Erratum: In the answer to the question "What are the strongest negatives about Autonomous under water collection?"

"The other challenge is pressure, which can be as high as 16,000 PSI at 5k depth. So this can damage instrumentation, or worse, hull integrity may be lost."

Correction: The pressure of 16,000 psi is at ~11,000m, the depth of Challenger Deep at the Mariana Trench, the deepest known location on our planet. Underwater, the pressure goes up by about 1 atmosphere (or 14.5 psi) every 10m, so at 5,000m, the pressure will be about 7250 psi, still very extreme, posing difficult engineering challenges for underwater vehicles and submersibles designed for those depths.

Student questions: Jnaneshwar Das colloquium on "Robots in the Wild: Collaborative Exploration and Mapping"

3/13/19

Question 1: How do you know what kind of neural network to use for training? (mask RCNN vs others)
A1: The choice of deep neural networks is driven our prediction goals. If it is segmentation, i.e., pixel-level class membership, then U-net and retinanet are a couple of examples of latest segmentation networks. For detecting objects, that is, not just segmenting rock pixels, but putting bounding boxes around individual rocks (or fruits), YOLO, and Faster RCNN are a couple of choices, that use region proposals. Mask RCNN is a reasonably new hybrid, that combines features of segmentation and object detection, so this is best suited for our rock classification work where we care about the shape, and identity of rocks.

Question 2: Is the annotation tool you have created available to the rest of SESE?
A2: Yes, we are making some improvements to our map-tiles and image annotation tool that is built over Django, and Leaflet, and Postgresql DB, and this will be released to SESE community soon.

Question 1: Is there a varient of the drones that can map lunar lava tubes considering there is no atmosphere?
A1: There has been work on jet-propelled multi'rotors' on Earth, and given much lower lunar gravity, we will not need a lot of thrust, especially for very low size weight and power (SWaP) solutions. So a very small quad-drone with a jet pack at the tip of each arm, and lightweight sensor suite with IMU, LiDAR, and camera with onboard lighting. The controllers will have to be designed to handle the different dynamics compared to a quad rotor, though the basic principles of guidance, navigation, and control will be similar to those of Earth-based drone.
PS: here is a company named Jetoptera building futuristic jet-propelled drones _and_ taxis.
http://www.jetoptera.com/products/

Question 2: Can the same process used to track red algae be used to track jellyfish? In the Gulf red algae and jellyfish were the two main reasons you were not suppose to go into the water.
A2: Perhaps with onboard cameras, and the same technology we are using to detect fruits. In case of algae, we assumed that water currents advect algae. For detection, we were using point sampling using fluorometers and other sensors.

Question 1: How often do poor samples get collected when collecting the 10 samples?
A1: The hiring problem's solution prescribes collecting the last candidate sample if no sample is found that is better than the ones in the "observation window". So this sample has a higher probability of being suboptimal. Fortunately, these serve as good control samples. Details of this work can be found in the following publication,
J. Das, F. Py, H. Harvey, J. Ryan, A. Gellene, R. Graham, D. Caron, K. Rajan, and G. Sukhatme, "Data-driven Robotic Sampling for Marine Ecosystem Monitoring", in International Journal of Robotics Research, 34(12):1435-1452, Oct 2015.
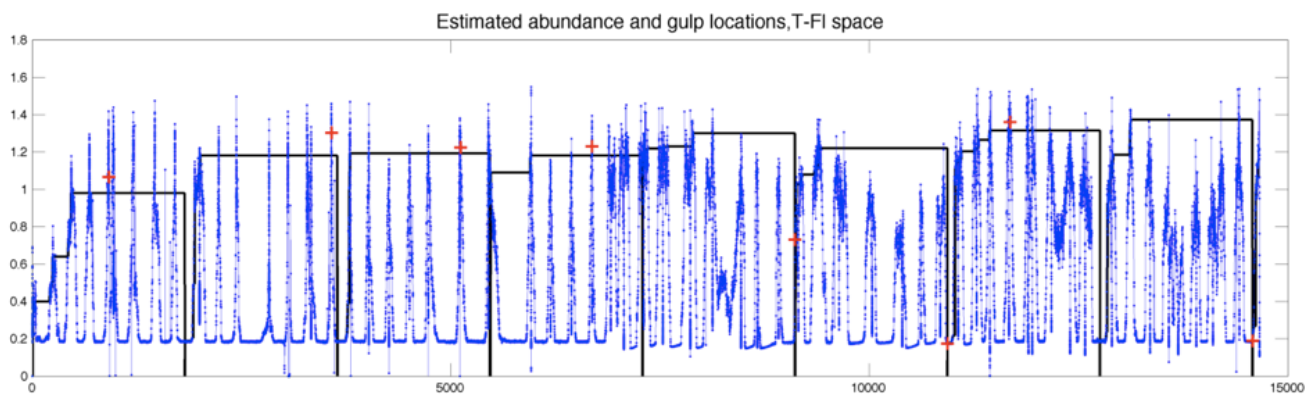https://dornsife.usc.edu/assets/sites/378/docs/Caron_pdfs/2015_Das_IJRR.pdf

Fig 1: The hiring problem is applied to each of the 8 gulping decisions in this mission, based on the utility peaks. Black lines show current best score for each window, which attains the highest value within ⅓ of the total window size, and subsequently, the next best sample is chosen.
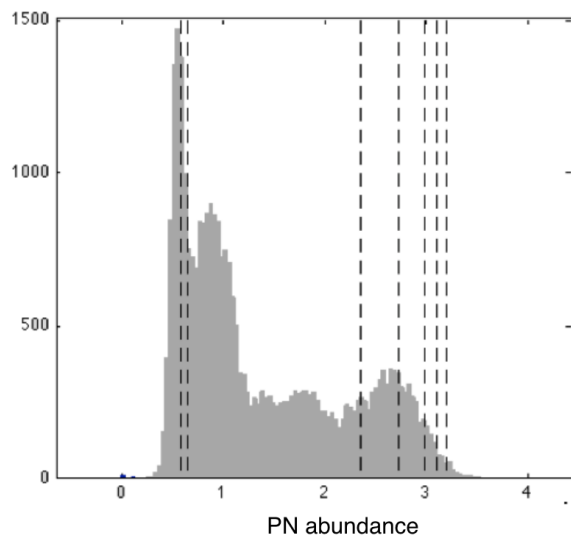

Fig 2: Note two sample collections that have low organism abundance, these were two of the three samples collected at window ends.

Question 2: How many repeats of an organism is needed to indicate a model should be made?
A2: Modeling is always carried out after a mission on newly collected and lab-annotated samples. As for filtering during the mission, as the AUV sees measurement of environmental context (temperature, salinity, pH, etc.) multiple times a second while passing through the water column in a saw tooth pattern, it chooses the peak of each dive as the candidate for observation, or sampling. This is shown in Fig 1.
If you are interested in learning more about the sampling algorithm, evaluation, and field trial, please check out the following article.
J. Das, F. Py, H. Harvey, J. Ryan, A. Gellene, R. Graham, D. Caron, K. Rajan, and G. Sukhatme, "Data-driven Robotic Sampling for Marine Ecosystem Monitoring", in International Journal of Robotics Research, 34(12):1435-1452, Oct 2015.
https://dornsife.usc.edu/assets/sites/378/docs/Caron_pdfs/2015_Das_IJRR.pdf

Question 1: Are there any planned projects for drones to fly and conduct in-situ extraction on the Moon?
A1: Not yet, though Prof. Meenakshi Wadhwa has ideas for lunar sample collection. Drones with propellors will not work on the moon of course due to lack of atmosphere, so we will have to develop a jet-propelled system, with appropriate onboard imaging instrumentation to help guide sample extraction decisions.

Question 2: There is an ArcGIS tool called Crater Counting, which counts craters on the surface of body to get relative age, and do you know of any machine learning programs that are in the works to automate this tool to recognize craters?
A2: Thanks for pointing this out. Upon looking at the ArcGIS crater counting tool, it seems it uses image processing techniques, and not deep learning, to extract contours of craters, and analysis of contours is carried out there after. Using deep neural networks may improve contour detection, especially by training on already available crater datasets. So, this is a good area for trying out data-driven (expert annotation + AI predictions + expert correction) mapping approaches.

Question 1: From the video of counting apples, it seems that there is a limitation to visualizing only the exterior of the plant canopy. Can you overcome this by different imaging modes to look through the canopy without looking through the fruits?

A1: Good question, occlusion is always a challenge. Fruits that are not visible in a frame, may become visible in subsequent frames since the camera is moving. We keep track of detected fruits across multiple frames, and can get around occlusion to a large extent. Still, there will be fruits that cannot be seen at all. For that, we have explored backscatter X-ray imaging that can "see through" leaves and detect fruits -- round bags of sugared water that are very visible to this mode of imaging.

backscatter X-ray image of the apple trees
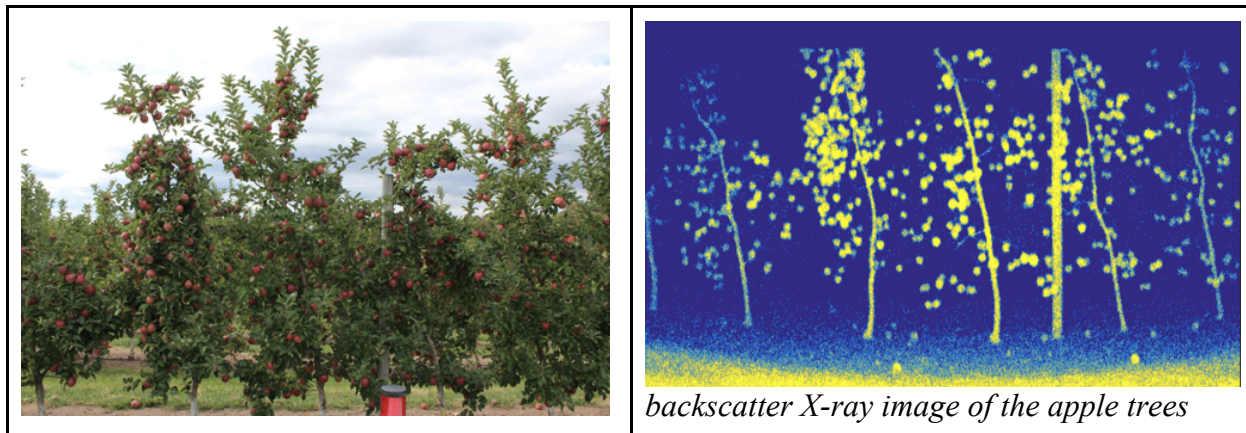
Question 2: The rock imaging work showed that you had picked up some polygons counted as boulders which were actually only shadows of the rock next to the polygon. You mentioned that a lot of this work may be very new, but have you made any progress on delineating the rock from the shadows of rocks?

A2: Great observation! Sometimes our network detects shadows as rocks, and we think this is because the training dataset does not have any instances of rock-shadows annotated as negative examples. Typically, when we draw a polygon around a rock , all the other pixels outside the polygon are marked as negative pixels. We have observed that the network detects shadows of rocks that are comparatively tall, and away from neighboring rock clusters. We will be retraining with human correction of the shadow predictions (2-3 examples), and then observe what happens. We expect the false positives to go away.

Another approach we are considering is adding the elevation of each pixel to the network as a fourth channel, since we already have this from the DEM orthomaps. This may eliminate the need for annotating shadow-rocks as negative examples, since these do not have any vertical extent (shadow falls flat on the ground).

This is a great example for why we have to always be improving by sampling, and sometimes correcting, the neural network predictions.

Question 1: Does aerial imagery have a high enough resolution to record something as small as mica flash to determine mineral orientation?
A1: This will depend on the camera resolution, and distance from target object. Currently, we are resolving features that are a couple of centimeters for rocks, and less than a centimeter for fruits, where we fly much closer to target. Also, if the drone is hovering, and imaging a fixed location on the ground, we could get mm resolution by fusing the imagery. Lot of options, reach out to discuss!

Question 2: What is the flight range of an average UAV?
A2: Multirotors have about 10-30 minute flight time. Newer Hybrid VTOL drones have wings as well, so it takes off and lands like a multirotor, can uses its wings during transit. Hence the energy usage is much lower during cruise, and they can fly for an hour or more. Purely winged drones such as SenseFly eBee also have very good endurance of 90 min. Note that for gas-engine winged drones, the range can be much higher, we do not use that format.

Question 1: You mentioned moving to Unreal engine for simulation. Would this require getting new people on the team or is this something you're well versed in?
A1: Unreal engine has a learning curve, and although we have people in our team using the Unity engine for VR work, we are looking for game developers well versed with Unreal engine.

Question 2: Where do you hope AI will help most critically in terms of volcanology?
A2: I believe the main role of AI and Machine Learning will be to facilitate safe navigation of the drone(s) close to a volcano (obstacle avoidance, terrain relative navigation, probe deployment). Our ongoing experimentation with spectroscopy and tomography may also benefit from deep neural networks, think inverse-modeling aided by AI.

Question 1: With fruit counting, how do you prevent inconsistencies, like counting fruit twice, in the data?
A1: Since we track fruits, double-counts are eliminated across frames. However, say we are mapping a single orange tree by going round and round. We would want our count to stop after one revolution around the tree. This requires 3D mapping of the fruits so that we know we are looking at the same fruit after coming around. Challenge here is if we are looking at fruit clusters, or fruits very close to each other. With good mapping methods, we can ensure no double counting, even in extreme examples such as this.

Question 2: What information can be interpreted from the AUV's water collection system data?
A2: The water samples are usually processed using molecular probes to measure abundance of specific genus of zooplankton or phytoplankton. During the Gulf of Mexico oil spill, the same AUV was deployed at the gulf by another team, where they used colored dissolved organic matter (CDOM), backscatter, and other in-situ measurements as proxy to collect water samples potentially rich in hydrocarbon, and then, from lab-analysis of samples, the specific signature of hydrocarbons were determined (did this water mix with hydrocarbon from well-head, or is this just from some hydrothermal vent?)
https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.20399

Question 1: What would it mean if a team were to successfully spot the drone in the competition?

A1: It would suggest that the barrier to entry for such tasks have now been lowered. Autonomous object search and recovery was earlier accessible to senior PhD students in sophisticated labs, or defense labs. Successful, and reliable recovery of lost objects would hint that even undergrad and junior grad students can tackle the challenge. The team will also receive an award at the next NSF Cyber-Physical System PI meeting in November, and receive a reasonable cash prize.

Question 2: What if we don't have training data for some of the rocks we are trying to identify – would that mean we would miss it?

A2: Great question. To use the fruit counting as an example, let's say we can learn a generalized model of a fruit (round colored things in canopy), but we cannot precisely detect the type of fruit. Similarly, for rocks, we could train on a generic rock dataset and be able to detect rocks, but not identify particular type of rocks. So a sampling strategy can be used to optimally annotate a _few_ rocks of a particular type, and learn a subclass for the "species" of rock. So, I would say that there will be more false positives of the target class of rocks, though we should still be able to detect rocks broadly.

In our case, bedrocks were not annotated, and surprisingly, also not detected. This hints that the network is skipping rocks that are too "continuous and big". Such scenarios can be troublesome, for instance, if you were interested in bedrocks.

Question 1: How are these methods expected to be applied in the future to various bodies in space?

A1: A low hanging fruit will be to apply our techniques for optimal sampling with ground-based and space-based telescopes -- where should we point our telescope next, based on all the data we have so far, to maximize some objective, for example finding a new exoplanet.

For next-generation autonomy on spacecrafts to make optimal-decision making, we need to ensure we have addressed the high energy and thermal footprint of training deep neural networks. Perhaps training will have to happen on Earth or resources external to the spacecraft, based on new potentially-useful imagery sent back by the spacecraft. Inference on the other hand (predicting using a trained model), is not very expensive, so once model parameters have been sent to the spacecraft, it can start "seeing clear" detecting features that were previously not identifiable.

Question 2: How do the different abiotic conditions of different planetary bodies affect the different technical approaches necessary?

A2: To cite one factor for drones, lack of atmosphere (or very thin atmosphere) will make both navigation and thermal dissipation somewhat challenging. For underwater, other challenges will be present. How well will multibeam SONAR on a AUV work in methane lakes? However, I am just getting

Question 1: Is it possible to repurpose the sensors used for fruit counting and use them to collect real time data of local temperature, humidity, and wind to better predict local weather forecasts?
A1: Fruit counting is involves extracting semantic information from imagery (where are the individual fruits), and tracking them, to acquire a 3D "semantic map" of fruits. If there are parallels in meterology where we need to detect visual features in multi-spectral images then fruit counting will be a good analogue. Otherwise, we can carry out optimal point sensing with sensor networks using techniques that connect with the use of Gaussian processes. As an example, JPL's Regional Ocean Modeling System (ROMS) fuses data from ocean gliders, and moorings, to provide ocean current predictions. https://science.jpl.nasa.gov/projects/ROMS/
Also check out the following paper for a flavor of how optimal sensor placement (or data collection) can be carried out for better modeling.
http://www.cs.cmu.edu/~./ajit/pubs/Guestrin2005.pdf

Question 2: What kind of academic background does one need to get involved in this field of study?
A2: I would say, equal interest in science and engineering, with a good foundation in statistics and probability theory. Also, although we do not have to be master coders, we must not be afraid of coding. So find a way to enjoy coding, perhaps through interesting projects.

Question 1: When detecting fruit on the trees, is the scanner able to differentiate the difference between ripened and unripened?
A1: We have not addressed ripeness yet extensively, there has been related work in this arena that use color cues to do so. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5017387/

A better approach will be sensing volatile organic compound (VOC) in-situ, that would serve as "nose and tongue" for the robots. Additionally, we are also looking at spectroscopy and multi-spectral imaging.

Question 2: There was a short video showing a UAV attempting to remove a leaf sample from a tree, and having some difficulty; why not install a pair of clippers on the arm?
A2: My apologies for not mentioning this during my talk, the gripper (end-effector) is composed of three clippers, and a tongue to 'grab' down on the leaf during the cutting process. The drone crashed in that trial in the video, because of control issues. We need to improve how we control the drone when the arm is in contact with the branch. Alternatively, we have to let go when the drone is becoming unstable due to the swinging branches etc, and reattempt. Here is a publication on this work, and a longer video clip if you are interested in learning more about the system.
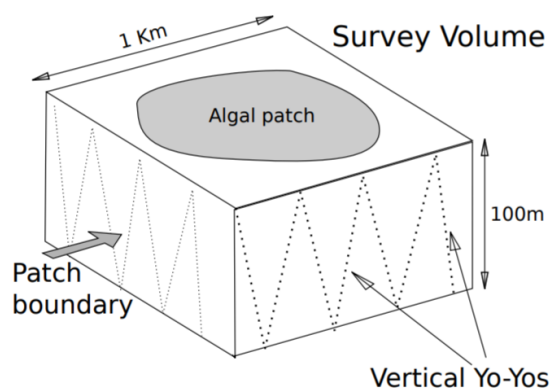https://label.ag/Aerial_phytobiopsy_ICUAS17_0187_FI.pdf
https://www.youtube.com/watch?v=JmjJXrCwpbU

Question 1: Why does the underwater robot go around the algae bloom in a box pattern?
A1: Good question, this was to characterize what's entering and leaving the volume encapsulating the patch. In particular, marine scientists involved with the experiment were interested in nutrient flux (we also measure nitrate concentration with our AUV).

Figure 5: The box survey pattern of an AUV which circumscribes a patch volume being sampled.
http://oceandatacenter.ucsc.edu/MBHAB/hotspots/publications/Das%20et%20al%202012.pdf

Question 2: Do your training sets of rocks or fruit have to contain the specific rock to prevent biases toward other shapes, colors, etc?
A2: Yes, that is our goal, though it is hard to identify in advance _all_ the possible scenarios. So, we introduce diversity by augmenting out training datasets by flipping, stretching, and distorting images, and soon, we will be using game engines to play with lighting and other properties, to an extent. The best way to improve the system further, in my opinion, is to explore predictions, and correct false negatives and false positives, and retrain. AI can also help in identifying potential candidates for further labeling by using unsupervised learning. Hannah Kerner, a graduate student at SESE wrote a nice paper on a related problem of filtering very large sets of imagery from Mars, to much smaller image sets that may be relevant for scientists.
https://www.sciencedirect.com/science/article/pii/S0098300417309688

Question 1: What are some key physical parameters that are used as training sets for the deep neural network to get information on rock boundaries?

A1: We primarily focus on having annotated examples of diverse size, shape, and texture of target rocks, and enough examples of what are not rocks. The latter is usually extracted from pixels that are outside rock annotations, and marked as negative. You may have noticed that we sometimes land up detecting rock-shadows as rocks, so now we are working on interactively correcting AI predictions, and retraining.
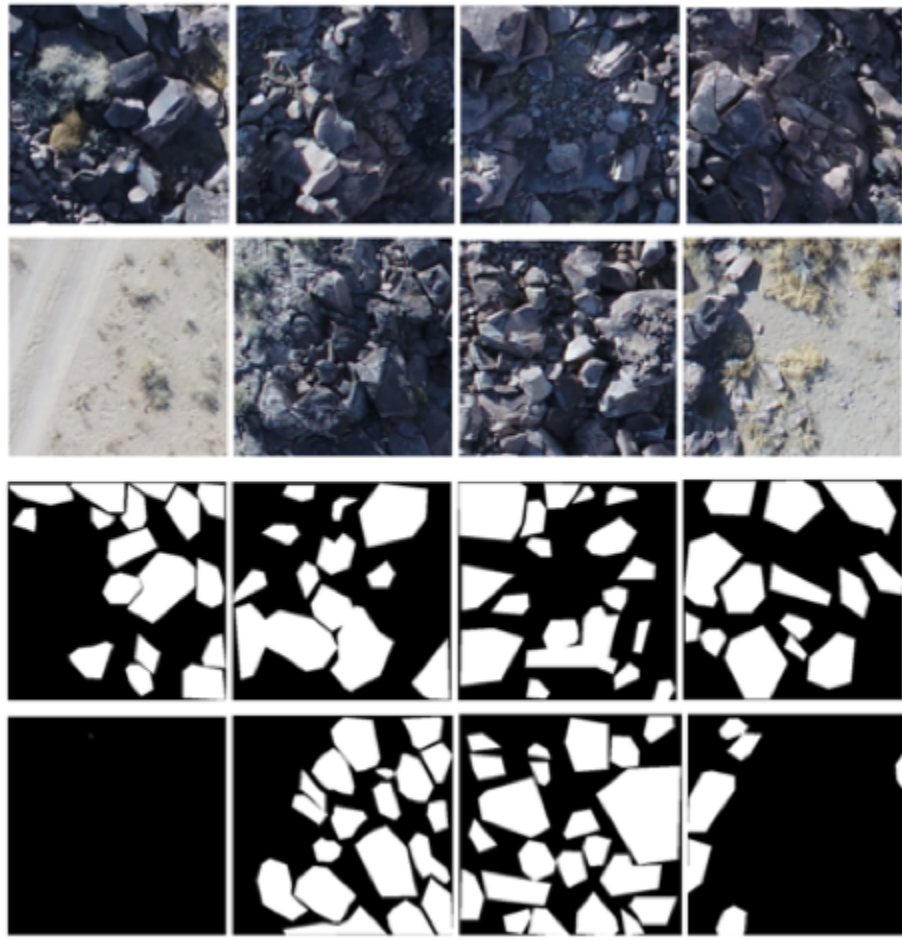
(raw image tiles on top, annotation masks on bottom)

Question 2: What is the typical range of view for the drones used for crop health monitoring?

A2: When flying above trees to carry out NDVI mapping (crop stress), we are usually about 20m above ground, apple and orange trees are usually 2-4m high. If you want to map pecan trees that can be very tall, upto 20m tall, you will have to fly higher. With onboard lighting, we fly pretty close to objects, about 4-5m. When flying between rows, looking sideways (fruit counting examples), we have to be between the two rows ideally. Picking the appropriate camera lens (wide angle-ness) helps to an extent in mitigating challenges when having to fly close, in cluttered scenes.

Question 1: Have you patented any of the technologies you've worked on?

A1: Yes, we have two patents filed or pending. Code/sources and various datasets are open, and available for academic and non-commercial use.

1. Systems, devices, and methods for robotic remote sensing for precision agriculture
   V Kumar, GB Cross, C Qu, J Das, A Makineni, YS Mulgaonkar
   US Patent App. 15/545,266
2. Systems, devices, and methods for agricultural sample collection
   D Orol, L Vacek, DV Kaufman, J Das, V Kumar
   US Patent App. 15/974,243

Question 2: What is the next step in your plume tracking project with Prof. Clarke?

A2: We want to carry out an autonomous air sample collection mission in our small-scale setting with coalfire plume, and thereafter, we may travel to a site that has fumaroles (WA state has come up in conversations). Lot of engineering has to be done before we can travel to an actual volcano. First, we have to develop good in-situ sensing approaches to measure SO_2 and CO_2 for the whole or part of plume. Other challenges are, drone design to withstand heat, operations that are beyond line of sight (will need FAA permission), 3D mapping of plume (for full-scale experiments), and finally, failure recovery, for example, if we lose one motor due to exposure to plume chemistry, can we still recover the drone? Also, we are thinking that perhaps a winged airplane, or a hybrid VTOL, may be able to ride the thermal currents from the volcano. We will have to identify the best way to navigate around and in the volcano plume.

Question 1: With AUV sampling, do you have specialized ones that take samples at deeper depths?

A1: The sample collection depth depends on the AUV depth rating, and the sample collection system is a bay that can be moved to a higher depth-rated AUV. MBARI in fact carried out water sample collection in Gulf of Mexico with the same AUV I used, at much higher depths of 1200m. https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.20399 . Some of the other MBARI AUVs are rated to 5k depth, the question then is, whether the spring-loaded piston based system will operate reliably at those high pressures.

Question 2: With payload recovery, could you use radio transmitters or radar to triangulate longterm probes if autonomous flight has problems?

A2: Excellent suggestion. Indeed, the probes may be in a valley or a region with very poor GPS localization, and triangulation through beacons would be the only way to locate those, especially if the locations are not known in advance. In fact, since radio waves don't work underwater (beyond a few inches or a meter), localization underwater is carried out often using acoustic triangulation. You may enjoy this paper, https://onlinelibrary.wiley.com/doi/full/10.1111/1365-2478.12670

Note that since the probes are not going to move, unlike in the ocean water-column where things get advected, the drone can also "remember" how the surrounding topographical map looks based on when the sensor was deploymeyed, and locate the probe again purely based on localization with topographical landmarks.

Question 1: What are the strongest negatives about Autonomous under water collection?

A1: First challenge is communication. Radio waves don't work in water, so you are restricted to acoustic communication to and from the AUV. Sound waves bend in water, and the speed of sound can also change, due to temperature and density variations.

The other challenge is pressure, which can be as high as 16,000 PSI at 5k depth. So this can damage instrumentation, or worse, hull integrity may be lost.

Then these is salinity of the water, biofouling that can render sensors ineffective after few weeks of underwater operation.

Question 2: What is the future of the spatial distribution of rock traits?

A2: If there are any events that fracture rocks, or move them (example, flooding), then the spatial distribution of rock traits may change.

Question 1: Can the data collected be used to create a manipulatable VR envorenviro with the unreal engine?

A1: yes, in fact we are working on this as a part of ASU ASURE VR Challenge, for which we are a finalist. Our project is called "Virtual Reality and Machine learning tools for Characterizing Rocks in 3D for Geological Applications", and I encourage you to attend our final demo mid April. Please reach out if you need more info on this!

Question 2: Can the robots detect crop readiness as well as health and number?

A2: Yes, crop ripeness is something we have been interested in, and are exploring with spectroscopy and multi-spectral imaging. The future, I believe, will involve visual-olfactory sensing of ripeness, perhaps through sensing of volatile organic compounds (VOCs), that may be better cues for ripeness (think nose and tongue of a robot). VOC technology is evolving, and I believe low size, weight and power (SWaP) affordable VOC sensors will hit the market within the next 2-3 years.

Question 1: How long do you foresee this taking to impact interplanetary exploration in a significant way?

A1: GPUs will have to become more energy efficient to allow in-situ adaptation and learning. For navigation on planetary surfaces, a lot has already been learned on Mars, and we will probably learn a lot from the Mars 2020 rover missions. Also, we are already learning tremendously through development of new robotics and AI tools in emerging markets such as agriculture and mineral prospecting. I foresee some of these technologies translate to aiding interplanetary missions.

Question 2: Are there any foreseeable problems with the massive influx of data that could arise from widespread use of AI controlled research robots

A2: The general wisdom is that more data is good, in fact deep neural networks are meant for domains with a LOT of data. The bigger challenge will be data storage, and more importantly, data transfer.

Das_Answers to student questions with Erratum 319.pdf

Question 1: In your opinion, what do you think the biggest difference between Earth-based marine robots, and space-base? For example, a space robot sent to search for life in extra-terrestiral oceans.

A1: Since you asked biggest difference, perhaps it is that the robot may be in Titan, with liquid methane for example. Then, we will have to understand how the sensors we use in water will behave in something that is not water. Will multibeam SONAR work effectively to map terrain? How effective will acoustic communication be?

Question 2: Do you think that, for planetary science, air based geologic and seismic monitoring more viable than land based?

A2: Air based mapping will be agile/maneuverable, fast, and can scale rapidly through use of swarms that work collaboratively. Challenges are, FAA regulations on Earth, limited flight time, and in general the danger of flying (fairly heavy) objects above people and things. Advantage of ground robots are that you can pack a lot of sensors and energy. I would imagine the future being hybrid, where planetary, geologic, and seismic monitoring is carried out by drones, that are based off of ground bases (rovers). The drones can then come back to the rover to recharge, transmit data, retrain models, and then take off and reach far away places with improved models, and bring back even better data, and enable improvement of maps and models.

Question 1: How effective is the use of the optimal stopping theory in you data analysis?

A1: Overall, we saw _much_ tighter variances on cumulative regret (our performance metric), compared to other strategies. The following publication has a discussion of results. Here is the publication on this work, fig 5 & 7 have some info, though you may have to read the experimental design, and the results discussion a bit to make sense of these. https://dornsife.usc.edu/assets/sites/378/docs/Caron_pdfs/2015_Das_IJRR.pdf

Question 2: At what level of precision do your cameras perform the detection and counting of all the fruits from any given tree?

A2: We are able to detect 99% of all _visible_ fruits on canopy, and even fruits that are somewhat occluded, but visible in some of the image frames. Fruits not showing up in some of the frames for the image stream is OK, our filtering techniques can still track. As for fruits deeply occluded, we will need technologies such as backscatter X-ray imagine to detect and count those. Overall, we can assume that the visible fruit is a good proxy for total fruit count of a tree, assuming the canopy conditions are not changing rapidly. So, you can see that this requires some ground-truthing post harvesting. See the following papers on some discussion of challenges and results for fruit counting.

http://label.ag/sys-prec-ag-case-2015.pdf
http://label.ag/counting-apples-oranges-R-AL-2017.pdf
http://label.ag/fruit-counting-iros18-inreview.pdf

Question 1: What is the most hazardous environment Robotics and AI have explored?
A1: Perhaps Space! Next will be _into_ a volcano, if we can do it :)

Question 2: How much have drones and sensors increased the crop yields of farms where they have been used?
A2: The short answer is, we do not know yet, but we will know in a few years! Growers have seen enough evidence and feels strongly about the potential to fund field trials. This is on top of funding from federal agencies such as USDA that would typically fund high-risk high-impact research.
Crop yield can be increased by better harvest planning through improved crop yield estimation, reduced fruit/crop loss, and better crop health monitoring. All three areas are active research topics, and working with growers, we are testing and evaluating which cost-effective methods can be incorporated into agriculture. Here is a recent paper on a work involving plant pathologist, crop consulting company, and roboticist (me!) to test out our ideas, in this instance, for crop disease. https://apsjournals.apsnet.org/doi/abs/10.1094/PDIS-08-18-1373-RE

Question 1: To what extent can/do everyday farmers utilize precision agriculture technology?
A1: Fruit counting is the low hanging fruit, and we are working with our collaborators to apply this technology to existing farm equipment first, for example , sensor suite mounted on all-terrain vehicles (ATVs) or tractors. Drones are being incorporated into crop scouting, here is a recent publication. https://apsjournals.apsnet.org/doi/abs/10.1094/PDIS-08-18-1373-RE

Question 2: How much human annotation is needed for rock trait mapping?
A2:We trained our network on a set of images with a total of 1688 annotated rocks of diverse size and appearance (excluding bedrock). This model was then used to detect about 82,000 rocks. Training took 17 hours on an NVIDIA RTX 2080 Ti GPU.

Question 1: Will geologists be able to depend on drones and robots to do field work in the future?
A1: Absolutely! We are already doing this, and the autonomy and ease of use of systems is only becoming better. Oceanographers have been using AUVs (underwater drones) for a while now for their research given the inaccessibility of sites.

Question 2: What extreme environments will we not be able to explore with this technology at first?
A2: Into the crater of an active volcano! We will probably have drones lurk around far above in the plume. We have thought of winch based methods to lower a metal probe into the lava from a drone high in the air, though this will require a lot of considerations including can a cable withstand intense heat etc.

Das_Answers to student questions with Erratum 319.pdf

Question 1: Are there any environments, whether they be hazardous, crowded, etc., that you or others have had difficulties using robotics/AI in?

A1: Indeed, if the environment is dynamic (cars or people moving, strong wind that makes plant move a lot), then the localization and mapping algorithms will suffer. Then, there is the challenge of predicting when AI will fail (you may remember the Tesla crash where AI labeled a semi as an underpass). This last bit is particularly tricky, since you cannot guarantee you have considered all situations, so you could try a few things a) train on very diverse data b) exploit simulations and other techniques to augment training data c) interactively correct AI predictions often since our surroundings may have new things, especially in built environments, and most importantly d) develop methods for formal verification of AI systems. The last point on verification is challenging since it is still hard to "peek" inside a network to determine what it has learned, given its complexity (about 100 million parameters). New tools are coming out though to enable us to look at the layers of the network. In the meantime, we will be constrained to a clinical approach of testing the network on variety of test datasets to see how it is performing. So if an environment has a lot of detail and nuances, it may be challenging since we may have missed out on annotating and training some important features.

Question 2: Have you been approached to potentially use these same type of target identification methods in-situ on telescopes to actively identify objects while observing?

A2: It has come up in conversations, however, I have not actively started looking into the problem yes. Please reach out to me if you are interested! :)
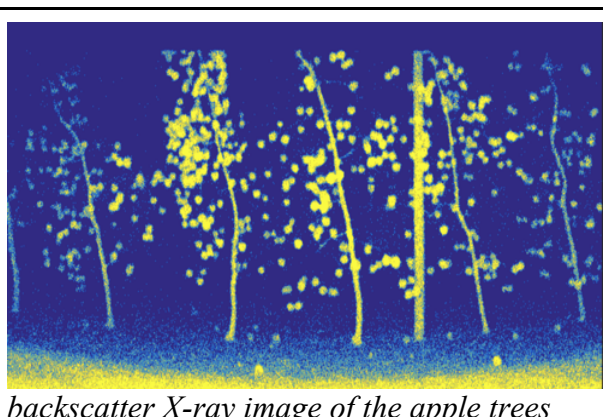
Question 1: Can we use irreversible Monte-Carlo method to make robots sample quicker?
A1: Thanks for the suggestion. I have not explored irreversible Monte-Carlo for my work, it looks useful, I will check it out! The Gaussian Process (GP) based upper confidence bound (UCB) sampling strategy involves model update after every stage, so any alternative approach will have to be able to capture that. Also, some of the performance guarantees based on cumulative regret that I did not get into, are derived from properties of GPs, so it will be interesting to find analogues with other approaches.

Question 2: In the fruit counting case, what will you do if many fruits are hidden by leaves?
A2: Great question. We assume we are only counting visible fruits, and the canopy occlusion is assumed to be a species/location/season specific factor. Occlusion is somewhat reduced since the camera is moving, and can see occluded fruits can be visible in some frames, and our tracking and filtering pipeline can count them. For deeper embeddings, backscatter X-ray imagine is something we have experimented with, and I believe it has a lot of promise both for fruit counting and automated pruning.

backscatter X-ray image of the apple trees