An Automated Microsurgery in *Drosophila* for High Throughput Brain Imaging

Matthew Golub and Bikiran Goswami
CS229 Machine Learning Final Project

1. Abstract

We demonstrate a supervised learning approach to tracking the location of live, immobilized flies for purposes of performing automated laser microsurgery to prepare flies for *in-vivo* two-photon imaging. We explored least-squares estimation and multi-class logistic regression for classifying pixels in a digital image as either fly eye, fly skin or background. The multi-class logistic regression classifier outperformed the least-squares estimator across nine different experiments conducted with incremental training and testing image sets. We used k-means clustering of machine-labeled image pixels to estimate the location of fly eyes in test images. When clustering pixels classified as eye point by the multi-class logistic classifier, k-means correctly located the fly eyes in 100% of our largest test set, which consisted of sixteen images with varying amounts of defocusing.

2. Introduction

2.1 Background

Drosophila melanogaster, better known as the fruit fly, has long been studied as a model organism. With a well-characterized genome, *Drosophila* genetics can be manipulated in the laboratory with relative ease. Understanding the fruit fly is tremendously relevant to solving problems in human health, since roughly fifty-percent of a fly's protein sequences have an analog in humans. From a neuroscience perspective, *Drosophila* are particularly interesting due to the simplicity of their neural circuits, which perform computation to guide behavior. A major challenge to the *Drosophila* brain research community has been in acquiring statistically significant sets of neural data.

Fly brain imaging requires an unobstructed optical path between brain and microscope. In the past, highly trained biologists have manually dissected the fly head, removing areas of tissue that protect the brain. This painstakingly slow process has severely limited the process of decoding the circuitry of the fly brain in recent years. A high throughput imaging system must utilize a fast, accurate and automated microsurgery to overcome the limitations associated with manual dissection.

We have demonstrated a laser ablation strategy for removing protective tissues surrounding the fly brain. A human operator mounts a fly in an acrylic restraint and uses a commercially available laser microdissection device to ablate a 300x100 micron region above the fly brain. We aim to automate this entire procedure, enabling imaging of tens or hundreds of flies in a single session.

2.2 An Automated Laser Microsurgery

To automate this microsurgery, we must 'teach' our laser microdissection device to intelligently locate the desired ablation region on each fly. From a low magnification image of a mounted fly, the computer must first identify the position of the fly and navigate a motorized stage such that the fly is placed under a set of microscope objectives. Next, the computer must analyze a 10X magnification image of the fly to locate the 300x100 micron region that we are

interested in ablating. Once this region is located, the computer uses a 20X magnification image to autofocus the camera and laser on the fly head.

2.3 Computer Vision Implementation and Drawbacks

We have implemented computer vision techniques to carry out this three-stage positioning process. From the low magnification image, we are able to detect the approximate location of the fly head using geometric cues from the acrylic mount. For more precise tracking of the fly head, we use color information to identify the fly's eyes. Specifically, we apply color intensity thresholding to choose eye seed points from a 10X magnification image, region grow over the seed points, and then dilate the region-grown image to remove false positives. K-means clustering is used to identify the centroid location of each eye. Finally, we apply edge-based autofocusing on a 20X image centered between the fly eyes.

The low magnification detection and high magnification autofocus steps are computationally inexpensive, and they are robust against color and lighting variation since they operate only on grayscale intensity images. The computer vision-based 10X eye detection algorithm is very computationally expensive, however, and thus not well-suited for real time application. Further, the eye detection is highly dependent on color information, which must be hard-coded into the algorithm. These major drawbacks motivated a machine learning-based approach to detecting eyes.

3. Machine Learning Approach to Fly Eye Detection

3.1 Task Definition

We employed supervised learning techniques toward classifying image pixels as belonging to one of three classes: fly eye, fly skin or background. Two independent methods were compared for classifying image pixels: Least-squares (LS) and multi-class logistic regression (MCLR). Each classifier was trained using human-labeled images—regions of eye, skin and background were processed into training examples.

3.2 Algorithm Definition

Each pixel's red/green/blue (RGB) and hue/saturation/value (HSV) color intensities were used as training features for the LS and MCLR classifiers. We implemented a one-versus-all classification scheme for MCLR. Prediction of a new image's pixel labels are made by multiplying each pixel's $[1\ R\ G\ B\ H\ S\ V]^T$ vector by the weights vector computed by LS or MCLR, and thresholding the result accordingly. Across an entire image, this prediction requires far fewer computations compared to the region growing method previously described. After pixels are classified, k-means clustering is used to find the centroid location of detected fly eyes. Example input and output images for LS and MCLR are shown in figure 1.

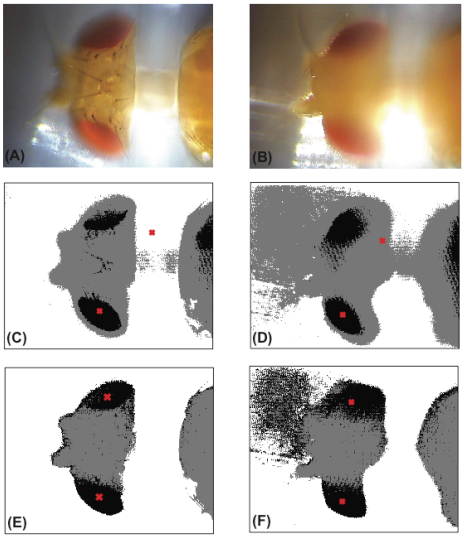

Figure 1. Example classification performance. In-focus (A) and 400um defocused (B) images of a fly head under 10X magnification are classified using LS (C and D, respectively) and MCLR (E and F, respectively). Black points have been classified as eye, gray points as skin, and white as background. K-means estimated fly eye locations are shown in red. In both of these examples, LS fails to correctly locate the eyes, while MCLR succeeds in both examples.

4. Experimental Evaluation

4.1 Construction of Training Image Sets

To evaluate the effect of image focus on training efficiency, we constructed three training image sets consisting of combinations of in-focus, 200um defocused, and 400um defocused images. Each training image set consisted of 10X magnification images with rectangular regions hand-labeled as either eye, skin or background. Only the pixels within these hand-labeled rectangular regions were used as training data. An example training image is shown in figure 2. We chose a small training set, X_{train}^{S}, to include four in-focus images, a medium training set, X_{train}^{M}, to include X_{train}^{S} and three additional 200um defocused images, and a large training set, X_{train}^{L}, to include X_{train}^{M} and three additional 400um defocused images. Note that X_{train}^{L} includes in-focus, 200um defocused, and 400um defocused images. We trained each classifier, {LS, MCLR} on each of the training sets, $\left\{X_{train}^{S}, X_{train}^{M}, X_{train}^{L}\right\}$, resulting in six weight vectors.

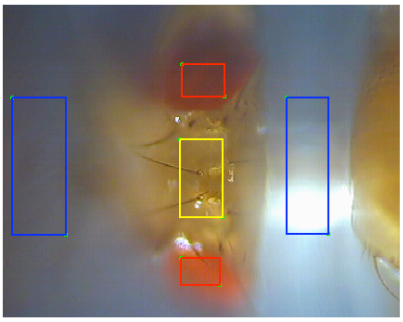

Figure 2. Example hand-labeled training image. Pixels within the red, yellow and blue rectangles are designated as eye, skin, and background, respectively. Pixels outside of these rectangles are not used for training.

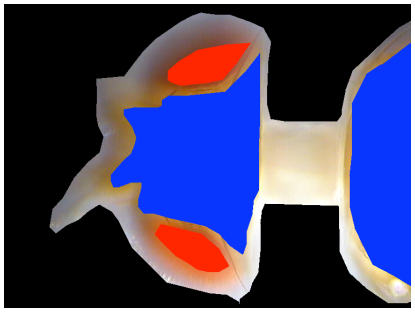

Figure 3. Example hand-labeled testing image. Red, blue, and black regions correspond to eye, skin, and background labels, respectively. Some points remain unlabeled due to the ambiguity of their true class labels. These unlabeled points are not considered while assessing machine-labeling performance.

4.2 Construction of Testing Image Sets

We evaluated machine-labeling performance on three testing image sets consisting of untrained images with identical distribution to the training sets with respect to image focus. We denote the test sets as $\left\{X_{test}^{S}, X_{test}^{M}, X_{test}^{L}\right\}$, for the small, medium, and large training sets, respectively. We applied each of the nine weight vectors from the training phase to each of the three testing image sets, resulting in eighteen machine labeled images.

Machine labeled images were compared against hand-labeled versions of the testing images. Hand-labeled testing images differed from the hand-labeled training images, in that the testing-labels were applied to *every* pixel that could be unambiguously identified as eye, skin or background. Some ambiguous pixels were left unlabeled, to prevent algorithm performance penalization in regions that do not clearly belong to one class or another. An example hand-labeled testing image is shown in figure 3.

4.3 Results

We use a performance index (PI) as a metric for quantifying machine-labeling performance with respect to hand-labeled testing images. We calculate PI according to equation 1.

Equation 1.

$$PI = \frac{\left(\#\,correct\ labelings - \#\,incorrect\ labelings\right)}{total\ \#\ hand\ labeled\ pixels}$$

We calculate a three PI's for each machine-labeled image, corresponding to the three classes, eye, skin, and background. The performance results, summarized in figure 4, show MCLR outperforming LS in all nine experiments. For all classifiers, performance degraded monotonically as the testing set grew from X^{S}_{test} to X^{L}_{test}. Physically, this means that the classifiers had more trouble labeling pixels from out-of-focus images than pixels from in-focus images. For practical purposes, we are most interested in each classifier's performance on X^{L}_{test}, since we would like our algorithm to work regardless of image focus. On X^{L}_{test}, the MCLR classifier performed optimally when trained on X^{S}_{train}. Physically, this means that the in-focus images provided the best information for classifying pixels in variable focus images.

Following machine-labeling, we apply k-means clustering to eye pixels to determine the centroid location of each eye. In all experiments, k-means clustering on MCLR-labeled correctly detected the location of the fly eyes. The LS classifier did not prove nearly as successful; on X^{L}_{test}, k-means succeeded on 72%, 55%, and 45% of images when trained on X^{S}_{train}, X^{M}_{train}, and X^{L}_{train}, respectively.

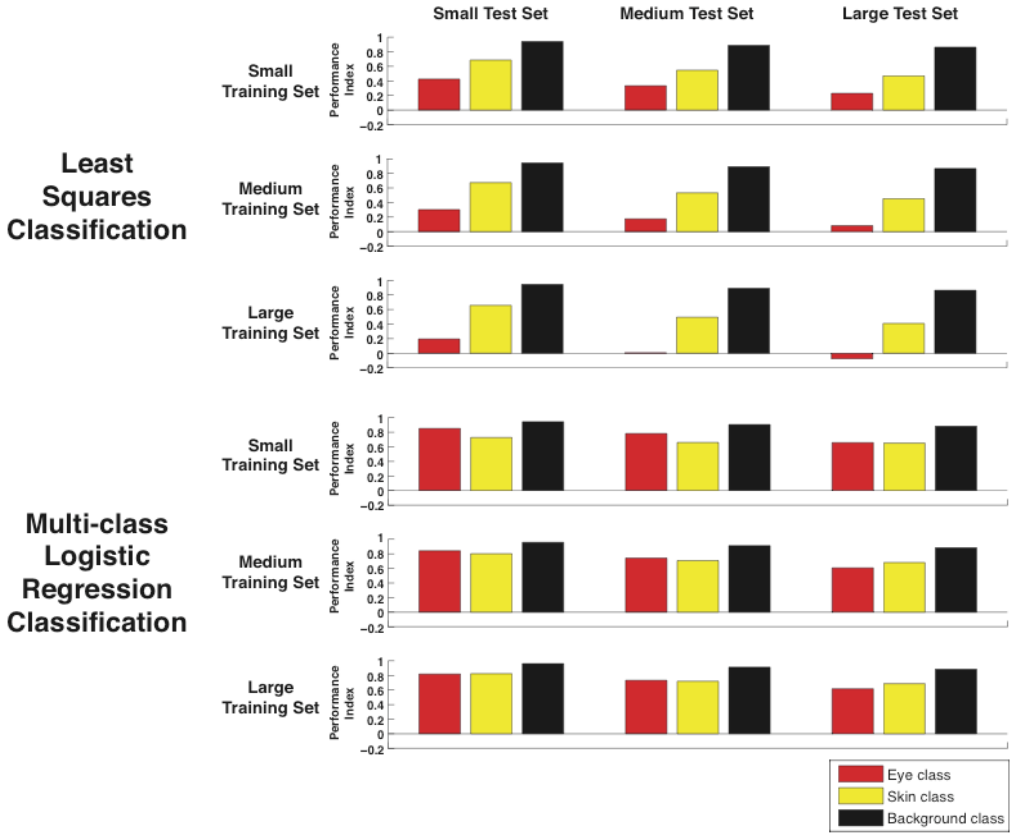

Figure 4. Summary of performance results of the least-squares (LS) and multi-class logistic regression (MCLR) classifiers. Plots across rows correspond to classifier performance when trained on the same image set, but evaluated with respect to different testing image sets. Plots down columns correspond to classifier performance on the same testing image set, after training on different image sets.

5. Future Work

While we are not ready to present results, we are currently working on a multi-class support vector machine implementation (MCSVM) to solve our pixel classification problem. We are interested in using a quadratic kernel according to a $\Re^6 \rightarrow \Re^{36}$ feature mapping. From our experimental results, we found MCLR to yield the best performance. Depending on the success of our MCSVM progress, we will integrate either MCLR or MCSVM directly into the software that controls our laser microdissection device.

We would like to allow for online training, to enable researchers to bring in fly lines that may have radically different appearances compared to the OK107 flies that our classifiers were trained on. Other fly lines may have genetic modifications, such as white eyes—clearly the classifiers would need to be retrained at the beginning of a microsurgery session.

6. Conclusion

We have successfully demonstrated the use of supervised machine learning algorithms to classify regions in high-magnification images of the fly head. We characterized least-squares estimation to multi-class logistic regression classification for our fly tracking purposes. We have shown that MCLR trained on in-focus images results in optimal classification performance on in-focus and out-of-focus images. Using k-means clustering of MCLR-labeled points, we achieved a 100% success rate in locating the position of fly eyes. These results suggest that we will be able to perform high throughput microsurgeries to prepare live flies for two-photon brain imaging.