Word Count: 1440

Light Pollution Mapping of Pittsburgh using Cessna Imagery

Brian Zhang, *Carnegie Mellon University School of Computer Science*
John Wenkovitch, *Virginia Polytechnic Institute and State University*
Diane Turnshek, *Carnegie Mellon University*

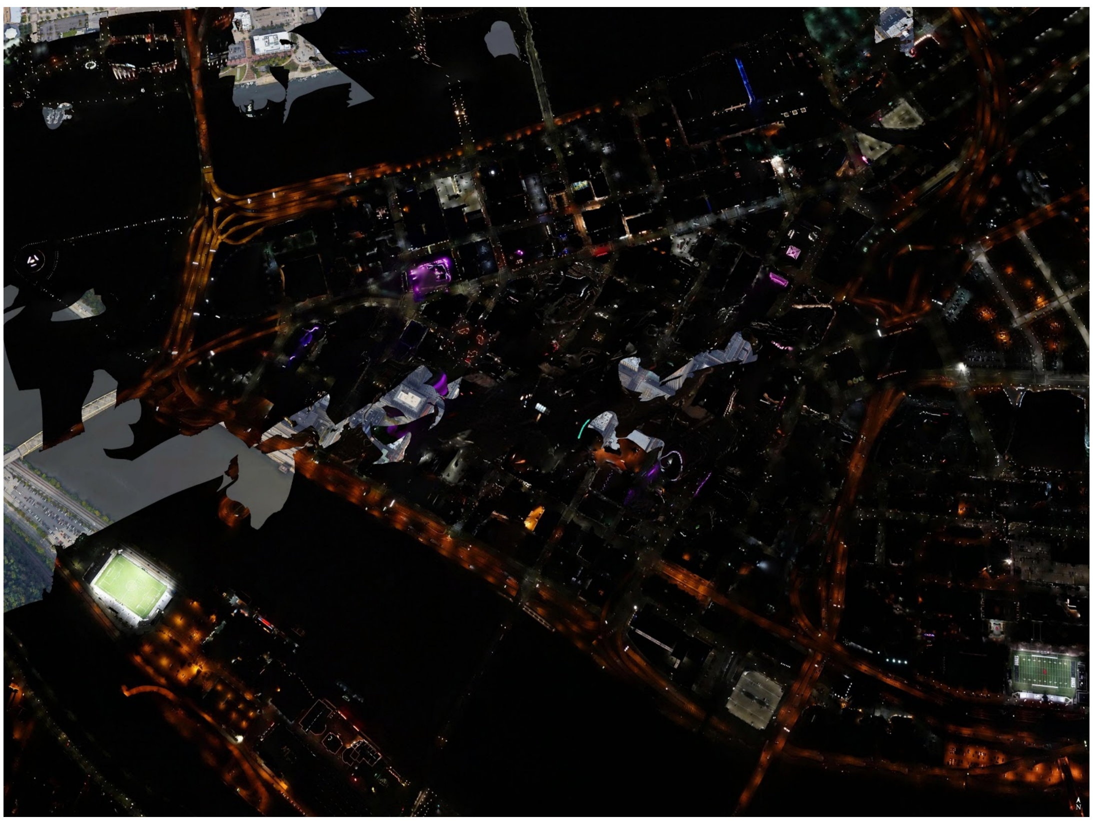

Generated light pollution map of Pittsburgh at night

Brian Zhang
bez@andrew.cmu.edu

1. Project Information

Over the course of the past century, the rapid urbanization of cities has resulted in an exponential rise of light pollution as a result of streetlights and other artificial light fixtures. These sources of artificial light at night (ALAN) have been shown to have deleterious effects on astronomy, natural ecosystems, and even human health (Cho et al., 2015). As a result of this, more and more people have been beginning to rethink the way we illuminate our cities at night with an emphasis on minimizing the extent of light trespass, skyglow, and other unwanted and excess light.

In an effort to reduce ALAN in 2014, the City of Pittsburgh has already begun plans for phasing out their 40,000 sodium vapor streetlights for more modern, energy efficient LED bulbs (Sanserino, 2014). Despite these many benefits of LED lights over sodium vapor, it must be noted that an irresponsible undertaking can exacerbate the issue of light pollution as LED lights have a much higher color temperature, and thereby more light pollution, than sodium vapor bulbs. On the other hand, an environmentally conscious mindset which takes advantage of the controllability of these smart LEDs to adjust the intensity and temperature of the lights based on time or motion sensors can be extremely effective at minimizing ALAN (Smart Streetlights, n.d.).

Our project hopes to create a light pollution map of Pittsburgh prior to and after the replacement of streetlights with LEDs. This allows us to understand how light pollution was impacted in Pittsburgh by this switch to LED streetlights and potentially shed light on what can be done better in both Pittsburgh and other cities to successfully develop a light pollution conscious streetlight infrastructure. As of this year, Pittsburgh has only installed LED bulbs in a

Brian Zhang
bez@andrew.cmu.edu

small percentage of its neighborhoods due to significant delays in the program. Thus, this paper only describes work done on light pollution in Pittsburgh prior to the conversion.

For this project, we used night-time photos taken of Pittsburgh from a Cessna flight by Sebastian Scherer of Carnegie Mellon University's Air Lab. The flight took place on October 29, 2019 beginning at 11:32 pm and lasting roughly 40 minutes, taking off from Allegheny Airport. The flight covers most of downtown Pittsburgh and attempts to cover the entire city via following a zig-zag like pattern. The aircraft cruised at an altitude of around 750 - 825 meters for the majority of the flight.

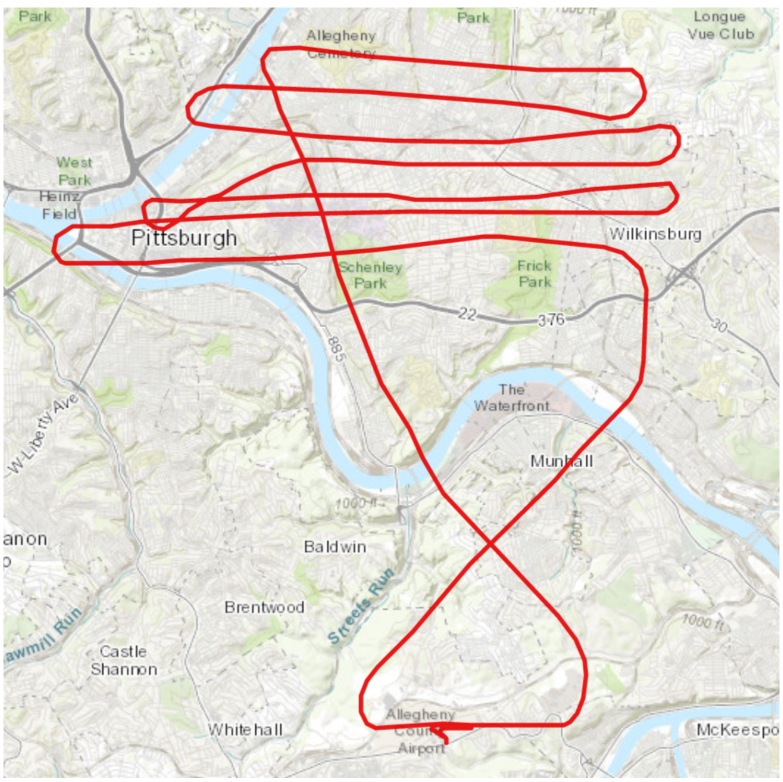

Fig. 1 Path of Cessna flight over Pittsburgh taking off from Allegheny Airport

2. Internship Role and Experience

My role in this project was in creating the actual light pollution map from the Cessna photos. There were ~2800 photos along with some tracking data which was in total around 22GBs. Our team had a subscription for the DroneDeploy service which uses computer vision to stitch together aerial images and generate 3d models and maps of a region. As such, much of my time as an intern was spent learning to work with the DroneDeploy software and creating an image processing pipeline for our raw images into a more DroneDeploy friendly input.

During the start of the project, I was waiting for the Cessna photos to download and read up on DroneDeploy. I expected the map generation to be a very simple process since DroneDeploy was handling the complicated computer vision algorithms. I quickly realized this to not be the case when, after the download finished, I uploaded all of the images to DroneDeploy and the map generation failed after around 8 hours of waiting. DroneDeploy complained that the images we were uploading were not the best quality and lacked a lot of the geopositional metadata (longitude/latitude/altitude) necessary for processing.

“Good” Photos	“Bad” Photos

Fig. 2 Sample raw Cessna photos.

After that attempt, I filtered out the “bad” photos that lacked the necessary metadata and the images that were excessively blurry, did not depict much, or at a low altitude. The remaining 229 photos had minimal blur, clearly visible structures, and were optimal for DroneDeploy according to their FAQ. However, this time the DroneDeploy processing failed for a new reason: low coverage. By filtering out so many images, there simply were not enough photos and insufficient overlap to reconstruct all but two disappointingly tiny portions of Pittsburgh despite the range of photos spanning the entirety of the city.

Fig. 3 The two dark regions was all DroneDeploy was able to generate with the images. The blue dots represent geolocation of photos used for rendering.

After some more failed attempts, I began looking at the GPS tracking data from the Cessna I had to see if I could do anything with it. It consisted of longitude, latitude, altitude, and other readings taken every second or so from the plane. Additionally, I found that while many images lacked GIS metadata, they all had information about the date/time captured. I initially overlooked this \`date\time captured\` tag due to the times not matching up to the actual flight times. Nevertheless, I eventually realized that the \`time captured\` tag was offset by exactly 4 hours from the actual flight times (i.e. a photo with the \`time captured\` tag of 6:30pm was actually taken at 11:30pm). I suspect that this was a result of the timezone settings of the camera.

With this, I was able to corroborate the \`time captured\` tag of each photo with the GPS tracking data to determine the exact longitude, latitude, and altitude data for that image and update its metadata for DroneDeploy. I converted the GPX tracklog file into a CSV spreadsheet that read into a python script which automated this process for all of the images. I struggled some with working with datetime objects, and converting from decimal coordinates to hours/minutes/seconds but in the end it was not too difficult. Additionally, the python script also compressed the 1607 images, with no noticeable loss in quality, to just 1.5GBs which was a much more manageable size to work with.

Finally, after all of this, DroneDeploy was able to generate a more complete light pollution map of Pittsburgh. There are still some significant gaps in the map due to the imperfect Cessna flight.

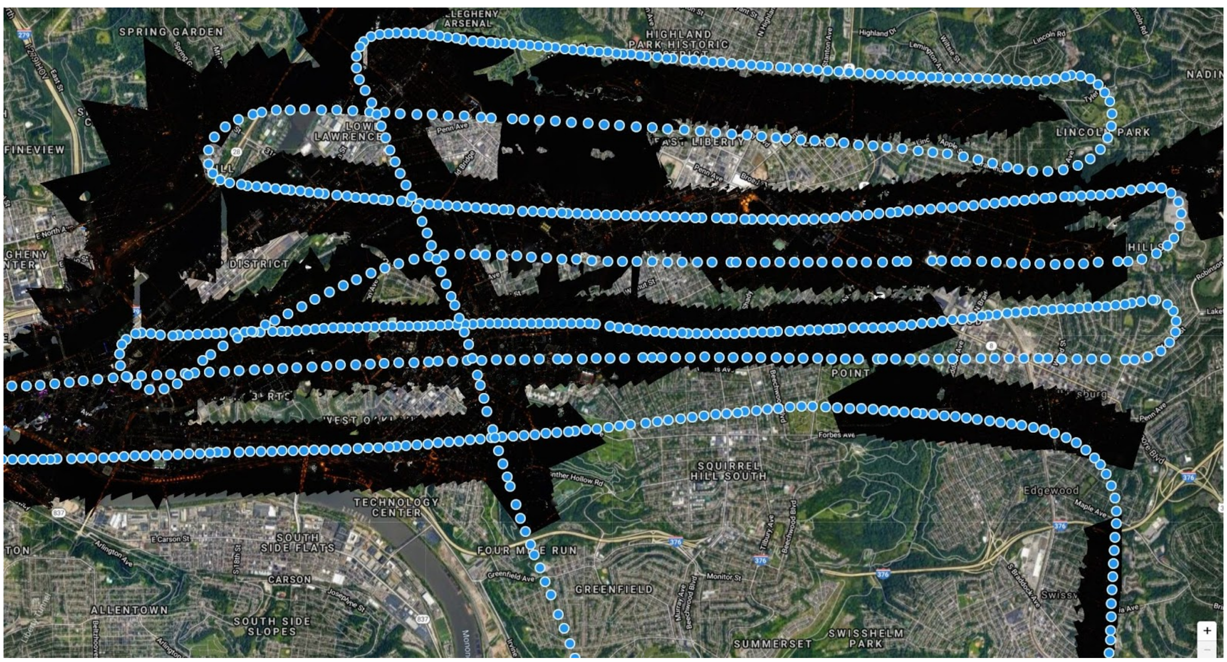

Fig. 4 DroneDeploy map generated from using python processed imageset. The blue dots represent geolocation of photos used for rendering.

After some minor adjustments and tweaks, we have the final map with an astonishing max resolution of 5.68in/px. However, DroneDeploy limited me to exporting the map files as a

Geotiff at a resolution of 20in/px due to file size limitations. Using Google Earth, I was able to generate a few more visualizations (See Appendix).

The python code for preprocessing the photos can be found here.

3. Reflection

The project was pretty interesting and exciting for me as I have always been interested in working with computer vision. I was astonished by how useful and versatile DroneDeploy is and getting to play around with it was really cool. I initially imagined this project to involve a lot of complicated math, matrix rotations, and ML for imagery but it turns out that DroneDeploy took care of all that for me. I was both relieved that the project wasn't going to be as difficult as I imagined it to be and disappointed as I thought those concepts would be fun to learn. I still got to learn a lot about what kind of amazing work was possible with computer vision through working with DroneDeploy.

Both Mr. Wenkovitch and Mrs. Turnshek gave me almost complete freedom on my task. I had autonomy to decide for myself how to approach the problem they tasked me with and the time to experiment, test, fail, and eventually succeed in generating my map. Outside of our first meeting which was greatly beneficial for me in understanding the goals of the project, we interacted relatively little outside of my periodic project updates.

All in all, I was able to successfully reconstruct a map of a large portion of Pittsburgh. However, there is still work to be done in analyzing the map to uncover geospatial light pollution levels, luminance, and more. Future research can explore the spread and types of streetlights as classified by their light emission on our map as well. Also, a parallel project still needs to be done in the future to map Pittsburgh after the transition to LEDs are complete.

4. Works Cited

Cho, Y., Ryu, S. H., Lee, B. R., Kim, K. H., Lee, E., & Choi, J. (2015). Effects of artificial light at night on human health: A literature review of observational and experimental studies applied to exposure assessment. *Chronobiology international*, *32*(9), 1294-1310.

Sanserino, M. (2014). Switching on: Installing LED streetlights to save Pennsylvania taxpayers money. Retrieved from https://www.post-gazette.com/business/powersource/2014/03/25/Switching-on-Installing-LED-streetlights-to-save-Pennsylvania-taxpayers-money/stories/201403250025

Smart Streetlights. (n.d.). Retrieved from http://smartpittsburgh.org/programs/smart-streetlights

5. Timeline

Date	Hours	Tasks Performed
7/22/2020	3	Preparation for meeting with Diane and John by reading up on various materials related to the project. Talked to Diane and John about the goals of the project and where we were at.
7/23/2020	3	Read up on DroneDeploy and made several attempts at generating maps after finally downloading GDrive Cessna image files. Failed due to lack of metadata and bad parameters (selected “structure” setting rather than “terrain” on DroneDeploy).
7/24/2020	2	More failed attempts at generating maps.
7/25/2020	2	Filtered bad images from dataset. More attempts to upload to DroneDeploy. First successful map generation. However, the map was only generated in two tiny portions of the city.
7/26	3	Created python script to add altitude data to images which lacked this metadata. Also compressed images to a workable size.
7/27	4	Converted GPX tracklog file into a csv file. Also added missing Long/Lat metadata to all images based on tracklog after adjusting for the time offset between the flight and \`date/time captured\` metadata on images.
7/28	2	Successful generation of Pittsburgh map with processed dataset. Attempted to fill some empty gaps by uploading images again. There were some marginal improvements but ultimately not worth it to continue to optimize map.
7/31	2	Exported final map as Geotiffs and imported into Google Maps to generate presentable images of final map for Diane. Also located the shapefile of Pittsburgh which was used to add borders around the city on the map.

Brian Zhang
bez@andrew.cmu.edu

6. Appendix

Fig. 5 Full DroneDeploy map of Pittsburgh with city limits in blue.

Brian Zhang
bez@andrew.cmu.edu

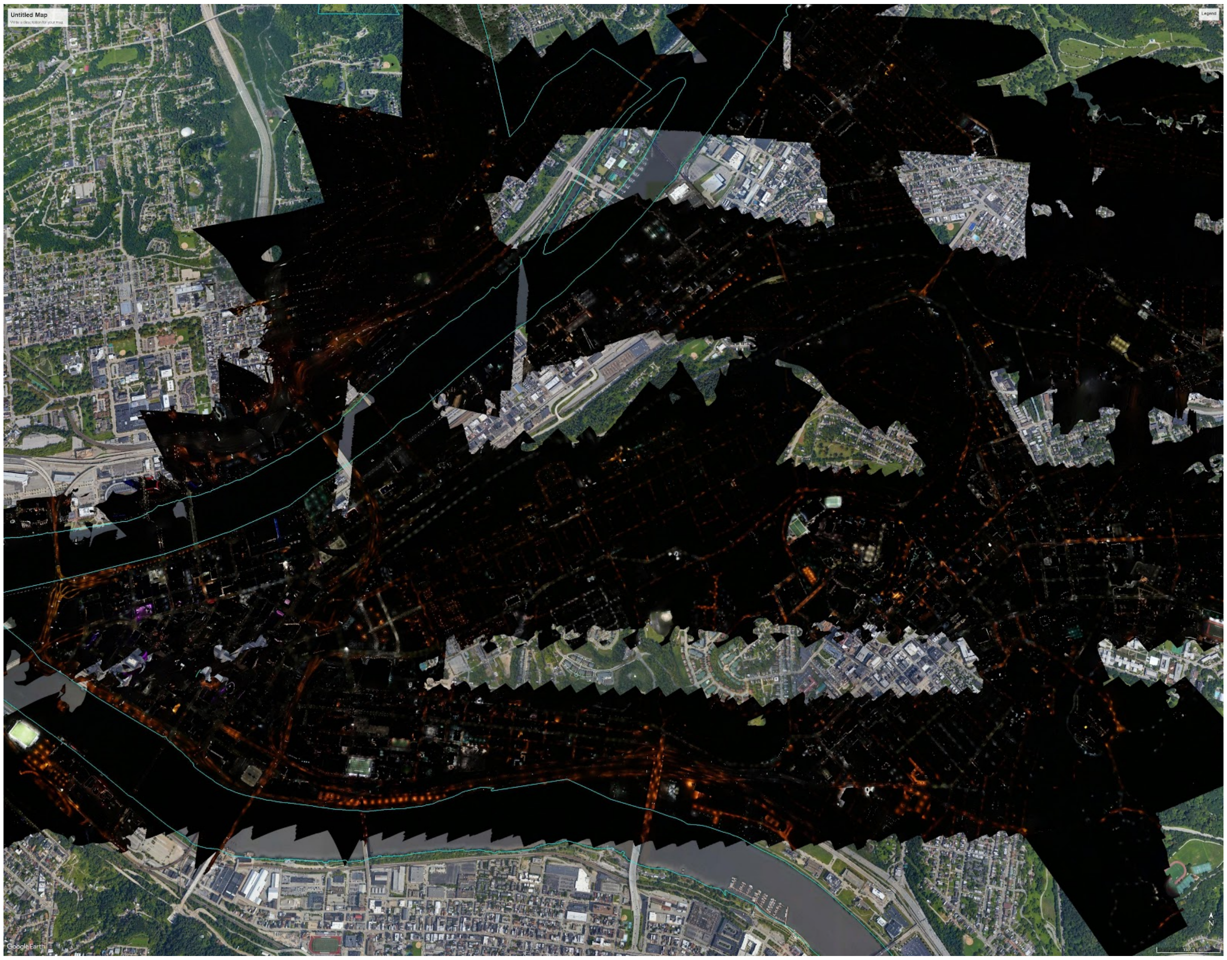

Fig. 6 DroneDeploy map of Pittsburgh downtown with city limits in blue.

Brian Zhang
bez@andrew.cmu.edu

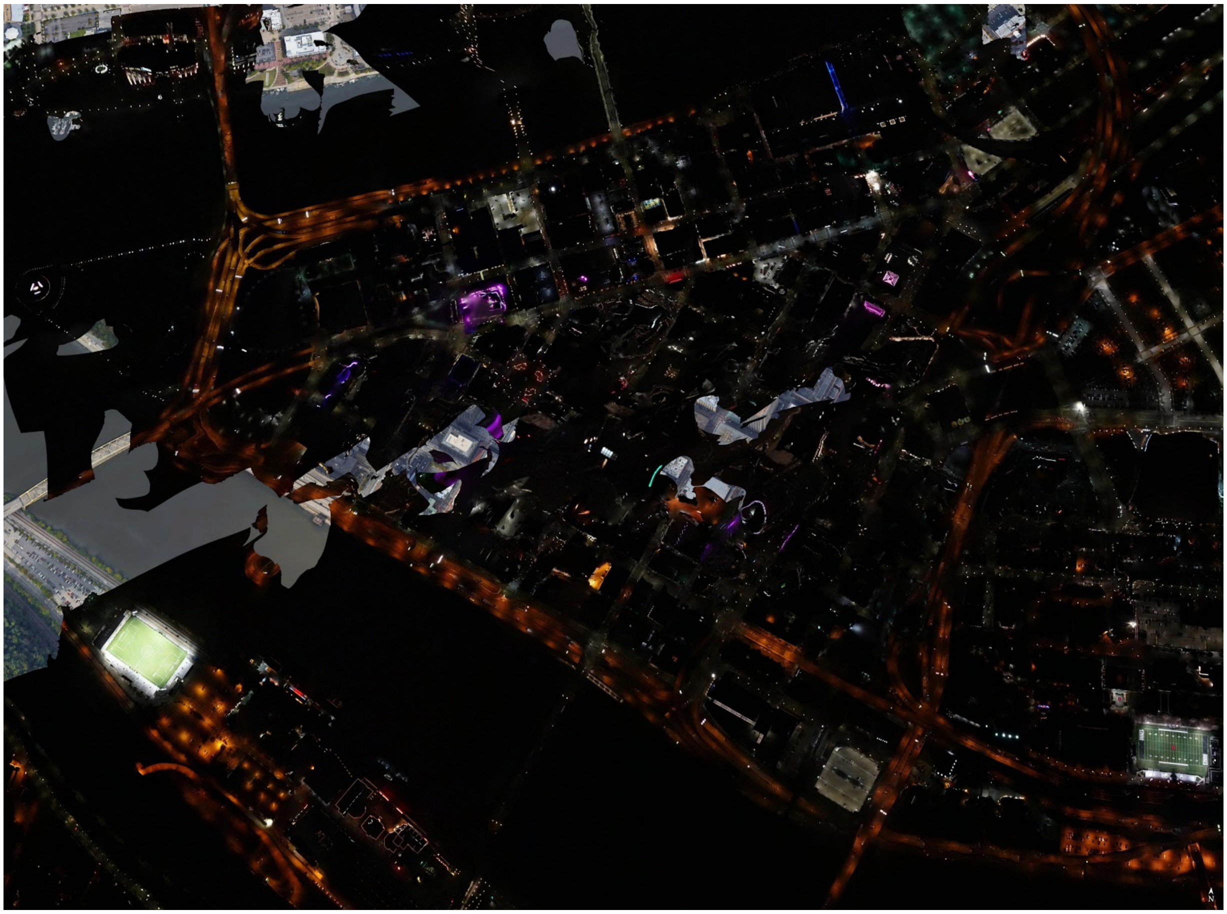

Fig. 7 DroneDeploy map of Pittsburgh downtown zoomed.

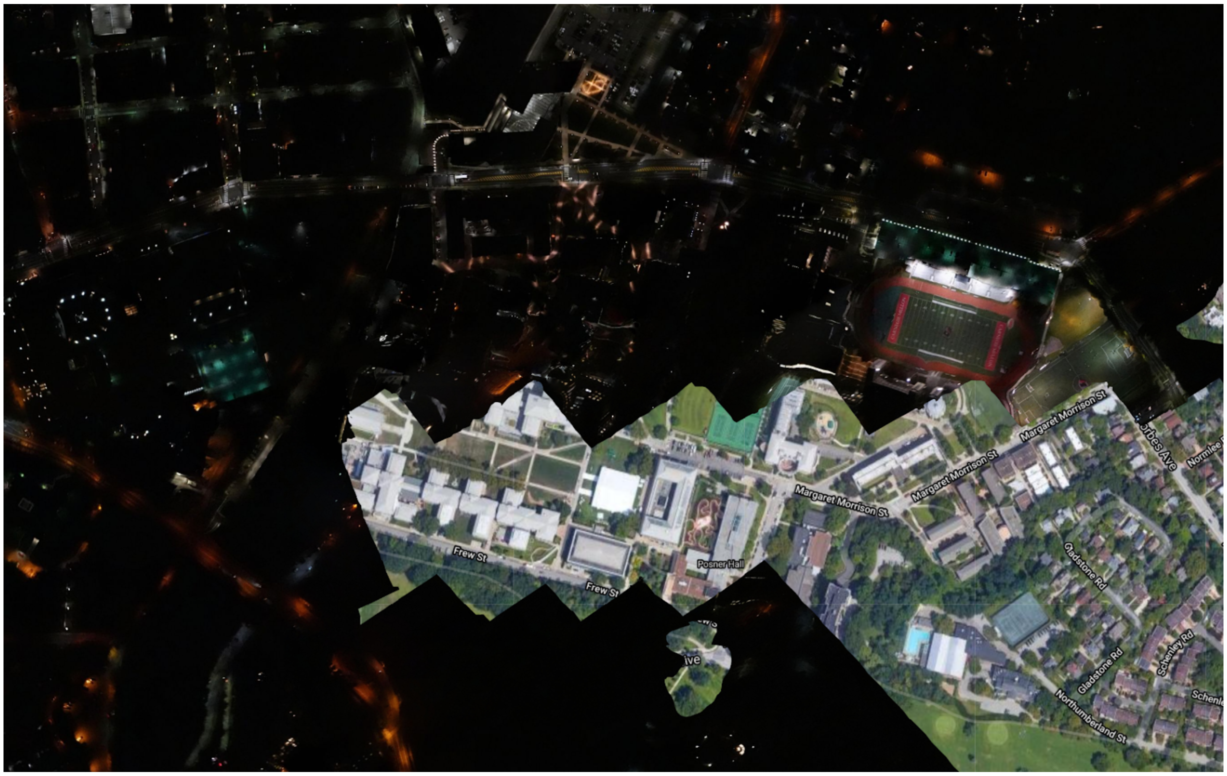

Fig. 8 DroneDeploy map of Carnegie Mellon University and surrounding areas. Much of the campus was not rendered due to lack of coverage.