

# THE CAVEATRON

## QUICK AND EASY 3D CAVE MAPPING FOR CAVERS

By Joe Mitchell



What if a device existed you could carry into a cave to make a 3D map of it that could be viewed on your computer or even printed as a 3D Model? What if it could also record traditional survey data with an accuracy similar to conventional instruments but far more quickly and repeatable? What if this device was handheld and designed specifically for cave survey, reasonably simple to the average caver surveyor and based on existing techniques? And what if it was affordable for the average caver?

Meet the Caveatron, a new electronic cave survey tool. You can use it to take your distance, azimuth and inclination measurements in a few seconds and then view your line plot. You can travel down a passage with it and scan the walls to produce a 3D point cloud. You can also scan rooms and alcoves by rotating it in place. Since it produces accurate and detailed measurements of the cave

walls, it can greatly reduce the amount of hand sketching required and the data can be used to produce maps with a level of precision previously unseen except in the most detailed of surveys. Handheld and relatively compact, it is designed to be easy and practical to use in a typical cave environment and operate for many hours at a time.

### Development

A few years ago, I started to get interested in how you might use electronics to improve cave mapping. I participated in the LIDAR survey at Devil's Sinkhole, which was impressive but the equipment was big, heavy, slow and very expensive. I also saw some small survey equipment cavers were making such as the DistoX and the Shetland Attack Pony, which were also very nice and can speed up a survey by taking shot data, but were not

exactly what I was looking for. The time consuming part is sketching, and I wanted to find a way to make it faster and more accurate. Rather than waiting for someone to make the cave survey device I wanted, I realized that my job as an electrical engineer meant that I had some of the knowledge and skills to do it myself. I got two of my fellow co-workers and cavers Steve Gutting and Gregg Williams involved since Steve has lots of soldering and electronics assembly experience and Gregg has extensive machining experience and could help with mechanical design.

My first idea was a basic device that produced LRUDs (left, right, up, and down) measurements continuously along a passage. Since a laser solution seemed too big and costly at the time, I thought ultrasonics might be the way to go. I had worked with some small ultrasonic sensors previously that were relatively inexpensive, so I tried them out in a cave and they seemed to work well. The idea for the first prototype of the Caveatron was to use four of these to get the LRUDs and to somehow figure out the position of the ultrasonic sensors as you move. Getting compass angle and tilt was not too hard since small electronic modules with a magnetometer and accelerometer were not too hard to come by. Distance was another matter. Laser distos (rangefinders) would be the best solution, but they are almost all made with proprietary interfaces and there is no real way to communicate with them (except for some very expensive industrial Models). As an alternative, I came up with the idea of using a remote transponder with its own ultrasonic emitter that sent a pulse to a receiver on the main unit. An infrared LED in the main unit was used to send a signal to activate the transponder's ultrasonic emitter. The time from sending the IR pulse until receiving the ultrasonic pulse gave a pretty accurate distance measurement (to a centimeter or so) at up to 15 m away.

The first prototype had a very basic interface with a numeric membrane keypad and a 2-line text LCD. It stored the LRUD data on a built-in SD card. After a few cave tests, it became obvious that the Caveatron had a lot of potential, including being used to take the regular distance, azimuth and inclination shot data. It also seemed that if it was going to do all these things, a better way of controlling it was needed. I came across some low cost touchscreen displays which would not only allow me to make whatever user interface I needed, but also allow data to be displayed to the user instead of just being stored. So I set out to build the first complete prototype including the display, the LRUD ultrasonic sensors, and added a laser pointer to accurately align it for taking shots and LRUDs. This prototype ran on 4 AA batteries. During testing, it worked

well and was it pretty cool to watch it build up the LRUDs as you walked through the cave. This version was tested in several caves and presented at the 2013 TSA convention. However, with further testing some problems became apparent. First the compass calibration was erratic. It would stay relatively good for a time but then suddenly change for no apparently reason (found to occur when replacing the batteries). Another problem, which became apparent on a longer trip in Deep Cave, was that after a few hours the ultrasonic sensors would start to behave erratically and eventually just quit altogether. After coming out of the cave, they would then go back to normal. I concluded that humidity was the cause, which was a major problem since it would only allow use for short periods.

About this time, I came across a low cost LIDAR scanner that was used in robotic vacuum cleaners and discovered that some people had figured out how to hack it and use it for other purposes. This LIDAR scans a full 360 degrees at 1 point per degree at 4.5 rotations per second. This sounded very interesting, so I bought one to try it out. Sure enough it seemed to work, so we built a new prototype based on previous version – still using ultrasonics for distance and a transponder – to test out the LIDAR in a cave. We tried it out in Robber Baron and after looking at the data, I was blown away by how great the data looked. You could view the cave in 3D and when looked at from inside the cloud, it was very recognizable! At this point in time it was only a point cloud, so still crude, but after working out more sophisticated post-processing of the data to get rid of noisy points and add surface normal, I discovered that it was fairly simple to render this point cloud into an actual surface to make solid walls and get a truly realistic looking 3D Model that you could “explore” both inside and out. This was far better than what I had originally set out to do and nothing else like this existed at such a low cost and small size.

The final piece of the puzzle was finding a laser rangefinder to replace the ultrasonic transponder. I came across a blog where someone had found a disto with a hidden serial port that seemed to work. I ordered one and sure enough, it was possible to command it and get distance measurements back over that connection. It worked well and reliably acquired distance data out to more than 45 m at up to 3 measurements per second.

One issue with incorporating the LIDAR is that it generates a lot of data. The original processor I was using (an Arduino Mega 2560) only had a 16 MHz 8-bit processor with 8kB of memory and 256kB of code storage. This was clearly not enough since LIDAR data was being lost and the touchscreen I was using exhibited slow

refreshes. Since I had become familiar with the Arduino platform, I wanted to stick with it and fortunately they had just come out with a much faster design, the Arduino Due with an 84 MHz 32-bit ARM processor with 96kB of RAM and 512kB of code storage. The downside was that it used a different voltage and the different type of processor would require modifying portions of the code. Between all the new components, and the increasing complexity of the software, I concluded that I needed to redesign the system and rewrite the code from the beginning to be more flexible and easier to work with. The user interface components were written into a separate library, the GUI code was redesigned to be interrupt based and the code was re-written to divide it up into smaller re-useable modules.

Since I was essentially starting over, I wanted to use what I learned from the early versions to make this one even better. First, switching to a built-in rechargeable battery was necessary to avoid the calibration challenges I encountered with replaceable batteries. A 4400 mAh Lithium-ion battery pack was selected to provide more than 8 hours of run time. The next piece was the touchscreen. The previous one had a low resolution and the two fonts it provided were too limiting for the GUI and hard to read. I found a new LCD that was slightly bigger with a brighter screen and with almost double the resolution (480 x 320 pixels) for about the same price. It also included an SD card slot that was more than 8x faster and a font chip with a wide range of sizes that were much more readable. It could also save images on the chip, so I could use it to store the on-screen button icons that Jill Orr had so generously created.

Other changes included a real-time clock with a small replaceable button cell that keeps time regardless of whether the Caveatron is powered on or off. A battery monitor board measures the charge of the Li-Ion pack and another board safety regulates charging of the battery. Other electronics provide the various voltages needed throughout the system. A piezoelectric buzzer is used to provide audio tones while taking shots and passage scans to reduce the need to look at the screen for feedback. About the only thing that didn't change was the compass module, although a newer version

was chosen with double the accelerometer resolution.

## Design

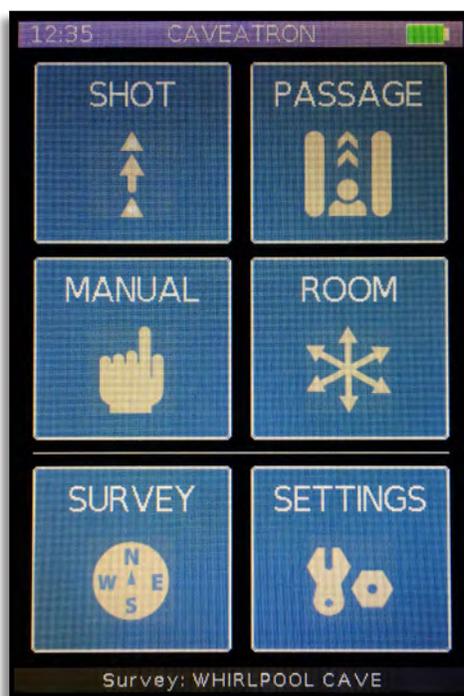
The Caveatron has two components: the main unit and the LIDAR module. On the right hand side of the main unit is the recessed On/Off push button switch with a red ring LED built into it to indicate when battery charging is in progress. On the left side is a capped mini-USB port that is used to charge the battery, connect to a computer to download data, and update the firmware. In the front is a small window for the laser rangefinder and in the rear are two rings for the neckstrap. On top is a metal rail for mounting the LIDAR module and a connector with a rubber cap for connecting the LIDAR module cable. Finally, the top has a cutout for the recessed LCD touchscreen.

The LIDAR module has two rail mounts on the base to attach it to the main unit, which are tightened by hand with small thumbscrews. Near the front are windows around all four sides through which the LIDAR scans. A curved window would have been nice to avoid having to have corners (which obstruct a small part of the scan area) but curved windows severely distorted the distance measurement.

Separating the LIDAR from the main unit was done for several reasons. Since the LIDAR unit is somewhat sizeable in diameter (about 4.5") and must protrude in front of the system to enable 360° scanning, the overall shape of the system is not particularly streamlined. By making the LIDAR removable, the system can be

made more compact for transport, or alternatively the Caveatron can be used without the LIDAR unit at all, making it a more compact conventional survey instrument. Finally, a more advanced (but far more expensive) LIDAR could be used if more range or resolution is required.

Both parts use off-the-shelf heavy-duty electronics enclosures made from ABS plastic with integrated O-rings to protect against dust and water penetration. Each of the connectors also has a sealing gasket and another is used around the touchscreen. The windows are glued into place with silicone sealant and, along with the touchscreen, are recessed to reduce the risk of scratches or impact damage. The touchscreen



Main Menu

is also covered with a screen protector. Although the Caveatron is probably not suitable for water caves, it should do fine in Moderately wet and muddy environments.

## Operation

Having surveyed caves for quite a few years, I wanted to make its operation familiar by using and building on existing survey techniques. The basis of a cave survey is the line plot, which is also used by the Caveatron as the skeleton on which the LIDAR point cloud scans are built. Just like conventional surveys, stations are set and the Caveatron is used to measure the distance, azimuth and inclination between them. All the LIDAR scans are then referenced to these stations. For instance, if you are doing a scan as you move down a passage (a traverse), you walk toward a particular station while the Caveatron continually measures its location off of that station. The coordinates of the point cloud for that section are based off the position of that station as determined from the line plot. Since only a single station is used to position each scan, you can make as many different traverses as you want to a particular station to build up a larger area or get additional coverage for areas that may be shadowed in particular positions.

One thing that is a bit different from a conventional survey is the use of a retroreflective card that is held on the station to which you are shooting. This was needed as a way to solve the problem of being sure that you were actually measuring the distance and direction to the station when doing a passage scan with the LIDAR. Since it is impossible to hold the Caveatron perfectly on station while traversing with it, the scan point wanders around quite a bit. The laser rangefinder has been modified so that it only gets valid readings when it hits the retroreflective card and all other surfaces are ignored. So, as you move down a passage with the Caveatron, you don't need to be pointed exactly at the station continuously, just from time to time. The Caveatron does the rest, filtering out off-angle measurements and interpolating between valid measurements. Having the card also helps with the accuracy of station-to-station shots in that it ensures you are actually pointed exactly at the station and not measuring something next to or behind it.

Each survey is stored in a different file, so at the



Shot Screen

beginning of a new survey, you enter a name for it (typically the cave name) and the date is automatically added to make it distinct. To take a shot, you enter SHOT Mode and enter in the names for the “from” and “to” stations (the “from” station is pre-filled from the previous “to” station but can be easily changed). At this point the red laser comes on so you can align the Caveatron to the station you are shooting to. You have the choice on the screen of selecting which of the bottom corners to measure from, depending on which is more convenient. You can also hold it at an angle or upside down if your station is on the ceiling, and this is accounted for in the software. A large

button covering most of the screen is used to activate the shot and an audible tone indicates when the shot has started and finished. A shot takes about 3 seconds, during which 3 distance, measurements and 900 azimuth and inclination measurements are taken and averaged. At the end of the shot, the measurements are shown on the screen and the shot can be accepted or redone. If a problem occurred, such as the Caveatron moved during the measurement, the shot will fail which is indicated by a special low-pitched tone.

PASSAGE Mode is the main method for taking LIDAR scans. It is done by gradually traversing the Caveatron (by walking, crawling, climbing, etc) toward the reference station. As with SHOT Mode, PASSAGE Mode starts by entering the station to which the traverse will be made (again auto-filled from the last shot) and a traverse number. After pointing the laser at the station and getting a starting position, the LIDAR scan begins. As you move down the Passage, you attempt to keep the laser pointed at the card. It's not required to constantly hit the card as that would be nearly impossible while moving, but just to hit it periodically and not stray the laser pointing angle too far away. The Caveatron beeps whenever it gets a valid position reference off of the card, so you can tell when you hit it. If you move too fast either forward or backward, too much time passes between valid measurements, or if the angle strays too far from the card, the Caveatron will stop the traverse. When the traverse is finished, colored statistics are shown in the screen to tell you if the traverse data is good (green), so-so (yellow) or not acceptable (red). You can then chose to do another

traverse, redo the previous traverse or return to the main screen.

Although a fair amount of processing has been incorporated to filter out poor data, account for the jostling and motion that accompanies a traverse, and protect against blunders, the quality of a LIDAR scan does depend on some operator skill. Without trying too hard you can get at least reasonable data but it may have accuracy issues or be somewhat sparse in coverage. However, with more careful use, you can get higher precision and more detailed coverage. Through practice, I have found that you can learn to hold it more carefully to improve accuracy even while moving through rough terrain, travel at a good smooth pace to increase scan density, and hold the Caveatron in such a way as to get the best coverage.

The last survey mode is ROOM Mode which is used to capture LIDAR scans of areas that PASSAGE Mode might miss, such as alcoves, passage corners or ends, or rooms that don't continue. This Mode is probably the best way to safety scan pits. The idea is to descend to a certain point, stop and scan, descend a bit further and repeat until you get to the bottom. In ROOM Mode you don't move while scanning, you just rotate the Caveatron around an arc or complete circle while holding it at a fixed location. This is best done on a rock but can be done hand-held in the air if you can rotate it steadily enough. This Mode does require a station to reference its position, so you start by entering the reference station and a splay number and getting a reference shot to the station. The LIDAR scan begins and you slowly spin the Caveatron on the reference point until the desired area is covered. As with other Modes, statistics and redo options are displayed after the scan is complete.

A fourth measurement Mode, MANUAL Mode, is also provided to obtain quick distance, azimuth or inclination readings without having to enter station codes. Both shots and continuous readings can be taken. However, in Manual Mode, no data is saved, only displayed on the screen.

Besides the measurement Modes there is a SURVEY menu and a SETTINGS menu. The SURVEY menu is used for initial survey setup, to view statistics for the current survey (such as length, depth and number of shots) and the line plot of the current survey. Plan and profile



Shot Results



Line Plot

views as well as station labels can be displayed. Loop closures are detected and indicated on the screen so you can quickly see the quality of the closures by how close they are. The current station is also highlighted to quickly orient your position. In the SETTING menu there are a variety of functions to set the date and time and view information about the system and SD card, and perform re-calibration. When the LIDAR is attached, a screen is available to show a live view of the scan, which helps in determining if you are getting the coverage you want.

## Calibration

Proper calibration is important when dealing with electronic accelerometers and magnetometers and it can be challenging to get reliable and reasonably accurate readings with these types of devices. A great deal of development effort and code was put into handling the calibration and make it as transparent to the user as possible. Most of the calibration is done once, at the time of assembly, and does not require any user intervention. This is the case for the accelerometer, which is calibrated by carefully leveling the Caveatron, taking readings in all six orientations and computing the nine calibration parameters.

For the magnetometer, things are more complex since measurements can be affected not just by physical misalignment but also by the onboard electrical currents (soft iron) and metal inside the enclosure (hard iron). Only the physical alignment can be calibrated during assembly. This is accomplished by placing the Caveatron on a

rotating, non-metallic platform, outdoors away from other influences, and rotating it to a number of known angles measured with a Suunto. Unfortunately, the hard and soft iron values vary over time so need to be occasionally recalibrated by the user. This posed a challenge to make recalibration as simple as possible and able to be done in the field without any special equipment. Computing all 15 of the calibration values requires some rather sophisticated math involving eigenvectors. Although more simplistic computational approaches exist (usually ignoring soft iron effects), after experimentation these did not seem to produce accurate enough values. Thus, the more complex method was implemented.

To perform the user calibration, one only needs to find a location away from magnetic influences and rotate the Caveatron through as many angles as one can over a full sphere of directions. Although simpler methods ignore this, for accuracy it is important to get a reasonably uniform distribution of points. While rotating, the software automatically figures out if enough data has been collected over each section of angles and a progress bar is shown on the screen indicating how much of the total required data has been captured. The process typically takes about a minute or so to complete and when the system determines enough data has been collected, the calibration values are automatically calculated, displayed and stored in onboard memory. These can be updated as often as you wish, but it has been found through testing that the calibration is relatively stable over many days or even weeks.

## Post-Processing

Although the onboard processor can handle some basic filtering of the incoming LIDAR and position data during scanning, it can't generate the point cloud. For that I needed to write custom software that can run on a Windows or Mac computer with one program to view, download, and delete the files stored on the Caveatron and another program to post-process them. The first step of post-processing is to generate an X,Y,Z coordinate file using David McKenzie's Walls cave survey software. The shot file generated by the Caveatron is already in Walls survey file format, so is easy to import, view and process a line plot.

The output station list file from Walls is then loaded into the point cloud processing software along with the LIDAR data file from the Caveatron. This software brings up a list of the traverses and room scans, one or more or all of which can be selected to begin the point cloud processing. Alternatively, selecting one of these also allows you to manually step through it so you can view the data in graphical form as single cross-sections. The software automatically applies additional filtering to get rid of bad data and noise, interpolate the Caveatron position between valid measurements, remove excessive motion, compute the directional normal for each point, and then stitch all the traverses together into a single file with the absolute X, Y, Z coordinate computed for each point. The program saves a text file that can then be easily loaded into one of several freely available point-cloud viewing programs such as Cloudcompare or Meshlab. To take the

## SPECIFICATIONS

	Distance	Azimuth	Inclination	3D Scans
Sensor Type	Laser rangefinder	3-axis magnetometer	3-axis accelerometer	Single-axis rotating LIDAR
Measurement Range	0.13 – 45 m	360°	+/- 90°	0.05 – 4 m over 360°
Resolution	1 mm	0.1°	0.1°	1 mm 1° (angular)*
Accuracy	, ± 5 mm	, ± 1.5°	, ± 1.0°	several cm (depends on many factors)
Measurement Rate	1/sec	300/sec	300/sec	4.5 rotations/sec (1620 points/sec)

data processing one step further, Meshlab has a function to turn the point cloud into a rendered mesh through a surface reconstruction algorithm. That creates a 3D file that shows the walls of the cave as a solid surface allowing you to fly around or through the cave on your computer. As a final step, Meshlab can export a file suitable for 3D printing so you can make your own physical cave Models!

## Testing

The final version of the Caveatron has been tested in several caves. At Robber Baron, over 400 m of the cave has been surveyed on multiple trips. At Deep Cave the entrance area was surveyed and at Devil's River State Natural Area, two small caves and one shelter were completely mapped. A comprehensive test was performed at Whirlpool Cave, where 240 m from the entrance to the Travis County Room, including both routes, were surveyed on a single trip that was about 7 hours in duration. On this trip, we attempted to be as thorough as possible, scanning all the alcoves and side rooms along the way. Discounting breaks and lunch, the average survey rate was about 40-50 meters per hour. One of the crawlways was quite wet including plenty of mud, standing water and ceiling drips, giving the Caveatron a good test of that type of environment. At the end of the day, the battery was still holding up and the Caveatron was working normally. All of the data from each cave has been rendered into solid surface Models revealing the caves like they have never been seen before. Robber Baron, Whirlpool and one of the Devil's River caves have also been 3D printed with the Whirlpool Model being more than 30 inches long. 2D maps are great for showing you the general configuration of the cave, but there is a lot they don't capture. Even though I've been in Robber Baron hundreds of times, being able to hold the cave in your hands or view it in 3D on the computer shows you things about the cave like slope, bedding planes, passage shapes and relationships between parts that were not obvious before. Combining the speed at which surveys can be done with a whole new dimension of measured detail has the potential to dramatically change the way we can survey caves.

## What's Next

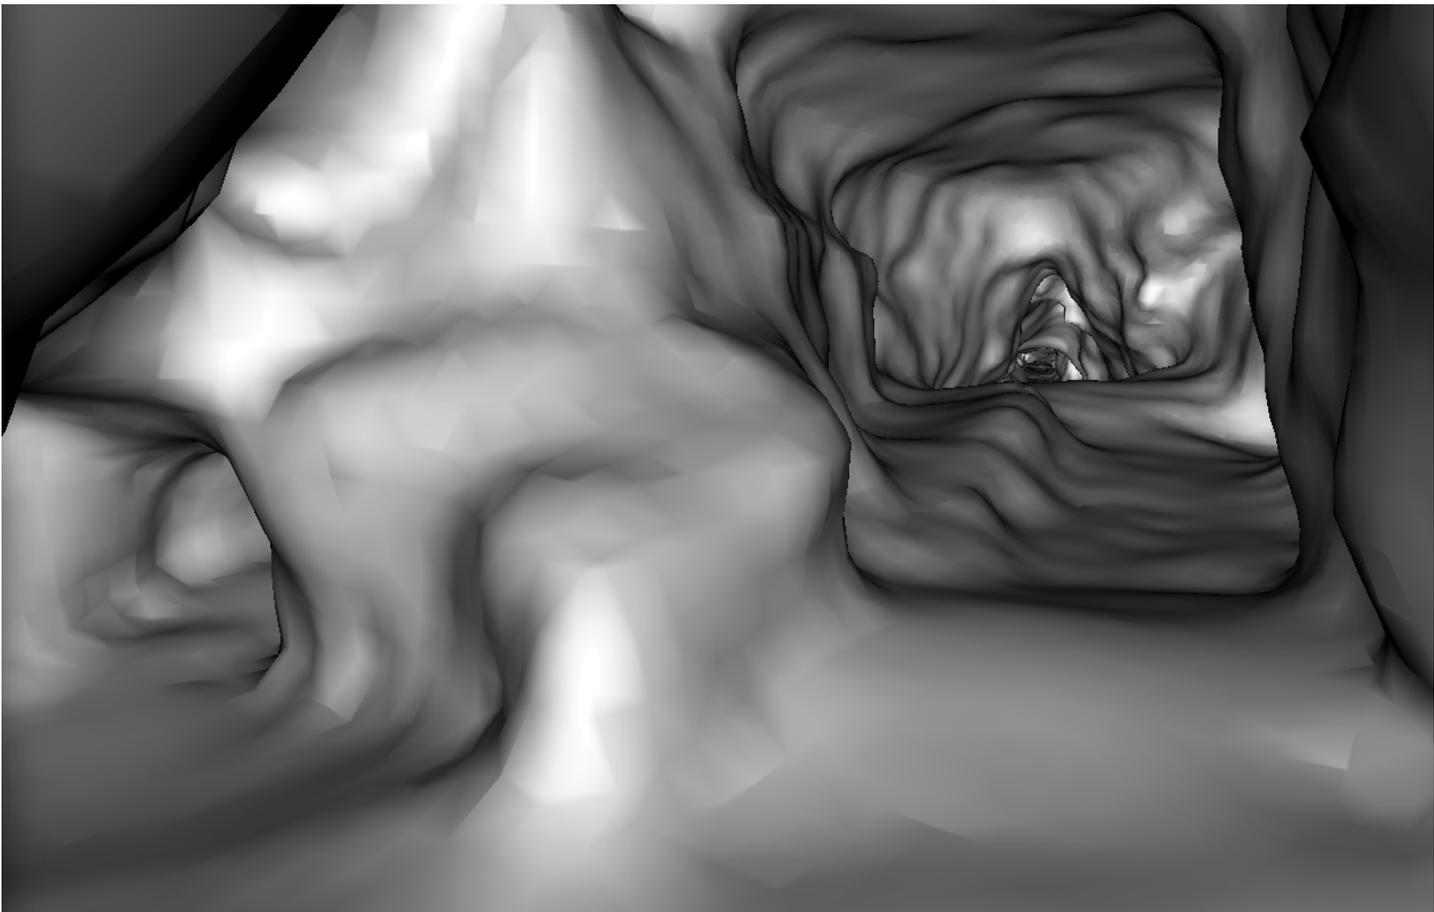
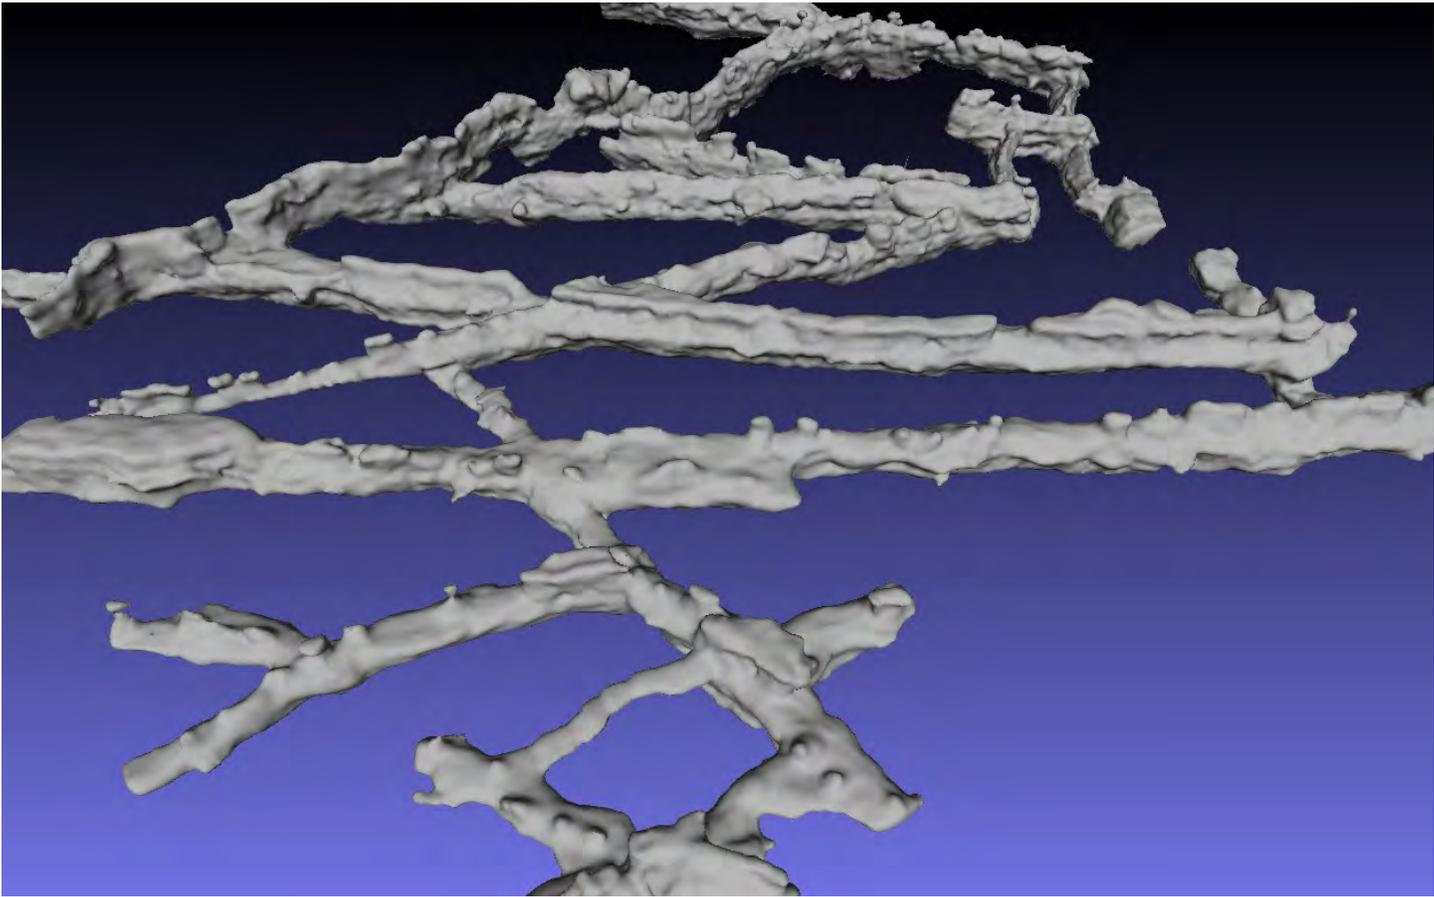
Although the design and software are now mostly complete, a lot of manual work is involved in assembling the Caveatron with some of the most time consuming labor being the wiring. We are currently preparing a printed circuit board layout that will combine most of the individual electronics modules onto one board, eliminating most of the wires. One more challenge

remaining to solve is to create a better front cover for the LIDAR that doesn't cut off as much of the corners as the current one. Five additional Caveatrons are currently being assembled to distribute for field trials. A webpage will be setup soon on the Texas Speleological Survey website with more information on the Caveatron including downloadable cave Model files. And of course there are lots of caves to get out and survey! My eventual goal is to not only figure out how to make these available to sell but also to open source the design so that anyone with the right skills can build their own.

Lots of people have helped out with this project so far. Steve Gutting has helped with the electronics design, has done most of the wiring and electronics component assembly, and is working on the circuit layout. Gregg Williams has helped with the enclosure design, has done most of the machining work and is currently preparing drawings for producing more. Gregg also printed the 3D cave Models. Jill Orr developed the icons for the user interface and Ellie Falgout developed a logo and originally came up with the Caveatron name. Many cavers have been involved in testing, most notably Evelyn Mitchell and Mike Harris. This project is for cavers, and I would like more cavers to be involved, so if you are interested in helping, please contact me at [joemitchell@satx.rr.com](mailto:joemitchell@satx.rr.com).



Plastic model of Whirlpool Cave



Top: Oblique print of Robber Baron Cave. Bottom: 3D rendering of the Entrance Room in Robber Baron Cave