Project Proposal

Vision Guided Mosquito Dissection for the Production of Malaria Vaccine

Nicholas Greene - ngreen29

Mentors: Balazs Vagvolgyi, Alan Lai, Parth Vora

EN.601.656 CIS II Spring 2021

2/25/21

Contents

1 Introduction

The goal of this project is to develop several vision based methods which are necessary to further the progress of developing a robotic mosquito dissection system. This system automatically extracts the salivary glands from mosquitoes which is an important step required in the production of malaria vaccine.

2 Motivation

Malaria is a major global human health issue. In 2019, there were 229 million cases of malaria, with 411 000 deaths [7]. A successful effort to prevent the spread of malaria would have a dramatic impact on the lives of millions of people. Currently, the most successful method for combating malaria on a large scale is by preventing mosquito-human contact through various means such as mosquito nets and insecticides [6]. These have proved remarkably effective, however there has been an apparent observable plateau in their effectiveness in recent years.

Sanaria is a company which has developed the first malaria vaccine which is effective and feasible to produce [4]. Malaria is a parasite which spreads among humans by inhabiting the salivary glands of a host mosquito, which spreads the malaria parasite when it bites humans. The production and distribution of Sanaria's vaccine involves the extraction of these infected mosquito salivary glands. This step during production is a severe bottleneck and is preventing large scale production and deployment of the Sanaria vaccine.

A robotic mosquito dissection system for automating the extraction of infected salivary glands is being developed to make salivary gland extraction significantly more efficient which will enable large scale production of the malaria vaccine [5].

3 Prior Work

The prior work, as outlined in [5], is the mechanical design and software architecture of a robot system which is able to position, decapitate, and extract the salivary glands of mosquitoes in parallel. In addition, a handful of vision algorithms based on both image processing and deep learning have been implemented and are fully integrated with the ROS system. Figure 1 shows a rendering of the robot hardware and the results one of the vision algorithms. The existing hardware and software architecture can accommodate the addition of new camera sensors and vision algorithms without requiring any additional development.

4 Goals

The goal of this project is to create vision algorithms which will further the development of the robotic mosquito dissection system which was described in section 3 by providing automatic evaluation for steps in the system. In addition, these algorithms will also be useful for reducing the operating costs and waste. The specific vision algorithms are: classification of the turntable cleaning station outcome, classification of the gripper cleaning outcome, and regression of the mosquito exudate volume which is the final product of each mosquito dissection. Furthermore, each of these algorithms will be implemented twice, once

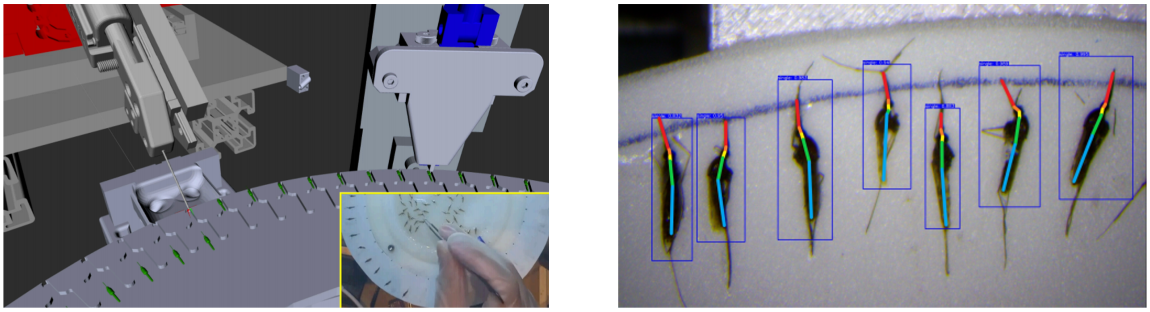

Figure 1: Activities and Deliverables

with traditional image processing, and once with deep learning for a total of six algorithms. The existing software architecture allows for the focus of this project to be purely on the six vision algorithms which is critical to finishing the project given its short time-frame.

5 Technical Approach

This section outlines the technical approach. An important consideration is that the cleaner classification tasks each will receive an image from a single mono RGB camera. The regression task will receive two mono RGB images from cameras which are positioned to provide approximately orthogonal views.

5.1 Image Processing Approach

For the cleaner classification tasks, a combination of image filtering and thresholding will be used to determine if a mosquito body or some other debris is located in the image or not. This approach is very likely to succeed due to the consistency of the lighting and the fixed camera position in the robot setup. For the exudate volume regression task, filtering and thresholding will be used to segment the blob of exudate in the two provided images. Subsequently, the volume estimation will consist of two different methods. The first will perform an elliptical Hough transform in each view and then use the approximately orthogonal geometry of the views to estimate an ellipsoid around the blob. The second method will involve counting the number of pixels of exudate in each view and then scaling that value appropriately.

5.2 Machine Learning Approach

Considering the small amount of ground truth data that will be available, transfer learning is the most feasible approach to creating machine learning models. Learning will be initialized with pre-trained ResNet [3] and YOLO [2] networks for the classification tasks. Depending on the results, more models may be used. All but the last few initialized layers will be frozen during training. For the regression task, there are two sub-tasks: segmentation and regression. Transfer learning on a segmentation network such as SegNet [1] will be attempted to segment the mosquito exudate in both views. This may not work well due to the difficulty of segmentation coupled with the small number of ground truth images. If so, the image processing method for segmentation will be used instead. These segmented image pairs can then be used to train a relatively simple model to regress the volume.

5.3 System Integration

In the current layout of the system, there are two computers. The main robot computer interfaces with the actuators and cameras of the robot, while the vision computer, which contains a GPU and which performs the vision tasks, is remote. Integration will be done by creating a ROS Service wrapper around each image processing and machine learning function. This will allow the robot main computer to call all of the remote vision functions as needed. The flow of data though the system can be seen in figure 2

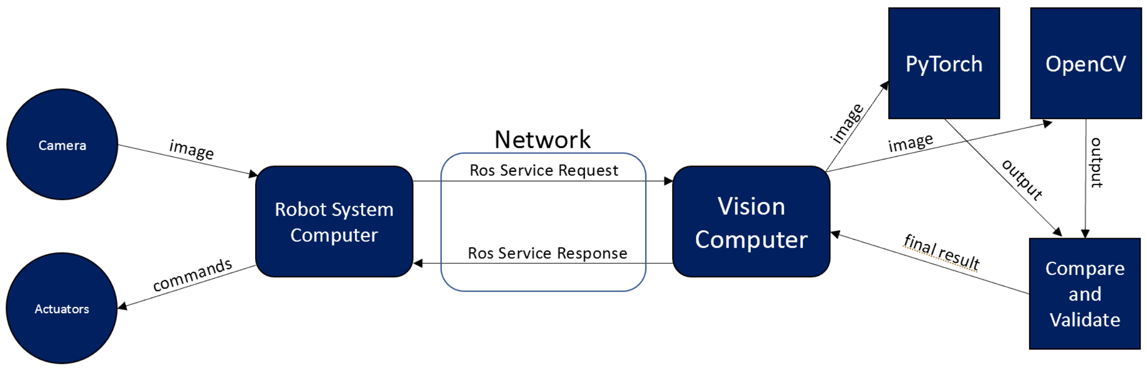

Figure 2: System Overview

6 Key Activities and Deliverables

The key activities and deliverables are outlined below in Figure 3. Note that they are categorized into minimum, expected, and maximum tiers. For brevity the chart is condensed; the contents of the minimum tier will be repeated for the respective tasks of the expected and maximum tiers.

For each of the three tasks, cleaning station classification, gripper cleaning classification, and exudate estimation, the deliverables consist of the following components:

Machine Learning Approach

- Dataset of 300 hand annotated images
- PyTorch based python code for training a neural network for the task
- PyTorch based python function for quickly evaluating the network's prediction on an image
- ROS service python wrapper for integrating the neural network predition with ROS
- Wiki documentation detailing the network architecture, training procedure, and ROS integration in python

Image Processing Approach

	Activity	Deliverable	Deadline
Min	Collect images of cleaning station and annotate	Annotated cleaning station dataset	3/1
	Implement image processing method for cleaning station	Functioning image processing based code and high quality documentation in a wiki	3/15
	Develop and train a neural network on the cleaning station dataset	Trained parameters of a neural network classifier, along with code and high quality documentation in a wiki	3/18
	Integrate both with the rest of the system using ROS	Working ROS services which can be successfully interfaced with from the robot control computer, and thorough documentation for how to use them in a wiki page	3/22
Exp.	Repeat for gripper cleaning	Same deliverables as for turntable cleaning station.	4/7
Max	Discuss with hardware team about collecting exudate data	Plan for collecting exudate ground truth data	4/1
	Repeat for exudate volume estimation	Same deliverables as for previous two tasks.	5/1

Figure 3: Activities and Deliverables

- OpenCV based C++ code of an image processing algorithm for the task
- C++ ROS service wrapper for integrating the algorithm with ROS
- Wiki documentation detailing the image processing algorithm and the ROS integration in C++.

.

A more detailed breakdown and discussion of the estimated completion time for each task can be seen in the Timeline section and in figure 5.

7 Dependencies

The dependencies, and their respective contingency plans are explained in detail in figure 4.

8 Timeline

The Gantt Chart in figure 5 shows the expected working period for each task. There is significant overlap in many of the tasks due the opportunities to work on tasks in parallel. Namely, it is possible to work on the image processing tasks with only a few images, and it is to work on any other task while waiting for the neural networks to train. Additionally, the period to complete the gripper cleaning classification task is expected to be significantly shorter than that of the turntable cleaning classification because it is very similar, and most of the code from one will be reusable with the other. A last point of discussion is that the discussions for collecting ground truth exudate volume data must begin early to allow time for the implementation of the plan.

Dependency	Need	Status	Contingency Plan	Planned Deadline	Hard Deadline
Continued access to GPU	GPU for training neural network	Currently have access to the Diva computer	I have a very capable personal GPU	Ongoing	Ongoing
Cameras for collection of images of each task need to be mounted and integrated	To collect images for annotation	Cameras mounted for turntable cleaning station and squeezing station but not for gripper cleaning	Use existing but less desirable views from the other existing cameras.	3/14	3/22 (for remaining camera)
Hardware team	300+ images of turntable cleaning station in progress with mosquitos.	Hardware team is currently running experiments and collecting these images on a near daily basis while they run their experiments	Begin working on the image processing methods for the other tasks while waiting for additional training data collection.	3/1	3/7
Hardware team	300+ images of gripper cleaning in progress with mosquitos.	The gripper cleaning water jet has not yet been installed, so there are no ongoing gripper cleaning experiments	Small Delay: Begin working "out of order" on the image annotation and image processing for the exudate estimation task Long Delay or Expected Failure: Change from the gripper cleaning task to one of the are many remaining vision tasks	3/14	3/22
Hardware team	Method for ground truth collection of mosquito exudate volume for training images	I will bring this up during the weekly meetings	Abandon neural network approach and only do image processing for this task	3/15	4/1

Figure 4: Dependencies

9 Roles and Management Plan

I am the sole group member for this project and my mentors are Balazs Vagvolgyi, Alan Lai, and Parth Vora. Balazs will be my main contact for help with the robot system and with the ROS interfacing. Alan and Parth previously implemented similar machine learning and image processing based vision algorithms which are currently in use. They will be a valuable resource for consulting about practical implementations.

The management plan is the following. I am attending weekly meetings with the full mosquito project team where I will be able to discuss with the hardware team, receive feedback, and/or bring up concerns. I will additionally meet with Balazs on a somewhat regular basis to keep in sync and to ask for help. These meetings will be as needed, with approximately weekly or biweekly frequency. I will also correspond with everyone else over email.

10 Reading List

H. Phalen, P. Vagdargi, M. Schrum, S. Chakravarty, A. Canezin, M. Pozin, S. Coemert, I. Iordachita, S. Hoffman, G. Chirikjian, and R. Taylor, "A mosquito pick-and-place system for pfspz-based malaria vaccine production," IEEE Transactions on Automation Science and Engineering, 2020, issn:1545-5955, doi:10.1109/TASE.2020.2992131.

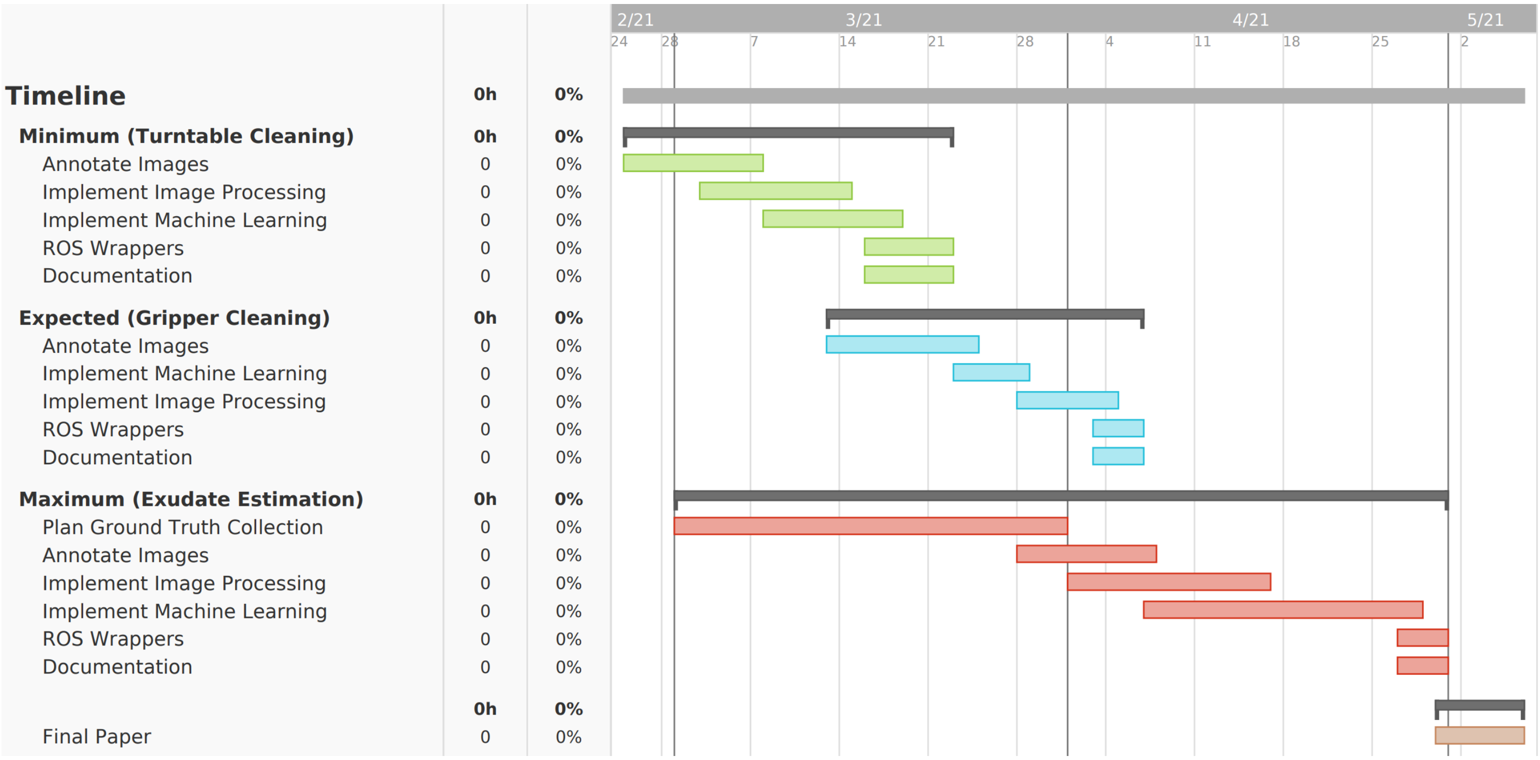

Figure 5: Gantt Chart of the estimated work timeline

S. Lathuilière, P. Mesejo, X. Alameda-Pineda and R. Horaud, ”A Comprehensive Analysis of Deep Regression,” in IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 42, no. 9, pp. 2065-2081, 1 Sept. 2020, doi: 10.1109/TPAMI.2019.2910523.

Lo, Frank Po Wen, Yingnan Sun, Jianing Qiu, and Benny Lo. ”Image-Based Food Classification and Volume Estimation for Dietary Assessment: A Review.” IEEE journal of biomedical and health informatics 24, no. 7 (2020): 1926-1939.

Belagiannis, Vasileios, Christian Rupprecht, Gustavo Carneiro, and Nassir Navab. ”Robust optimization for deep regression.” In Proceedings of the IEEE international conference on computer vision, pp. 2830-2838. 2015.

Chandrarathne, Gayani, Kokul Thanikasalam, and Amalka Pinidiyaarachchi. ”A comprehensive study on deep image classification with small datasets.” In Advances in Electronics Engineering, pp. 93-106. Springer, Singapore, 2020.

Hurtik, Petr, Vojtech Molek, Jan Hula, Marek Vajgl, Pavel Vlasanek, and Tomas Nejezchleba. ”Poly-YOLO: higher speed, more precise detection and instance segmentation for YOLOv3.” arXiv preprint arXiv:2005.13243 (2020).

References

[1] Badrinarayanan, V., A. Kendall, and R. Cipolla (2017). Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence 39*(12), 2481–2495.

[2] Bochkovskiy, A., C.-Y. Wang, and H.-Y. M. Liao (2020). Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*.

[3] He, K., X. Zhang, S. Ren, and J. Sun (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778.

[4] Jongo, S. A., S. A. Shekalaghe, L. P. Church, A. J. Ruben, T. Schindler, I. Zenklusen, T. Rutishauser, J. Rothen, A. Tumbo, C. Mkindi, et al. (2018). Safety, immunogenicity, and protective efficacy against controlled human malaria infection of plasmodium falciparum sporozoite vaccine in tanzanian adults. *The American journal of tropical medicine and hygiene 99*(2), 338–349.

[5] Li, W., Z. He, P. Vora, Y. Wang, B. Vagvolgyi, S. Leonard, A. Goodridge, I. Iordachita, S. L. Hoffman, S. Chakravarty, and R. H. Taylor (2021). Automated mosquito salivary gland extractor for pfspz-based malaria vaccine production. *preprint*.

[6] US Centers for Disease Control and Prevention (2021). “Essential Services for Malaria”. [Online]. Available: `https://www.cdc.gov/coronavirus/2019-ncov/global-covid-19/maintain-essential-services-malaria.html`.

[7] World Health Organization (2020). “World Malaria Report 2020”. [Online]. Available: `https://www.who.int/publications/i/item/9789240015791`.