Wearable device recognizing activities done by the arm using machine learning

Endre Yllö

Building a wearable device that senses the arms movements and identifies what action it is performing, in order to aid rehabilitation and human activity recognition. By combining inertial measurement units and radar with a convolutional neural network.

Human activity recognition is an active field of study that concerns itself with identifying what a person is doing and how they are doing it. Utilizing various sensors combined with artificial intelligence enables a wearable device to peer into the wearer's everyday life in an effort to gain insight as to what actions are being performed. This capability can prove useful when implemented for the purpose of rehabilitation. The device consists of inertial measurement units capable of sensing acceleration and rotation along three principal axes, one of them placed by the elbow and one on the wrist. It also features a radar module placed on the wrist directed towards the back of the hand. By analyzing the data produced by the sensors it becomes possible to determine the position of the arm and how it changes. The radar is intended to both approximate the flexion of the wrist and collect some useful information about what is in front of the hand. Combining the information gathered by the sensors and feeding it to a convolutional neural network tasked with classifying it according to a set of tasks, allows the device to guess as to what action is being performed by the user, as well as providing information on how it is being performed.

Part of the challenge is to collect a sufficient amount of high quality training data for the neural network, this was done by letting volunteers wear the device while executing tasks according to a protocol consisting of 19 tasks. Even using limited training data collected on 10 people, it was possible to achieve accuracies in excess of 68%. Further examination of the performance revealed that the neural network had a low reliance on the radar when classifying tasks. Although some of the 19 tasks should be indistinguishable based on radar data alone, this hints that the radar might not have been optimally configured for this purpose and further development of the device should make an effort to remedy this. The radar module could also be omitted from the device completely in order to avoid certain trade-offs induced by its inclusion.

In conclusion, the device performed well considering the limitations of the training data. Further improvements to the system were suggested based on the discoveries