Determining Gap-Crossing Decision Making Processes in *Drosophila melanogaster* Using Machine Learning

Vinay Bhaip

Abstract

The human brain contains roughly one hundred billion neurons, meaning that even simple human decision making mechanisms are difficult to process at a neuron-based level. To understand the complex systems in the human brain, *Drosophilia melanogaster*, the common fruit fly, which has roughly one hundred thousand neurons, presents a way to map mechanisms and neural circuitry to decision making processes. *D. melanogaster* displays numerous behaviors in which it must make a decision, one of which is called "gap-crossing behavior" when approaching a gap larger than its stride length. This gap-crossing behavior is helpful to analyze since it requires a fly to engage in complex decisions while still using mechanisms common to both humans and flies. This paper presents a system comprised of two machine learning models to automatically extract information, such as the length and number of attempts of gap-crossing behaviors, since manually computing this information can be both time-consuming and subjective. The system achieves an optimal accuracy of 84% on identifying frames with attempts and is used to analyze the decision making processes of flies across various genotypes. This paper identifies neurons in layers 3, 4, and 5 of the fan-shaped body, an area in the fruit fly's brain, as having an important role in allowing for gap-crossing behaviors to occur. This therefore demonstrates the viability of applying machine learning approaches to behavior identification and linked changes in behavior to disruptions in specific types of neurons. This approach can be applied across behaviors and across species to link brain structure to function.

1 Introduction

1.1 Decision Making Processes

When faced with a choice, the human brain has numerous electrical signals passing through its neurons, a process known as decision making, the ability for a brain to discern which action it wants to do. A decision is much more complicated than a normal reflex, a quick response to external stimuli; a decision uses cognitive awareness of the environment to inform a choice. While it is generally understood what different areas of the human brain contribute to decision making, there remains a lack of understanding of the neural circuitry, the underlying connections between individual neurons that determine which actions occur. The human brain's prefrontal cortex, the front-most area of the frontal lobe, is known to play a significant role in decision making. The actual processes and connections in this area that control decision making, however, are largely unknown [1].

Decision making in the human brain has been found to use recurrent excitations of synapses, the connections between neurons [2]. Understanding the specific neurons that engage in this process in the human brain is crucial for advances in neuroscience. Comprehending the decision making processes at a lower neuron-based level is key toward understanding complex diseases like Attention-Deficit/Hyperactivity Disorder (ADHD) and Parkinson's Disease, both partly characterized by impaired abilities to carry out decisions [3].

However, human brains contain nearly one hundred billion neurons, so understanding the underlying processes becomes challenging. Moreover, each of these neurons have tens of thousands of synapses, so creating a connectome, a map of all the connections, becomes difficult [4]. A connectome is crucial to see the pathways and mechanisms that occur in the brain at a neuron-based level in decision making.

To address this issue of an inability to analyze the complex human brain, researchers have looked toward the fly brain, which contains a significantly less number of neurons: about one hundred thousand. The neurons and the synapses in this brain have been almost completely charted out, which means there is a 3D model connectome that allows researchers to see the connections between various neurons [5]. This allows for insights in seeing which neurons are activated in the neural circuitry and how that affects different pathways.

Previous work has found that in the fly brain's central complex, the center of decision making processes in *D. melanogaster*, there exists an internal compass that allows the fly to keep track of its relationship with its environment [6]. This indicates that a fly, similar to a human, is able to understand temporal context from its past decisions to inform its future decisions. The central complex is also associated with motor modulation, the ability for the brain to control body parts like its limbs [6]. This becomes important as it shows that this one region that controls decisions has access to control body parts as it deems it necessary.

1.2 Gap-Crossing Behavior

In daily tasks, the common fruit fly often encounters situations where in its path, it must cross a gap, a larger than normal spacing between two platforms. When *D. melanogaster* approaches a gap, it has three main options to get past it. The first option is for it to simply walk across the gap, but this only works when the gap length is below a specific threshold length. The second option is for it to fly over the gap. Interestingly, it also has a third option: initiating what is called a "gap-crossing behavior."

When the length of the gap is beyond its stride length but not long enough to merit flight, *D. melanogaster* can cross the gap by performing a unique behavior of extending its two frontal limbs, swaying them vertically until they make contact with the opposing edge. Once the fly links onto the opposite side, it swings itself over and proceeds to lift itself up till it is on the other side.

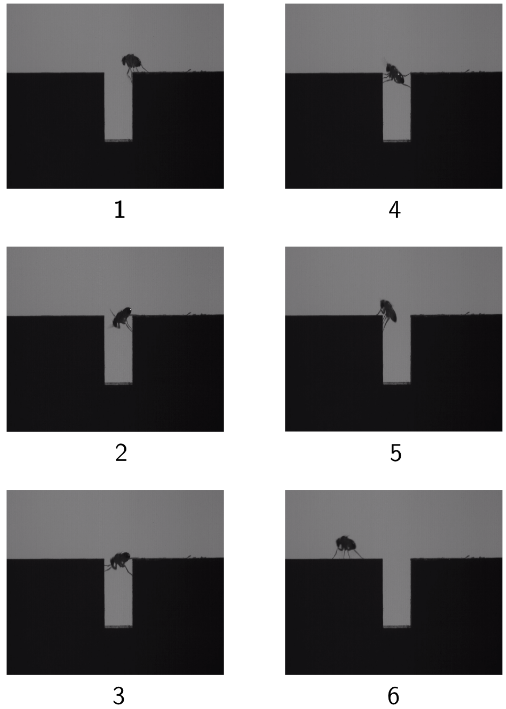

Figure 1: Successful Gap-Crossing Example.

This gap-crossing behavior prompts the question of under which conditions does a fly initiate this behavior. Gap-crossing behavior is interesting to look at it, because a fly has to cognitively assess its environment and determine whether or not to cross. For example, a fly might have to assess whether it has sufficient energy to perform the motor action of moving its limbs.

Previous studies have found that *D. melanogaster* processes the distance of the gap using parallax motion. As flies move around, they store information about their body size in long-term memory. This allows them to compare their stored body size to the gap size and determine whether or not to initiate the gap-crossing behavior [7].

Flies are generally able to cross gaps ranging from 2.5 mm to 4.3 mm, though they cross gaps of smaller distances at much greater rates [8]. In a given attempt, there are several factors which contribute to whether or not a fly will successfully cross. The most obvious factor is the size of the fly and the size of the gap – larger flies are capable of crossing smaller gaps more easily due to their long limbs. Another factor is the location of where the fly begins its attempt in relation to the edge; flies that attempt to gap-cross closer to the edge that they are reaching out from are closer to their target edge on the opposite side. Lastly, the angle θ that flies make with the perpendicular line to the edge determines the length of the actual distance they have to cross. A θ of 0 means that the fly is crossing the gap directly, whereas a θ between 0 and $\frac{\pi}{2}$ means that the fly is angled and therefore is taking a longer path to the opposite side.

Niven found that genetic altercations in a region known as the protocerebral bridge, an area in the aforementioned central complex, affect the ability of a fly to position itself at a minimal θ to cross the gap

[9]. This implicates that this area of the fruit fly brain affects decision making, though it does not shed light on the internal neural circuitry as the gene modifications were not on a neuron-based level.

Wolff et al. shows that the anatomy of the central complex is structured such that it likely integrates sensory information to inform motor movements [10]. The presence of a system that incorporates environmental information suggests that specific neurons play a role in regulating decisions involving moving limbs. Honkanen et al. further shows that since insects have small brains, yet are able to perform complex actions in navigation, it is likely that insects have evolutionarily developed a neural circuitry to process navigation in the central complex [11].

1.3 Genotypical Differences

Understanding what roles various parts of the fly brain contribute in gap-crossing behavior is critical to understanding decision making as a whole. To do this, it is evident there needs to be comparisons of gap-crossing behaviors across various genotypes. Slight modifications in the fly brain that result in specific changes to gap-crossing behaviors can reveal information regarding the purpose of various groups of neurons.

For example, a fly that has a significant proportion of gap-crossing attempts with non-zero θ values might mean that the specific fly's heading-direction neurons, which monitor the orientation of a fly, are off-centered. Additionally, a lack of attempts from a fly with a specific genotype might mean that its stored information about its body size is an underestimation of its true body size. Comparing these phenotypical differences to genotypical differences allows for a correlation between which parts of the brain influence which actions.

1.4 Machine Learning

With the rise of computational power in recent years, machine learning presents a viable solution to problems with large amounts of data. Neural networks are increasingly popular models used to uncover hidden relationships among data, such as images and videos, to approach complicated problems.

Until now, there did not exist an automatic classification system to detect when a fly attempted to cross a gap. Previous studies required researchers to go through thousands of videos to determine details about flies initiating gap-crossing behaviors. This presents an opportunity to speed up the process and create an accurate system that is able to extract even more information from these videos.

2 Methodology

2.1 Data Collection and Preprocessing

The data in this study were collected and used with permission by Tilman Triphan of Howard Hughes Medical Institute's Janelia Research Campus [12]. Each data point comprised of a ten second video filmed at sixty frames per second. Each video shows a fruit fly coming in from one side of a platform and approaching a gap in an attempt to cross it. The videos used in this paper were filmed from a top-view angle. The number of videos of flies with modified genotypes available for analysis (which is the data I would eventually analyze, not the training/testing data my model used) totaled to 730 videos with some videos having multiple attempts of gap-crossing and other videos having no attempts of gap-crossing.

The videos were taken of flies with varying genotypes, which had specific neurons disabled, and gaps that varied between 3.5 mm and 5.0 mm, at 0.5 mm intervals. The experimental setup varied slightly across videos as the original experiment was not designed with the pipeline proposed in this study in mind. The experimental setup also varied within videos themselves due to environmental perturbations, which meant that simply subtracting a static background frame from all the frames would not work.

The data was already labeled on a per-video-level basis, which meant that each ten second video had a label attached as to whether it had successfully crossed the gap, unsuccessfully attempted to cross the gap, or had not attempted to cross the gap. This study required the knowledge of the specific frame where a fly initiates gap-crossing behavior to allow for analysis of a fly making multiple attempts in a single given video.

I labeled 351 videos of the specific frames when flies attempted to cross gaps (with each video comprising of one or more attempts) using the image analysis software Fiji [13]. I noted the success or failure of each attempt, as well as whether the fly was making its attempt on target or not (meaning that if it extended its legs fully, whether or not it would reach the other side). Each video was reflected horizontally to generate more data points, as the orientation of the dataset videos on the horizontal axis was already varying. I generated data points with flies not gap-crossing by picking frames of the gap-crossing videos where the fly was not attempting to cross the gap. Picking the non-attempt data points from the same video as the attempt data points was intentional so that the model does not simply memorize the background but rather generalizes what actions a fly performs in gap-crossing.

The total dataset that I labeled was as follows:

	Attempt	**Nothing**	**Total**
Train	644	644	1288
Validation	160	160	320
Test	200	200	400
Total	1004	1004	2008

Table 1: Breakdown of Data Splits. Each number refer to the number of samples.

The training/validation dataset and the testing dataset was randomly split so that the training/validation dataset used 80% of the data and the testing dataset used 20% of the data. Within the training/validation dataset, I used another 80-20 split so that the training data ultimately was roughly 64% of the total dataset and the validation data was roughly 16% of the total dataset.

2.2 Contextual Data Representation

Given the nature of the problem, it is valuable to have temporal information to discern whether at a given time a fly is crossing a gap or not. To address this, I wrote a script that, given a frame number, will return a list of x frames, with a gap of y frames between each of those consecutive x frames. For the purposes of this paper, I define x as the *"frame context"* and y as the *"frame gap."*

I generated the data points with a frame gap of five (meaning I appended frames to a list skipping every five frames) and a frame context of eleven (meaning there were eleven total frames, centered at the frame number that the event occurred). The value in this is that now each data point did not only contain spacial information, since each frame is an image, but each data point also contained temporal information, since each data point contains multiple frames over time.

2.3 Feature Extraction

While passing in the original images into a model may present a successful system, the model might not be able to understand the underlying motion within the frames. Even with a temporal contextual representation, there exists a need to visualize the velocity of the fly at any moment, as gap-crossing behavior comprises of a fly performing specific motions.

A way to visualize this motion in images is to use optical flow. Optical flow is a vector map between two consecutive frames, showing where each pixel moved from the first frame to the second frame. This vector map is represented as an image, with different magnitudes and directions of vectors corresponding to various colors with variable opacity. Using optical flow allows for a more rich understanding of the motion of the flies, rather than just a static view of a frame. The model will ideally not have to learn about the motion of the fly from scratch, but rather focus on learning unique aspects of the motion that lend themselves to classification of gap-crossing attempts. I wrote a program using OpenCV to extract the optical flow of the datapoints [14].

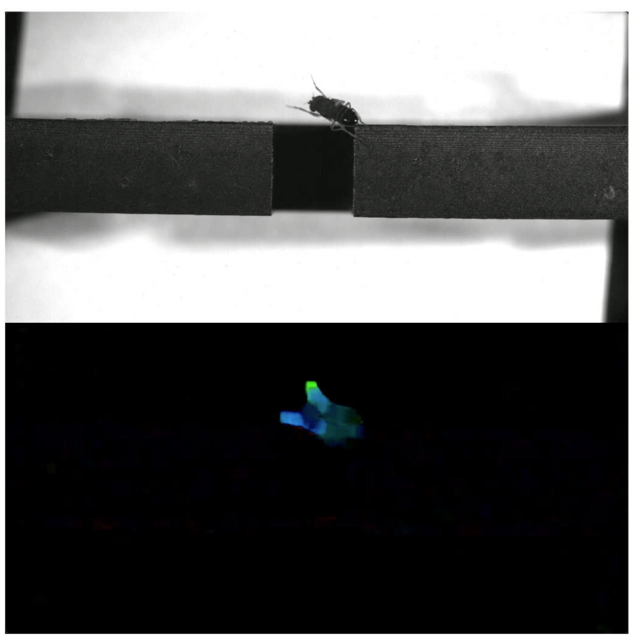

Figure 2: Optical Flow Representation of Frame for Gap-Crossing Fly.

2.4 Model

Since the datapoints had three dimensions (two given by the image, and the third given by the presence of multiple images), I accordingly used a 3D Convolutional Neural Network (3D CNN). The model takes in the data and outputs a two element vector, one element corresponding to the probability that the datapoint is an attempt and the other element corresponding to the probability that the datapoint is not an attempt.

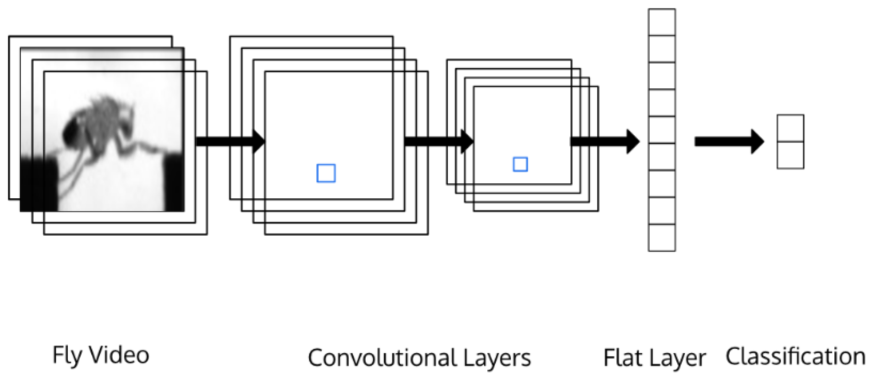

Figure 3: Convolutional Neural Network Model Example. My project uses 3D convolutions to process the data.

I generated two models to train on the data. One model trained on the datapoints corresponding to videos filmed normally. The second model trained on the optical flow representation of the normal videos. Training two separate models allows for a use of both of them in a system such that each model attempts to account for the weaknesses of the other. While the optical flow model might have a richer view of motion, it might miss important details by overgeneralizing the data, for example. To account for this, both models are used in the final system. To make a prediction on a given data point, the datapoint is passed into both models. The prediction is the label of the model with more confidence.

This model design was based off the current state-of-the-art I3D Two Stream Convolutional Neural Network that combines two models, one that takes in the original data and another that takes in the optical flow data [15]. The I3D model itself was not used for this study as its network structure is complex and contains numerous adjustable parameters to the point where it would just memorize the gap-crossing dataset. Instead, I based the architecture of my model off the I3D Network, but significantly simplified the model in hopes that the system would be able to learn underlying information from the data and be quick enough to use in practice.

I built my model in Python using Keras, a common library used for machine learning [16]. The model consisted of two 3D convolutional layers, followed by three standard fully-connected layers. The purposely-chosen simplistic architecture ensures that this model is feasible as a tool in the future, since it would not take too long to run on videos. If a more complex model was chosen, even if there would be an increased accuracy, it would not be practical to use at a large-scale level for analyzing numerous videos since it would be too slow.

2.5 Frame-Level Behavior Identification

To find out in a single video when a fly is initiating gap-crossing behavior, I run the model at fifteen frame intervals to generate a list of predictions with probabilities that a fly is attempting to cross a gap. These probabilities can be graphed, and distinct peaks can be attributed to distinct attempts.

This generates two different graphs, since there are two different models. The probabilities are averaged to create one system of predictions, and then a rolling average with a window of five datapoints smoothens out the curve. Whenever the graph exceeds a 50% confidence in its prediction, the timestamp of the attempt as well as the length of the attempt is recorded. This information is then attributed to the genotype of the fly.

3 Results

3.1 Detection Results

	Attempt	**Nothing**
Attempt	605	199
Nothing	18	786

Table 2: Confusion Matrix for Normal Model on Training and Validation Data. True labels are on the vertical axis and predicted labels are on the horizontal axis.

	Attempt	**Nothing**
Attempt	668	136
Nothing	28	776

Table 3: Confusion Matrix for Optical Flow Model on Training and Validation Data. True labels are on the vertical axis and predicted labels are on the horizontal axis.

Tables 2 and 3 show how each model performed on the training and validation data and the breakdown of each model's performance for the specific classes. This comprised of roughly 80% of the total data available for making this model.

	Attempt	**Nothing**
Attempt	150	50
Nothing	6	194

Table 4: Confusion Matrix for Normal Model on Testing Data. True labels are on the vertical axis and predicted labels are on the horizontal axis.

	Attempt	**Nothing**
Attempt	140	60
Nothing	13	187

Table 5: Confusion Matrix for Optical Flow Model on Testing Data. True labels are on the vertical axis and predicted labels are on the horizontal axis.

Tables 4 and 5 show how each model performed on the testing data and the breakdown of each model's performance for the specific classes. This comprised of roughly 20% of the total data available for making this model.

On the validation data, the normal model achieved an 84% accuracy and the optical flow model achieved an 87% accuracy. On the testing data, the normal model achieved an 86% accuracy and the optical flow model achieved an 82% accuracy.

	Precision	**Recall**	**F1 Score**
Normal Model	0.962	0.750	0.843
Optical Flow Model	0.915	0.700	0.793
Average	0.939	0.725	0.818

Table 6: Classification Statistics for Normal Model and Optical Flow Model on Testing Data.

Table 6 shows the classification statistics, which includes the precision, the recall, and the F1 score, for each of the models. The average values of these statistics for the dual system that this paper proposes are listed as well.

3.2 Prediction System

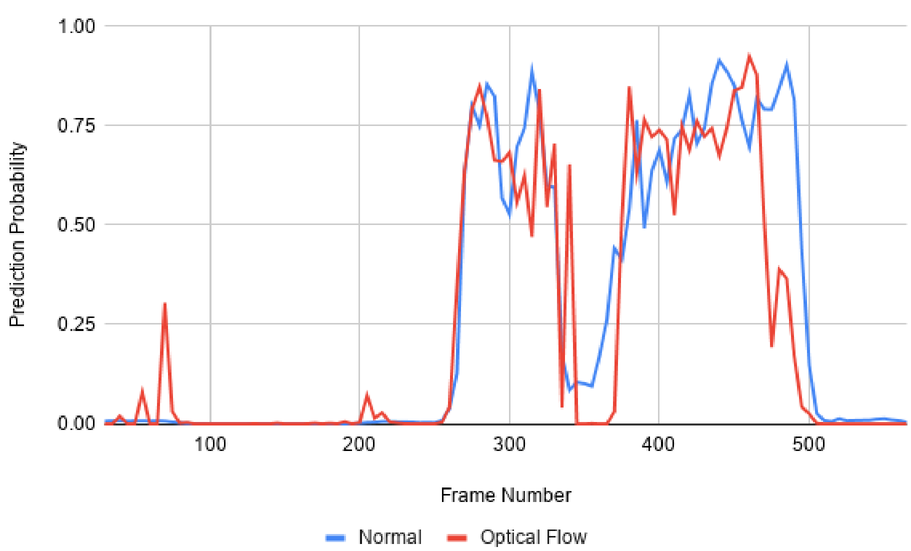

Figure 4: Raw Predictions on Sample Testing Video.

Figure 4 shows how each model can be used to produce a graph of the probabilities of gap-crossing occurring in a video. To illustrate each model's predictions, this graph was generated by running the model after every five frames in a sample testing video.

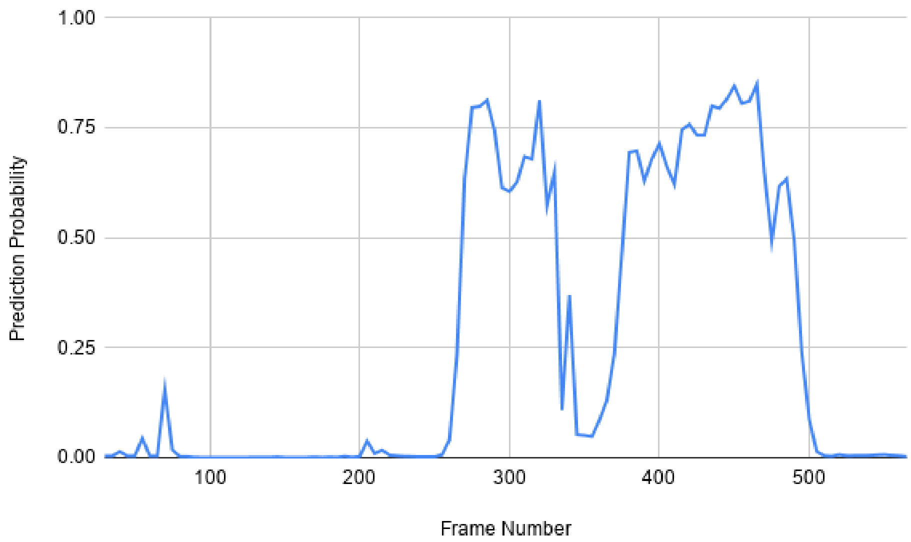

Figure 5: Averaged Predictions on Sample Testing Video.

Figure 5 shows the averaged predictions of both the normal model and the optical flow model for the same testing video in Figure 4.

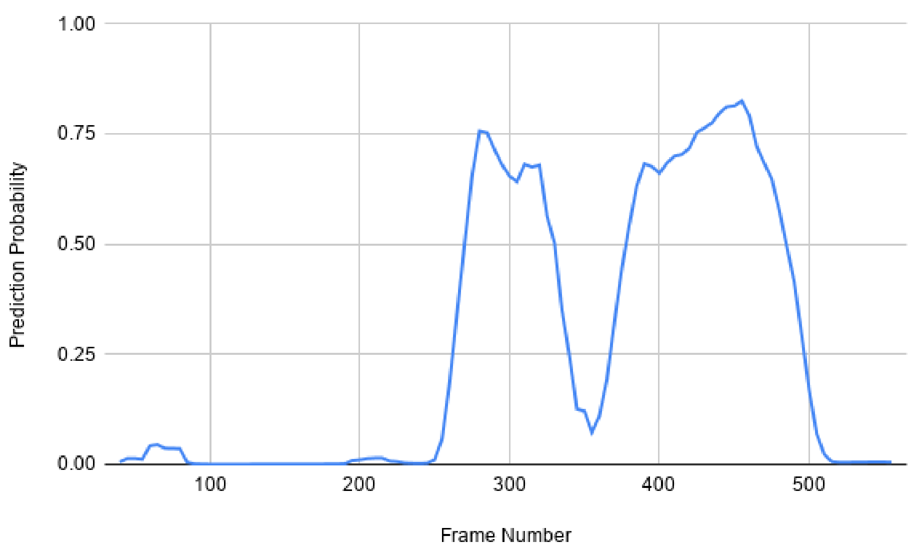

Figure 6: Smoothened Averaged Predictions with Window of Five on Sample Testing Video.

Figure 6 shows a rolling average implemented on the graph from Figure 5. This leads to a smoothening of the curve and allowing for distinct events of gap-crossing to be more visible.

3.3 Extracted Attempt Information

	Num Attempts Avg	Length Avg	Length SD	N
Control	1.079	106.622	74.698	38
Layer 1	1.071	62.308	45.893	14
Layer 3/4	1.000	50.357	33.831	20
Layer 4/5	1.000	57.500	30.516	13
Layer 6	1.200	76.000	50.636	15
Layer 8	1.066	87.119	48.915	168

Table 7: Extracted Information from Attempts. Num Attempts Avg is the average number of times a fly attempts to cross a gap in each video. Length Avg and Length SD is the average length and standard deviation of an attempt that a fly makes, respectively. The samples analyzed only include videos marked as including a successful gap-crossing event. The labels on the vertical axis correspond to which neurons had been inhibited in the genotypes tested.

	Attempt Percentage	N
Control	65.766	111
Layer 1	75.676	37
Layer 3/4	81.132	53
Layer 4/5	59.459	37
Layer 6	87.879	33
Layer 8	64.270	459

Table 8: Percentage of at Least One Attempt in a Video. The samples analyzed include both videos where flies attempted to gap-cross and also where flies did not do anything. The labels on the vertical axis correspond to which neurons had been inhibited in the genotypes tested.

4 Discussion

4.1 Model Discussion

The optical flow model did better on the validation data, gaining an 87% accuracy compared to the 84% accuracy of the normal model; however, on the testing data, the normal model outperformed the optical flow model with an accuracy of 86% compared to an accuracy of 82% (for an average accuracy of both models of 84%). An explanation for this could be that the optical flow model was able to understand the motion of the videos better in the training dataset, but was not able to extrapolate that to the testing dataset. The optical flow model might not have performed as well as the normal model did on the testing dataset because the optical flow representation loses some information about the fly's relationship to its environment, which can be important for discerning gap-crossing behavior.

The similar accuracy of each model on their own testing dataset and their respective validation dataset indicates that both models learned the underlying features that indicate an attempt is occurring at a given frame. If both models had a significantly lower accuracy on the testing dataset than the validation dataset, there would be reason to believe that the models had just memorized insignificant noise from their training data.

Both models have a high precision, with an average precision of 0.939, indicating that whenever the models predict that there is a gap-crossing attempt, there's a high probability that the models are correct. This is especially valuable because it means that the system is not going to mark that many false-positives that would likely inflate the number of times that it predicts a fly is crossing a gap.

On the other hand, the models have a relatively low recall, with an average recall of 0.725, so the models might be missing more attempts when they occur. While this is not ideal, for this paper, it is better to have a higher precision rather than a higher recall because it ensures that the predicted number of attempts is not exaggerated.

While the normal model outperforms the optical flow model, it is still valuable to combine both models into the prediction system, as the optical flow model might be able to pick up on certain nuances that the normal model cannot pick up. Figure 4 indicates that the optical flow model has a more erratic prediction system than that of the normal model, but it also is more confident about its predictions at times. For example, at around frame 350, the normal model has a prediction of roughly 0.1 whereas the optical flow model has a prediction of roughly 0. The confidence of the optical flow model can be an advantage to the system in the way that if the optical flow model is extremely confident, it can sway the prediction one way, whereas if the optical flow model is not as confident, then the more stable normal model can take control of the prediction.

Figure 6 indicates that the implementing a rolling average with a window of five provides a successful method to smooth out any erratic predictions. It is important to note that between the two peaks that correlate to attempts in Figure 6, the prediction is not as close to 0 as it is before the attempts. This might be attributed to the lack of training data points of gap-crossing occurring between attempts.

4.2 Genotypical Analysis

The genotypes analyzed each inhibited different layers of neurons in the fan-shaped body, a component of the central complex area of the brain that researchers believe is involved with decision making. These layers of neurons, numbered 1 through 9, possibly each receive unique environmental cues as feedback in the neural circuitry [17].

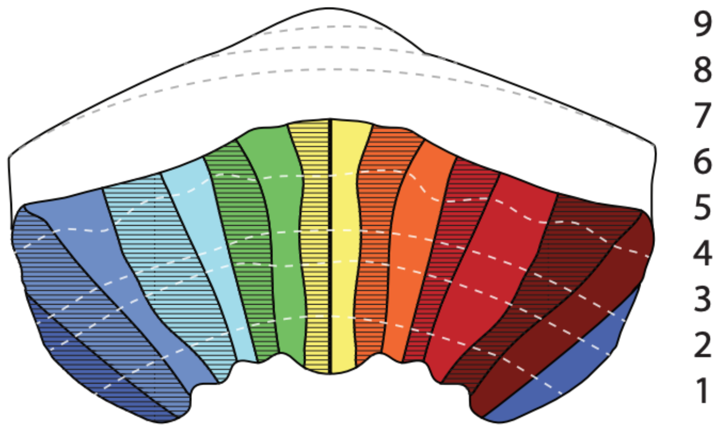

Figure 7: Fan-Shaped Body Diagram. Different contextual environmental information is received from each layer of neurons. Used with permission from Daniel Turner-Evans.

Figure 7 shows the different layers of the neurons that can be modified. This paper analyzed the modification of layers 1, 3, 4, 5, 6, and 8, as the flies with these layers of neurons modified had a sufficient number of videos of gap-crossing flies (at least ten). The videos that were available for me to test had modified lines that inhibited different layer neurons. Pan et al. found that layers 1 and 5 neurons were used for storing patterns it saw in flight [18].

Table 7 shows that the flies with layer 4 and 5 neurons inhibited had an attempt length of roughly 55 frames compared to the control's average attempt length of roughly 105 frames, meaning that flies with this genotype have an attempt length of roughly $\frac{1}{2}$ the attempt length of the control flies. This supports the idea that layer 4 and 5 neurons are heavily involved in receiving environmental cues and processing those for gap-crossing behavior. Since the flies with those neurons inhibited have attempts with shorter lengths, this could mean that layer 4 and 5 neurons play an important role in gap-crossing, like comparing the gap size to their body size. As noted before, flies are known to store their body size from a young age [7], so one possibility is that the layer 4 and 5 neurons could be processing the gap size and comparing it to this stored body size to discern whether or not to cross the gap.

A two-sample t-test comparing the average lengths of attempts between the control flies and the flies with layer 4 and 5 neurons inhibited gives a two-tailed p-value of 0.0261. This p-value is less than 0.05, meaning that the average length of attempts for this genotype being less than that for the control flies is statistically significant.

Table 7 also shows that flies with layer 3 and 4 neurons inhibited have short lengths of attempts just as the flies with layer 4 and 5 neurons do. Moreover, both types of these flies attempt to cross the gap only once on average in the videos analyzed. This seems to indicate that these flies are unlikely to make multiple attempts to cross a gap. In contrast to the flies with layer 4 and 5 neurons modified, the flies with layer 3 and 4 neurons modified attempt to cross the gap at a much higher percentage, according to Table 8. A possible explanation for this could be that layer 3 and 5 neurons have opposite functions of discerning whether a fly should attempt to gap-cross. In this case, the layer 4 neurons are what discern how long the fly should attempt before giving up.

Another two-sample t-test for the flies with layer 3 and 4 neurons modified and the control flies gives a two-tailed p-value of 0.0023, which is statistically significant. This p-value is also much less than the p-value from the t-test involving the layer 4 and 5 neurons, suggesting a stronger connection between layer 3 and 4 neurons to gap-crossing.

The flies with the layer 8 neurons modified have similar attempt statistics to those of the control flies. Table 8 shows that the attempt percentage of these flies is very similar to the attempt percentage of the control flies with both percentages around 65%. This indicates that layer 8 neurons play a very small role, if not no role, when it comes to the decision making of a fly for whether or not to initiate gap-crossing behavior.

The clear difference between how some layers have an evident affect on gap-crossing abilities and how other layers have a very little affect on gap-crossing abilities indicates that decision making processes are localized to specific layers of neurons, and not something the entire region takes part in.

4.3 Future Work

The model proposed in this paper is successful, but in order to extract more meaningful analysis, requires more videos of flies with different neurons inhibited to analyze. The t-tests show statistical significance, but a larger N would help to corroborate these findings. This paper was limited by the relatively small number of videos available for each type of fly. In the future, researchers would use this model to test more samples for the types of flies that I tested to confirm the results. This paper also does not analyze flies with neurons in layers 2, 7, and 9 inhibited due to a lack of data available, so the next step in this work would be to run my model on videos with other flies.

Moreover, there are other regions in the central complex fly brain, like the protocerebral bridge and the ellipsoid body that are postulated to be involved in decision making [17]. Future work consists of analyzing videos of flies with neurons in these regions modified and seeing if there is an effect on the ability of gap-crossing.

5 Conclusion

This paper proposes a successful system that is able to discern whether or not a fly is attempting to cross a gap at a given frame. The system comprises of a dual model setup that uses machine learning. One model takes in preprocessed images from the video of the fly, whereas the other model takes in optical flow representations from the preprocessed images. Both models are used to predict whether a fly is gap-crossing at a given frame.

This system can be run on videos to output graphs of the probability that a fly is gap crossing over the frames in a video. These graphs by themselves provide valuable information as to what a fly is doing throughout a video. The peaks of these graphs are interpreted as when the fly is attempting to cross a gap. Information about each gap-crossing attempt, like the average number of attempts in a video and the average length of each attempt, can also be extracted.

Applying this system to a collection of videos of flies with varying genotypes that have specific layers of neurons inhibited reveals that neurons in layers 3, 4, and 5 of the fan-shaped body play a significant role in decision making processes. One possible explanation of what these neurons do is that they take in the environmental context to discern what is the length of the gap in relationship to their body size. When these neurons are inhibited, they may be unable to process the gap size's true length and have its ability to make comparisons altered.

The work done in this paper helps pave the way for processing the neural circuitry of the brain. This paper shows that within general areas of the brain, specific neurons can have vastly different functions in decision making. Understanding the functions of all these neurons provides for a better way to process how diseases that affect decision making work. The work in this paper sets up a framework for future testing of *D. melanogaster* to find out more functions of different neurons in the fruit fly brain. This allows for a richer understanding of the specific neurons and their functions in the brain's neural circuitry.

References

[1] Coutlee, C. G., & Huettel, S. A. (2012). The functional neuroanatomy of decision making: prefrontal control of thought and action. Brain research, 1428, 3-12.

[2] Wang, X. J. (2008). Decision making in recurrent neuronal circuits. Neuron, 60(2), 215-234.

[3] Lee, D. (2013). Decision making: from neuroscience to psychiatry. Neuron, 78(2), 233-248.

[4] Mitchum, R. (2018). Neuroscientist leads unprecedented research to map billions of brain cells. Retrieved November 11, 2019.

[5] Zheng, Z., Lauritzen, J. S., Perlman, E., Robinson, C. G., Nichols, M., Milkie, D., ... & Calle-Schuler, S. A. (2018). A complete electron microscopy volume of the brain of adult Drosophila melanogaster. Cell, 174(3), 730-743.

[6] Turner-Evans, D. B., & Jayaraman, V. (2016). The insect central complex. Current Biology, 26(11), R453-R457.

[7] Krause, T., Spindler, L., Poeck, B., & Strauss, R. (2019). Drosophila Acquires a Long-Lasting Body-Size Memory from Visual Feedback. Current Biology, 29(11), 1833–1841.e3.

[8] Pick, S., & Strauss, R. (2005). Goal-driven behavioral adaptations in gap-climbing Drosophila. Current Biology, 15(16), 1473-1478.

[9] Niven, J. E. (2010). Visuomotor control: Drosophila bridges the gap. Current Biology, 20(7), R309-R311.

[10] Wolff, T., & Rubin, G. M. (2018). Neuroarchitecture of the Drosophila central complex: A catalog of nodulus and asymmetrical body neurons and a revision of the protocerebral bridge catalog. The Journal of comparative neurology, 526(16), 2585.

[11] Honkanen, A., Adden, A., da Silva Freitas, J., & Heinze, S. (2019). The insect central complex and the neural basis of navigational strategies. Journal of Experimental Biology, 222(Suppl 1), jeb188854.

[12] Triphan, T., Nern, A., Roberts, S. F., Korff, W., Naiman, D. Q., & Strauss, R. (2016). A screen for constituents of motor control and decision making in Drosophila reveals visual distance-estimation neurons. Scientific reports, 6, 27000.

[13] Schindelin, J., Arganda-Carreras, I., Frise, E., Kaynig, V., Longair, M., Pietzsch, T., ... & Tinevez, J. Y. (2012). Fiji: an open-source platform for biological-image analysis. Nature methods, 9(7), 676.

[14] Mordvintsev, A., & K, A. R. (2013). Optical Flow. Open-CV Python Tutorials.

[15] Carreira, J., & Zisserman, A. (2017). Quo vadis, action recognition? a new model and the kinetics dataset. In proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 6299-6308).

[16] F. Chollet et al. (2015). "Keras."

[17] Pfeiffer, K., & Homberg, U. (2014). Organization and functional roles of the central complex in the insect brain. Annual review of entomology, 59, 165-184.

[18] Pan, Y., Zhou, Y., Guo, C., Gong, H., Gong, Z., & Liu, L. (2009). Differential roles of the fan-shaped body and the ellipsoid body in Drosophila visual pattern memory. Learning & Memory, 16(5), 289-295.