BrainCreators
Research Internship 2021-2022
Advanced Object Detections for Robotic Conveyor Belt Applications

---

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
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:

- **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.

- **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.

- **Fashion & Retail**: this year with a focus on generative models for Virtual Try-on of clothing items.

- **Road surface inspection**, combining Deep Learning Object Detection with Geo-information (and possibly 3D data).

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

- **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-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
Or visit our website: https://www.braincreators.com/contact

The Research Internship Position: Advanced Object Detections for Robotic Conveyor Belt Applications

Summary

Join our Robotics Vision team as a research intern. Work on R&D for visual Ai methods for advanced object detections on conveyor belt setups. The end product will be used for sorting and manipulation of various object types presented on conveyor belts. Describe your scientific results in a publishable paper.

Product

Classified, so no link yet, sorry.

This product is currently under NDA with our partners. All we can disclose at this point is that it has strong roots in considerations related to circular economy and environmental sustainability. Once you are on board as a research intern, all details can be disclosed. For now, we keep the description of the research general. This does not influence the relevance of the scientific motivations behind the project, as described below.

BrainCreators wants to spend 1 year of R&D effort on the challenges involved with object recognition and localization for sub-second response robotics, and aim to deliver working modules of the pipeline at the end of the internship. A publishable paper should also be the aim of the research intern and their internal/external supervisors.

Technology

For reasons of confidentiality, we will phrase the research topics in terms of general but challenging ML questions.

Annotation Efficient Pixel Segmentation

Visual segmentation at the level of pixels (combined with depth/laser scans) can greatly improve the 3D accuracy for our robot manipulators. Exactly where is the center of gravity of an object? This is harder to estimate from a mere bounding box than from a pixel mask. Next, imagine a horse-shoe shape? Solving the question of where to grasp the object involves seeing which pixels are part of the object and which are not.

However, pixel segmentation typically also requires pixel level annotated image masks, which are expensive to make. Novel methods allow a short cut, under the right conditions. See [1] for a promising approach that might form the starting point for our research. There is code available here [2]. Our first experiments with this method were very promising. We would like a research intern to develop and optimize this approach for our client's setup.

The Open Set Problem.

In our field of application, accurate (and fast) recognition of known object classes is only part of the problem. An equally important problem is the detection of out-of-distribution events, i.e., quickly deciding if the robot is looking at an individual object of an unknown class. Questions relating to this issues are often grouped under the name Open Set Problem.

In a deployment environment where the model might at any time be confronted with new object types, this can become increasingly important. Existing approaches to the Open Set problem focus primarily

on full image classification. Our challenge would be to develop this further into the field of object detection.

This is not an easy challenge. Depending on criticality of the application, and on expected vs actual visual variation of the objects, we are confronted with a difficult algorithmic tuning problem. Moreover, some cases might require development of new algorithms instead of tuning existing algorithms. Finally, on the implementation side, questions arise about efficiency, training times, and speed of inference. Being able to provide the client with the right spectrum of choices is part of the R&D effort. This will allow the client to decide how to approach different streams of objects, based on specific assumptions for different deployment scenarios.

Open Set has been studied for years. Yet, solutions often depend on constraints of the application domain. There are many useful literature surveys, e.g., [3] and [4]. The candidate is encouraged to explore additional literature on the topic.

One topic of particular interest is Open Set for Object Recognition, going beyond mere image level labeling. This is a relatively new topic. For example, see [5].

The Taxonomy Problem.

Related to Open Set, the Taxonomy problem instead asks how to optimally structure the set of known object classes. While a flat taxonomy is the simplest choice from a ML point of view, it may not scale very well into hundreds or thousands of classes. At these scales, a natural choice is to define sets of more generic parent-classes, each of which contains more specific sub-classes, into a tree-like structure: the taxonomy.

Choices for the overall taxonomy structure range from shallow and wide to narrow and deep. Ideally, the taxonomy structure reflects requirements from the use case, enabling us to sort objects into the most optimal way. However, two new issues arise.

First, the distribution of available training data over the taxonomy. Data will become more sparse for some classes as a price we pay for more specific object sorting. While this might be attractive from the point of view of application requirement, it can also be detrimental from an ML point of view. Research is needed to find the right trade-off here.

Second, the distribution of ML decision making over the taxonomy. Instead of having one master model recognizing all classes in the data taxonomy, it might be attractive to view the most generic levels of the taxonomy as one between ML models instead of just object classes. Multiple models, each specializing in different branches of the taxonomy, might train and perform better than one master model.

Again, there is a trade-off. Too little model specialization may put too much stress on the master model, as different levels of visual variation between classes is too high (or low). Too much specialization may run us into compounding errors, as multiple levels of specialization each propagate their errors to their submodels. Research is needed to understand how to approach these trade-offs in practice. New algorithms will have to be developed and tuned to specific properties of the data.

Problems like this, concerning taxonomy and engineering choices have been around ever since ML started using big data. Solutions often depend on the particular constraints of the application domain, and the data that is available. An interesting starting point in the literature might be the group that wrote [7] and [8]. For an impression of how industry players have approached related problems, [6] is an interesting read.

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.

Some interesting research questions may concern the subtle interactions between the Open Set problem and the Taxonomy Problem, like the following:

- Perhaps we want some branches of the taxonomy to be more "open" than others?
- Perhaps different levels of visual variation are allowed in different classes?
- Perhaps the decision making should be allowed to backtrack through the taxonomy of specialization, depending on open set recognition?
- Perhaps we want algorithmic specialization between unknown vs. known object classes to be before, after, or independent of specialization between known classes only?
- Perhaps we would like some form of automated taxonomy restructuring to take place, depending on the frequency of open set detection and particularities of the data stream?
- How can existing scientific literature and software help us with these issues?
- How much of this should we develop ourselves?
- How much of this is specific for the client's setup, or can be reused in other applications?

For pixel segmentation, research questions may include

- How far can we take the ideas in [1] and the code in [2] in our practical setup?
- If we want to go further, is the next step in algorithm development, innovative use of loss functions, or data preparation? Or all of these combined? How exactly?
- There have been debates on how to treat the relative importance of loss-terms recently, and if linear combinations of loss terms is the best way to force optimal solutions. Are the loss terms of [1] susceptible to these considerations? Can we improve on the results of [1] by going beyond a single linear combination of the two loss terms involved?

The candidate is encouraged to think of their own research questions. We are always open to good ideas, and want to foster an environment in which ideas can grow. Please let us know your own views on the questions above.

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.

Support for working with various types of robots and automatic conveyor systems will be provided by our partners. We also envision a modular character for the internship, where most functionality on our topics of interest can be developed somewhat independently of the mechanical engineering questions related to the robots.

Sources

[1] BoxInst: High-Performance Instance Segmentation with Box Annotations
https://arxiv.org/abs/2012.02310

[2] Github: BoxInst: High-Performance Instance Segmentation with Box Annotations
https://github.com/aim-uofa/AdelaiDet/blob/master/configs/BoxInst/README.md

[3] Recent Advances in Open Set Recognition: A Survey https://arxiv.org/abs/1811.08581

[4] Survey: Open-set Classification
http://caiocv.com/wp-content/uploads/2018/07/survey-open-set.pdf

[5] The Overlooked Elephant of Object Detection: Open Set
https://www.researchgate.net/publication/341405186_The_Overlooked_Elephant_of_Object_Detection_Open_Set

[6] How we Scale Machine Learning https://scale.com/blog/how-we-scale-machine-learning

[7] On Flat versus Hierarchical Classification in Large-Scale Taxonomies
https://papers.nips.cc/paper/2013/file/cbb6a3b884f4f88b3a8e3d44c636cbd8-Paper.pdf

[8] Learning Taxonomy Adaptation in Large-scale Classification
https://www.researchgate.net/publication/309104482_Learning_Taxonomy_Adaptation_in_Large-scale_Classification

---

For more information, please contact our head of research,
Maarten Stol: maarten.stol@braincreators.com
Or visit our website: https://www.braincreators.com/contact