



Stichting NIOC

Stichting NIOC en de NIOC kennisbank

Stichting NIOC (www.nioc.nl) stelt zich conform zijn statuten tot doel: het realiseren van congressen over informatica onderwijs en voorts al hetgeen met een en ander rechtstreeks of zijdelings verband houdt of daartoe bevorderlijk kan zijn, alles in de ruimste zin des woords.

De stichting NIOC neemt de archivering van de resultaten van de congressen voor zijn rekening. De website www.nioc.nl ontsluit onder "Eerdere congressen" de gearchiveerde websites van eerdere congressen. De vele afzonderlijke congresbijdragen zijn opgenomen in een kennisbank die via dezelfde website onder "NIOC kennisbank" ontsloten wordt.

Op dit moment bevat de NIOC kennisbank alle bijdragen, incl. die van het laatste congres (NIOC2025, gehouden op donderdag 27 maart 2025 jl. en georganiseerd door Hogeschool Windesheim). Bij elkaar zo'n 1500 bijdragen!

We roepen je op, na het lezen van het document dat door jou is gedownload, de auteur(s) feedback te geven. Dit kan door je te registreren als gebruiker van de NIOC kennisbank. Na registratie krijg je bericht hoe in te loggen op de NIOC kennisbank.

Het eerstvolgende NIOC vindt plaats in 2027 en wordt dan georganiseerd door HAN University of Applied Sciences. Zodra daarover meer informatie beschikbaar is, is deze hier te vinden.

Wil je op de hoogte blijven van de ontwikkeling rond Stichting NIOC en de NIOC kennisbank, schrijf je dan in op de nieuwsbrief via

www.nioc.nl/nioc-kennisbank/aanmelden-nieuwsbrief

Reacties over de NIOC kennisbank en de inhoud daarvan kun je richten aan de beheerder:

R. Smedinga kennisbank@nioc.nl.

Vermeld bij reacties jouw naam en telefoonnummer voor nader contact.


Actively Evolving Course

NIOC 2023

Dr. Klaas Dijkstra

Lector Computer Vision & Data Science
NHL Stenden

30 Maart 2023



Inhoud

De onderwijscontext



Het abstracte onderwijs concept

De concrete uitwerking






Vragen aan jullie!




Zelf aan de slag




Voorbeelden van Computer Vision & Data Science



Het Lectoraat Computer Vision & Data Science



Het Lectoraat Computer Vision & Data Science



Onderzoek

Werken in projecten




Semester 1 (30 ECs)

Evidence-based portfolio

Knowledge and skills (2 sprints)

Applied research project (5 sprints)


portfolio-based assessment



Leerweg onafhankelijk

Portfolio-gebaseerd,
programmatisch toetsen.

Holistisch beoordelen aan het
einde van ieder semester.



**“Ieder project is een avontuur,
welke leeruitkomsten horen daar
bij?”**

Hoe stem je verwachtingen af?

Gering aantal **leeruitkomsten**,
breed geformuleerd.

Educatie

Learning Outcomes



Educatie

Learning outcome 1

The student designs, develops and tests independently, within a team and methodologically correct, machine-learning algorithms that automate visual inspections that meet the customer's specifications.

Learning outcome 2

The student creates and manages, in collaboration with domain expert(s), a representative annotated and balanced dataset with the required quality to develop and test machine-learning algorithms.

Learning outcome 3

The student optimizes algorithms, independently and based on customer specifications, and makes them scalable so they can be applied in practice.

Learning outcome 4


The student develops himself proactively and with a high degree of responsibility, to guarantee his own sustainable employability and thus also contribute to the development of professional practice and the knowledge domain.

**“Design-based education is
geen one-size-fits-all”**


Hoe ziet DBE er voor jou uit?

Design Based Education binnen NHL Stenden

Onderwijsconcept




Design Based Education binnen NHL Stenden




Stepping into Design-Based Education:

Chapter: Applied research using design-based education in a technical context



Integratie van meerdere concepten:
DBE, Agile, Waterfall, Scientific method.




Ons onderwijsontwerp


Regelmatige feedback in
“progress meetings”.

Korte design thinking
sprints in een BaMa leerlijn.

Minor (Bachelor)



Master (Master of Science)



Increasing complexity and independence

Orientation

Final qualifications

“Eén keer toetsen per
semester in de feedback
factory”

Hoe schaal je dit op?

De topics

Themes	Computer Vision	Machine Learning	Deep Learning
Groups	Introduction	Fundamentals	Hardware
Topics	Math for Machine Learning 1	Linear Algebra 1	Spectral Imaging 1
Semester	Scientific Programming 1	Image Features 1	Vision Hardware 1 2
	Tools 1	Annotation 1	Camera Geometry 1 2
	My First Pipeline 1	Validation 1	Computing Hardware 1 2
		Data Exploration 1	3D Computer Vision 2
	21st Century Skills	Hyper Parameter Tuning 1 2	
	Scientific Writing 1	Dimensionality Reduction 2	Learning
	Professionalization 1 2	Transforms 2	Methodology 1
	Ethics 1 2		Probability and Statistics 1
		Tasks	Training 1
	Misc. and Future	Classification 1	Linear Models 1
	Reinforcement Learning 2	Segmentation 1	Neural Networks 1
	Big Data 2	Object Detection 1	Convolutional Neural Nets. 1
	CNNs for Other Modalities 2	Other tasks 1	Generative Adversarial Nets. 2
		Unsupervised Learning 2	Advanced Architectures 2
		Anomaly Detection 2	
		Explainable A.I. 2	

Een topicstructuur

Een topic bestaat uit:


- Een centraal onderwerp
- Een lijst met *learning resources*
 - (video's, boeken, papers, tutorials, etc.)
- *Exercises* die automatisch gecontroleerd worden.
- *Quiz* met theorievragen.
- Heeft een omvang van één dag.
- Iedere docent-onderzoeker kan ieder topic verzorgen.

Voorbeeld learning resources:

Main Books



Gonzales, Rafael C., and Richard E. Woods. *Digital image processing*.



Deisenroth, Marc Peter, Aldo Faisal, and Cheng Soon Ong. *Mathematics for machine learning*. Cambridge University Press (e-book version freely available at <https://mml-book.github.io/book/mml-book.pdf>)



Goodfellow, I., Bengio, Y., Courville, A., & Bengio, Y. *Deep learning*. Cambridge: MIT press. (e-book version freely available at <https://www.deeplearningbook.org/>)



Bishop, Christopher M. *Pattern recognition and machine learning*. Springer.

Main Online Courses



Stanford University Course [CS229](#) Machine Learning




Stanford University Course [CS230](#) Deep Learning



Stanford University Course [CS231](#) Convolutional Neural Networks for Visual Recognition

Een voorbeeld topic:

Advanced Classification



Classification is a common task in deep learning. The goal of classification models is usually to directly predict the class of an image. This topic contains example pipelines based on the classic LeNet architecture. The LeNet example runs on the CIFAR10 dataset. Below are a few example images from the dataset.

Resources

Before continuing please get familiar with classification using convolutional neural networks. To get you started, here you can find some landmark overviews about classification with deep learning. Read those carefully. There are plenty of deep learning architectures for classification.

Online Course

- Stanford University [Link](#)
 - Lecture 1: Introduction to Convolutional Neural Networks for Visual Recognition

Scientific Papers

A survey research paper that discusses some important classification architectures in greater detail:

- A Survey of the Recent Architectures of Deep Convolutional Neural Networks, [Link](#)

The following bullets contain links to the papers of several important classification architectures, starting with LeNet from 1998.

- Gradient-Based Learning Applied to Document Recognition, [Link](#)
- ImageNet Classification with Deep Convolutional Neural Networks, [Link](#)
- Going Deeper with Convolutions, [Link](#)

An important class of deep learning architectures are called Inception or GoogLeNet. Throughout the years multiple improvements have been done.

- Rethinking the Inception Architecture for Computer Vision, [Link](#)
- Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning, [Link](#)

Additionally an ever increasing amount of classification architectures exists:

- Deep Residual Learning for Image Recognition, [Link](#)
- Aggregated Residual Transformations for Deep Neural Networks, [Link](#)
- Densely Connected Convolutional Networks, [Link](#)
- Squeeze-and-Excitation Networks, [Link](#)

Additional Resources

Following are additional resources that might be helpful in the learning process of the classification with deep learning.. Below is a link to a Medium post which gives a nice overview of some important classification architectures.

Communicate sprint producten van de studenten



To the effect of dataset order in multiple object tracking

Moustafa Elhagaly

Supervisors: Klaas Dijkstra and Lucas Ramos

Abstract—In this paper, we explore the effect of domain adaptation and catastrophic forgetting in conjunction with multiple object tracking and its application of two-wheeler tracking to see if the order of datasets during training matters. We employ a Siamese Multiple Object Tracker (SiamMOT) and train it using different permutations on two public datasets (Multiple Object Tracking (MOT) and Specialized Cyclists Dataset (SCD)) and a proprietary dataset (Traffic Intersection Dataset (TID)). We ran experiments on the datasets using different permutations to test the performance when trained on a single dataset, combining datasets and sequence training. We also qualitatively tested the generalizability of the best model in dusk/night footage. Training exclusively on the TID dataset results in the highest IDF1-score, and combining datasets results in a lower IDF1-score compared to when training exclusively on TID. Catastrophic forgetting occurs when training the model with datasets in different orders, where swapping orders of datasets leads to a reduction of about 30% in performance. We have shown that the order of datasets during training plays an important role when adapting datasets from different domains. The best model shows promising results when testing the generalizability on data from different conditions. The qualitative results of the best model on crossing a red light detection show the possibilities of using tracking-by-detection models for other traffic safety indicators.

Index Terms—multiple object tracking, traffic analysis, siamese network, object detection, incremental learning, catastrophic forgetting

1 INTRODUCTION

According to the global status report on road safety published by the World Health Organisation (WHO), approximately 1.3 million people die each year as a result of road traffic accidents. On average, a cyclist dies every 12 minutes and 30 seconds somewhere in the world [1]. Traffic Safety Specialists are in charge of analyzing traffic footage to identify risk situations and identify possible areas for improvements, while, traffic psychologists analyze the behavior of road users identifying dangerous behavior. Moreover, expert knowledge is required during the analysis. This analysis is time-consuming since it requires the specialist to watch footage that corresponds to long periods of time; therefore, conclusions are often drawn based on the analysis of shorter video fragments. Such approaches can lead to biased results and a distorted vision of a certain traffic region's overall problems. Traffic analysis using Deep Learning could assist in identifying and measuring hazardous situations in traffic, reducing the time necessary for analysis and being less biased. Deep learning technology has seen major advances in recent years with detection algorithms including Faster R-CNN [2], SPPNet [3] and YOLO [4, 5, 6, 7]. Given the advancements in object detection, these models have been widely used in multiple




Fig. 1. Tracking-by-detection paradigm. First the object is detected and then the detections are associated across frames to form tracks.

trajectories by estimating object motion. Despite the recent advances in MOT, many challenges remain: 1) the tracker has to deal with multiple objects that need to be tracked from the moment they appear to the moment they disappear from the scene; 2) frequent object

Paper

Dataset – tiling

- Fixed tiling

- Each tile is a class

- Anomaly

- Training: normal

- Validation: normal


CutPaste framework

Self-supervised representation learning for anomaly detection [2]

- CutPaste aims to create irregular patterns on the normal data where the patterns resemble a possible anomaly

- CutPaste alteration: cutting a random patch from the image and pasting it back on

- Binary classification: normal & altered
- ResNet-18 [3]



Presentatie op het symposium Computer Vision & Data Science

Automatic quantification of traffic safety with multiple object tracking using deep learning

Moustafa Elhagaly
Supervisors: Klaas Dijkstra & Lucas Ramos

Winter 2022

Abstract

Abstract: In this paper, we explore the effect of domain adaptation and catastrophic forgetting in conjunction with multiple object tracking and its application of two-wheeler tracking to see if the order of datasets during training matters. We employ a Siamese Multiple Object Tracker (SiamMOT) and train it using different permutations on two public datasets (Multiple Object Tracking (MOT) and Specialized Cyclists Dataset (SCD)) and a proprietary dataset (Traffic Intersection Dataset (TID)). We ran experiments on the datasets using different permutations to test the performance when trained on a single dataset, combining datasets and sequence training. We also qualitatively tested the generalizability of the best model in dusk/night footage. Training exclusively on the TID dataset results in the highest IDF1-score, and combining datasets results in a lower IDF1-score compared to when training exclusively on TID. Catastrophic forgetting occurs when training the model with datasets in different orders, where swapping orders of datasets leads to a reduction of about 30% in performance. We have shown that the order of datasets during training plays an important role when adapting datasets from different domains. The best model shows promising results when testing the generalizability on data from different conditions. The qualitative results of the best model on crossing a red light detection show the possibilities of using tracking-by-detection models for other traffic safety indicators.

Introduction

According to the global status report on road safety published by the World Health Organisation (WHO), approximately 1.3 million people die each year as a result of road traffic accidents. On average, a cyclist dies every 12 minutes and 30 seconds somewhere in the world [1]. Traffic Safety Specialists are in charge of analyzing traffic footage to identify risk situations and identify possible areas for improvements, while, traffic psychologists analyze the behavior of road users identifying dangerous behavior. Moreover, expert knowledge is required during the analysis. This analysis is time-consuming since it requires the specialist to watch footage that corresponds to long periods of time; therefore, conclusions are often drawn based on the analysis of shorter video fragments. Such approaches can lead to biased results and a distorted vision of a certain traffic region's overall problems. Traffic analysis using Deep Learning could assist in identifying and measuring hazardous situations in traffic, reducing the time necessary for analysis and being less biased. Deep learning technology has seen major advances in recent years with detection algorithms including Faster R-CNN [2], SPPNet [3] and YOLO [4, 5, 6, 7]. Given the advancements in object detection, these models have been widely used in multiple

Experiments and Results

• **Tracker and dataset permutations**

- The best IDF1 score (82.8%) is achieved by training on the TID dataset.
- The effect of MOT is visible when comparing the single and combined experiments.
- The best result of the sequence experiment is when fine-tuning MOT → TID.
- The experience of dataset order when training is important as shown when training SCD → MOT → TID resulting in an IDF1 score of 81.8% and reversing the order of the MOT and TID results in an IDF1 score of 79.8%.

Dataset	Tracker	IDF1	AP
MOT	SiamMOT	77.8%	27.0%
SCD	SiamMOT	80.8%	27.0%
TID	SiamMOT	82.8%	27.0%
MOT + SCD	SiamMOT	79.8%	27.0%
MOT + TID	SiamMOT	80.8%	27.0%
SCD + TID	SiamMOT	81.8%	27.0%
MOT + SCD + TID	SiamMOT	80.8%	27.0%
MOT	DeepSORT	77.8%	27.0%
SCD	DeepSORT	80.8%	27.0%
TID	DeepSORT	82.8%	27.0%
MOT + SCD	DeepSORT	79.8%	27.0%
MOT + TID	DeepSORT	80.8%	27.0%
SCD + TID	DeepSORT	81.8%	27.0%
MOT + SCD + TID	DeepSORT	80.8%	27.0%

• **Tracker evaluation with traffic light detection**

- The best model (single SCD) is used and results in long tracks allowing for the tracker to track objects past the traffic light which in conjunction with traffic light state allows for crossing red light detection.

Materials and Methods

Network:

- SiamMOT [2] that uses Faster R-CNN for object detection and Explicit Motion Scheduling that explicitly learns a matching function between the same instance in sequential frames.

Datasets:


- MOT: Multiple Object Tracking
- SCD: Specialized Cyclists Dataset
- TID: Traffic Intersection Dataset

Poster op het symposium Computer Vision & Data Science

Progress and Assessment Cycle Master CV&DS

Portfolio-gebaseerd,
programmatisch toetsen.


Holistisch beoordelen aan het
einde van ieder semester.



**“Het proces centraal stellen
geeft ruimte voor flexibele
inhoud.”**


Hoe ziet jullie proces eruit?

De pipeline



1. Een pipeline biedt een gestructureerde en herkenbare methode om praktijkproblemen op te lossen.
2. Een pipeline bestaat uit elementen in categorieën zoals *load data*, *visualize* en *save results*.


Evolving course



First the network is trained and each 10th epoch the results are validated on a separate set. The model with the lowest loss on the validation set is saved to yolov5.pth (early stopping). The training and validation loss are shown in a tensorboard session. During validation one of the processed tiles for each minibatch is shown in tensorboard and after a few epochs. In our training session this looked like this ²:



The testing pipeline loads the model and estimated bounding boxes on the full image. The testing pipeline outputs the metrics Precision, Recall and F1-Score. The results are shown in a tensorboard session. In our run the results looked like this:



1. Nieuwe elementen of pipelines worden gemaakt binnen de projecten.
2. Studenten kunnen deze pipelines via het digitale systeem maken en hier *Exercises* mee doen.
3. Studenten kunnen deze pipelines gebruiken als basisoplossing voor hun project.
4. Nieuwe pipelines worden toegevoegd aan het digitale systeem voor toekomstige projecten en studenten.

Een Exercise vanuit Python (voor de ontwerper)

```
def create_internal_camera_matrix(f: float, s: float, u: float, v: float, a: float) -> np.ndarray:
    """
    This method generates an internal camera matrix for the pinhole camera model with
    the specified parameters.

    :param f: The focal length in mm
    :param s: The image scaling
    :param u: The x-translation in the camera coordinate system
    :param v: The y-translation in the camera coordinate system
    :return: A [3, 4] internal camera matrix with focal length f, scaling s, translation u,v and skew a.

    >>> create_internal_camera_matrix(8.0, 2.0, 10.0, 5.0, 0.1)
    array([[16. ,  0.1, 10. ,  0. ],
           [ 0. , 16. ,  5. ,  0. ],
           [ 0. ,  0. ,  1. ,  0. ]])
    >>> create_internal_camera_matrix(8.0, 2.0, -10.0, -5.0, 0.2)
    array([[ 16. ,   0.2, -10. ,   0. ],
           [  0. ,   16. ,  -5. ,   0. ],
           [  0. ,   0. ,   1. ,   0. ]])
    """
    mat = np.zeros([3, 4])
    mat[0, 0] = f * s
    mat[1, 1] = f * s
    mat[0, 2] = u
    mat[1, 2] = v
    mat[0, 1] = a
    mat[2, 2] = 1.
    return mat
```

Een *Exercise* vanuit Jupyter (voor de student)

Exercise 4

This method generates an internal camera matrix for the pinhole camera model with the specified parameters.

param f

The focal length in mm

param s

The image scaling

param u

The x-translation in the camera coordinate system

param v

The y-translation in the camera coordinate system

return

A [3, 4] internal camera matrix with focal length f, scaling s, translation u,v and skew a.

```
In [ ]: # Please add code to this function to complete the assignment.
# You can use the description and code cells below to check you answer.
# Tip: With print('hello') you can print debugging info.

def create_internal_camera_matrix(f: float, s: float, u: float, v: float, a: float) -> np.ndarray:
    # Add/adjust code.
    return
```

```
In [ ]: # This cell performs a doctest of the function.
newline = '\n'
your_answer = repr(create_internal_camera_matrix(8.0, 2.0, 10.0, 5.0, 0.1)).strip()
real_answer = str("array([[16. , 0.1, 10. , 0. ],\n          [ 0. , 16. , 5. , 0. ],\n          [ 0. , 0. , 1. , 0. ]])")
assert your_answer == real_answer, f'Wrong output, see details below. {newline}{newline}Your answer {newline}{newline}({your_answer})'
print('<Test succeeded!>')
```

```
In [ ]: # This cell performs a doctest of the function.
newline = '\n'
your_answer = repr(create_internal_camera_matrix(8.0, 2.0, -10.0, -5.0, 0.2)).strip()
real_answer = str("array([[ 16. , 0.2, -10. , 0. ],\n          [ 0. , 16. , -5. , 0. ],\n          [ 0. , 0. , 1. , 0. ]])")
assert your_answer == real_answer, f'Wrong output, see details below. {newline}{newline}Your answer {newline}{newline}({your_answer})'
print('<Test succeeded!>')
```

De student geeft
zelf de uitwerking.

Het systeem kijkt
deze na en geeft
de student
feedback.

Ons onderwijsontwerp

“You cannot get educated by this self-propagating system in which people study to pass exams, and teach others to pass exams, but nobody knows anything. You learn something by doing it yourself, by asking questions, by thinking, and by experimenting.”
– R. Feynman

Student ervaringen

- **“Je weet waarvoor je het doet door de directe koppeling met de praktijk” – Master Student**
- **“Wat super fijn is de mogelijkheid om voor- en achteruit te kunnen kijken in de topics” – Master Student**
- **“I liked the DBE way of teaching since we got to work on a project with real-life applications instead of just a 'school' project. Additionally, we also collaborated with people from other disciplines during our project, which I found quite interesting.” – Minor Student**
- **“I think it was a very interesting experience as an Erasmus student. I had to develop my skills and research capability in order to make what was expected.” – Erasmus Student**

Zelf aan de slag

- **Master-apprentice** werken aan échte problemen.
- Het **werkproces centraal stellen** in “*dataverwerkings pipelines*”.
- **Geen vakken**, maar **compacte “topics”**.
- Altijd **actueel cursusmateriaal** door een “*evolving course*”.
- **Korte** design thinking **sprints** in een BaMa leerlijn.
- **Regelmatige feedback** in “*progress meetings*”.
- Gering aantal **leeruitkomsten**, breed geformuleerd.
- **Portfolio-gebaseerd**, programmatisch toetsen.
- **Holistisch beoordelen** aan het einde van ieder semester.