WE WANT YOU!!

INTERNSHIPS & THESIS PROJECTS

Clockworks is a Rotterdam based startup specialized in computer vision & A.I., developing a deep-learning solution for real-time, high-performance, digitization of analog meters and sensors under any condition.

- Are you eager to apply advanced A.I. techniques to real-world problems?

- Are you interested in deep-learning for object-detection, segmentation and classification?

- Are you interested in optimizing neural nets for mobile devices?

- Would you like to work with pixels (trillions of them) and computer vision?

- Then Clockworks is looking for you!

Interested in joining our team?

please send your CV and motivation to careers@clockworks.co

clockworks

Research projects & Internships

General "scene" text detection
- Scene text detection/recognition is the art of reading any text anywhere. Whilst we can currently read some of the most vital text on the meter already, we are always looking to expand this capability to more data and more text.

Viability of capsule networks in real world settings
- Capsule networks have recently been proposed and produce promising results on toy problems. However, few practical applications are known. Are capsule networks suitable for object detection, and how could we interpret their learned features? Does routing-by-agreement ensure a logical arrangement of multiple components?

Reading barcodes
- While QR codes are designed to be easily read from images, barcodes are surprisingly difficult to read; as a little bit of noise in the image is introduced, they quickly become completely illegible. We would like to find a way to read them in a more robust way, which will require some creative solutions, for instance using recurrent CNN's.

Extending training data
- Every AI student knows that current deep learning methods thrive on lots of data. How can we augment our currently annotated data sets such that we can improve our deep learning models? Possibilities include generating new data using a Generative Adversarial Network (GAN) or by implementing a 3D simulator, or augmenting existing data by simulating flashes and/or dirt.

Real-time inference to capture the perfect frame from a video stream
- We want to make our product as fool proof as possible. To achieve this it would be great to assist in capturing that perfect shot of the data. To accomplish this, some part of our image-to-data pipeline has to run locally (on either a smartphone, or in the browser (using tensorflow.js?)) or at the very least fast enough to work on a video stream to the server. Will you be our speed-monster?

Transferring object detection to image segmentation
- Segmentation data contains less noise within the label than bounding boxes do. Therefore segmentation could allow for more accurate detections and easier extraction of rotation of the segmented object. Can you repurpose our bounding box-annotated dataset for segmentation?

Evaluation of object detection methods
- Recent research has produced several object detection methods based on deep learning. Some emphasize speed, while others specialize in precision. Can you objectively compare, combine and apply these methods and determine which networks are most suitable for our specific use-case?

Irregular object detection
- Object detection networks usually assume rectangularly shaped objects. Although this allows for efficient and conceptually easier implementations, objects are rarely exactly rectangular and without rotation. Can we alter the current object detection network to be able to handle rotated bounding boxes or even irregularly-shaped polygons?

More information? contact us at careers@clockworks.co