In biomechanics we use kinematic analysis to examine changes in joint angles during a wide range of movements. Kinematic analysis is often performed with a camera system combined with a set of reflective markers placed over bony landmarks. For example, at JYU we often use a set of Vicon cameras for this purpose. However, sometimes it is not feasible to use reflective markers in practice, such as when collecting data outside of a laboratory setting. Another problem with traditional motion analysis systems is their high cost and steep learning curve.

We recently performed a study where we aimed to overcome some of these difficulties. We used a very cheap camera (GoPro) to record people running underwater. This is a good example of a setting where reflective markers are not feasible, so an alternative approach is needed. Also, some cameras are not well suited to being used underwater, so this is a challenging environment to perform motion analysis.

In our study, we took an existing algorithm called DeepLabCut (Insafutdinov et al., 2016; Mathis et al., 2018; Pishchulin et al., 2015) that is designed for object tracking, and we modified it to track the locations of lower limb joint centres. Using this information we were able to go further and calculate joint angles, as well as various other parameters that are useful in typical gait analysis applications.

The DeepLabCut method is an open-source algorithm that uses deep neural networks to predict the locations of individual points in an image. First, we trained the neural network to identify the points we were interested in, such as the hip, knee and ankle joint centres. To do this, we showed it several labeled examples of these points, allowing it to gradually 'learn' how to identify each one. Neural networks are normally very data hungry. For example, in order to reliably train a network to identify a cat in an image, it usually needs to see many thousands of photos of cats before it can accurately identify a cat in a previously unseen image. With our approach we instead take advantage of something called transfer learning, where we take a model that has already been trained to identify objects based on many thousands of images (or even millions). We then modify part of the model to re-train it to perform our particular task. Using this approach, we typically need far fewer image examples in order to train the neural network to accurately identify our target object. Below are 2 examples of images that have been labelled by a human and a neural network:

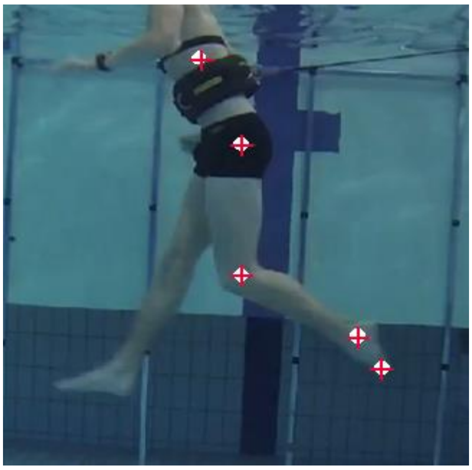

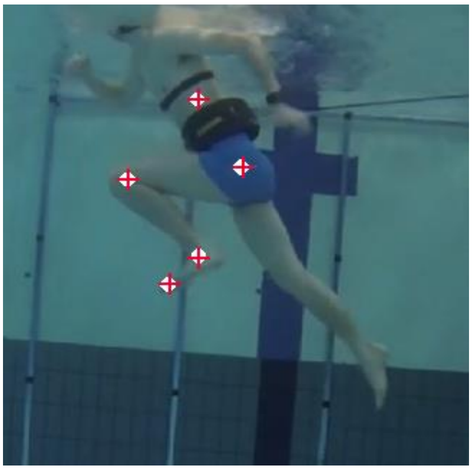

In these examples from 2 different people, labels placed by a human are shown by red crosses, and those placed by the neural network are shown by white circles.

One of the major advantages of our approach is that once a neural network has been trained for a particular task, the trained model can be used to analyse new image/video examples very quickly (in a few seconds, or potentially even in real-time). In a paper we published recently, you can see a few examples of this process.

We found that 300-400 labelled images were enough to train the network to be able to position joint markers with an accuracy level similar to that of a human labeler, with a mean difference of around 1cm. It is likely that we can further improve this accuracy, because there are many ways our approach could be modified, which we are just beginning to explore. The level of accuracy that we achieved is sufficient for many 2D applications, such as sports biomechanics and coaching, and rehabilitation/training scenarios, and we are now exploring various scenarios where our method could be very useful, such as clinical cases and sports performance.

As well as joint kinematics, we also used our method to compute some other relevant parameters such as joint range of motion and cadence on a stride by stride basis, and we found good test-retest reliability of kinematics measured with this method. Some of the results from our recent paper are shown in the figure below.

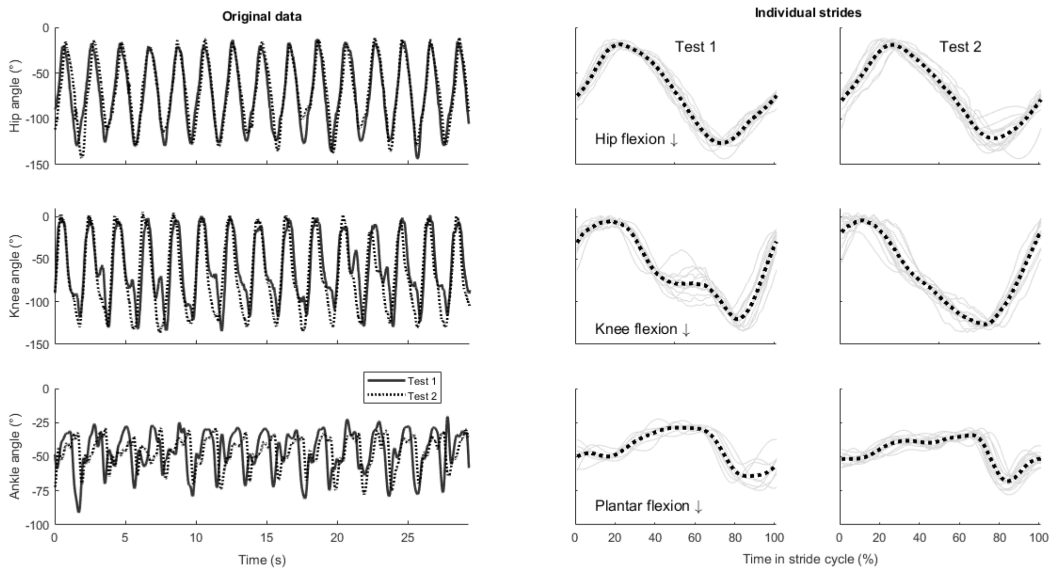

In this figure we present joint kinematics measured from one subject measured on two different days. The raw traces are shown on the left. On the right, you see individual strides that were segmented from the raw data, and it is clear that this person ran in a kinematically similar way between test sessions.

Our approach is very low-cost, since it works well with cheap cameras and there are no other major expenses. After training a neural network to work with a given type of data, it would be possible to simply go and record the videos, and then run the analysis. I am currently streamlining the whole analysis process so that this can be done as efficiently as possible. There are many exciting projects coming up where we will use this and other artificial intelligence based methods, and eventually our goal is to provide near real-time feedback, making this a valuable tool for coaches, athletes, clinicians, and whoever else has a need for biomechanical analysis.

Our method only requires a small amount of manual labelling of image frames, and a well-trained network can potentially be re-used for many projects. Given the challenges associated with imaging deepwater running, I firmly believe that this approach could easily be modified to analyse kinematics in other human movements and measurement settings, simply by re-training the network using a suitable dataset. In this study we only used a single camera to record the running, but by using additional cameras, this approach could be used to perform 3D analyses, as already demonstrated in various animals (Nath et al., 2018).

We recently published an article about this work in the Journal of Biomechanics. For more information see this link, and be sure to watch the supplementary videos, which give a much better idea of the value of our method.