

AllCloud Helped EyeControl Increase Accuracy in an Algorithm for a Wearable Personal Communication Device for Patients Who Can't Communicate Verbally

About EyeControl

EyeControl offers a wearable personal communication device to people whose medical condition prevents them from communicating effectively. The technology tracks eye movement and converts it into communication.

Qualified patients are reimbursed for the device by the Israel Ministry of Health and also by the NHS in the United Kingdom. The technology is currently being adopted by intensive care units and trialed by COVID-19 patients in Israel.

Summary

The device EyeControl offers a head-mounted infrared camera that tracks eye movements, the camera takes a picture of the eye and sends it to a deep learning model that classifies each movement into 1 of 6 possibilities: up, down, right, left, straight (no movement), or blink. Before turning to AllCloud data expert, EyeControl's machine learning (ML) model had achieved a 92% accuracy rate but the EyeControl team knew they could do better.

EyeControl decided to replace their model with a 2-step flow:

1. Use an object detection model to detect 5 key zones in the pictures taken by the camera: among those key zones are the pupil and the eyeball.
2. Based on the coordinates of these zones, use supervised learning to output a prediction for the eye position.

AllCloud's role was to develop a complete ML pipeline, an object detection model, and then a classifier.

Challenge

EyeControl's existing model was already highly performant but developing just another "good" model would not have been enough for them. Moreover, it was important for EyeControl to achieve quick inference time. It is well known that even with the current state-of-the-art object detection, there is a tradeoff between accuracy and inference time, and it is not guaranteed to get great results on both fronts.

In addition, in efforts to improve accuracy, replacing an image classification model with a pipeline of object detection and an additional supervised learning algorithm means adding a level of complexity, uncertainty, and error margin.

Lastly, each training job for object detection took hours even on very powerful GPUs. Blind tuning and training could get extremely long and costly, therefore EyeControl searched for a way to control their timeline and budget.

The Algorithms

The two combined requirements for high accuracy and low inference time steered us towards developing the SSD algorithm. The AWS built-in object detection algorithm is SSD with and is running on two basic networks: vgg16 or resnet50. Although it is well known that the resnet architecture provides much greater speed from the two, we decided to test both in case the vgg's precision would be significantly higher. The test set was tested on a t2.medium machine to emulate the device's allocated computational power.

The results from the tests only solved part of the pipeline work. AllCloud and EyeControl would need to then understand which machine learning classifier would be appropriate for the remaining part. We decided to explore a few possibilities, such as:

- A logistic regression
- A neural network
- An SVM
- A random forest

The confidential classifier ended up providing far better accuracy and very low inference time, it was clear which classifier should be selected.

Project Management

At AllCloud, one of our key project management principles is transparency: keeping the customer highly involved throughout the entire process.

In this case, the challenge was to provide insights throughout our very technical research process without the need to make the customer a machine learning expert. The requirement led AllCloud to define simplified KPIs based on ML measurements and standard reporting. For example, 'precision' was defined as the percentage of success of the machine predictions. This set of simplified KPIs along with classic ones like 'accuracy' and standard graphs were used to share the progress of each algorithm and iteration, providing EyeControl with a good indication of which direction they should continue testing.

Together with EyeControl, we conducted weekly milestones and demonstrated continuous improvement in overall accuracy. This provided the customer with the confidence and required visibility they needed.

"Working with AllCloud's team was a great experience, the daily communication and quick response to every question and suggestion allowed EyeControl to achieve all its goals on time with continuous improvement to our data pipeline and trained models."

Osama Gerai, Team leader at EyeControl

Research

An additional challenge was the need to conduct a research project that was experimental in nature within a strict budget. We took advantage of the fact that time constraints were not an issue and planned the project with variable resource allocation of 200% to 10% according to the type of the task, percentage of machine time expected to be used during this task, and the required deliverable to hit the next milestone on the project roadmap. Using this method, AllCloud was able to deliver significant value, utilizing 98.2% of the allotted budget.

An important milestone after completing object detection of the first model was whether we should continue to develop a second model. We asked ourselves questions like: Do we have a good chance at overachieving EyeControl's legacy model? Object detection results are not comparable to image classification model results, of which the customer already had. We did manage to overachieve their legacy model by comparing two ground truth datasets, which demonstrated that object detection results are superior to human-generated results in terms of locating the center of an object. This was a key insight, as it is commonly assumed that the object's location is the most important parameter in generating overall accurate end-results. Following this demonstration, EyeControl and AllCloud decided to continue and develop a second model which later proved to be the right decision.

Solution

In order to reduce the cost of training the object detection model and allow for many iterations, AllCloud significantly reduced the cost of training by using Spot Instances. This only marginally increased the time of training and allowed for easy iteration. Quick iterations then allowed us to tune our model effectively and reach very high evaluation metrics.

The precision of the object detection model was confirmed as soon as we started applying the selected classifier. Without fine-tuning, we had already reached reasonable accuracy scores, attesting to our model's strength.

This reasonable accuracy was still below 92% so we, therefore, started tuning the hyper-parameters of our classification model. Through fine-tuning, we eventually found the best combination that gave us a very high accuracy score.

The Results

AllCloud's contributions to EyeControl's device accuracy proved to be successful. The new pipeline implemented achieved a 98.58% accuracy rate on the test set with an average of 266ms inference time. The accuracy is therefore much higher than Eyecontrol's legacy model of 92%. The model weighs 100 MB and has an inference time corresponding to 4 frames per second; both results are considered good and suitable for the customer at the first stage of the project.