

# BrainCreators

## Research Internship 2021-2022: Understanding Activities in Anonymized Video Surveillance

---

### [General Information](#)

[Contact and Interviews](#)

[Introduction and overview](#)

[Research internships in our applied vertical teams](#)

[Research internships on other activities](#)

[What we offer, what we expect](#)

[The Research Internship Position: Understanding Activities in Anonymized Video Surveillance](#)

[Summary](#)

[Product](#)

[Technology](#)

[Research Questions](#)

[Engineering & MLOPs](#)

[Sources](#)

---

## General Information

### Contact and Interviews

**Before you read on, .. we encourage interested candidates to contact us as soon as possible for an intake interview!**

Please contact our head of research, Maarten Stol: [maarten.stol@braincreators.com](mailto:maarten.stol@braincreators.com)

Or visit our website: <https://www.braincreators.com/contact>

Like previous years, we have a limited number of available positions, and expect another rise in the number of candidates. Interviews will take place in September & October, and decisions to hire will be made for a timely start in November 2021.

To some degree, and given equal skills, this will be a first-come-first-serve selection process. (there is a chance of new positions opening up later in the year, so if you read this after November 2021, the door is not fully closed yet)

## Introduction and overview

Welcome! You are looking at the research internships BrainCreators has on offer in 2021-2022.

BrainCreators is at the forefront of applied AI, with many years of successful research internship projects that combine cutting edge science with the challenges of applying AI in the real world. Located at Amsterdam's Prinsengracht and Science Park, we are a growing team of AI experts, software developers, MLOps & DevOps specialists and researchers.

## Research internships in our applied vertical teams

The 4 business verticals that offer a research internship position this year are:

- **Road surface inspection**, combining Deep Learning Object Detection with Geo-information (and possibly 3D data).
- **Video surveillance**, based on, and extending our anonymization tooling. The focus is on understanding person and crowd behavior, anomaly detection, and video retrieval, all based on video representation Deep Learning and self-supervision.
- **Conveyor belt applications**: recognition, localization, and manipulation by robot of objects on a conveyor belt. Challenges concern high variance of object shape and visuals, and detection of out-of-distribution imagery.
- **Fashion & Retail**: this year with a focus on generative models for Virtual Try-on of clothing items.

## Research internships on other activities

In addition to our business verticals, there are research topics that are more general, or concern pure research which is not immediately related to our commercial activities.

If you would like more information on topics like these, please contact our head of research, Maarten Stol: [maarten.stol@braincreators.com](mailto:maarten.stol@braincreators.com)

- **MLOps** is an essential part of every product we roll out live. Topics include data unit tests, live evaluations, deployment monitoring, handling shifting data, containerization, building KubeFlow pipelines, and scaling deployments.
- **Symbolic/Subsymbolic Hybrid Ai** In particular we are interested in compensating a lack of annotated training data with symbolically encoded background knowledge about the application domain. If valuable explicit background knowledge is available in the form of rule-based information, then we are interested in e.g., imposing this knowledge as regularizers on our object detection models, or in other ways to exploit relational information.
- **Astronomy** A position working in tandem with our partners on the Cortex Consortium in the field of astronomy. BrainCreators is an industrial partner in this 6 year project, providing research and development with a focus on topics like neural network compression and autotuning of real-time ML pipelines. For a general impression see:
  - <https://www.uva.nl/en/shared-content/faculteiten/en/faculteit-der-natuurwetenschappen-n-wiskunde-en-informatica/news/2019/06/self-learning-machines-hunt-for-explosions-in-the-universe.html?cb>
  - <https://www.esciencecenter.nl/projects/cortex/>

## What we offer, what we expect

### We offer:

- Be part of a growing company with a proven track record in applied Ai
- A research internship position on one of our vertical teams
- Interaction with research interns from our other vertical teams, in a science oriented horizontal research team.
- A protected environment for your research, without distraction by commercial deadlines of the team
- Opportunities to contribute to the team by developing dual-use software: for your own research and the team's products.
- Weekly supervision on scientific progress, experimental design, and thesis text
- Weekly supervision on software development and code reviews
- Daily contact with the vertical team, and morning stand-up meetings
- Weekly participation in internal ML workshops, sharing ideas with others
- Access to compute resources (in addition to University resources)
- Opportunity to work from home, or work from our HQs at Prinsengracht or Science Park Amsterdam.
- A financial compensation of 300 euros per month
- Learn all the essential things a Master program typically does not offer, e.g.,
  - onboarding with software development skills,
  - MLOps skills,
  - optimal use of compute resources,
  - versioning of ML and datasets,
  - collaboration software,

- and communication skills.
- Be the eyes and ears of your team, looking for promising academic developments that might be relevant to the vertical
- Opportunities to become a permanent team member, and join as ML engineer after the research internship.

**We expect:**

- Workload contribution of 40h per week, 6-8 months (all activities related to your MSc program are included in this 40h, other jobs and classes are not)
- Capable to work independently on your own research questions and experiments
- Active participation in team effort when needed
- Solid control of spoken and written English language
  
- A strong opinion on ML research and how to apply it in practice
- Solid fundamental knowledge of ML theory and practice
- Overall knowledge level of a graduating Ai MSc student
- Good PyTorch skills
- Good understanding of the required mathematics
- Good software development skills
  
- Active participation in internal workshops, presenting your progress, and discussing your experimental design choices with your team and other verticals in the company
- Willingness to rewrite the thesis as a publishable paper
- Co-authorship for your thesis supervisors on publications derived from the thesis.

Please contact our head of research, Maarten Stol: [maarten.stol@braincreators.com](mailto:maarten.stol@braincreators.com)

Or visit our website: <https://www.braincreators.com/contact>

---

## The Research Internship Position: Understanding Activities in Anonymized Video Surveillance

### Summary

Join our Video Deep Learning team as a research intern. Work on challenging video technologies for understanding surveillance footage, all based on our anonymization product ARA. Applications range from understanding Human-Object interactions, recognition of actions involving more than one object/person, understanding crowd behavior, or making sense of 3D information. The end product will be used for our products in the Ai video field. Describe your scientific results in a publishable paper.

## Product

<https://www.braincreators.com/spatial-inspectors/ara>

Used in airports, retail, and other large spaces, our video surveillance anonymizer, ARA, is a keystone technology. It enables GDPR sensitive footage to be anonymized with great recall and precision. This allows third parties to process the data on the cloud for purposes other than security alone.

The base functionality of ARA has been online for about one year. We are currently working to extend the product with extra functionality. Some are pure downstream processing, like e.g., people counting in user defined areas and integrating output into geo-information systems, or anomaly detection based purely on ARA like movement vectors and other output statistics. Other new functionalities involve the deep learning stack behind all this, e.g., better action recognition based on RGB & optical flow data, or video representation learning for retrieval purposes.

BrainCreators wants to spend 1 year on deep learning based video R&D, and work with a research intern on one of the new applications. We aim to deliver a working solution at the end of the internship, and write a publishable paper.

## Technology

For an example of what our video team did last year, please have a look at [1]. This was part of our effort for better Self-Supervised action recognition and video retrieval.

Two examples of topics we would like to work on this year are

1. Human-Object interactions (HOIs)
2. Understanding depth information

HOIs are a challenging topic for standard convolutional neural architectures. The notion of locality that is hardwired in convolutional neural networks is often not the right one for understanding HOIs. One way forward is the QPIC method in [3], with its transformer architecture. Here, attention is used for aggregating image wide context that is relevant for HOIs and the query mechanism captures individual HOI pairs to avoid mixing up multiple interactions in the same image.

Reproducing and improving upon the results of QPIC seems a promising approach for our extended ARA functionality. For other resources besides [3], see the HOI detection overview on Papers with Code [2].

Understanding depth information with standard video surveillance hardware is another important challenge, for instance, when we want to plot a person's position on a map given a single video stream, or understand how two moving objects are related in real space.

The best approach often depends on the (type of) available data and constraints of the deployment environment. Self-supervision seems to play an important role in modern approaches, as do transformer based models.

In [9,10] vision transformers are used for dense prediction tasks in a way that may be relevant for us. Their dense vision transformer provides finer-grained predictions compared to fully convolutional networks, useful for e.g. monocular depth estimation. Also see github pages [4,5,6,7,8] as an example of methods we may consider using.

The research intern will have to help select and/or implement any methods to be used during the project. Starting points may include any of the provided sources below, or other software and literature sources found by the research intern.

## Research Questions

At BrainCreators, research interns have considerable freedom to define their own research questions. We do, however, provide scope and direction, and maintain the possibility to veto ideas that are too far removed from our commercial interests. That said, part of the internship should have a strong scientific orientation, and aim to result in a publishable paper. Another part of the internship is the development of software modules to be integrated into our product stack.

For this project we want to create extra functionality for our video deep learning applications. Examples of research questions include, but are not limited to:

- Which neural architectures are best suited for the Human-Object Interactions the ARA team is interested in?
- How can we best leverage available trained models, and fine tune on customer data to achieve new levels of performance?
- After establishing a baseline of 3D understanding using available camera footage, how can we improve this by using methods like [4,5,6,7,8] ?
- Are dense vision transformers a feasible solution in our industrial settings? If not, how can we compromise best? Will real-world compromises lead to innovative ideas concerning neural architectures?
- What are innovative ways to collect and use real world data from our video surveillance locations, for use in our own R&D efforts?
- The research intern candidate is encouraged to bring their own ideas to the table (even during the selection rounds!) and help shape the goals for the project.

## Engineering & MLOPs

The research intern will be partly responsible for integration of developed technologies into our product stacks, to facilitate deployment and scaling of the solutions with MLOPs.

While this requires a substantial amount of skills that are often different from typical Ai research, we hope to provide the research intern the opportunity to learn as much as possible, and implement the solution together with our team.

## Sources

- [1] Self-supervised Video Representation Learning with Cross-Stream Prototypical Contrasting <https://arxiv.org/abs/2106.10137>
  - [2] Papers with Code: Human-Object Interaction Detection on V-COCO <https://paperswithcode.com/sota/human-object-interaction-detection-on-v-coco>
  - [3] QPIC: Query-Based Pairwise Human-Object Interaction Detection with Image-Wide Contextual Information <https://arxiv.org/pdf/2103.05399v1.pdf>
  - [4] Github: Visual 3D Detection Package <https://github.com/Owen-Liuyuxuan/visualDet3D>
  - [5] Github: Stereo3D <https://github.com/Owen-Liuyuxuan/visualDet3D/blob/master/docs/stereo3d.md>
  - [6] Github: MonoDepth Prediction <https://github.com/Owen-Liuyuxuan/visualDet3D/blob/master/docs/monoDepth.md>
  - [7] Github: Mono3D <https://github.com/Owen-Liuyuxuan/visualDet3D/blob/master/docs/mono3d.md>
  - [8] Monocular 3D Object Detection <https://paperswithcode.com/task/monocular-3d-object-detection>
  - [9] Github: Vision Transformers for Dense Prediction <https://github.com/intel-isl/DPT>
  - [10] Vision Transformers for Dense Prediction <https://arxiv.org/abs/2103.13413>
-