Machine learning for BIOMEDICAL ENGINEERING

There's a very active machine learning community at Stellenbosch University, to the extent that a Maties Machine Learning event is held every second Friday lunchtime, with short talks or open discussions to bring people together and strengthen machine learning research at the university.

Recently, it was the turn of Professor Dawie van den Heever, head of the Biomedical Engineering Research Group (BERG) in the Department of Mechanical and Mechatronic Engineering, to share information on projects involving machine learning. Here, *Quest* gives a brief summary of a few of these.

Prof. Dawie van den Heever with Sophia, the humanoid robot.

Attention-Deficit/Hyperactivity Disorder

An estimated 5% of learners in South Africa have Attention-Deficit/Hyperactivity Disorder (ADHD), which occurs as two sub-types – inattentive and hyperactive – and is associated with a lack of concentration, no sense of time, poor memory, low self-esteem and poor social skills. Typically, teachers or parents notice these problems in a learner, who is then referred to a specialist, such as a child psychiatrist or educational psychologist, for a proper diagnosis. This involves interviews with the family and the completion by the parents and teachers of questionnaires called rating scales, as well as a series of psychometric tests for the learner.

Many learners, especially in rural settings, do not have access to advanced healthcare services, so they may never get the diagnosis and hence the treatment and support they need. To address this, researchers at BERG developed a tablet-based game that can be used as a rapid screening tool for ADHD. Final diagnosis would still need to be done by a specialist, but the tool could be used by teachers or nurses before referral, as a diagnostic aid by specialists, or even to monitor response to treatment.

The PANDAS game developed by BERG researchers as a screening tool for ADHD.

Drawing on the acronym of its name 'Paediatric Attention-Deficit/Hyperactivity Disorder Application Software' – PANDAS – the game uses a panda bear as its main character, which has to perform tasks and deal with various obstacles at increasing speed. Two different games to cater for each of the ADHD sub-types were developed. Data is recorded while the learner plays the game, and machine learning is then used to identify patterns in the data that discriminate between players with and without the condition.

Knee replacements

The knee is the largest joint in the human body. It is prone to injury, especially through rigorous sport, and to disease, particularly osteoarthritis in the elderly, when cartilage wears away and the underlying bones rub against one another, causing pain and a loss of function. Patients requiring a 'knee replacement' usually have to rely on an artificial implant that comes in standard sizes from a variety of manufacturers. Where an implant does not match the geometry of the patient's healthy knee, the surgeon tries to make the implant 'fit' by shaving away more bone. Often, the mismatch means that normal motion is not fully restored, or that the patient experiences more pain and discomfort during the rehabilitation process.

An X-ray showing a knee implant.

The BERG researchers obtained a large database of CT scans of normal knees from the United States, and used it to model knee geometry – examining, for example, the relationship between the width of the femur and tibia, and the size and curvature of the condyles. They could then use machine learning to estimate the knee geometry of a patient's healthy knee from CT scans of the damaged knee, well before surgery. This would help in ordering the 'best fit' standard implant, but could also be used to custom-make patient-specific knee parts by metal 3D printing, also known as additive manufacturing.

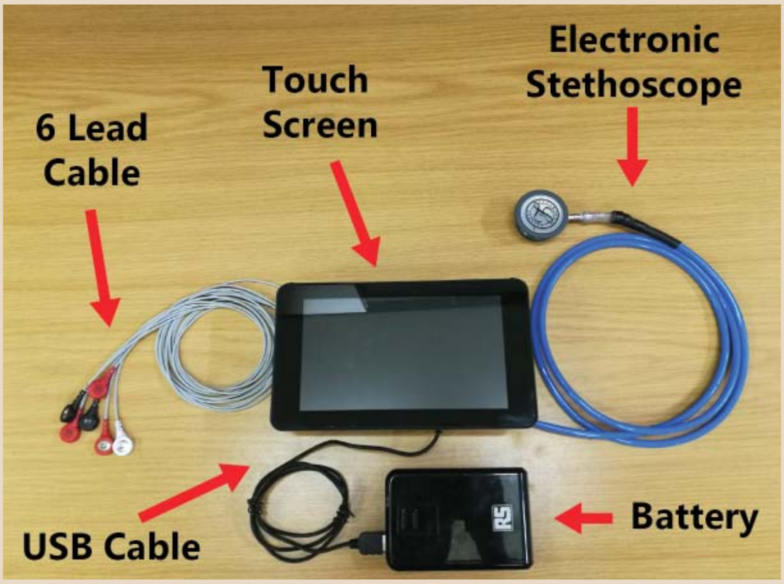

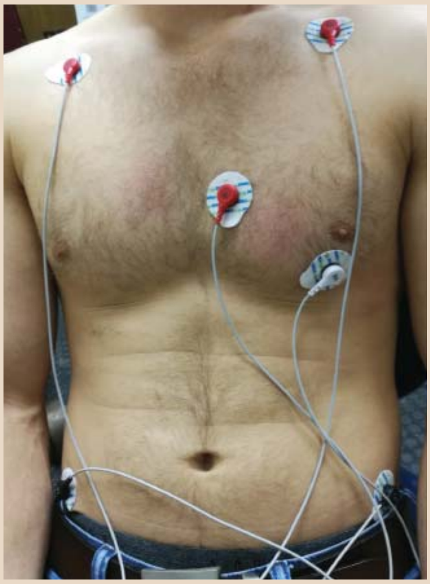

The portable ECG device developed by BERG researchers could be used in rural areas to detect heart problems.

Cardiovascular disease

Worldwide, cardiovascular disease is a leading cause of death, but diagnosing it requires expensive medical equipment, operated by trained technicians. For this reason, an electrocardiogram (ECG) can only be done at larger hospitals in major cities in South Africa. A standard 12-lead ECG makes use of 10 electrodes attached in specific places on the chest and limbs to record the heart's electrical activity, represented as a graph of voltage versus time.

The BERG researchers developed a portable ECG device that could be used in rural areas to screen for cardiovascular disease, allowing patients with potential heart problems to be referred to hospital for further examination if necessary. The device has a reduced number of leads and electrodes, as well as an electronic stethoscope that records heart sounds, with all data saved to a tablet for instant display and replay.

A student wears an EEG cap to measure patterns of electrical activity in the brain.

Machine learning was used in two ways. Firstly, it allowed 'lead reconstruction' – in other words, the device is able to reproduce the results of a standard 12-lead ECG that doctors are accustomed to using when making a diagnosis. Secondly, by training the neural network on 'normal' ECGs and ECGs representing different heart abnormalities, the device was able to diagnose abnormal heart activity with 90% accuracy when tested on 70 subjects in a clinical study at Tygerberg Hospital. It is expected that this would improve with an increased number of training sets.

Brain waves

An electroencephalogram (EEG) is a test usually used by doctors to diagnose problems with the electrical activity of the brain. Electrodes placed on the scalp allow patterns in electrical activity, or brain waves, to be recorded. The main use of EEGs is to investigate epileptic seizures, but also head injuries, sleep disorders, coma, dementia, and brain inflammation or tumours.

Researchers at BERG have used EEGs and machine learning to do 'mind-reading'! Test subjects were shown pictures, words or sentences on a computer screen while brain activity was recorded, and neural networks were then trained to identify what the subject was thinking about. Good results were obtained for training on the EEG data of a specific person, but the neural network did not generalise well to other people, especially to people it did not train upon at all.

"But this is still cool," concludes Dawie. "It tells us that people's brains work differently, and that it is possible to train AI to know what object or concept someone is thinking about, based just on brain activity".

More information can be found on the websites of BERG (https://berg.sun.ac.za/) or Maties Machine Learning (https://mml.sun.ac.za/).