diseases, and identify issues such as standing water, irrigation malfunctions, and persistent automated machinery errors in planting or cultivation. The technology proposed by ViSUS, Utah State University, and the University of Utah in collaboration with Crop Solutions LLC is part of a broad initiative in agriculture addressing the need for a 70% increase in food production by 2050 in response to the projected growth of the world's population.

ACKNOWLEDGMENTS

This work was funded in part by NSF Innovation I-Corps award 1542353, NSF SBIR awards 1549187 and 1738448, and AggieAir's work is funded by Utah MLF 2006-2021.

REFERENCES

[1] "Openvisus." http://www.visus.org (2021).

[2] Venkat, A., Christensen, C., Gyulassy, A., Summa, B., Federer, F., Angelucci, A., and Pascucci, V., "A scalable cyberinfrastructure for interactive visualization of terascale microscopy data," in [*2016 New York Scientific Data Summit (NYSDS)*], 1–7 (Aug 2016).

[3] Kumar, S., Vishwanath, V., Carns, P., Summa, B., Scorzelli, G., Pascucci, V., Ross, R., Chen, J., Kolla, H., and Grout, R., "Pidx: Efficient parallel i/o for multi-resolution multi-dimensional scientific datasets," in [*Proceedings of IEEE Cluster 2011*], (September 2011).

[4] Kumar, S., Vishwanath, V., Carns, P., Summa, B., Scorzelli, G., Pascucci, V., Ross, R., Chen, J., Kolla, H., and Grout, R., "PIDX: Efficient parallel I/O for multi-resolution multi-dimensional scientific datasets," in [*IEEE International Conference on Cluster Computing*], (2011).

[5] Kumar, S., Vishwanath, V., Carns, P., Levine, J., Latham, R., Scorzelli, G., Kolla, H., Grout, R., Ross, R., Papka, M., Chen, J., and Pascucci, V., "Efficient data restructuring and aggregation for I/O acceleration in PIDX," in [*High Performance Computing, Networking, Storage and Analysis (SC), 2012 International Conference for*], 1–11 (Nov 2012).

[6] Pascucci, V., Scorzelli, G., Summa, B., Bremer, P.-T., Gyulassy, A., Christensen, C., Philip, S., and Kumar, S., "The visus visualization framework," in [*High Performance Visualization: Enabling Extreme-Scale Scientific Insight*], Bethel, E. W., Childs, H., and Hansen, C., eds., ch. 19, 401–414, Chapman \& Hall/CRC Computational Science (2012).

[7] Summa, B., Scorzelli, G., Jiang, M., Bremer, P.-T., and Pascucci, V., "Interactive editing of massive imagery made simple: Turning atlanta into atlantis," *ACM Trans. Graph.* **30**, 7:1–7:13 (April 2011).

[8] "AggieAir Lab at USU." http://aggieair.usu.edu/ (2021).

[9] Calvin Coopmans, James Humble, and Matthew Givens, "Software requirements, architecture, design, and results of a modular, multispectral, scientific quality unmanned aerial system payload for civilian remote sensing," in [*Proc.SPIE*], **11008** (May 2019).

[10] "Robot operating system (ros)." https://index.ros.org/doc/ros2/ (2021).

[11] Lawrence, M., "Big bots in little agriculture." https://slate.com/technology/2012/06/automated-farm-equipment-and-small-scale-farmers.html (June 1, 2012).