Exploring the potential of computer vision and machine learning in enhancing the functionality of an EMG-controlled prosthetic hand

Jethro Odeyemi, M.Sc. Defense

Supervisor: W. J. (Chris Zhang)

Department of Biomedical Engineering

May 2023

Problems with existing Myo-prosthetics

Training time for users to gain accurate control of Electro-Myography (EMG) prosthetics is lengthy, taking several months on average

Prosthetics with automatic control methods negatively trade-off on

- Weight
- Cost
- Appearance
- Size

Knowledge gap

This research addresses the knowledge gap surrounding the absence of prosthetic arms capable of **automated object gripping** and **hand gesture performance**, without imposing the negative impacts of **extensive user training** and **excessive weight** in the prosthetics.

It specifically focuses on tackling this gap from a software perspective.

Overall Objective

The overall objective of this thesis is to explore the application of computer vision and machine learning technologies in prosthetic hands to automate the process of performing hand gestures and grip control of the prosthetic hand.

Specific objectives

- Train tree-based classification models in accurately interpreting EMG signals
- Train an agent using existing SAC, PPO and DQN algorithms to enable the prosthetic hand's end effector to automatically grasp objects and to examine how each of the mentioned models affects the hand’s ability to effectively grasp objects.
- Test the optimal model on varying objects to examine how object’s physical properties affect prosthetic’s hand ability to effectively grasp objects.

Experiment 1: EMG signal data analysis for hand gesture estimation

The goal of this experiment is to develop a model that can interpret EMG signals in predicting specific hand gestures.

Procedure:

- ❑ EDA
- ❑ Data preparation
- ❑ Model training

Data explanation

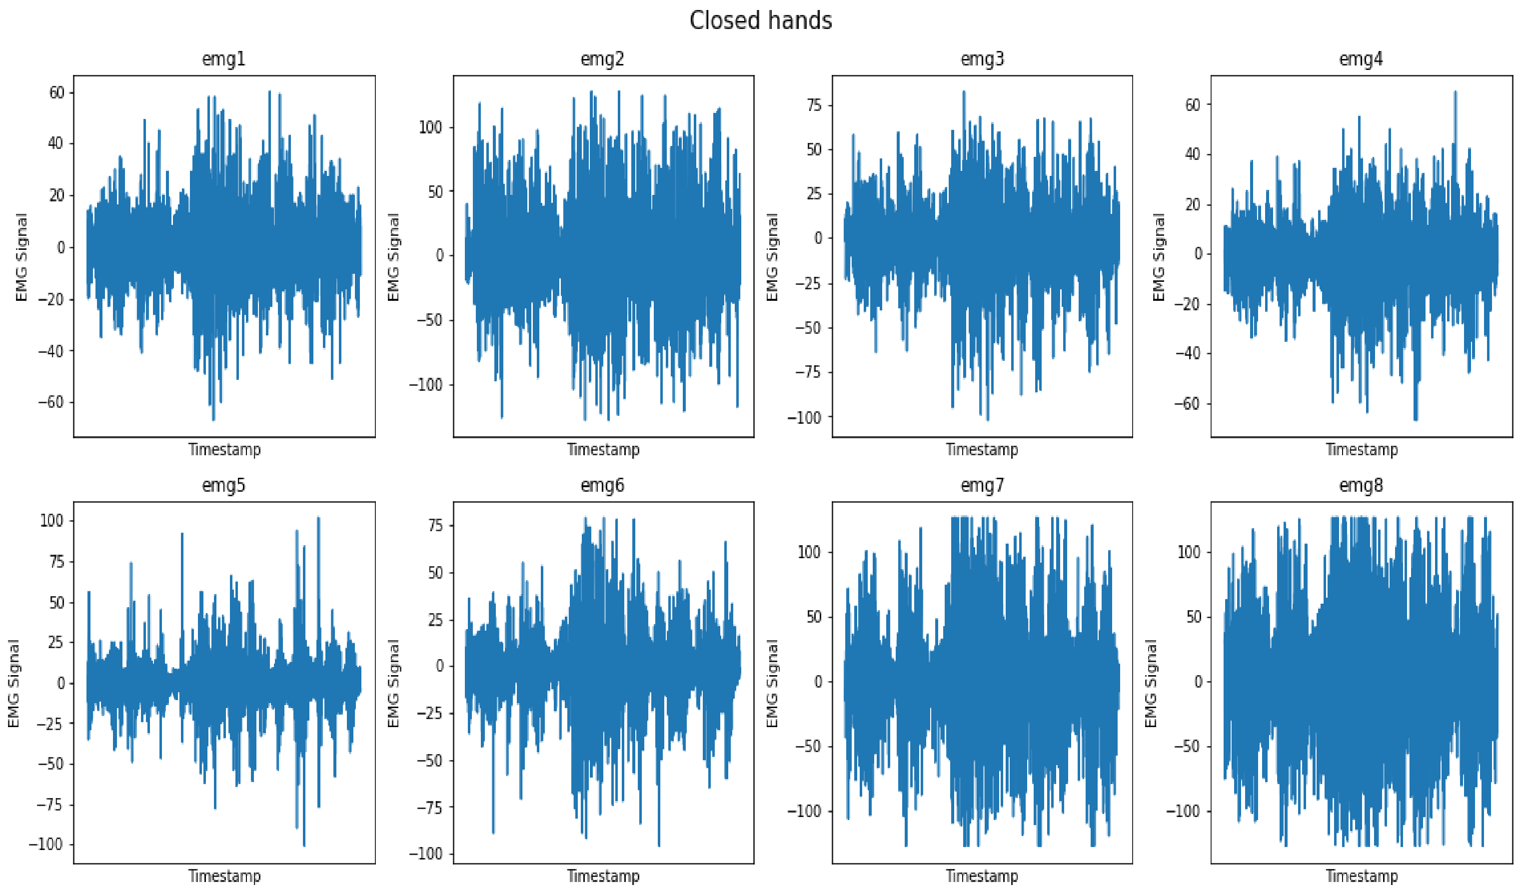

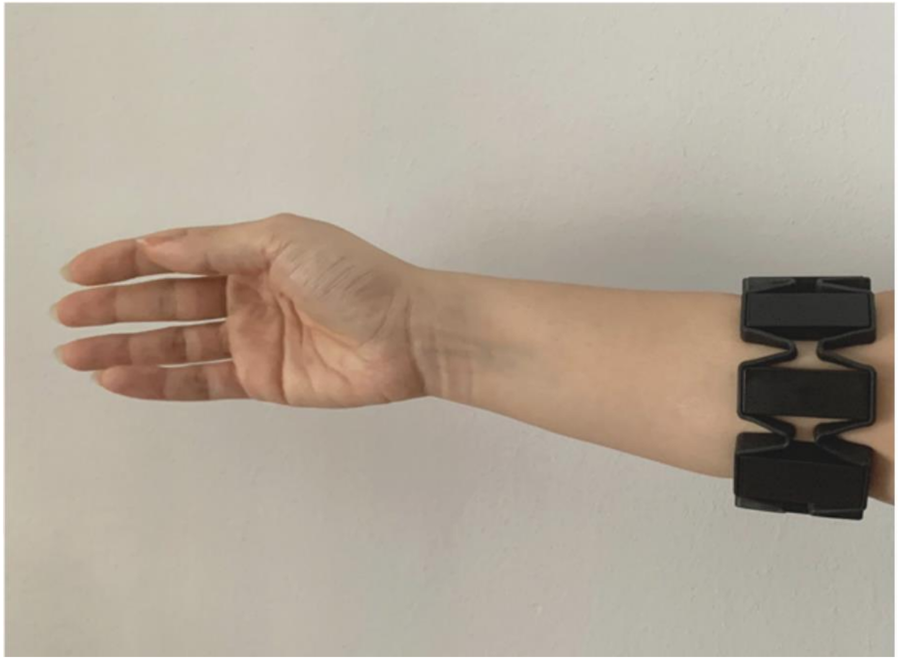

Source: Jiang, et al., 2022

Data preparation

Standard Scaling

- ❑ Improves algorithm performance and data interpretability
- ❑ Scale the mean to 0 and standard deviation to 1

Univariate Feature Selection

- ❑ Calculates F-value for features, converts to p-value

Hyperparameter Tuning

- ❑ Grid search cross-validation used to find optimal hyperparameters

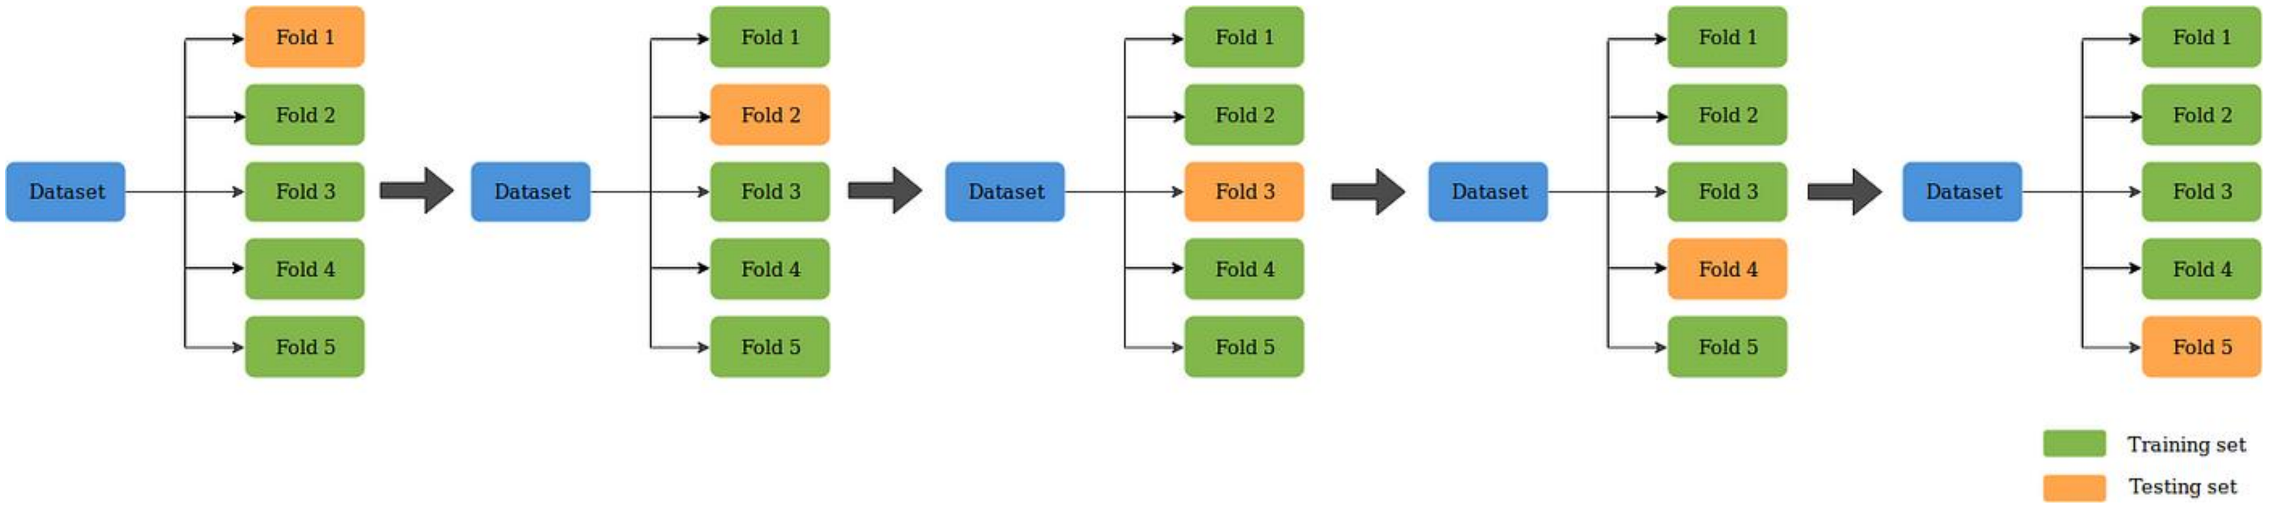

Model training – Decision Tree and Random Forest

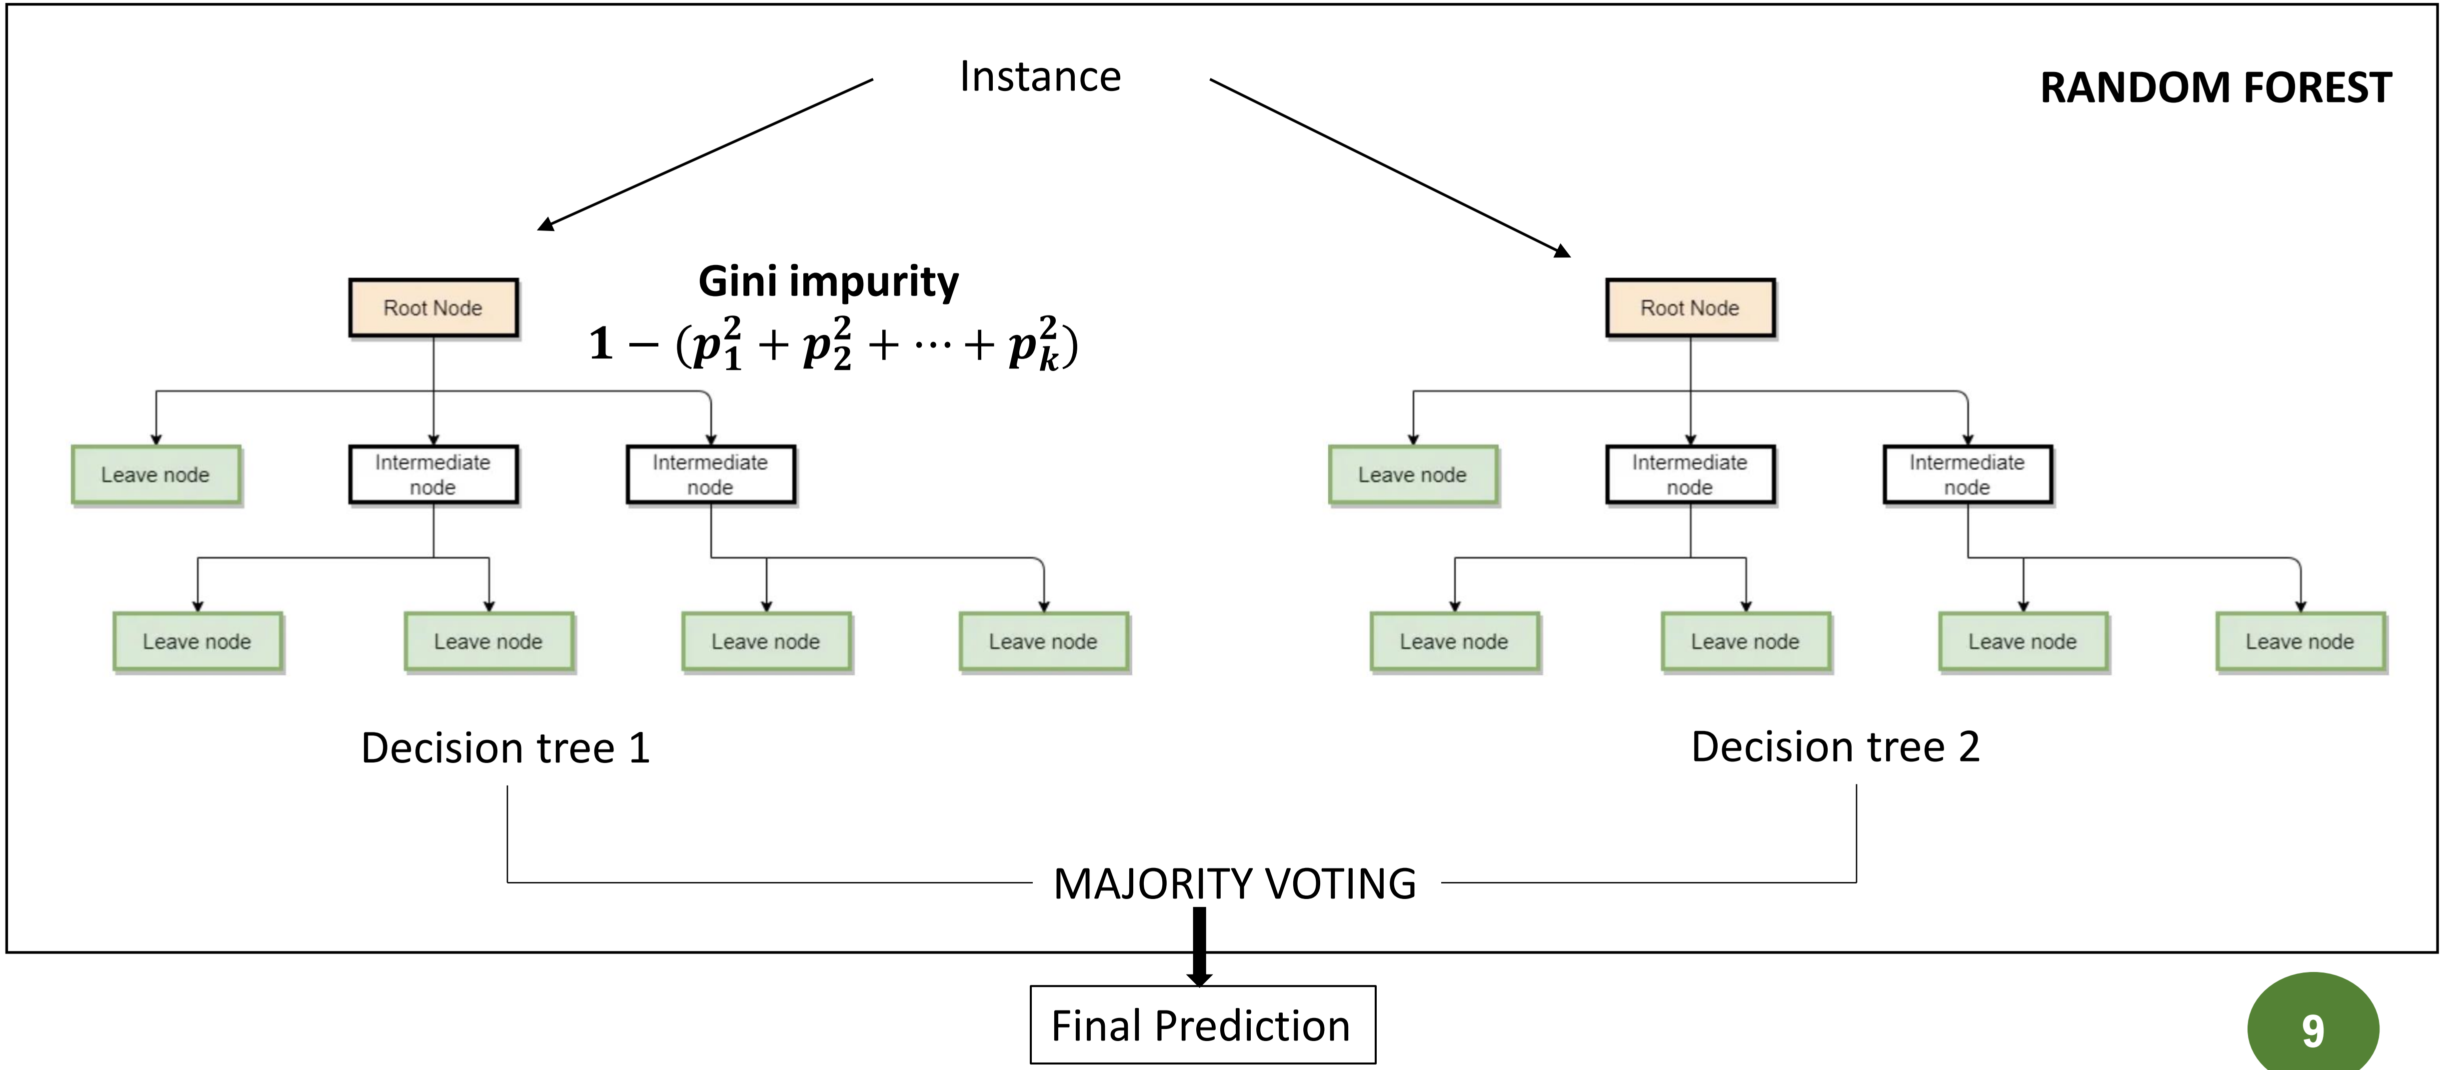

Experiment 2: Automated grip control

This experiment is setup to train the simulated prosthetic arm gripper to be able to grip an object in minimum time and with optimal force.

The training environment scene consists of a gripper, plane, and random objects spawned on a table.

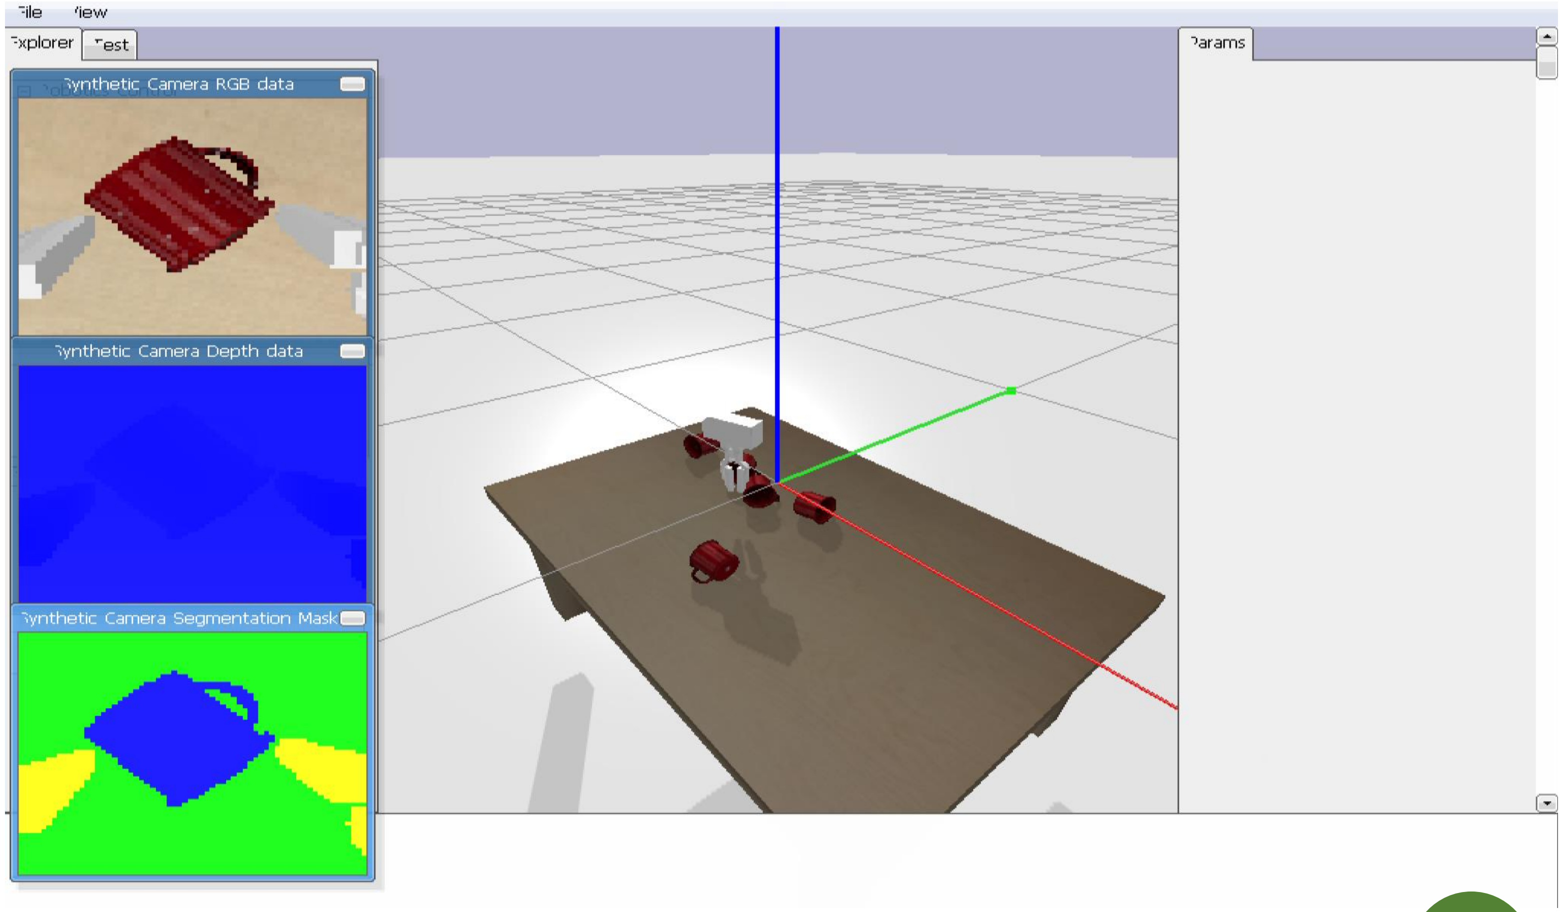

Environment Setup

- ❑ An RGBD camera mounted on the gripper provides RGB data, depth data, and segmented mask information.
- ❑ Gripper movements are determined by a collection of values representing translation, yaw rotation, and gripper open/close actions.
- ❑ Rewards or punishments are given in each episode based on task completion and time taken.
- ❑ A shaped reward function is used, specifically designed for the gripping task.

Sensor

The sensor used in the task is a camera mounted at the midpoint of the gripper's base.

Observations captured by the sensor represent the state space in the Markov decision process.

A perception pipeline is implemented using RGBD observations, which combine RGB color data with depth information.

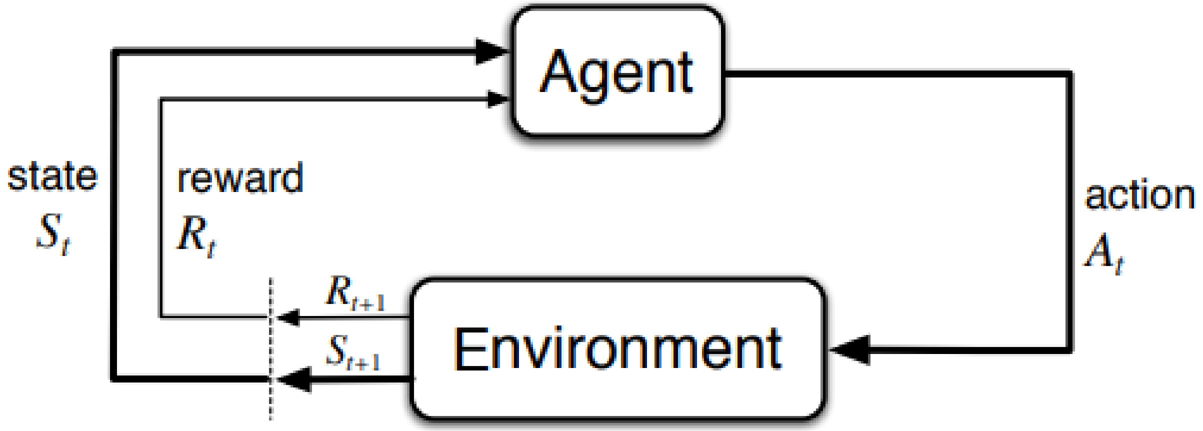

Sensor data processing

- ❑ Gripper width information is padded into a three-dimensional array as an extra channel.
- ❑ During training, the extra channel is removed, allowing the robot to learn how to use the sensor data and gripper width.
- ❑ Convolutional neural networks are used to learn features from the RGBD data.
- ❑ Gripper width information is concatenated with the RGBD tensor to form the observation vector.

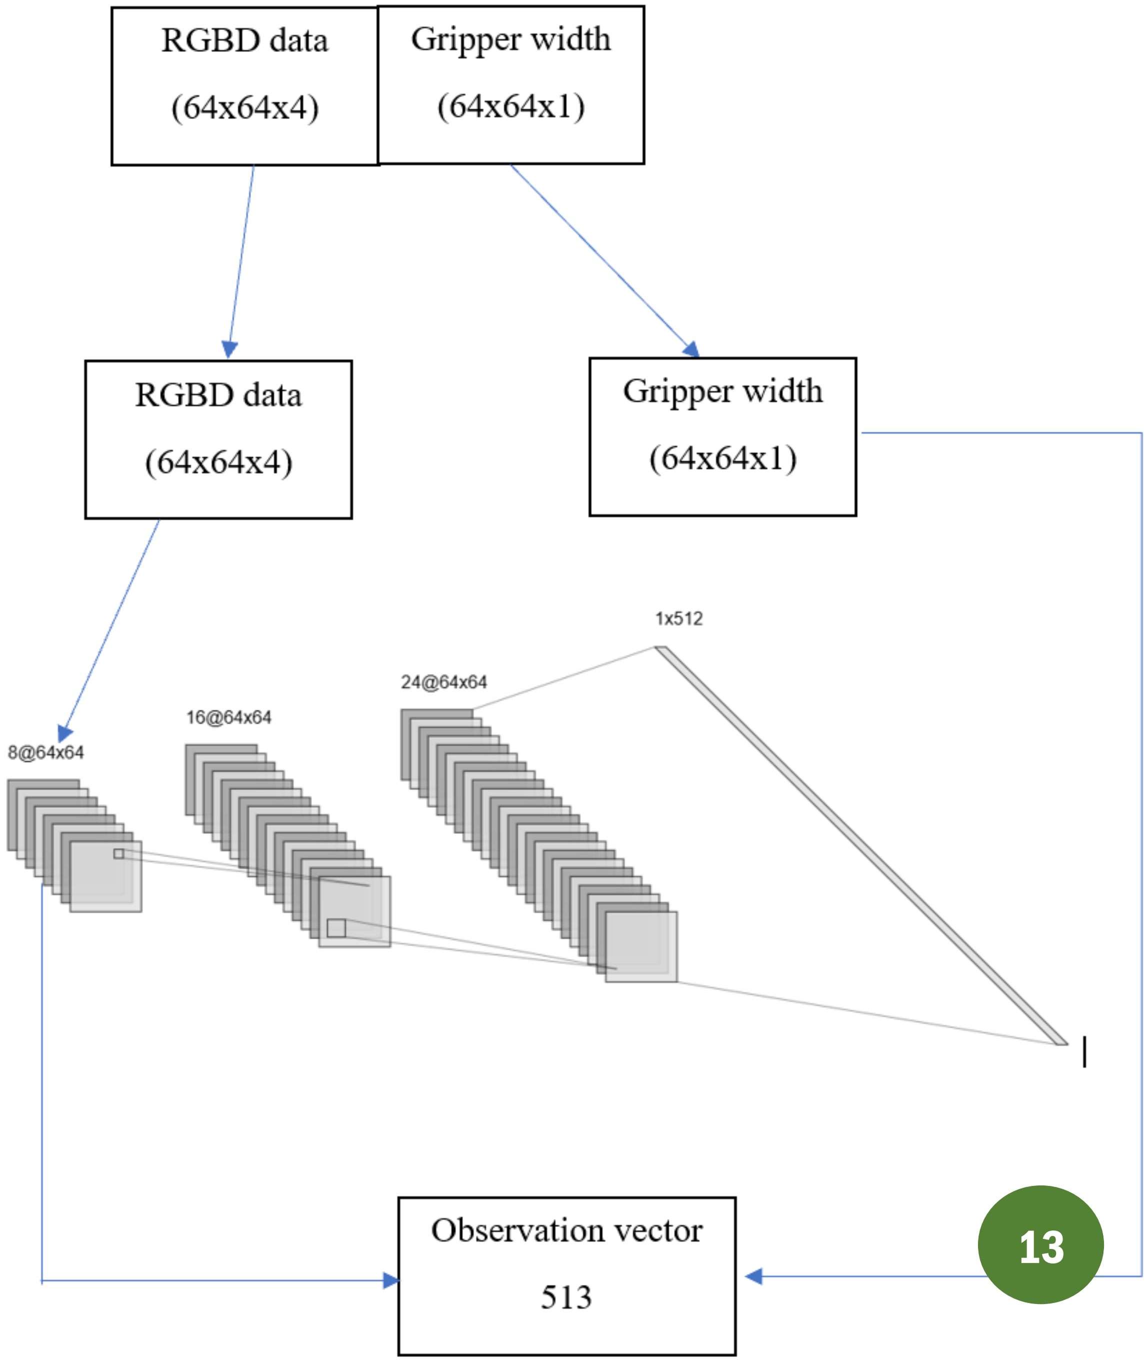

Reward Function

A reward function was used to encourage efficient grasping and quick lifting of the object while penalizing the agent for taking too long.

	Non-terminal state	Terminal state
Object grasped?	$r_g - r_{tdp}$	$r_t - r_{tdp}$
Not grasped?	$-r_{tdp}$	@timeout $-r_{tdp}$

Where r_t = terminal reward, r_g = grasping reward, r_{tp} = time and distance penalty

Training Process

- Training starts from a reset state with objects on the table and the gripper mid-air.
- Reward functions are defined to provide feedback based on the agent's actions.
- Episodes consist of the agent taking actions, receiving rewards, and aiming to maximize cumulative reward.
- The agent updates its policy based on observed rewards and the current state, exploring the environment to learn the optimal policy.

Experiment 1 – Results and Discussion

Model Evaluation Report:

- The decision tree classifier reports a higher precision and recall for most classes, resulting in an overall accuracy of 94% compared to the 88% accuracy of the random forest classifier.
- The similarity in macro average and weighted macro average suggests that the models perform similarly across all the different EMG channels.

	Decision Tree		Random Forest	
Channels	precision	recall	precision	recall
1	0.94	0.95	0.84	0.89
2	0.93	0.94	0.93	0.94
3	0.92	0.93	0.92	0.93
4	0.98	0.98	0.98	0.98
5	0.93	0.93	0.93	0.93
6	0.91	0.91	0.91	0.91
7	0.94	0.94	0.94	0.94
8	0.92	0.92	0.92	0.92
9	0.94	0.94	0.94	0.94
10	0.94	0.93	0.94	0.93
Accuracy		0.94		0.88
Macro avg	0.94	0.94	0.94	0.94
Weighted avg	0.94	0.94	0.94	0.94

Experiment 1 – Results and Discussion

EMG signal channels importance:

- For both models, the EMG signal from channel 2 was identified as the most important feature.
- The importance of the timestamp was relatively high in both models, indicating that constant use of the device can significantly affect the model's performance, even more so than the EMG signal channels themselves.

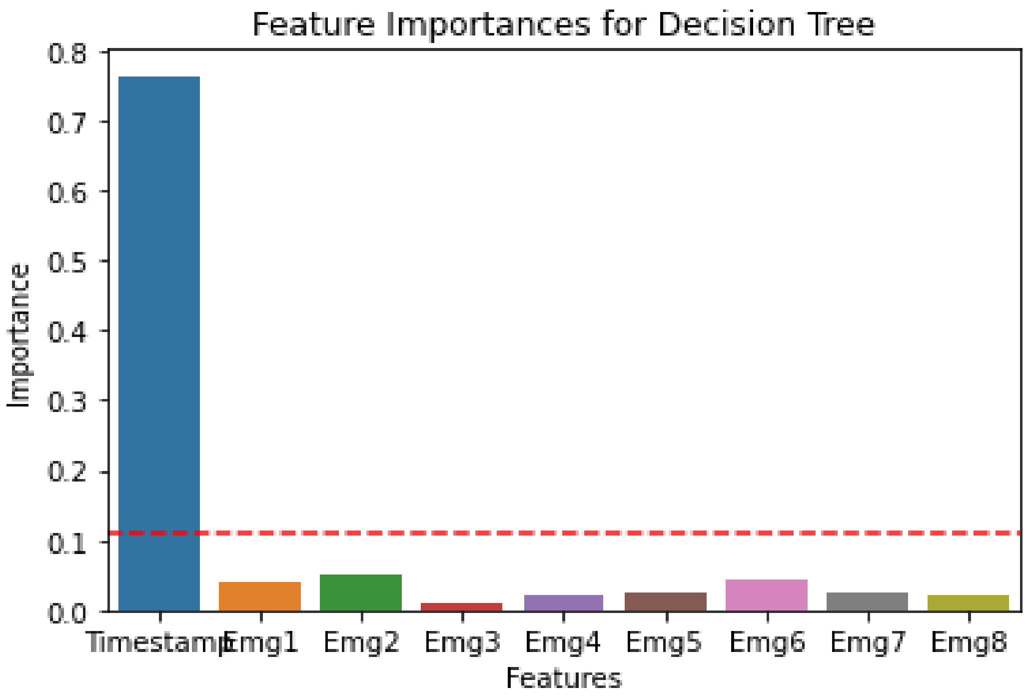

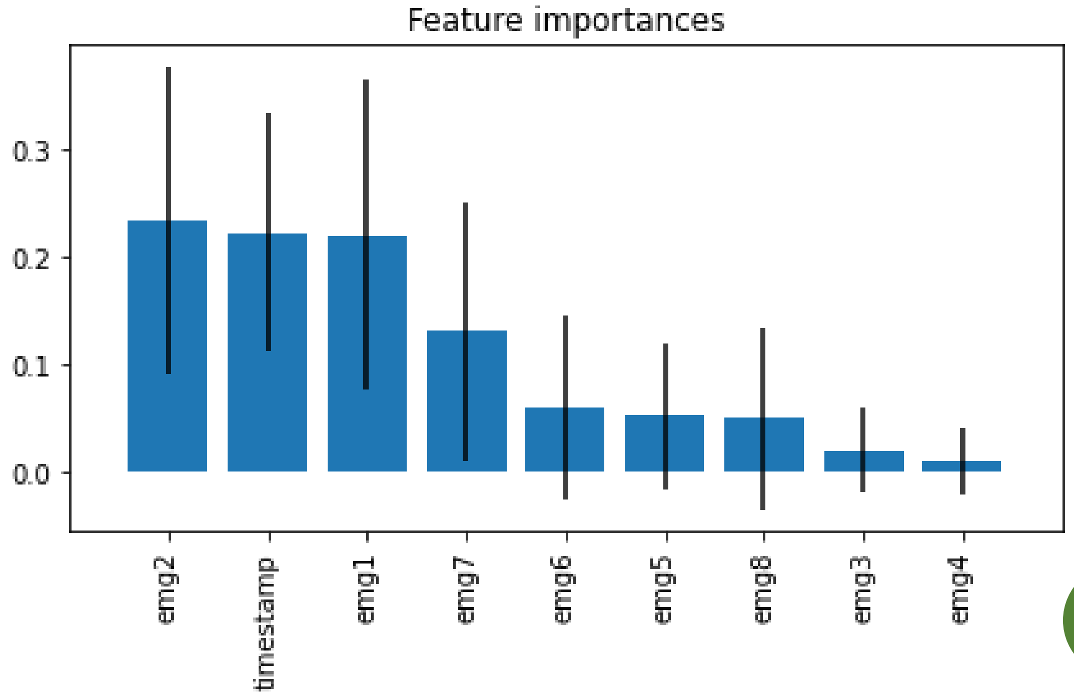

Experiment 2 – Results and Discussion

Algorithm-specific success rate:

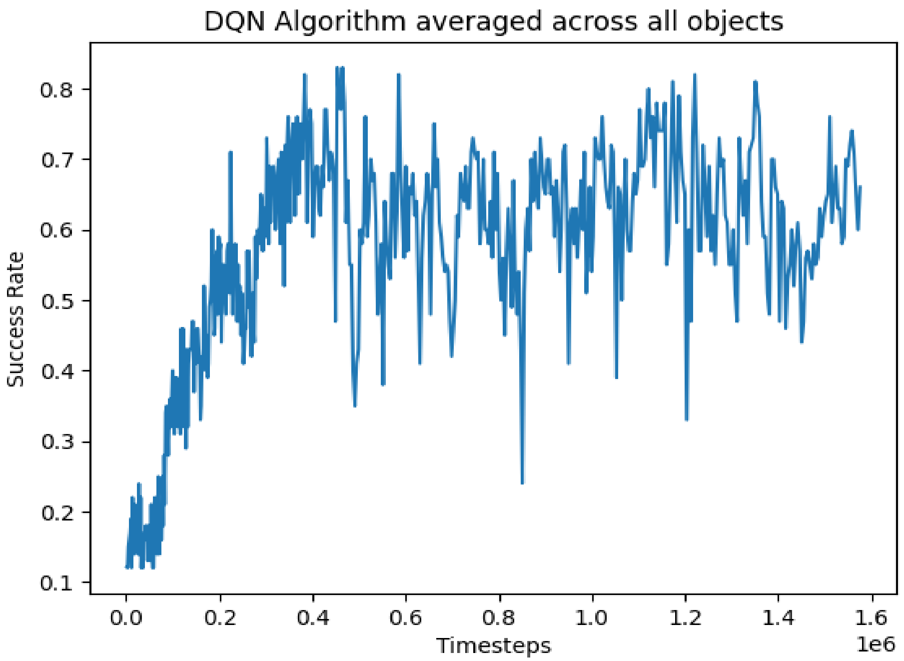

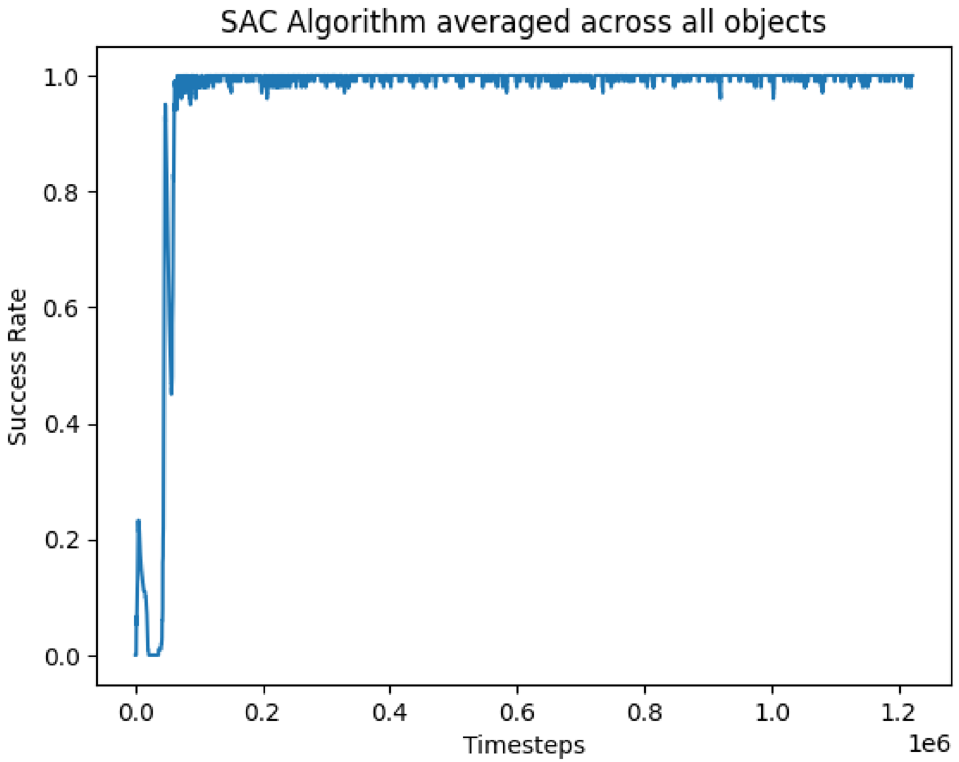

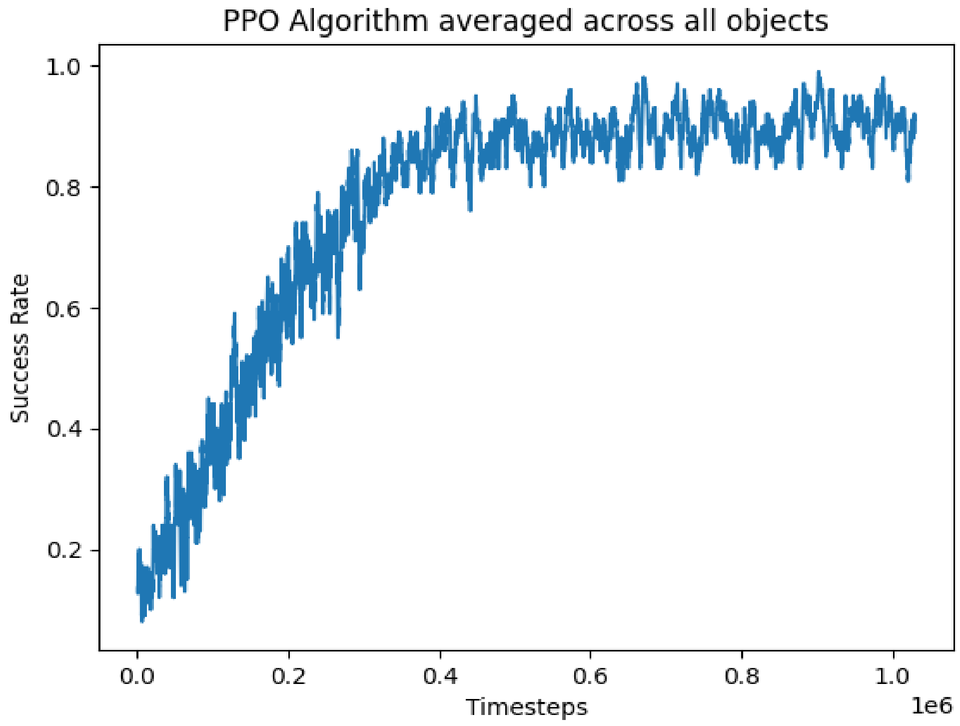

Experiment 2 – Results and Discussion

ANOVA test to examine the significant differences between the three algorithms:

At p-value (2.2e-16) < 0.05, reject the null hypothesis.

There is a significant difference in the performance of the three algorithms.

Summary	Result
DF	2
Pillai's trace statistic	0.57982
Approximate F-statistic	10148.93
P-value	2.2e-16 ***
Significance code: 0 '***'	

Experiment 2 – Results and Discussion

Object-specific success rate:

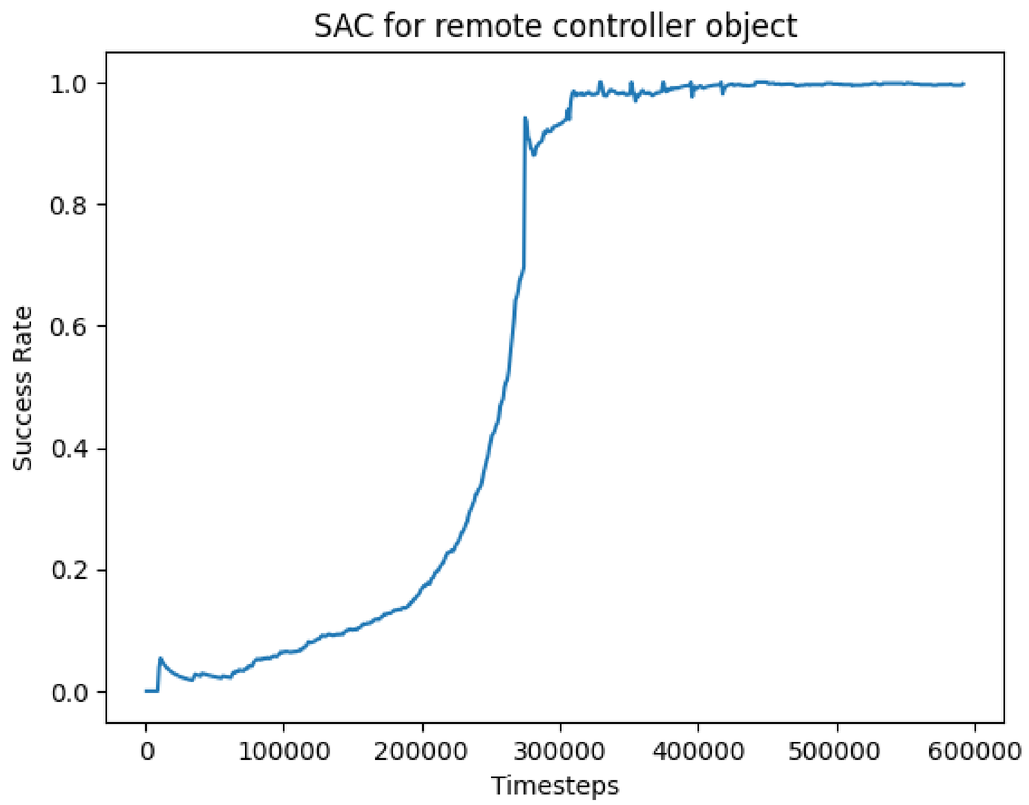

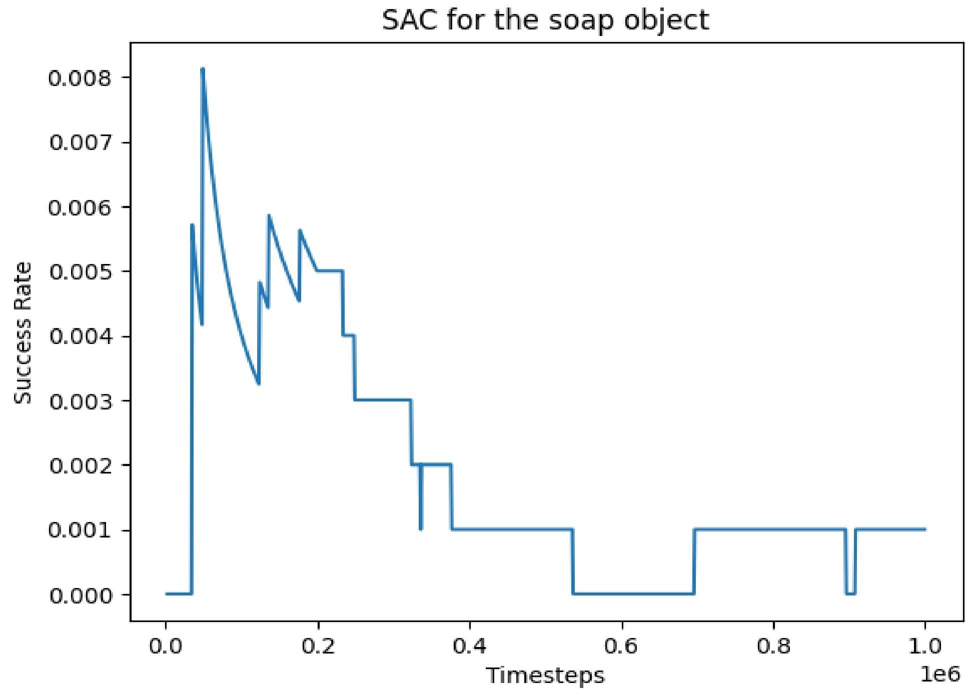

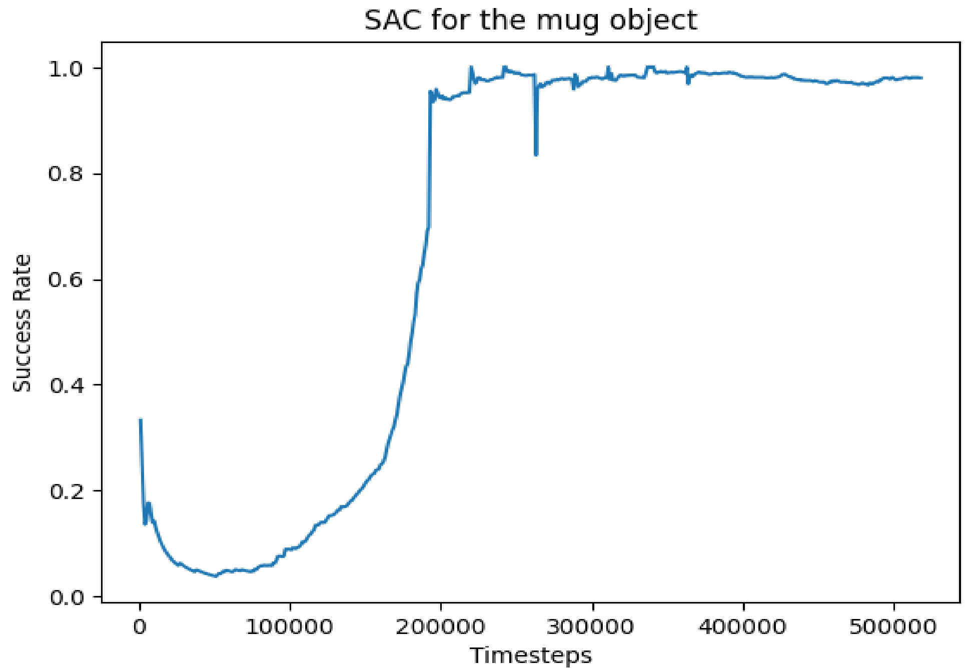

Experiment 2 – Results and Discussion

ANOVA test to examine the significant differences between the three objects using the SAC policy:

There is a significant difference in the performance of the SAC algorithm on the three different objects .

Summary	Result
DF	2
Pillai's trace statistic	0.6888
Approximate F-statistic	552.9
P-value	2.2e-16 ***
Significance code: 0 ‘***’	

Conclusion

- In Myo prosthetics with intelligent gesturing capabilities, frequency of use has an impact on the model’s performance in accurately interpreting the signals.
- In EMG devices, the different signal channels do not have the same importance in the overall signal interpretation. Specifically, for an 8-channel EMG device, the channel 2 is the most significant.
- SAC algorithm excels the most in comparison with DQN and PPO when training a prosthetic hand to automatically grasp objects with appropriate force and at the correct contact point.
- Object physical properties such as the shape and texture is a major determinant of how effectively the prosthetic hand can grip objects without it slipping or getting damaged.

Contribution

- For an 8-Channel EMG device, the identification of the second channel as a key driver in gesture recognition offers practical implications for designing and implementing robust and efficient gesture recognition systems, as it provides guidance for feature selection and prioritization in real-world applications.
- The research reveals that the success of prosthetic grasping is influenced by object physical properties, such as shape and texture. Understanding these factors can inform the development of robust prosthetic grasping systems capable of effectively gripping objects without slippage or damage.
- By scaling the learning rate for each parameter, the performance of the model used in sensor data preprocessing is significantly improved, particularly for sparse data. This enhancement builds upon prior work by Breyer, et al., (2019) and optimizes data handling in EMG-based applications.

References

- M. Breyer, F. Furrer, T. Novkovic, R. Siegwart and J. Nieto, "Comparing Task Simplifications to Learn Closed-Loop Object Picking Using Deep Reinforcement Learning," in IEEE Robotics and Automation Letters, vol. 4, no. 2, pp. 1549-1556, April 2019, doi: 10.1109/LRA.2019.2896467.Britannica, T. E. (2018, June 20). Metacarpal. Retrieved from Encyclopedia Britannica: https://www.britannica.com/science/metacarpal
- Jiang, N., Pradhan, A., & He, J. (2022). Gesture Recognition and Biometrics ElectroMyogram (GRABMyo). PhysioNet. doi:https://doi.org/10.13026/dzhr-nj10.
- Haarnoja, T., Zhou, A., Abbeel, P., & Levine, S. (2018). Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. *arXiv*.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal Policy Optimization Algorithms. *arXiv*.
- Roderick, M., MacGlashan, J., & Tellex, S. (2017). Implementing the Deep Q-Network. *arXiv*.

Thank you