BrainCreators
Research Internship 2021-2022: Advanced Object Detections for Road Quality Maintenance with Geo-Information

---

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:

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

- **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.

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 Road Quality Maintenance with Geo-Information

Summary

Join our Deep Learning Road Quality team as a research intern. Work on R&D with visual- and geo-data for advanced object detection tasks to ensure road quality on an international level. Describe your scientific results in a publishable paper.

The end product of your research will be used for optimizing the performance of our Ai road inspectors. We are currently active in Europe, India, and South-America, and are always looking to expand further.

Product

Inspech: Automated Road Inspection
https://www.braincreators.com/inspech-digital-road-inspections?hsLang=en

Together with our road inspection partners, BrainCreators develops solutions for Ai driven inspection of road quality. Advanced imagery, RGB, video, lidar/laser, and pointcloud, are integrated into a geo-spatial interface, making new insights possible and repairs more efficient. Under the hood, many Ai related questions still deserve closer research attention.

BrainCreators wants to spend 1 year of R&D effort on the challenges involved in deep learning for object recognition integrated with suitable geo-information systems. Your team will aim to deliver a working module for the pipeline at the end of the research internship. A publishable paper should also be the aim of the research intern and their internal/external supervisors.

Technology

Examples of interesting technology include the following topics. These are examples we know we are interested in. An internship does not need to cover all of these, but combinations of topics like this are perhaps of higher scientific and practical relevance.

Data set design

We are interested in the impact of dataset design on our models for road inspection. The typical real-world class imbalance of majority and minority classes is ever present. But there are more issues here. The work in [3] provides an interesting start that might be relevant for our work on road infrastructure too. Short video introductions can be found in [1] and [2]. Key questions asked by the authors include:

- What is the influence of the similarity between base and test classes?
- Given a fixed annotation budget, what is the optimal trade-off between the number of images per class and the number of classes?
- Given a fixed dataset, can features be improved by splitting or combining different classes?
- Should simple or diverse classes be annotated?

Note that your work would include lifting these questions, and their proposed solutions, from image classification to the task of object recognition in ML.

Taxonomy design

Another issue that we would like to combine with questions around data set design is the challenge of variations in taxonomy. We have clients in different countries, and each have their own taxonomy for road defects. Some defects may be unique to one country, but many are not. However, even the shared defect types may in some cases still result in a taxonomy of a different structure. T

aking a model trained in one country to another country, some classes may be split, others merged, and yet others may get fragmented to multiple classes in the new taxonomy. We want to avoid having to start labeling from scratch every time we move to a new taxonomy. But how do we re-use what we have in an optimal way. Some ideas can be found in [4], in particular the method of Taxonomy Loss Masking.

Model Calibration for Object Detection

A third issue we have is the difficulty of translating our ML metrics into financial concepts such as how much income the model generates, and the financial risks associated with deployment at scale.

Ever since the now classic paper on neural network calibration [5], we know that modern neural networks have a tendency to sacrifice calibration for performance. Not a problem when all you want is to win a scientific competition measured by accuracy. But in the real world, end users need calibrated confidence scores as indicators of model correctness. For example, for deciding which inputs should be reviewed by humans. In that case, high confidence mistakes cost money!

More generally, calculating the financial gains of deployment, and finding the optimal scale of adoption are all problems that can only be solved by well calibrated models. The challenge for us is to take the next steps for ML calibration in the road inspection domain.

Some good starting points exist, like [6], which plots calibration metrics as a function of the x-y grid of the input. However, for roads the camera is moving, so the two dimensions have different roles in the data, and road scans can have overlap. This might influence visualizations of the type proposed in [6].

Finally, where [6] mostly aims at calibration post-training, other approaches exist that attempt to deliver well calibrated models during training. One example might be the work in [7], but others exist. The candidate is encouraged to look for more sources, and propose their own ideas to us.

The Geo-information aspect

All of our datasets live on maps, and every pixel has gps coordinates of a certain accuracy. Currently, labeling is done on the image level. We are moving towards labeling on the map. This has a number of consequences for our models and the training cycle.

One consequence of data living on a map is data sample dependency. For example, if two scans are adjacent in geo-space, should their inferences be assumed to be independent of each other? Basic ML theory says yes, as all data is assumed to be independent (and identically distributed). But common sense says no. A defect that extends to the edge of one image, should be assumed to continue on the next adjacent image. If there is overlap between the images, even more so.

Other background knowledge might also be incorporated into the deployment pipeline. A product is more than an ML model. For example, when a road bends or changes elevation, this might be associated with certain defects being more likely. Do we wait until the model discovers this, or do we add background knowledge to the pipeline?

Integration of geo-information and related background knowledge is an important challenge when taking ML products to the next level. You are invited to help us think about new solutions!

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 combinations of the technology described above:

- Given a fixed annotation budget, what is the optimal way to approach data set design and taxonomy design?
- How would we evaluate the proposed ideas?
- How should we set up an experimentation pipeline to test an idea like taxonomy masking the optimal designs of our road inspection data set?
- How should we prepare for ingestions of a new client's data, a new country, a new continent?

- What are the best ways to adapt existing ideas for model calibration for object detection to applications in road inspection?
- How far can we optimize with post-training calibration methods?
- Is it worth the time and effort to adopt more complex methods that yield calibrated models by design? How much performance gain do we get, and at what price?

- How should geo-information background knowledge be added to the ML pipeline?
- How should data sample dependencies be treated by the pipeline? Only in post-processing, or during training. If so, then how should the independence assumption be broken?

- Finally, what is the impact of our semi-supervised / self-supervised training loop on the ideas from these research questions?
- Our semi-supervision methods create some feedback in the training cycle. How can we avoid unwanted effects that are introduced as a result of innovations in data set and taxonomy design?

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 data from image capture hardware will be provided by our partners and us. We also envision a modular character for the internship, where most of our topics of interest can be developed independently from the engineering questions related to image capture.

Sources

[1] (ECCV 2020) Impact of base dataset design on few-shot image classification (1 minute)
https://www.youtube.com/watch?v=6EL9oeUDYAU&t=58s

[2] (ECCV 2020) Impact of base dataset design on few-shot image classification - Long Presentation
https://www.youtube.com/watch?v=M9LOiJY4ZO4

[3] Impact of base dataset design on few-shot image classification
http://imagine.enpc.fr/~sbaio/publications/fewshot_dataset_design/index.html

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

[5]On Calibration of Modern Neural Networks
https://paperswithcode.com/paper/on-calibration-of-modern-neural-networks

[6] Multivariate Confidence Calibration for Object Detection https://arxiv.org/abs/2004.13546

[7] Learning for Single-Shot Confidence Calibration in Deep Neural Networks through Stochastic Inferences https://arxiv.org/abs/1809.10877

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