Graduation Assignments Spotr 2022-09

Introduction

What does Spotr do? We Inspect millions of buildings in seconds

Our vision: To create a sustainable world, it is essential to have proper insights in buildings

Our mission: We pro-actively inspect every building in the world

Who is Spotr? A team of 30 highly motivated young enthusiasts

Where are we located? We regularly meet in our office in The Hague and often work from home

What do we make? Our web-based platform provides user-friendly access to the real estate of our customers. Our customers are mostly housing associations and they use the platform to get a quick view of their real estate to schedule maintenance.

Which problem do we solve? We create detailed data about buildings in order to support the transition to a sustainable world

What makes AI research at Spotr fun? We have the biggest database of building information in the Netherlands, containing close to a million images of buildings. Many are labeled and are used to train our recognition algorithms, but there is still an impressive amount of information hidden in the image data. Do you want to help us find it? All our data and algorithms run in the cloud, making both training and inference incredibly scalable. Our highly skilled team of Data Scientists and engineers will support you with all tooling so that you can focus on the AI research for improved building recognition. You will get a chance to experience how AI is put into production, from research to large-scale deployment. Your supervisor has many years or experienced in academic research and supervising (MSc and PhD students).

What are the challenges that we face for ML?

We have currently defined 5 assignments for MSc graduation projects:

1. Unsupervised or semi-supervised learning to extract info for unbalanced building object classes
2. Few-shot model adaptation for internationalization: from our Netherlands ML model towards other countries
3. Multi-modal data sources: exploitation of multiple other non-image data sources
4. Building semantic segmentation: extend the pure pixel-based labeling with application information
5. Change detection over time or different buildings

Are you a match for our team? You will be warmly welcomed in our super pro-active R&D team. Together with our computer vision experts and data scientists, we will take machine learning to the next level. We will prepare large amounts of data to support the research that you will be doing. We have extensive experience with academic supervision (both MSc and PhD projects).

More info? www.spotr.ai

Interested in an assignment? Send your CV and a short motivation letter to Rob Wijnhoven, rob@spotr.ai

The assignments

1. Unsupervised or semi-supervised learning to extract info for unbalanced building object classes

Problem definition

We have annotated several thousands of images with pixel-accurate information about facades and roofs. In each picture we have labeled several object classes for building materials (brick, concrete, wood, frame, glass panel, rain pipe, etc.). On the one hand we suffer from large class imbalances in the training data. There are plenty of windows, but rain pipes occur much less frequently. Other classes like ladders or ventilation openings are very rare and the current labels are insufficient to train a model at all for these classes. How do we find more instances of these classes in our very large set of building images? More interestingly, how to we find possible new object classes in an unsupervised way?

Work environment

You will be warmly welcomed in our super pro-active R&D team. Together with our computer vision experts, data scientists and engineers, we will take machine learning to the next level. You will work on our existing large-scale building dataset of ~700k images of which ~30k images are annotated at pixel accuracy. You will train your models using PyTorch in the cloud, so that training time is unlimited and you can run many experiments to support your research. Our existing production training pipeline is available, but you can also choose to use different tools of your choice. After researching related available solutions for the problem in academic literature, we jointly decide on a solution direction.

At Spotr, everyone decides to work from home or to come to our office in The Hague (10 minutes walking from the central train station). Especially in the beginning of the project, we expect you to be in the office for 2 days, you can work from home the rest of the week if desired. Joint time with our team in the office is important to facilitate discussions and brainstorm sessions, and to support you technically once you want to start training models in the cloud. We will track progress in a weekly meeting. We expect you to work independently, but highly support frequent technical discussions on creative ideas. We have extensive experience with academic supervision (both MSc and PhD projects).

Expectations

- Literature review to find related solutions
- Implementation of a new or existing ML model
- Together we will select a set of classes for investigation
- You will propose a process to find more instances of these classes in our large unannotated dataset
- Depending on the experience and interest of the student we can focus on finding rare classes or finding new object classes in an unsupervised way

2. Few-shot model adaptation for internationalization: from our Netherlands ML model towards other countries

Problem definition

Although houses in the Netherlands, Germany and Belgium might look quite similar at first sight, there are many differences in their construction that lead to different ways of managing the properties from a maintenance standpoint. From a civil engineering perspective this is a huge challenge. Our current ML models work well on Dutch houses, and in general will give decent results on Belgian and German houses, but we expect that houses in certain less developed parts of the world will lead to challenges. The typical ML approach is: collect images, label them and train your model, but this does not scale if we want to inspect every building on the planet in the coming years. We need ML models to adapt from the Netherlands to Brazil or India. Note that we have extensive data sources of buildings in the Netherlands that enables us to build a super model for Dutch houses. We want you to think of how we can adapt this model to other countries. As an example, we have an initial image dataset with a few hundred houses in the UK in different regions that can be used for experiments and validation.

Work environment

You will be warmly welcomed in our super pro-active R&D team. Together with our computer vision experts, data scientists and engineers, we will take machine learning to the next level. You will work on our existing large-scale building dataset of ~700k images of which ~30k images are annotated at pixel accuracy. You will train your models using PyTorch in the cloud, so that training time is unlimited and you can run many experiments to support your research. Our existing production training pipeline is available, but you can also choose to use different tools of your choice. After researching related available solutions for the problem in academic literature, we jointly decide on a solution direction.

At Spotr, everyone decides to work from home or to come to our office in The Hague (10 minutes walking from the central train station). Especially in the beginning of the project, we expect you to be in the office for 2 days, you can work from home the rest of the week if desired. Joint time with our team in the office is important to facilitate discussions and brainstorm sessions, and to support you technically once you want to start training models in the cloud. We will track progress in a weekly meeting. We expect you to work independently, but highly support frequent technical discussions on creative ideas. We have extensive experience with academic supervision (both MSc and PhD projects).

Expectations

- Literature review to find related solutions for domain adaptation
- Evaluation of an existing model and benchmark SotA performance
- Propose improvement to the applied model
- Benchmark on both public dataset and the Spotr dataset
- Explore how novel object classes behave without any prior labels (example: some UK houses have rain pipes above the roof, used for ventilation of toilets)

3. Multi-modal data sources: exploitation of multiple other non-image data sources

Problem definition

Dutch buildings are a great start for global building inspection because lots of data are available. Apart from the large image database that we have at Spotr, we have access to various other data sources, such as construction year, energy label, WOZ value (tax value of the building), rebuild value (for insurance), geographical location, social neighborhood score, etc. We think that it would be great to couple all this information to the images of each building. This can be performed both at the input side of the ML model (the information is fed as additional input data) and help the model to recognize the façade pixels better, or it can be added to the estimation of the output of the ML model, so that the model can also estimate energy labels and rebuild value from just an image as input. We believe that some data sources might be used as input, while others should be used as output. Your task is to investigate how we can integrate these data sources and how (much) it can help ML models.

Work environment

You will be warmly welcomed in our super pro-active R&D team. Together with our computer vision experts, data scientists and engineers, we will take machine learning to the next level. You will work on our existing large-scale building dataset of ~700k images of which ~30k images are annotated at pixel accuracy. You will train your models using PyTorch in the cloud, so that training time is unlimited and you can run many experiments to support your research. Our existing production training pipeline is available, but you can also choose to use different tools of your choice. After researching related available solutions for the problem in academic literature, we jointly decide on a solution direction.

At Spotr, everyone decides to work from home or to come to our office in The Hague (10 minutes walking from the central train station). Especially in the beginning of the project, we expect you to be in the office for 2 days, you can work from home the rest of the week if desired. Joint time with our team in the office is important to facilitate discussions and brainstorm sessions, and to support you technically once you want to start training models in the cloud. We will track progress in a weekly meeting. We expect you to work independently, but highly support frequent technical discussions on creative ideas. We have extensive experience with academic supervision (both MSc and PhD projects).

Expectations

- Literature review to find related solutions for extending purely image-based models to additional data sources
- Explore novel data sources
- Integrate additional data sources into the existing semantic segmentation model
- Benchmark on both public dataset and the Spotr dataset

4. Building semantic segmentation: extend the pure pixel-based labeling with application information

Problem definition

Semantic segmentation now ignores the fact that we are inspecting buildings. Other cues such as the corner points of the wall and the orientation of lines in the façade can help the segmentation to infer more accurate boundaries around building wall elements (windows are typically rectangular so why try to estimate each pixel independently). As an example of the DeeplabV3 segmentation model, the CRF post-processing module could be extended with this additional wall information.

Work environment

You will be warmly welcomed in our super pro-active R&D team. Together with our computer vision experts, data scientists and engineers, we will take machine learning to the next level. You will work on our existing large-scale building dataset of ~700k images of which ~30k images are annotated at pixel accuracy. You will train your models using PyTorch in the cloud, so that training time is unlimited and you can run many experiments to support your research. Our existing production training pipeline is available, but you can also choose to use different tools of your choice. After researching related available solutions for the problem in academic literature, we jointly decide on a solution direction.

At Spotr, everyone decides to work from home or to come to our office in The Hague (10 minutes walking from the central train station). Especially in the beginning of the project, we expect you to be in the office for 2 days, you can work from home the rest of the week if desired. Joint time with our team in the office is important to facilitate discussions and brainstorm sessions, and to support you technically once you want to start training models in the cloud. We will track progress in a weekly meeting. We expect you to work independently, but highly support frequent technical discussions on creative ideas. We have extensive experience with academic supervision (both MSc and PhD projects).

Expectations

- Literature review to find related solutions for integrating domain knowledge into the semantic segmentation model
- Propose and implement novel solution for domain knowledge integration
- Benchmark on both public dataset and the Spotr dataset

5. Change detection over time or different buildings

Problem definition

Detect (a) modifications in construction of buildings or (b) differences between different buildings in a street. The algorithm should react to real-world changes such as dormers added to the roof or enlarged windows, and not to changed vegetation or photogrammetric or light/shadow differences. A dataset will be made available by acquiring images from multiple years from our data provider. The challenge here is that the picture quality is much lower for older pictures, so that is another challenge that the network should consider.

Work environment

You will be warmly welcomed in our super pro-active R&D team. Together with our computer vision experts, data scientists and engineers, we will take machine learning to the next level. You will work on our existing large-scale building dataset of ~700k images of which ~30k images are annotated at pixel accuracy. You will train your models using PyTorch in the cloud, so that training time is unlimited and you can run many experiments to support your research. Our existing production training pipeline is available, but you can also choose to use different tools of your choice. After researching related available solutions for the problem in academic literature, we jointly decide on a solution direction.

At Spotr, everyone decides to work from home or to come to our office in The Hague (10 minutes walking from the central train station). Especially in the beginning of the project, we expect you to be in the office for 2 days, you can work from home the rest of the week if desired. Joint time with our team in the office is important to facilitate discussions and brainstorm sessions, and to support you technically once you want to start training models in the cloud. We will track progress in a weekly meeting. We expect you to work independently, but highly support frequent technical discussions on creative ideas. We have extensive experience with academic supervision (both MSc and PhD projects).

Expectations

- Literature review to find related solutions for change detection for buildings
- Initial focus of the project will be on either (a) change detection of single building over time, or (b) multiple buildings in a street of the same time
- Propose and implement novel solution for change detection
- Benchmark on both public dataset and the Spotr dataset
- Depending on the interest and experience of the student, the focus can be extended to the other change detection case (a/b) once a working solution is available for the first case