

HUMBOLDT-UNIVERSITÄT ZU BERLIN
MATHEMATISCH-NATURWISSENSCHAFTLICHE FAKULTÄT
INSTITUT FÜR INFORMATIK

Autonomous Racing with Deep Learning

Bachelorarbeit

zur Erlangung des akademischen Grades
Bachelor of Science (B. Sc.)

eingereicht von: Philipp Badenhoop

geboren am: 7.7.1996

geboren in: Berlin Neukölln

Gutachter/innen: Prof. Dr. Holger Schlingloff

Prof. Dr. Verena Hafner

eingereicht am:

verteidigt am:

Abstract

This thesis covers the implementation of an autonomous racing car in scale 1:8 using deep learning. Based on the current camera frame, our method employs a neural network which predicts the vehicle's future path. The network is trained in a supervised fashion using data which is recorded while a human drives the car manually around the racing track. We propose a simple model to control the steering angle and speed using the network's predictions. Furthermore, we show that our method is both easier to train and to scale up to higher velocities than the well known end-to-end approach which uses a neural network to output the steering commands directly. With our autonomous racing car, we were able to win the third place in the Deep Berlin Robocars Challenge.

Contents

1. Introduction	6
1.1. Related Work	6
1.2. Structure of the Thesis	7
2. System Architecture	8
2.1. Chassis	8
2.2. Hardware	9
2.3. Camera	10
2.4. Wheel Encoders	10
2.5. Software Architecture	10
2.5.1. Odroid	11
2.5.2. STM32	15
3. Deep Learning	16
3.1. Motivation	16
3.2. Artificial Neural Networks	16
3.3. Learning	19
3.3.1. Loss Function	20
3.3.2. Optimization	21
3.3.3. Backpropagation	25
3.3.4. Overfitting	28
3.4. Convolutional Neural Networks	29
4. Autonomous Driving	31
4.1. End-to-End Approach	31
4.1.1. Preprocessing	31
4.1.2. Neural Network Model	32
4.1.3. Training	33
4.1.4. Dataset	34
4.1.5. Test Results	35
4.2. Trajectory Prediction Approach	39
4.2.1. Overview	39
4.2.2. Odometry	41
4.2.3. Neural Network Model	45
4.2.4. Driving Model	45
5. Results	48
5.1. Wheel Encoder Anomaly	48
5.2. Dataset	49
5.3. Final Race	49
5.4. Limits	51
5.5. Possible Improvements	51

6. Conclusion	53
A. Code Snippets	58
A.1. End-to-End Model Implementation using Keras	58
A.2. Trajectory Prediction Model Implementation using Keras	59
B. Trajectory Prediction Approach Parameters	60
C. Source Code	61
C.1. Odroid	61
C.1.1. C++ Source Code	61
C.1.2. Python Source Code	76
C.2. STM32	98
C.2.1. C Source Code	98

1. Introduction

The development of fully self-driving cars is currently one of the most exciting and demanding challenges in modern vehicle industry. Traveling on the road should become safer and more comfortable in the absence of the human driver. However, teaching a computer how to navigate the vehicle in complex areas such as cities appears to be a difficult problem. Therefore, it makes sense to tackle autonomous racing first since one can focus on fewer and more simplified aspects.

The goal of this thesis is to implement an autonomous racing system in scale 1:8 which is able to compete at the Deep Berlin Robocars Challenge on 12 September 2018 [14]. During the competition, each team let their car drive 10 rounds around the track. Only the fastest lap time counted and was compared across the teams. There weren't any obstacles or other racing cars on the track during driving. When all four tires left the track which was delimited by white lane markings, a two seconds penalty was added to the lap. Obviously, this gave some tolerance which could be exploited in terms of driving behavior. Over the last couple of months before the challenge, the teams were able to visit the final racing track once a week to collect data and tweak their algorithms.

Our car has been originally designed for the semester project "Hochautomatisiertes Fahren" [4] to develop an autonomous platooning system. However, that system was neither designed for racing nor for driving on the final track at all. In terms of sensoring, the car was equipped with a single monocular camera, a wheel encoder on each rear tire as well as an ultrasonic sensor which was not actually used for racing.

Based on the camera's current frame, our autonomous racing algorithm uses a neural network that predicts a trajectory which the car should follow. Using these trajectories, we can employ a driving model to control the steering angle and speed. This method differs from the traditional deep learning end-to-end approach which uses a neural network to predict the steering angle from the image directly. Our neural network tries to solve the problem of *where* the car should drive to instead of *how* it should drive. Therefore, we gain more control over the system. For instance, our car is able to recognize straight lines and turns to accelerate and brake respectively. With this approach we were able to win the third place out of 12 teams in the Deep Berlin Robocars Challenge.

1.1. Related Work

Autonomous racing is quite a young sports and research topic. Popular international competitions such as *Roborace* [32] [9] and *Self Racing Cars* [33] [26] exist only since 2016. The participating teams build full-sized autonomous vehicles which drive on a real racing track. In [13], de la Iglesia Valls et al. introduce the design of a driverless racing car which won the *Formula Student Driverless* competition in 2017. The vehicle

was equipped with a variety of sensors including LIDAR, IMU, GPS and a stereo camera to perform *simultaneous localization and mapping* (SLAM). They employed the LIDAR to detect cones which were delimiting the track. Therefore, they didn't rely on visual information making it inapplicable to our use case where there only exists lane markings to define the track. Unfortunately, besides a few occasional blog posts, there are hardly any further scientific publications about other competitions yet.

Actually, most scientific work in the field of autonomous racing is primarily done either on miniature cars or in simulations. In [25], fairly complex control techniques are presented to show how to drive a 1:43 scale car at its physical limits. Rosolia et al. employ *model predictive control* to reach the same goal but in a simulated environment [30]. However, both approaches assume that the vehicle is able to localize itself on the track and do not focus on the visual aspects.

The rise of deep learning and convolutional neural networks revolutionized computer vision [15] and is excessively driving the development of autonomous vehicles (not specifically racing). In [7], *Nvidia* introduced a neural network architecture called *PilotNet* which acts as an end-to-end lane keeping assistant. The network takes an image as input and outputs a steering angle which is directly controlling the car. Ever since, several other advanced architectures have evolved using recurrent neural networks and residual connections for example [36].

There also exist some research on how to apply reinforcement learning to autonomous driving in simulations [31]. However, this method hasn't been successfully applied to the real world yet.

Very exiting work in autonomous racing including deep learning for visual perception has been done in [30] which uses a convolutional neural network to create a cost function for a model predictive control algorithm. It also employs a core idea of our trajectory prediction approach that is to use deep learning as a vision system and to split perception and control at the same time.

1.2. Structure of the Thesis

This thesis is structured as follows. Section 2 provides an overview on the car's hardware, chassis and sensors as well as the system's software architecture. In Section 3 we present an introduction to deep learning. Section 4 is devoted to the implementation of the autonomous racing algorithm. In Section 5 we describe our results and experiments with the system. Finally, Section 6 concludes the thesis by summarizing what has been done.

2. System Architecture

In this Section we give an overview of our racing car by presenting its chassis, sensors and hardware components. Furthermore, we provide the software architecture which outlines how to get from sensor data through steering commands to the final actuations on the mechanics.

2.1. Chassis

Since we adopted the car used in the semester project “Hochautomatisiertes Fahren”, there haven’t been many decisions regarding the chassis and mechanics of the system. The Losi Horizon Hobby 8IGHT-E 1/8-scale 4WD RTR Buggy serves as the car’s foundation. It includes a differential gear connecting the powerful 2500 kV motor to all four tires. The wheels are covered with rubber textures providing plenty of grip on the racing track. Since the track’s material was rubber too, there haven’t been any issues regarding slippage even under competitive speeds. Furthermore, a suspension system is build into the car, yielding a stable driving behavior in tight turns.



Figure 1: Photo of the racing car.

The acceleration of our motor turned out to be quite advantageous. In fact it’s specified to drive up to 50 mph. Although the top speed is far from reachable on such a small and curvy racing track, we definitely tried to exploit the car’s capabilities.

A servo is used to turn the front wheels, providing a maximum steering angle of

30° to each side. However, our initial servo appeared to be seriously slow. It took approximately half a second to turn the wheels from full left to full right and vice versa. This turned out to be a crucial problem and led to certain design decisions in the racing algorithm which are covered in more detail later on. Most of the competitors used the Donkey Car [3] which is roughly half the size of our model. They come with smaller wheels and weigh less so their servos act almost instantly. Fortunately, by replacing the servos, we could improve the steering delay a bit.

In terms of powering the car, we make use of two 7.4 V lipo batteries, each having a capacity of 4500 mA h. One battery is powering the motor exclusively whereas the other one powers the rest of the hardware.

Altogether, one can say that our car provided the most potential in terms of motor power and stability. However, compared to the majority of the competition, the vehicle was rather big and heavy, resulting in less tolerance of staying in the track during the turns.

2.2. Hardware

To compute the high level logic which outputs the steering commands, an Odroid XU4 is used running Ubuntu 16.04. The Odroid is preferred over the popular Raspberry Pi due to its superior processor. Its recommended operation voltage is 5 V and under full load its power consumption rises up to 5 A. Since our batteries output 7.4 V we installed a DC to DC step down converter. However, the board seems to be really power demanding, so we had to increase the output voltage of the converter to 5.7 V. Anything below that would lead the system to reboot when running the racing algorithm.




Figure 2: Photo of the car's hardware.

The STM32 microcontroller is responsible for running the low level realtime logic. This

includes handling the servo's and motor's PWM, getting data from the wheel encoders as well as implementing a PID controller which controls the desired speed. A UART connection is employed to communicate with the Odroid using a Mavlink protocol. To be able to connect each component with the microcontroller, the STM32 is plugged into a breakout board made by *Assystem GmbH*.

2.3. Camera

Our car is equipped with a DFM 22BUC03-ML camera [2] from *The Imaging Source*. It is capable of producing frames at a resolution of 744×480 pixel at 76 frames per second which is a much higher resolution than actually required in our case since we cropped and rescaled the image as described in more detail later on.

An important aspect proved to be the camera's mounting position. Initially, it was mounted at the front of the car, a couple centimeters above the wheels. We ended up positioning the camera at the rear, roughly 40 cm above the ground. Consequently, it gained a much better overview of the track making the predictions of the neural network more reliable.

2.4. Wheel Encoders

To be able to measure the car's motion regarding its speed and position, there are two wheel encoders, one on each rear tire. These wheel encoders consist of two components:

- A dice containing equally spaced wholes which is connected to its corresponding shaft.
- An optical sensor which emits an electrical impuls whenever a whole of the dice is passing through it. These impulses are referred to as *ticks* and trigger an interrupt in software which increments a simple counter variable.

In Section 4.2.2 we explain how this data is used to compute a path, showing where the vehicle drives along in metric coordinates which is called *odometry*.

2.5. Software Architecture

In this Section we explain how the aforementioned parts connect through software.

Our system is composed of three software components which are shown in Figure 4. Since we use a deep learning approach in our racing algorithm, our system has to perform two use cases:

1. Recording training data while being controlled remotely.
2. Driving autonomously.




Figure 3: Photo of the rear left wheel encoder.

Therefore, we consider the car being either in *recording mode* or in *autonomous mode*. In recording mode, a human is able to drive the car remotely. During this mode, camera images and data from the wheel encoders are written to an external storage device. A game controller is plugged into a Laptop which runs a program sending UDP packets to the odroid containing the steering commands. This way, we are able to use a wired controller which we did initially. One may also choose to connect the sensor of a wireless controller directly to the odroid. However, the range of those devices may be quite weak, so one should definitely follow the car while driving remotely on larger tracks.

2.5.1. Odroid

To create the odroid's software component, our design is based on the *robot operating system* (ROS) [29] which is a popular collection of libraries and tools to build robot applications. One of its main concepts constitutes the messaging system. In ROS, messages are transported between *nodes* either over *topics* or *services*. Nodes are processes running in the OS just like any other application. This way, developers are able to write nodes in different programming languages. However, this implies that some overhead is involved in sharing data between nodes due to inter-process communication. Fortunately, ROS provides the *nodelet plugin* for C++. Instead of creating a process for each individual C++ node, only a single process is launched which is called the *nodelet manager*. The nodelet manager handles the execution and message transportation of each node it is responsible for in its own process.

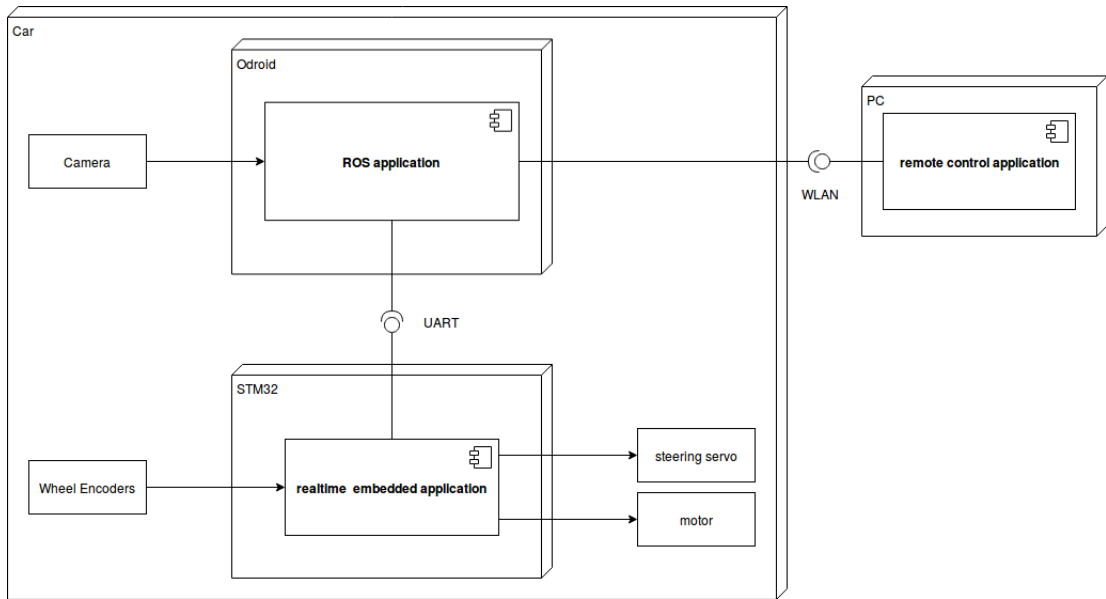


Figure 4: Component diagram of the system.

Messages and Topics

A node can publish messages to topics which other nodes can subscribe to. To avoid confusion, we point out that for each message in our particular system, there exists an equally named topic which the message is published to. For example, if there's a message *SetAngle*, then this message is transmitted over the topic */SetAngle*. These are the messages used in the system:

- *RemoteAngle* contains the steering angle transmitted over the network.
- *RemoteThrottle* contains the throttle value transmitted over the network.
- *RemoteState* encodes button presses so we can use the game controller to stop and continue recording.
- *SetAngle* contains the steering angle transmitted to the STM32 which should be applied to the servo.
- *SetThrottle* is interpreted as the speed the car should currently maintain and is transmitted to the STM32.
- *WheelTicks* provides an update from the STM32 containing the ticks for the left and right wheel encoders.

Note that there's no message containing camera information. We choose to read the camera images directly in the node where they are used. Thus, we avoid the overhead of serializing and deserializing the image data for transmission over a specific topic.

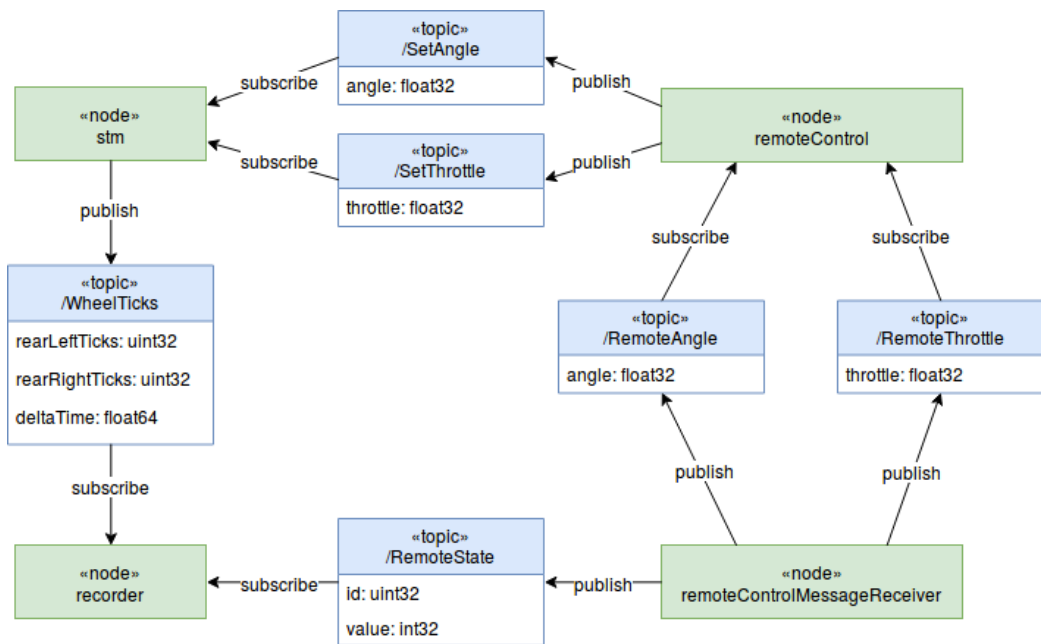
Nodes

In the following, we provide an overview of the nodes employed in the system.

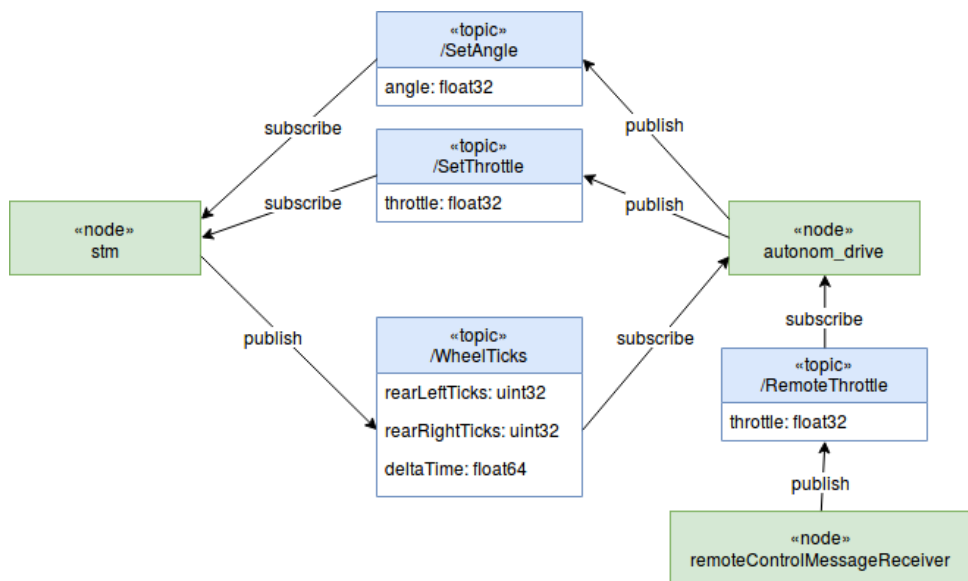
- *remoteControlMessageReceiver* receives UDP packets sent from the PC and publishes the contained information to the `/RemoteAngle`, `/RemoteThrottle` and `/RemoteState` topics respectively.
- *remoteControl* subscribes to `/RemoteAngle` and `/RemoteThrottle` and forwards that information by publishing it to the `/SetAngle` and `/SetThrottle` topics.
- *stm* subscribes to `/SetAngle` and `/SetThrottle` and transmits the data to the STM32 by handling the Mavlink communication over the UART. It also publishes the incoming wheel encoder updates to `/WheelTicks`.
- *recorder* is responsible for storing the training data which incorporates reading images from the camera as well as subscribing the `/WheelTicks` topic.
- *autonom_drive* runs the racing algorithm which publishes to `/SetAngle` and `/SetThrottle`. It subscribes the `/WheelTicks` and the `/RemoteThrottle` topic in order to still be able to get control over the vehicle's speed despite being in autonomous mode.

The ROS architectures for recording and autonomous modes are depicted in Figure 5(a) and Figure 5(b) respectively ¹.

¹Note that we don't provide any further lower level UML diagrams since we don't want to focus on the specific implementation details in this work.



(a) Active nodes and topics in recording mode.



(b) Active nodes and topics in autonomous mode.

Figure 5: ROS architectures.

2.5.2. STM32

Since we adopted the car from the semester project and there was already everything implemented that we needed, we left most parts of the STM32 unchanged. Its embedded architecture is based on a scheduler which periodically runs a set of tasks. A task is nothing more than a piece of C code. Global variables are used to share data between different tasks. In the following, we provide a list of the most important ones ²:

- *mavlink* is responsible to handle the Mavlink protocol using the UART.
- *vehspdctrl* implements a PID controller to control the target speed and updates the motor's PWM module accordingly.
- *stangproc* updates the servo's PWM module based on the target steering angle.
- *eict* configures and reads input capture timers to receive the ticks from the wheel encoders.

²Note that the poor naming stems from the fact that most of the STM software was left unchanged from the semester project.

3. Deep Learning

This Section provides an introduction into deep learning which is the core technology used in this work to solve the autonomous racing problem.

3.1. Motivation

Machine learning has gained rapid success in recent years. The era of deep learning began with *AlexNet* [21], a Deep Convolutional Neural Network which considerably outperformed its competition in the standard image classification challenge ImageNet. In the following six years, this technology has been applied to a wide variety of other problems such as natural language processing [12], handwriting recognition [40] and cancer detection [5]. However, these networks already existed in the 1980s [16] and were only able to succeed due to huge advances in high throughput computing.

History showed that problems which humans perform well at are typically hard to solve algorithmically by writing traditional, imperative programs. The classification of images is known for being such a phenomenon. Obviously, humans are able to instantly recognize objects like trees, cats and cars on an image. However there doesn't seem to exist any explicit rules how we actually perform this task. Therefore, new systems have been developed that are able to learn how to solve these problem by themselves, given some example data.

In machine learning, the most well researched learning paradigm is known as *supervised learning*. On an abstract level, a supervised learning model learns a function by fitting it with examples of input-output mappings. Taking the image classification as an example, the model gets a vector of pixels as input and outputs a binary vector specifying which classes it detected. The model is trained with large sets of image-to-class-label pairs. If trained properly, it is able to learn the specific features of each class in order to *generalize* to new data making it perform well on images it has never seen before.

The most sophisticated deep learning models for supervised learning are *artificial neural networks* (ANNs).

3.2. Artificial Neural Networks

Despite their name, modern ANNs don't really aim to model their biological counterpart which they were inspired from [15]. An ANN can be described as a graph consisting of nodes which are connected through edges. A node gets one or multiple signals from its input edges, computes a weighted sum of its inputs, applies a threshold function and outputs a new signal through its output edges. ANNs are structured in consecutive *layers*, consisting of one input layer followed by at least one hidden layer and finally the output layer. This structure is also known as the *multilayer perceptron* [11]. An example ANN is depicted in Figure 6. Each layer computes a function which can be




Figure 6: Visual representation of a basic multilayer perceptron consisting of a single hidden layer. Each node passes its output as weighted input to every node in the next layer.

represented as

$$f(x) = \sigma(Wx + b) \tag{1}$$

where

- x is the input vector of dimension m ,
- W is the $n \times m$ *weight matrix* where n is the number of nodes in the layer,
- b is the *bias* vector of dimension n and
- σ is the *activation function*.

The output of f is a n -dimensional vector which expresses the *activation* of each node in the layer. Each activation is given as input to every node in the next layer.

As one can see, the input edges are weighted. Therefore, an edge can be of high significance if its weight is large enough. On the other hand, setting it to zero prevents information to pass through the next layer. The original intuition behind the activation function was to define for which summed input the neuron actually *fires*, which means that its output is greater than zero. For example, using a neuron as a binary classifier, one may take the binary activation function

$$\sigma_i(x) = \begin{cases} 1 & \text{if } x_i > 0 \\ 0 & \text{otherwise} \end{cases} \quad \forall i \in \{1, \dots, n\} \tag{2}$$

so the i -th neuron becomes active as long as the term $(Wx + b)_i$ is greater than zero³. The bias can be viewed as a threshold of the weighted sum of inputs Wx which has to be reached so that the activation function lets the neuron fire.

³In fact, activation functions act always elementwise with the same operation for all i so one often omits the indexed notation.


Another purpose of the activation function is to introduce nonlinearity into the system. Let $\sigma(x) = x$ be the identity function. Now, consider a three layer model

$$\begin{aligned} y &= W_1x + b_1 \\ z &= W_2y + b_2 \end{aligned} \tag{3}$$


where x denotes the input layer, y the activations of the hidden layer and z the activations of the output layer. Thus, the input and hidden layer are connected through weights W_1 and the hidden and output layer are connected through weights W_2 . We can observe that the third layer is actually redundant since two composed affine transformations can be expressed as a single affine transformation:

$$W_2(W_1x + b_1) + b_2 = W_2W_1x + (W_2b_1 + b_2). \tag{4}$$

In practice, we want our models to approximate highly non-linear functions and therefore, a non-linear activation function should be applied to our weighted sum. A very common choice is the ReLU function [34] depicted in Figure 7(b). There are a couple of desired properties the activation function may fulfill however, those can only be motivated by understanding how an ANN learns.



(a) The binary function.



(b) The ReLU function: $\sigma(x) = \max(0, x)$.

Figure 7: Different activation functions.

3.3. Learning

Solving a supervised learning problem involves the following steps [22].

1. Collect example data and partition it into training, validation and test data. There is no general way to measure the quality of the data set so this always depends on the specific problem and has to be figured out by exploration.
2. Create (or refine) a model.
3. *Training step*. Optimize the model's parameters to perform well on the training data.
4. *Validation step*. Check how the model behaves on the validation data. If the model doesn't perform well, go back to step 2. Otherwise, continue with step 5.
5. *Testing step*. Once the model works reasonably on both training and validation data, we run it on the final test data right before deployment to see whether it is able to perform well in real life applications. It is crucial that this data is kept isolated during the previous steps since any further adaption to the model prevents us from objectively telling how the model generalizes to data it has never seen before.

When building neural network models, one has to deal with a wide variety of parameters. For structure, we partition them into two types:

- a) *Learnable parameters*. Parameters which are set during the training step. In neural networks, these are the weights and biases.
- b) *Non-learnable parameters* or *hyperparameters*. One may group these ones even further into:
 - *Topological parameters*. Regard the number and types of layers, the number of nodes on each layer as well as the choice of activation functions, to name a few.
 - *Training parameters*. Correspond to specific parameters of the training algorithm, the loss function and the method of weight initialization for example.

The difference between these types is that non-learnable parameters are selected by the human developer (through exploration and repetition of step 2) whereas learnable parameters are discovered by a training algorithm. For that, we need a metric which represents how well a model predicts on a given input and a *ground truth* output (or simply *label*). In literature, this metric is known as the *loss function*.

3.3.1. Loss Function

One may think of training or learning as an optimization problem. Given n samples of training data as inputs X_i and their corresponding labels Y_i for $1 \leq i \leq n$ as well as a model $\Phi(x, \Theta)$ where x is the model's input and Θ are the learnable parameters, our goal is to find Θ_{opt} to minimize the error \mathcal{L} between the model's predicted output and the labels

$$\Theta_{opt} = \arg \min_{\Theta} \mathcal{L}_{\Phi, X, Y}(\Theta) \quad (5)$$

with

$$\mathcal{L}_{\Phi, X, Y}(\Theta) = \frac{1}{n} \sum_{i=1}^n \mathcal{E}(\Phi(X_i, \Theta), Y_i). \quad (6)$$

Here, $\mathcal{E}(\cdot, \cdot)$ measures the error between the model's prediction $\Phi(X_i, \Theta)$ and a label Y_i so we call it the *output error measure*. \mathcal{L} is the average error over n training examples X and Y and is referred to as the *loss function* or simply the *loss* [22].

To provide a more concrete understanding, we show how this could be applied to autonomous driving using basic ANN models. A simple approach would be to create an artificial neural network model which takes the current camera image as input and directly outputs the target steering angle. Since there's no other processing involved between the sensory data and the steering commands, this is known as the *end-to-end* approach [7]. Data is collected by a human driving the car while recording the current steering angle for each new frame. The objective is that the network should mimic the human behavior. Therefore, this technique is referred to as *behavioral cloning*.

First, we have to define the neural network model. Its input layer takes a vector representation of an image which is done by stacking all the pixel values. The output layer is simply a single node which gives the steering angle. Determining the network architecture and the non-learnable parameters is done by many iterations of the training and validation steps.

In order to train the model, we have to define an appropriate output error measure. Optimally, our network always predicts the exact same steering angle that the human driver had entered based on the collected data. Given a model's predicted steering angle y and a ground truth steering angle \hat{y} , the error measure could be as simple as the squared distance

$$\mathcal{E}(y, \hat{y}) = (y - \hat{y})^2. \quad (7)$$

By squaring the difference between y and \hat{y} , we remove the sign and increase the penalty on larger errors. A widely used generalization to n -dimensional outputs is

known as the *mean squared error* (MSE) defined as

$$MSE = \mathcal{E}(y, \hat{y}) = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2. \quad (8)$$

There exist several other output error measures which are mainly divided into those that are suitable for regression problems (e.g. MSE or *mean absolute error* (MAE)) and the ones that are intended to be used for classification problems (e.g. *cross entropy* or *hinge*) [22].

Since we want to minimize the model's loss, we have to determine the appropriate learnable parameters Θ_{opt} as depicted in Equation 5. However, neural networks are extremely complex functions due to their possibly huge numbers of layers, neurons, weights and other parameters. Therefore, it's not feasible to get Θ_{opt} by trying to find an analytical solution. Thus, we need an optimization algorithm to compute the learnable parameters which minimize the loss function.

3.3.2. Optimization

To train our neural network, the first-order iterative optimization algorithm *gradient descent* [42] is used to find the appropriate learnable parameters Θ_{opt} . The *gradient* of a multi-variable function $f(x_1, \dots, x_n)$ is defined as

$$grad(f) = \nabla f = \begin{pmatrix} \frac{\partial f}{\partial x_1} \\ \vdots \\ \frac{\partial f}{\partial x_n} \end{pmatrix} \quad (9)$$

where $\frac{\partial f}{\partial x_i}$ is the partial derivative of f with respect to x_i . The important property is that for any given point x , the gradient $\nabla f(x)$ points to the direction of the greatest rate of increase of f . Likewise, $-\nabla f(x)$ points to the direction of the greatest rate of decrease of f .

Gradient descent uses this observation to find a local minimum. Starting at some initial point x_0 , we iteratively perform the following updates

$$x_{n+1} = x_n - \gamma \nabla f(x_n) \quad (10)$$

until the sequence (x_n) converges at a local minimum. However, to fulfill the desired condition $f(x_0) \geq f(x_1) \geq f(x_2) \geq \dots$ the update rate $\gamma \in R_+$ has to be chosen small enough.

In the following, we want to show how gradient descent is used to minimize the loss function $\mathcal{L}(\Theta)$. In practice, Θ is going to be high dimensional. For simplicity, let's assume a model consisting of only two learnable parameters $\Theta = (\theta_0, \theta_1)$ which we further assume to be both weights. Now, we can visualize the loss function as a




Figure 8: Visualization of a fictional error surface. The model contains only two weights as learnable parameters. Given some initial values for weight 1 and 2 (red cross), gradient descent is used to walk down to the nearest valley representing a local minimum in loss.

landscape where each configuration of θ_0 and θ_1 results in a position with height $\mathcal{L}(\Theta)$ as depicted in Figure 8. By initializing these weights randomly, we could potentially land on a peak of the error surface (red cross in Figure 8). To get the minimal local error, we use gradient descent to iteratively update the weights to walk down to a surrounding valley.

It should be noted that this method does not guarantee to find the *global* minimum. Hence, different weight initializations may lead to different results.

What's left is to determine the partial derivatives of \mathcal{L} with respect to each θ_i to compute the gradient $\nabla\mathcal{L}$. We can then simply update the learnable parameters using Equation 11

$$\Theta_{n+1} = \Theta_n - \gamma \nabla \mathcal{L}_{\Phi, X, Y}(\Theta_n) \quad (11)$$

until (Θ_n) hopefully converges at a local minimum. γ is referred to as the *learning rate* which is an important hyperparameter. The partial derivatives with respect to each learnable parameter can be computed using the backpropagation algorithm [41] (see Section 3.3.3).

When computing the loss, it matters to choose an appropriate quantity of training samples, given as X and Y . Let's assume to select the entire training set to calculate

the loss over. Consequently, we have to evaluate the model's prediction for each sample in the possibly huge training set in order to perform just a single update to the weights and biases. However, in practice we need to perform multiple updates since we do not know how far to walk into the direction of the gradient to arrive at the local minimum (see Figure 8 again). Since computing the loss function over the entire training set is usually very expensive, instead we can randomly subsample the data into smaller chunks called *mini-batches*. These mini-batches serve as an approximation of the full training set when used to compute the model's loss. The corresponding method is then known as *mini-batch stochastic gradient descent*.

To summarize the optimization procedure, Figure 9 gives high-level pseudocode of a mini-batch stochastic gradient descent algorithm used for training a neural network. The three most important steps are as follows:

1. *Forward pass* (line 4). We compute the model's predictions by processing the input X^{batch} . Then we can calculate the error between the network's output and the labels Y^{batch} . During the forward pass, we keep certain data in memory that is needed in order to perform the backward pass. For illustration, this data is returned as the function's result.
2. *Backward pass* (line 5). Based on the result of the forward pass, this function computes the gradient using backpropagation.
3. *Update* (lines 6 and 7). Finally updates the weights and biases according to Equation 11.

Algorithmus 1 : Mini-batch stochastic gradient descent algorithm

Input : $num_samples$ training inputs X .
Input : $num_samples$ training labels Y .
Input : $batch_size$.
Input : Neural network model Φ with weights \mathcal{W} and biases \mathcal{B} .
Input : Output error measure \mathcal{E} .
Input : Learning rate γ .
Output : Optimized weights \mathcal{W} and biases \mathcal{B} .

```
1  $num\_batches \leftarrow \lfloor num\_samples / batch\_size \rfloor$ 
2 for  $i \leftarrow 1, \dots, num\_batches$  do
    // Retrieve batches  $X^{batch}$  and  $Y^{batch}$  by randomly sampling
    //  $batch\_size$  samples from  $X$  and  $Y$  respectively.
3  $X^{batch}, Y^{batch} \leftarrow sample\_batch(X, Y, batch\_size)$ 
    // Forward pass: Compute the model's predictions and the
    // resulting error across the current batch.
4  $result = forward\_pass(\Phi, \mathcal{W}, \mathcal{B}, \mathcal{E}, X^{batch}, Y^{batch})$ 
    // Backward pass: Use the result from the forward pass to compute
    // the gradient vectors for  $\mathcal{W}$  and  $\mathcal{B}$ .
5  $\Delta\mathcal{W}, \Delta\mathcal{B} \leftarrow backward\_pass(result)$ 
    // Update: Use gradient descent to update the weights and biases.
6  $\mathcal{W} \leftarrow \mathcal{W} - \gamma\Delta\mathcal{W}$ 
7  $\mathcal{B} \leftarrow \mathcal{B} - \gamma\Delta\mathcal{B}$ 
    // Remove the samples used in the current batch.
8  $X \leftarrow X - X^{batch}$ 
9  $Y \leftarrow Y - Y^{batch}$ 
10 end
11 return  $(\mathcal{W}, \mathcal{B})$ 
```

Figure 9: High-level pseudocode showing the mini-batch stochastic gradient descent algorithm.

3.3.3. Backpropagation

In order to optimize the learnable parameters in a neural network, we still need a way to compute the gradient of the loss function. In other words, we have to get the partial derivatives of \mathcal{L} with respect to each learnable parameter in $\Theta = (\theta_1, \dots, \theta_m)$

$$\nabla \mathcal{L}_{\Phi, X, Y}(\Theta) = \begin{pmatrix} \frac{\partial \mathcal{L}_{\Phi, X, Y}}{\partial \theta_1} \\ \vdots \\ \frac{\partial \mathcal{L}_{\Phi, X, Y}}{\partial \theta_m} \end{pmatrix}. \quad (12)$$

The *backpropagation* algorithm allows us to evaluate these derivatives efficiently on arbitrarily complex neural networks.

First, we need to transform our model into a *computational graph*. Computational graphs represent functions where nodes are either input parameters or operations. Edges indicate the data flow. An example is depicted in Figure 10(a) which computes the function $f(x, w, b) = (wx + b)^2$. We could imagine that f represents a node in a neural network with a single input x and its corresponding weight w , bias b as well as the activation $\sigma(x) = x^2$. Note that we label each intermediate result of the operations, that is

$$p = wx \qquad q = p + b \qquad r = q^2.$$

Given some concrete input values to f , backpropagation computes the partial derivatives of the final output r with respect to each input parameter w , x and b which are $\frac{\partial r}{\partial w}$, $\frac{\partial r}{\partial x}$ and $\frac{\partial r}{\partial b}$. To refer to the neural network analogy again, the concrete inputs would be training input data x and randomly initialized weight w and bias b . In reality, one would actually need a fourth parameter y representing the label and some extra nodes in the graph in order to compute the error between the network's output and the label. However, we left that out for simplicity.

The algorithm consists of two steps, a *forward pass* (see Figure 10(b)) and a *backward pass* (see Figure 10(c)). During the forward pass, we evaluate each node's operation by forwarding the results through the graph in topological order. By doing this, we are able to compute the *local gradients* which are the partial derivatives of each node's output with respect to its own inputs. In our example, we start by calculating the local gradient for node p . It's inputs are $w = 2$ and $x = 4$ so we want to get $\frac{\partial p}{\partial w}$ and $\frac{\partial p}{\partial x}$

$$\frac{\partial p}{\partial w} = \frac{\partial wx}{\partial w} = x = 4 \qquad \frac{\partial p}{\partial x} = \frac{\partial wx}{\partial x} = w = 2.$$

Next, we propagate p 's output to the input of q . The partial derivatives of node q with

respect to its inputs $p = 8$ and $b = 3$ are given as

$$\frac{\partial q}{\partial p} = \frac{\partial(p+b)}{\partial p} = 1 \qquad \frac{\partial q}{\partial b} = \frac{\partial(p+b)}{\partial b} = 1.$$

Finally, we pass $q = 11$ to r and get $r = 121$ as well as

$$\frac{\partial r}{\partial q} = \frac{\partial q^2}{\partial q} = 2q = 22.$$

Using the results from the forward pass, the backward pass traverses the graph in reverse topological order and consecutively computes the partial derivatives of the graph's output with respect to each node's input. In other words, it measures the influence each node has on the final output (including the input nodes). Therefore, backpropagation heavily exploits the chain rule

$$\frac{\partial z}{\partial x} = \frac{\partial z}{\partial y} \frac{\partial y}{\partial x}. \tag{13}$$

For example, if we want to compute $\frac{\partial r}{\partial p}$, we can also write

$$\frac{\partial r}{\partial p} = \frac{\partial r}{\partial q} \frac{\partial q}{\partial p} = 22 \cdot 1 = 22.$$

Note that in this equation, $\frac{\partial q}{\partial p}$ is part of the local gradient of q which we already evaluated in the forward pass. On the other side, the first factor $\frac{\partial r}{\partial q}$ is the *backpropagated result* from q 's output node which is node r . Having determined $\frac{\partial r}{\partial p}$, we backpropagate this result to node p to compute $\frac{\partial r}{\partial w}$ and $\frac{\partial r}{\partial x}$


$$\frac{\partial r}{\partial w} = \frac{\partial r}{\partial p} \frac{\partial p}{\partial w} = 22 \cdot 4 = 88 \qquad \frac{\partial r}{\partial x} = \frac{\partial r}{\partial p} \frac{\partial p}{\partial x} = 22 \cdot 2 = 44.$$

Again, the second factor stems from the local gradient of p . By getting $\frac{\partial r}{\partial b}$ with


$$\frac{\partial r}{\partial b} = \frac{\partial r}{\partial q} \frac{\partial q}{\partial b} = 22 \cdot 1 = 22$$

the backpropagation algorithm finishes and outputs the gradient vector


$$\nabla f(2, 4, 3) = \begin{pmatrix} 88 \\ 44 \\ 22 \end{pmatrix}.$$



(a) Computational graph representing the function $f(w, x, b) = (wx + b)^2$.



(b) Forward pass. We compute each node's output (green numbers) as well as their local gradients in topological order using a single pass. (blue numbers).



(c) Backward pass. By using the chain rule, we evaluate the partial derivative of the output r with respect to each node's input (red numbers) in reversed topological order.

Figure 10: Application of the Backpropagation algorithm on a computational graph.

Although the provided example is quite small, the described method remains the same even for very large computational graphs which represent complex neural network architectures with millions of weights. The algorithm is extremely efficient since it goes over each node exactly twice and performs rather simple calculations on them.

We finally presented all the theoretical tools required to train neural networks. However, we don't want our models to only perform well on the training data. Referring to the autonomous racing problem, we could build an end-to-end neural network as described in Section 3.3.1, train it with lots of images and steering angles and very likely, it still won't be able to drive as expected in the real world. In order to get desired results on the validation and test process, some more concepts are involved.

3.3.4. Overfitting

It often occurs that the model performs very well during training but miserably during the validation and testing steps. We say that the model *overfits* to the training data. Overfitting appears when the training loss decreases while the validation loss increases. *L2 regularization*, *dropout* and *data augmentation* are common strategies to mitigate this problem.

Regularization

The idea behind regularization stems from the observation that overly complex models don't generalize well [18] [23]. Therefore, we can add another term \mathcal{R} to the loss function to penalize the model's complexity

$$\mathcal{L}_{\Phi, X, Y}(\Theta) = \frac{1}{n} \sum_{i=1}^n \mathcal{E}(\Phi(X_i, \Theta), Y_i) + \mathcal{R}(\Theta). \quad (14)$$

L2 regularization is a commonly used technique to penalize large weights which is defined as

$$\mathcal{R}(\Theta) = \lambda \sum_{i=1}^m \Theta_i^2. \quad (15)$$

Here, λ is another hyperparameter which represents the regularization strength. Intuitively, this technique could result in weights which are set so close to zero that they cancel out. As a consequence, large parts of the network may not be really in use anymore which virtually ⁴ shrinks the size of the model and potentially prevents overfitting [23].

⁴We say "virtually" since the actual computational costs are still the same as without regularization.

Dropout

Dropout is another regularization technique which randomly sets the activation of some neurons to zero (see Figure 11). The probability p of keeping neurons is another hyperparameter. We only drop activations during training. In order to maintain the same expected value across both phases, we multiply the output of each neuron by p during testing.

The intuition behind dropout is that it forces the neural network to have distributed representations since it can not rely on any single neuron to be reliably active for the representation of a state [24].




Figure 11: Visualization of dropout.

Data Augmentation

To further improve the model's generalization ability, we may expand the training set by applying some transformations. For example, we can mirror or rotate images or change the brightness and contrast. Thus we incorporate variations into the dataset which potentially makes predictions more invariant to different lighting conditions or perspectives not covered by the training data.

3.4. Convolutional Neural Networks

Neural networks are widely applied in image processing. However, traditional networks (as presented in Section 3.2) do not perform well on images. A 640×480 pixel image in RGB-colorspace results in an input layer with about 9.2×10^5 nodes, thus already 9.2×10^5 weights per node in the second layer. The second problem is that the network learns position-dependent features. Suppose we're training a neural network to recognize cats. On our training images, cats only appear in the upper left corner, so only neurons connected to this part get to know the specific features. Given an image where the cat sits in the bottom right corner, our network wouldn't be able to reliably detect it.

Convolutional layers are designed to alleviate these downsides. In a convolutional layer, each neuron has its own small region of the image it is looking at which is called the *receptive field*. Hence, neurons are not connected to every single pixel anymore. The




Figure 12: A convolutional layer, taking a $5 \times 5 \times 3$ image (left) as input. The third dimension represents the number of color channels. This convolutional layer has five 3×3 filters, resulting in five separate layers of neurons (right). Note that filters always operate on the full depth of the input.

receptive field is a rectangular filter containing weights which learn to recognize a specific feature of the input. These weights are shared across each neuron and therefore, they're able to recognize features position-independently. One filter learns just one feature but a convolutional layer has multiple filters which results in multiple layers of neurons. Thus, convolutional layers can be thought of as three-dimensional volumes, visualized in Figure 12.

Filters can be represented as blocks, sliding across the image to compute the activations of its corresponding neurons. Thereby, they take the dot product of the filter values w and the input volume of their current receptive field x and add a bias b to it. Suppose a convolutional layer whose input layer has a depth of 3, the filter size is 5×5 and the number of filters is 10. Consequently, each filter has $5 \times 5 \times 3 + 1 = 76$ parameters where the $+1$ represents the bias. Since there are 10 filters, the layer has exactly 760 learnable parameters, regardless of the input's width and height.

While sliding the filter across the input, we may also choose an interval to skip certain pixels. This interval is referred to as the *stride*. For example, a stride of 3×2 means each time we move the filter horizontally, we skip two pixels, and each time we move the filter vertically, we skip one pixel. In addition, a *pad* can be used to add additional zero valued pixel to the border of the input in order to preserve the spatial size.

Powerful modern convolutional neural networks typically consists of lots of stacked convolutional layers. The idea is that the features in each layer become increasingly more abstract as we go deeper into the network. Filters in the first convolutional layer typically detect small characteristics such as edges. Then, the second layer may combine the features from the first layer into curves or other shapes. Eventually in some deeper layer, the composition of higher level features starts to represent meaningful objects.

4. Autonomous Driving

In this Section we present two implementations of autonomous racing using deep learning: the established end-to-end approach and our new trajectory prediction approach.

4.1. End-to-End Approach

In the autonomous driving terms, end-to-end means that a neural network model directly outputs steering commands based on sensory inputs such as camera images. *Nvidia* developed a neural network called *PilotNet* [7] that performs lane keeping on real cars. PilotNet receives a single image of a front-mounted camera as input and outputs a steering angle. The model is trained based on collected data from human driver sessions. Since the network seemed to be very promising, our idea was to adapt it for our application and to see whether it can be trained to drive the car fast, reliably and competitive on the racing track.

4.1.1. Preprocessing

Before the image is passed to the network, we involve some preprocessing. The original image resolution produced by the camera is 744×480 . We select the region of interest by cropping the upper 744×200 pixels since these do not reveal any useful information about the track. Then, we resize the cropped image to have a resolution of 200×100 . Note that our network implementation differs in this aspect from PilotNet which resizes its image to 200×66 pixels. However, the lane markings are quite thin so by further reducing the resolution, some markings in curves would simply disappear.

As proposed by *Nvidia*, we also convert the RGB image into the YUV color space. Since the environment is mostly black and white, we would expect the RGB channels to be the same anyway. The Y channel may yield more variation across the channels, since it represents luminance and not color value.

We apply batch normalization to the input layer which makes each individual pixel value's distribution zero mean and unit variance. Input normalization is a commonly used technique as it helps the gradient descent algorithm to improve convergence time. Ioffe and Szegedy provide an in-depth discussion about the underlying theory which goes beyond the scope of this work [17].

During training, each sample has a 60% chance to be augmented. When being augmented, we apply a random change of brightness to the image. Furthermore, there's another 50% chance to flip the image horizontally. Of course, in this case we also change the sign of the steering angle. The idea behind this augmentation is that the model may learn how to drive in both directions based on the same image.

4.1.2. Neural Network Model

Figure 13 shows our model’s architecture which is an adaption of PilotNet. There are five convolutional layers which are intended to be feature extractors and their parameters were chosen empirically by *Nvidia*.

	Output Shape	Learnable Parameters
Steering Angle Throttle	2	22
Dense 10	10	510
Dense 50	50	5050
Dense 100	100	576100
Convolution 64 filters 3x3 kernel 1x1 stride	5x18x64	36928
Convolution 64 filters 3x3 kernel 1x1 stride	7x20x64	27712
Convolution 48 filters 5x5 kernel 2x2 stride	9x22x48	43248
Convolution 32 filters 5x5 kernel 2x2 stride	22x47x36	21636
Convolution 24 filters 5x5 kernel 2x2 stride	48x98x24	1824
BatchNormalization	100x200x3	12
Input	100x200x3	0

Figure 13: End-to-end model architecture based on PilotNet.

The first three convolutional layers feature a kernel size of 5×5 and a 2×2 stride whereas the two last convolutional layers only have a 3×3 kernel and a 1×1 stride. This is due to the fact that the output surface becomes increasingly smaller throughout the network. Furthermore, the filter size increases from 24 in the first layer up to 64 in the fourth and fifth layer which compensates the decrease in output volume size and adds more higher level features.

The three-dimensional output volume of the last convolutional layer is flattened into a one-dimensional vector of activations. This means that there are $5 \times 18 \times 64 = 5760$ neurons which are connected to the following 100 neurons of the first dense layer ⁵. There follow another two dense layers of 50 and 10 neurons respectively.

The final output layer consists of two neurons which represent the steering angle and the throttle. PilotNet only outputs the steering angle so this is another modification

⁵A dense (or fully connected) layer is just another name for a layer of stacked neurons as presented in Section 3.2.

as we wanted our model not only to be a lane keeping system but also to control the speed.

We use ReLU activation functions throughout the network except for the two output nodes which have linear activation functions. Unfortunately, *Nvidia* doesn't state explicitly which activation functions they use in PilotNet. However, it's quite likely that ReLU functions are employed due to their enormous popularity [39]. For computing the loss, we take the mean squared error measure and trained the network using the ADAM optimization algorithm [20]. ADAM is similar to the gradient descent algorithm but adds some improvements to get to a local minimum more quickly. Our learning rate is set to $\gamma = 5 \times 10^{-4}$ which was discovered by experimenting.

We wrote the model in Python using the high-level neural network library *Keras* which uses *TensorFlow* for the backend. This makes it easy to implement the network by simply using the Keras layers API. The full code for implementing the end-to-end model is given in Appendix A.1.

4.1.3. Training

We split the recorded data into a training set (80% of all samples) and a validation set (20% of all samples). An extra test set is not required since testing happens directly on the racing track. The two sets are randomly sampled into batches of size 40. Furthermore, the model is trained in epochs. In each epoch, the model processes the entire training set.

Training was very slow on the laptop we used initially which was a *Lenovo* Thinkpad L450 with an *Intel* Core i5 5200U and 8 GB memory. It took hours to train the neural network reasonably well on our datasets which made the process of model tuning very time consuming. Therefore, we wanted to train our model on a powerful GPU which speeds up learning due to its parallelization capabilities. The new training system contains an *Intel* Core i7 2600K with 4 cores, 16 GB memory and an *Nvidia* GTX 1080. However, to take advantage of the GPU, we had to adapt and optimize the data loading and preprocessing procedure.

To provide the training algorithm with data, we implemented batch generators which are callbacks to generate the next training or validation samples. Whenever the batch generator was called, it was reading 40 images from the hard drive into memory which was then sent over to the GPU. Since reading from hard drive is slow, most of the time wasn't spent on training but on loading data. The solution was to simply cache as many images as possible into the 16 GB memory. As a result, we could reduce the training time by a factor of 12 in our particular case.

Of course, when increasing the size of the dataset, we would also have to increase the memory size on the machine in order to maintain the same level of speedups. However, since our largest dataset only contained about 40000 images, this wasn't a big issue.

4.1.4. Dataset

To make the system move autonomously, we had to collect data by driving the car ourselves on the racing track. Therefore, we created two datasets. In the first dataset, we were trying to drive the car as close as possible on the track's center line which is depicted in Figure 14(a). The second one was produced by driving the car on an ideal line which is illustrated in Figure 14(b). By ideal line, we refer to the strategy of keeping the wheels mostly straight in order to accelerate as long as possible. The idea is to exploit our excessive motor power to outperform the competition.




Figure 14: Illustration of the driving strategies performed on the racing track to record two datasets.

Both datasets contained about 10000 images of driving. For the center line session, we were driving 12 laps with relatively slow speed. During the ideal line session, we recorded 28 laps but at the maximum speed which we could still drive without leaving the track. It should be noted here, that even though we recorded more laps in the second run, the datasets contain roughly the same number of images. This is due to the fact that we were simply driving faster while recording the ideal line.

We like to mention here that we actually recorded more data than 20000 images. In addition, we built our own test tracks in our office and performed several tests on them but we think that the two described datasets are the most representative and decision leading ones. Furthermore, since we changed the camera mounting position a couple of weeks before the final racing, we couldn't use the data we had previously recorded anymore. Nevertheless, this was the correct decision as Figure 15 shows the difference between the two setups. The car gained a much cleaner, wider and undistorted overview of the track. In fact, this becomes even more crucial for the trajectory prediction approach.



(a) Low-mounted camera position.

(b) High-mounted camera position.

Figure 15: Different camera positions.

4.1.5. Test Results

We trained our neural network model with the center line and ideal line datasets over 50 epochs using the same parameters.

Loss Analysis

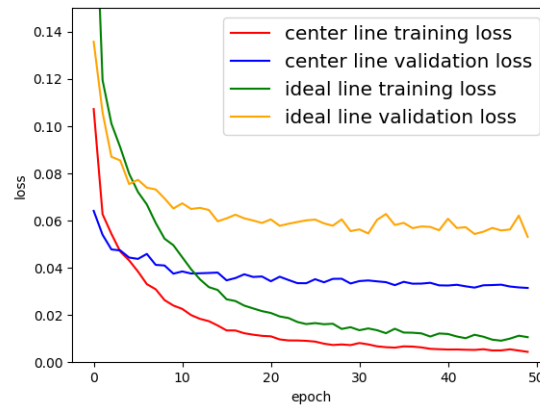


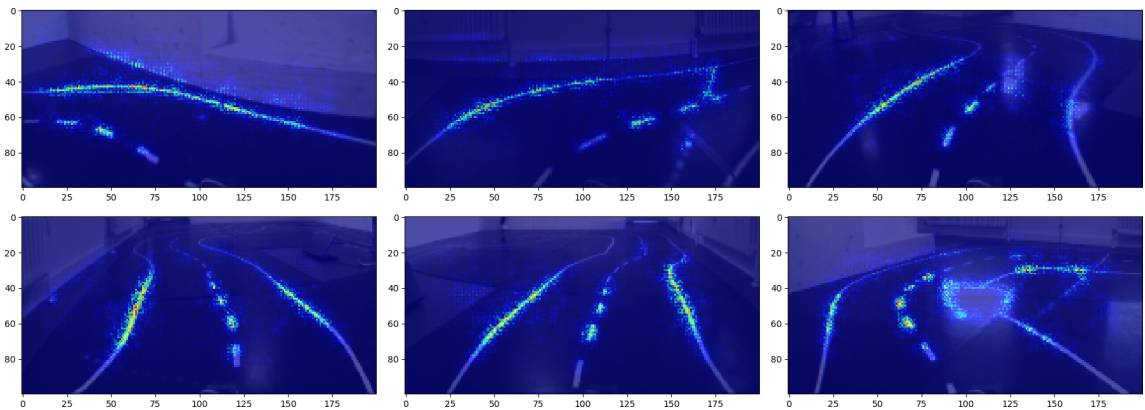
Figure 16: Training and validation loss of the model when trained with the center line and ideal line datasets over 50 epochs.

Figure 16 shows the loss diagram during training. We can observe that during both runs, the model didn't overfit to the training data since the validation loss does not tend to increase. The model seems to have more problems learning the ideal line dataset since the validation loss of the center line dataset is about 33% lower. Since we recorded 28 laps in the ideal line set, it contains much more variation compared to the 12 laps of the center line run. Furthermore, in the center line dataset, the steering angle heavily corresponds to the curvyness of the center lane markings directly in front


of the car. This is not the case when driving the ideal line since there are cases when the car drives straight even though the track makes a slight curve.

Visualization

To have some clue about which features of the image the neural network takes into account, Figure 17 illustrates *saliency maps* [35] based on some samples of the track. Saliency maps visualize the gradients from the network’s output with respect to the input pixels. The intuition is that if the gradient of a certain pixel input has a high value, it has much influence on the output and is important to the network. In Figure 17, we computed the gradients from the output of the fifth convolutional layer using *guided backpropagation* [37] and layered them over the original image.



(a) Saliency maps of the model trained with the center line dataset.



(b) Saliency maps of the model trained with the ideal line dataset.

Figure 17: Saliency maps.

On the center line saliency maps in Figure 17(a), we can clearly see that the model primarily looks at the lane markings and even ignores the reflections on the image at the top right. However, on the ideal line saliency maps in Figure 17(b), the markings are not nearly as much highlighted. It also seems to be the case that the model is triggered by noise and reflections seen at the top middle and top right images.

Driving Behavior

To test our car, we were driving it in the same environmental conditions that prevailed during the recordings. Then we performed another test in the evening to see how it reacts to different lighting conditions.

In the first test, the car was indeed behaving quite similar to the way we were driving it manually. It could both stay on the center line and also adapt the driving behavior of the ideal line. We don't actually have a robustness metric but the car would surely perform valid rounds in most of the ten available attempts under the given conditions.

However, this appeared to be quite different during the second test. Our car had severe problems to reliably stay on the track. The drastic change in behavior was quite unexpected since we augmented the dataset with different brightness settings which should actually lead to a better generalization ability. Maybe this problem can be fixed by extending the datasets with recordings in various other lighting conditions.

A much bigger problem turned out to be that we were not able to scale up the speed. By doing so with any dataset, the car was always leaving the track very soon. We believe that the main cause of this problem was produced by the slow steering of the servo. While the human driver adapts to this problem by intuitively steering earlier, the neural network is unable to take these dynamics into account. Replacing the servo improved the situation a little bit but we still didn't nearly reach superhuman level of racing.

We discovered another unpleasant property of the end-to-end approach: the car's ability to learn a certain driving behavior depended heavily on the way we performed steering inputs. Therefore, we created two recordings while driving the car on the center line. During the first run, we maxed out the steering input for a split of a second whenever we had to perform any course corrections. This means that there exists basically only three steering inputs: full left, full right and neutral. During the second run, we pushed very gently on the controls to perform smooth transitions to the input. While the car was driving well when fed with the data from the second recording, by training it with the first recording, it was barely able to drive a single lap.

This result might be inferred as follows. Imagine that we are creating a dataset by trying to keep the car as closely as possible on the center line but our inputs are only full left, full right and neutral. Next, consider driving a wide left curve. In order to keep the car on the center line, it's not possible to consistently maintain full left as input since the car would sheer off too far. Therefore, one has to quickly alternate between full left and neutral at a certain frequency to follow the line safely. Let's consider a sequence of consecutive images in this dataset. These images look quite similar to each other due to high recording frame rates. However, the sequence possesses a distribution of extremely dissimilar steering angles. If similar training inputs have strongly dissimilar labels, the network may either learn to approximate the mean of the label distribution or overfits by learning the subtle and irrelevant differences in these inputs. This is due to the nature of the learning process being just an optimization algorithm which

tries to minimize the error between the network's output and the labels. In both cases, the car performs extremely poorly as the neural network is unable to reproduce the actual driving behavior. This emphasizes the fact that the success of the employed deep learning techniques depend heavily on the choice of training data.

Our conclusion is that the end-to-end approach might seem to be an elegant way to implement a system. However, it has some serious downsides when being applied to autonomous racing. Training the car to obtain a desired driving behavior requires profound manual driving skills. Therefore, the vehicle will not perform much better than the human steering the car during the recordings.

4.2. Trajectory Prediction Approach

Based on the experiences we gained using the end-to-end approach, we asked ourselves two questions:

- Can we design a system that uses deep learning to process images into steering informations but being easy to train at the same time?
- Can we scale the speed such that the car is able to outperform any manually driven lap times?

In the following, we present the new design we came up with by considering these two requirements.

4.2.1. Overview

Our concept exploits the information retrieved by the wheel encoders to construct an odometry which gives a path of positions we were driving along. Based on some position p_0 in the odometry we can easily compute the position the car will be located at in a given distance. This allows us to sample the coordinates (p_1, p_2, p_3) which represent the car's future locations at 0.6, 1.2 and 1.8 meters ahead of p_0 . We then transform these coordinates into a local coordinate system where we assume that the car is currently located at $(0, 0)$ and pointing at direction $(0, 1)$. After applying the transformation, the new points (p'_1, p'_2, p'_3) represent a *local trajectory* which the car follows along. Figure 18 further illustrates this coordinate transformation which will be explained in more detail later on.

We modified the output layer of our end-to-end model to predict these trajectory coordinates instead of the steering angle and throttle values. Thus, we refer to this approach as *trajectory prediction*. For each image in our dataset, we compute the corresponding local trajectories the car was driving during the recording sessions. These image-to-trajectory mappings become the new training data which the model is fed with. Finally, during autonomous mode, we compute the steering angle and target speed based on the predicted trajectories using an appropriate driving model.

The core idea behind this approach is that we no longer train the neural network on *how* to drive the car but instead it only learns *where* it should be located at in the future. This has several advantages. First, the trajectories the network is trained with are computed independently of time. This means, during recording, we can drive the car around the track as slowly as required. Therefore, the car's behavior doesn't depend on manual driving skills anymore.


Secondly, we gain control over the driving dynamics. Assume the neural network would always predict the desired trajectories. This allows us to compute the speed and steering angle based on the car's physical limits. Obviously, a complex physical driving model would be required which was not possible to implement during the given time scope. However, we did make use of the fact that we're able to predict straight

lines and curves to adapt the speed accordingly.

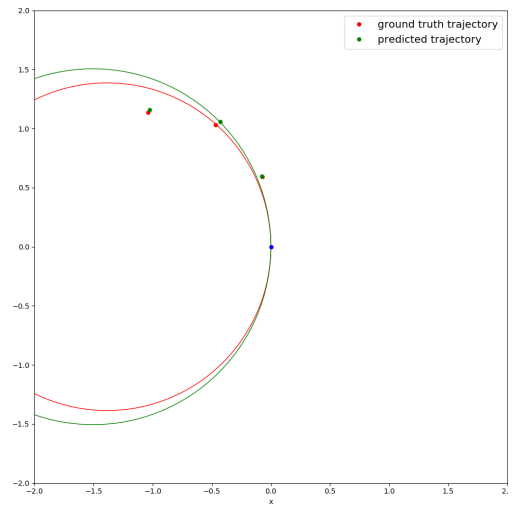
In the following sections we give an in-depth explanation about the implementation of our trajectory prediction approach.



(a) Current camera image of the car, driving along our test track.



(b) Recorded odometry data. The blue dot shows where the car is currently located. The red dots reveal the future locations in 0.6, 1.2 and 1.8 meters.



(c) The trajectory (red dots) computed by projecting the odometry coordinates into a local coordinate system where the center is the car's current position. The green dots represent the network's predicted trajectory given the current camera image as input. The circles give the turning radius which is then processed into a steering angle.

Figure 18: Illustration of the coordinate projection and the trajectory prediction.

4.2.2. Odometry

An odometry is a sequence of positions which are computed by estimating the car's movement based on the wheel encoder ticks. It's a relative positioning model which means that position p_t is evaluated based on p_{t-1} and the estimated change of position obtained by the current sensor output. Such models are sensitive to errors since any slight deviation from the actual position is accumulated over the sequence.

Positioning Model

Given the vehicle's current position $p_{t-1} = (x_{t-1}, y_{t-1})$ and orientation ω_{t-1} , we compute its new position p_t and orientation ω_t using the following *dead-reckoning* equations [19]. Initially, we choose $p_0 = (0, 0)$ and $\omega_0 = 0$.

Let d_l and d_r be the diameters of the rear left and right wheels respectively and let η be the number of encoder ticks per wheel revolution. The travelled distance per tick on each side is given by

$$c_l = \frac{\pi d_l}{\eta} \qquad c_r = \frac{\pi d_r}{\eta}. \quad (16)$$

Suppose that N_l and N_r are the number of ticks counted by the wheel encoders during the current time interval. Then, the travelled distance for each wheels is

$$\Delta s_l = N_l c_l \qquad \Delta s_r = N_r c_r. \quad (17)$$

We take the average to model the actual travelled distance of the car (or the imaginary center wheel):

$$\Delta s = \frac{\Delta s_l + \Delta s_r}{2}. \quad (18)$$

The change in orientation is approximated by

$$\Delta \omega = \frac{\Delta s_r - \Delta s_l}{T} \quad (19)$$

where T is the distance between the points where the rear wheels are touching the floor. We refer to this distance as the *track width*. Now, we can get the new orientation

$$\omega_t = \omega_{t-1} + \Delta \omega. \quad (20)$$

Finally, we update the position using simple trigonometric equations

$$\begin{aligned} x_t &= x_{t-1} + \cos(\omega_t) \Delta s \\ y_t &= y_{t-1} + \sin(\omega_t) \Delta s. \end{aligned} \quad (21)$$

Calibration

As mentioned already, odometry is very prone to errors. These errors can be classified into *systematic* and *non-systematic* errors [19]. Non-systematic errors are caused by the fact that the vehicle mostly doesn't drive in a perfect environment. For example, the wheels may be affected by slippage resulting in the registration of wheel revolutions that didn't actually move the car forward. Even small irregularities of the floor such as bumps lead to errors since Equation 17 doesn't model horizontal displacements.

Systematic errors are caused by measurement inaccuracies of the wheel diameters and the track width. For instance, by manufacturing rubber wheels, it's hard to avoid at least some slight deviations in size. Uncertainty is involved in measuring the track width since there doesn't exist a true single point where the wheels are in contact with the floor. In addition, by driving in a curve, the car might tilt to the side which shifts these points slightly along the wheels.

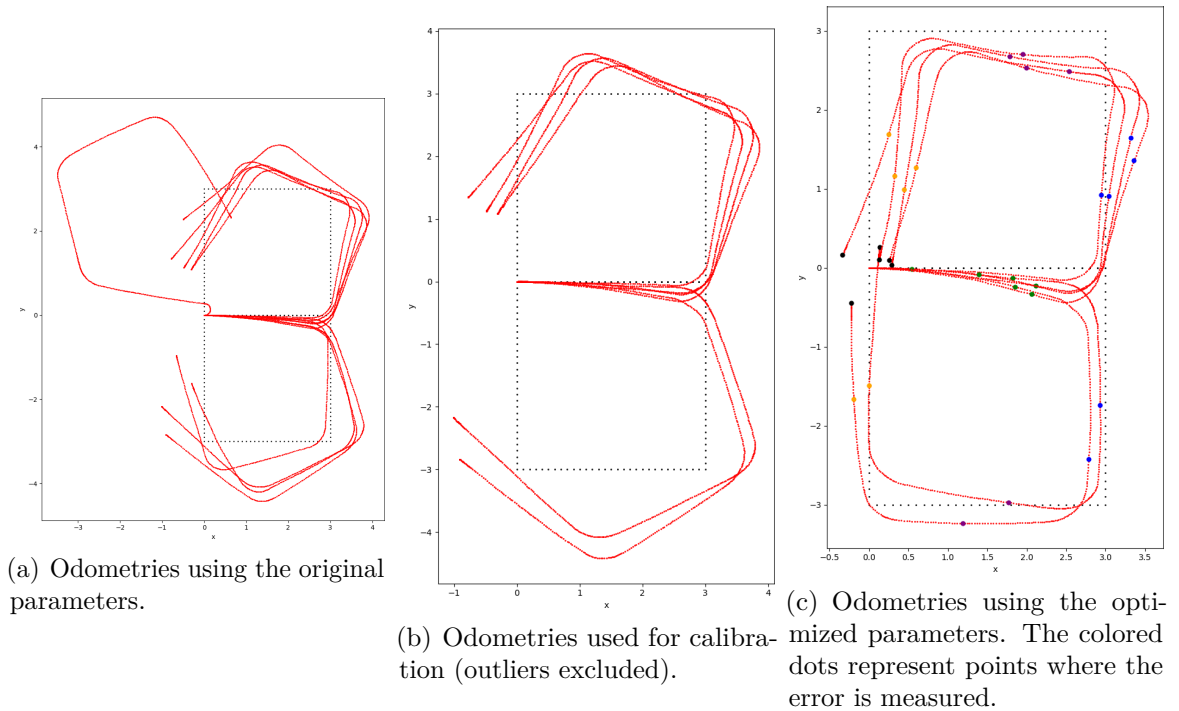


Figure 19: Reconstructed odometries using the ticks obtained by driving around a 3×3 meter square (black dots) in clockwise and counter-clockwise direction respectively.

Fortunately, compared to the inevitableness of the non-systematic error, we can minimize the systematic error by performing a calibration procedure. Therefore, the objective is to minimize the error between a computed odometry (using the equations above) and a corresponding ground-truth odometry by optimizing the parameters d_l , d_r and b .

In our case, we were driving the car along a 3×3 meter square by recording the

accumulated wheel encoder ticks N_l and N_r in a 20 millisecond time interval. We created five runs in clockwise and counter-clockwise direction respectively. The initial parameters were measured using a ruler which gave us the following dimensions:

$$\eta = 120 \qquad d_l = 11cm \qquad d_r = 11cm \qquad T = 27.5cm.$$

The resulting odometry is shown in Figure 19(a). Using these parameters, the car's position seems to drift heavily apart from the square. In addition, we observed that even though we were keeping the vehicle precisely on the line, there exist some substantial outliers which couldn't possibly be caused either by inaccurate parameterization or by incorrect driving. Therefore, during optimization, we only took those runs into account which seemed to be the most reasonable ones (see Figure 19(b)).

We measure the error between a reconstructed odometry and the ground truth odometry by meaning the squared distances between the final position at t_{final} and between positions at manually selected indices where we know the car was driving straight. At those selected indices, we only take either the distance in x- or y-coordinate into account, depending on whether the car was moving horizontally or vertically in the particular section.

For example, the blue dots in Figure 19(c) represent indices τ_i where we know for sure that the vehicle was heading vertically. Therefore, the x-coordinates of the reconstructed odometries at positions τ_i should be about 3.

In order to define the error function for the optimization problem, we need some formalism. An odometry can be described as a function $o(p_0, \Theta, N, t)$ [8] where

- p_0 is the start position (in our case $p_0 = (0, 0)$),
- $\Theta = (c_l, c_r, T)$ is the parameter vector,
- N represents the wheel encoder ticks and
- t is the current time or index.

Let N_i^{cw} and N_i^{ccw} be the number of ticks of the i -th run along the 3×3 meter square in clockwise (cw) and counter-clockwise (ccw) direction respectively. Furthermore, let $\tau_{j,i}^{cw}$ and $\tau_{j,i}^{ccw}$ be the manually selected indices for each individual run where

- $\tau_{1,i}^{cw}$ and $\tau_{1,i}^{ccw}$ is an index in the first horizontal section of run i ,
- $\tau_{2,i}^{cw}$ and $\tau_{2,i}^{ccw}$ is an index in the first vertical section of run i ,
- $\tau_{3,i}^{cw}$ and $\tau_{3,i}^{ccw}$ is an index in the second horizontal section of run i ,
- $\tau_{4,i}^{cw}$ and $\tau_{4,i}^{ccw}$ is an index in the second vertical section of run i and
- $\tau_{5,i}^{cw}$ and $\tau_{5,i}^{ccw}$ is the final index of the odometry of run i .

Given ticks data N with n runs in clockwise and m runs in counter-clockwise direction

and indices τ , our error function E is defined as:

$$\begin{aligned}
E_{p_0, N, \tau}(\Theta) = \frac{1}{2} \left(\frac{1}{5n} \sum_{i=1}^n \left[\right. \right. & o(p_0, \Theta, N_i^{cw}, \tau_{1,i}^{cw})_y^2 \\
& + (3 - o(p_0, \Theta, N_i^{cw}, \tau_{2,i}^{cw})_x)^2 \\
& + (-3 - o(p_0, \Theta, N_i^{cw}, \tau_{3,i}^{cw})_y)^2 \\
& + o(p_0, \Theta, N_i^{cw}, \tau_{4,i}^{cw})_x^2 \\
& \left. \left. + \|o(p_0, \Theta, N_i^{cw}, \tau_{5,i}^{cw})\|^2 \right] \right. \\
& + \frac{1}{5m} \sum_{i=1}^m \left[\right. & o(p_0, \Theta, N_i^{ccw}, \tau_{1,i}^{ccw})_y^2 \\
& + (3 - o(p_0, \Theta, N_i^{ccw}, \tau_{2,i}^{ccw})_x)^2 \\
& + (3 - o(p_0, \Theta, N_i^{ccw}, \tau_{3,i}^{ccw})_y)^2 \\
& + o(p_0, \Theta, N_i^{ccw}, \tau_{4,i}^{ccw})_x^2 \\
& \left. \left. + \|o(p_0, \Theta, N_i^{ccw}, \tau_{5,i}^{ccw})\|^2 \right] \right) \tag{22}
\end{aligned}$$

where $\|\cdot\|$ is the euclidean distance.

We approximate the gradient ∇E using finite differences and apply the *Broyden-Fletcher-Goldfarb-Shanno* algorithm (BFGS) [28] to compute the optimal parameters Θ . The resulted odometries are shown in Figure 19(c).

We can see that the final positions and the overall shape of the odometries heavily improved. However, given the fact that the runs were performed on a reasonably small 3×3 meter square, the error is still very large.

We tried out various other optimization techniques, including the popular *UMBmark* [19] method as well as other error functions [8] and even tried to manually calibrate the parameters. Nevertheless, the presented method provided the odometries which seemed to be closest to the actual square.

Trajectory Extraction

Let $f(p_t, s, t)$ be a function computing the position on the linearly interpolated path p_t which is reached by following this path for distance s while starting at index t . Given odometry (p_t, ω_t) , to compute the trajectory $Q' = (q'_1, \dots, q'_n)$ at index τ , we sample positions $Q = (q_1, \dots, q_n)$ at distances $\mathcal{D} = (D_1, \dots, D_n)$ apart from p_τ by computing

$$q_i = f(p_t, D_i, \tau) \tag{23}$$

To get Q' , we apply the transformation

$$q'_i = \mathcal{R} \left(\frac{\pi}{2} - \omega_\tau \right) (q_i - p_\tau) \tag{24}$$

where \mathcal{R} is the rotation matrix

$$\mathcal{R}(\alpha) = \begin{pmatrix} \cos \alpha & -\sin \alpha \\ \sin \alpha & \cos \alpha \end{pmatrix}.$$

In our experiments, choosing $n = 3$ positions with distances $d_1 = 0.6$ m, $d_2 = 1.2$ m and $d_3 = 1.8$ m away from the car’s current position p_τ turned out to work well. Larger distances may lead to bad model predictions since when entering a sharp turn, the camera cannot see much of the track. In this situation, the neural network may not be able to use the visual information of the track to predict a position which is far in the future and might learn different features instead during training.

We think that choosing more predicted positions to increase the resolution of the trajectory is unnecessary as three or even only two positions are enough to infer the steering angle in our case.

4.2.3. Neural Network Model

The only aspect in the neural network architecture that differs from the one presented in Section 4.1.2 is the output layer. We replaced the nodes for the steering angle and throttle with two nodes for each trajectory position which output their x- and y-coordinates respectively. Just as before, MSE is used as the error measure. The Keras implementation of the network is shown in Appendix A.2.

There are only slight changes in the preprocessing and training procedures. For example, during augmentation, instead of flipping the steering angle’s sign, we now flip the sign of the x-coordinate of the trajectory positions.

4.2.4. Driving Model

In the following, we describe how to get from a predicted trajectory to a final steering angle and throttle command.

Controlling the Steering Angle

We model the car’s motion using a bicycle model as shown in Figure 20. Assuming no slippage, the car follows along a circle with radius r when setting steering angle α . Given α the circle’s radius is modeled by r_α and given r the required steering angle is represented by α_r using the two equations

$$r_\alpha = \frac{b}{\tan(\alpha)} \qquad \alpha_r = \arctan\left(\frac{b}{r}\right) \qquad (25)$$

where b is the distance between the front and the rear axle which we refer to as the *wheel base*.




Figure 20: Illustration of the bicycle model.

Let $Q = (q_1, \dots, q_n)$ be a trajectory, our goal is to find steering angle α resulting in a turning circle of radius $|r_\alpha|$ and center $(r_\alpha, 0)$ which fits the trajectory best (see Figure 18(c) again). We can formulate this objective as another optimization problem where we want to find the scalar α which minimizes the error function

$$E_Q(\alpha) = \begin{cases} \frac{1}{n} \sum_{i=1}^n (q_{i,x})^2 & \text{if } \alpha = 0 \\ \frac{1}{n} \sum_{i=1}^n (\|q_i - (r_\alpha, 0)\| - |r_\alpha|)^2 & \text{otherwise} \end{cases} \quad (26)$$

given trajectory Q and the constraint $\alpha_{min} \leq \alpha \leq \alpha_{max}$. This formula computes the mean squared distance between point q_i in the trajectory and the point lying on the border of the circle constructed by r which is closest to q_i . Note that the special treatment of $\alpha = 0$ is required due to the fact that r_0 is not defined. In this case, we just take the distance between q_i and the x-axis.

Instead of applying a numerical approximation method to find the optimal α_{opt} , it turned out to be most efficient to simply compute

$$\alpha_{opt} = \arg \min_{\alpha} \{E_Q(\alpha) \mid \alpha \in \mathbb{Z}, \alpha_{min} \leq \alpha \leq \alpha_{max}\}. \quad (27)$$

So we just iterate over the integer set of steering angles⁶ and keep α where $E_Q(\alpha)$ is lowest. The loss in precision is negligible due to the low accuracy of the car's servo which exhibits a tolerance of about $\pm 2^\circ$.

It should be mentioned that we only considered the trajectory points $Q = (q_1, q_2)$ to fit the turning circle so we left out q_3 which is 1.8m apart from our current position. By leaving it in, the car starts to turn too early in front of curves and leaves the track as a consequence. In addition, small course corrections in q_1 and q_2 would be very likely to average out. Therefore, choosing Q carefully is crucial to fine tune the driving behavior.

⁶In our case, the maximum steering angle turning the wheels to the left is $\alpha_{min} = -30^\circ$ whereas $\alpha_{max} = 30^\circ$ turns the wheels most to the right.

Controlling the Speed

A desired driving behavior is to accelerate on straight lines and maintain a lower speed in curves to provide stability and avoid slippage. Due to the lack of time, we employed a very simple model to implement this behavior. Given trajectory $Q = (q_1, q_2, q_3)$, we just use the x-coordinate of q_3 as a curve detector and compute the target speed $w(Q)$ (in meter per second) with the formula

$$w(Q) = \begin{cases} \nu_1 & \text{if } |q_{3,x}| \leq \chi \\ \nu_2 & \text{otherwise} \end{cases} \quad (28)$$

where

- ν_1 is the target speed when driving straight,
- ν_2 is the target speed when detecting a curve and
- χ is the threshold in the x-coordinate of q_3 which decides when to use ν_1 or ν_2 .

Thus, there are only two possible target speeds ν_1 and ν_2 . The choice of χ determines how early the car should start to slow down to ν_2 in front of a curve. Obviously, these three constants have to be discovered by experimenting which is a downside of this approach.

After determining the target speed, we have to set the motor torque accordingly to maintain it. In each time step t we compute the error between the target speed w_t and the measured speed v_t

$$e_t = w_t - v_t \quad (29)$$

where v_t is computed by

$$v_t = \frac{\Delta s}{\Delta t} \quad (30)$$

and where Δs is taken from Equation 18. Our objective is to apply corrections u_t to the motor torque value $z_t \in [-1, 1]$ which minimizes the error e_t over time. Therefore, we utilize the popular *proportional-integral-derivative controller* (PID-controller) [6] [44] which calculates an appropriate correction with

$$u_t = K_p e_t + K_i \sum_{t'=0}^t (e_{t'} \Delta t) + K_d \frac{e_t - e_{t-1}}{\Delta t} \quad (31)$$

where $K_p, K_i, K_d \in \mathbb{R}_{\geq 0}$ are the coefficients for the proportional, integral and derivative terms respectively. We ended up choosing $K_p = 5 \cdot 10^{-3}$, $K_i = 0$ and $K_d = 1 \cdot 10^{-3}$ after tuning these coefficients manually. Using this configuration, we compromised a bit of error overshooting for a more rapid acceleration and deceleration.

5. Results

In this Section we present the results of our trajectory prediction approach. We do an offline analysis again and then show how it performed during the final race. In addition, we point out some issues of the system and propose possible improvements.

5.1. Wheel Encoder Anomaly

During our tests using the trajectory prediction approach, we noticed that each time the car performed a tight right turn, it was strongly accelerating. However, when driving left turns or straight, the controller was keeping the target speed just fine. It turned out that there exists an anomaly in the measurement of the wheel encoder ticks. We did three experiments where we were driving the car manually for a fixed distance s and recorded the wheel encoder ticks for each run. In the first experiment, we were driving a circle in clockwise direction while keeping steering angle $\alpha = 30^\circ$. We measured the radius r of that circle and determined s as the circumference $s = 2\pi r$. During the second experiment, the exact same circle was followed but in counter-clockwise direction with steering angle $\alpha = -30^\circ$. Finally, in the third experiment, we were keeping the vehicle on a straight line for distance s .

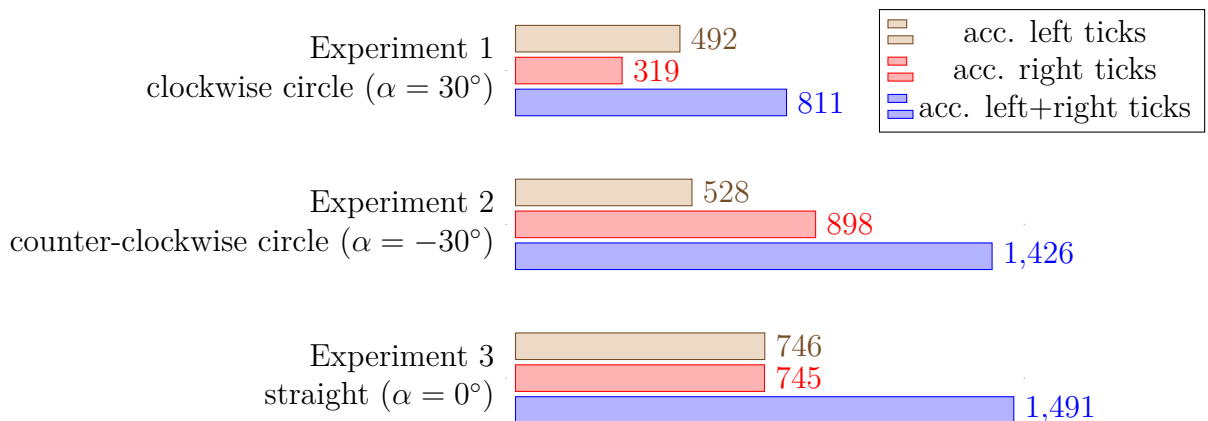


Figure 21: Results of the three experiments driving the car manually for a fixed distance. The diagrams show that the accumulated number of emitted encoder ticks are about 45% less when driving a clockwise circle as compared to when driving a counter-clockwise circle or a straight line with the same travelled distance.

Since the car covered the same distance across these runs, we would expect the accumulated number of recorded wheel ticks to be about the same too. For each experiment, Figure 21 shows the accumulated number of ticks recorded by the left and right wheel encoder. To make sure that our measurements are reliable, we repeated each experiment three times and averaged the results. We can see that our expectations

hold for the second and third experiments. However, when driving the clockwise circle, the total number of accumulated ticks is about 45% less compared to the other runs.

These results explain the misbehavior of the speed controller. When driving a clockwise turn, the number of emitted ticks reduces and consequently the measured distance over time decreases too even though the car maintains the same actual speed. Therefore, the PID controller thinks that the car slows down and misleadingly increases the motor torque which causes the aforementioned acceleration. In addition, the experiments may also explain the bad results of the odometry calibration. Unfortunately, we were not able to fix this problem since there wasn't enough time. In order to mitigate this problem, we decided to only drive counter-clockwise runs during the final race and during the recordings for our dataset.

5.2. Dataset

Our final dataset includes about 17000 samples of recordings while driving the car on the ideal line in the noon and in the evening. However, when only following the ideal line, the neural network doesn't learn how to behave in different situations such as when leaving the track accidentally. Therefore, we added the functionality to pause and resume the recording session using the buttons on our game controller. This allows us to hit pause and drive the car into a certain position that hasn't been covered by the dataset yet. Then we resume the recording session and lead the car back to the ideal line. We repeated this procedure to teach the car how to behave for various deviations from the ideal line.

Since we had no indication of whether we were currently pausing or resuming, after the final race it unfortunately turned out that we didn't include another 7000 samples by forgetting to hit the resume button.

5.3. Final Race

We can list the important parameters for the trajectory approach as

$$\Omega = (c_l, c_r, T, \mathcal{D}, b, \nu_1, \nu_2, \chi, K_p, K_i, K_d)$$

which were all introduced previously. The concrete values used in the final race are given in Appendix B.

We trained the neural network over 50 epochs with learning rate $\gamma = 1 \times 10^{-4}$. The training and validation loss were both converging so there didn't seem to be an indication for overfitting.

Figure 22 shows saliency maps using the trained trajectory prediction model. Although there seems to be much interest on the lanes (especially on the center line), especially the walls and reflections became salient objects which is not really desired.

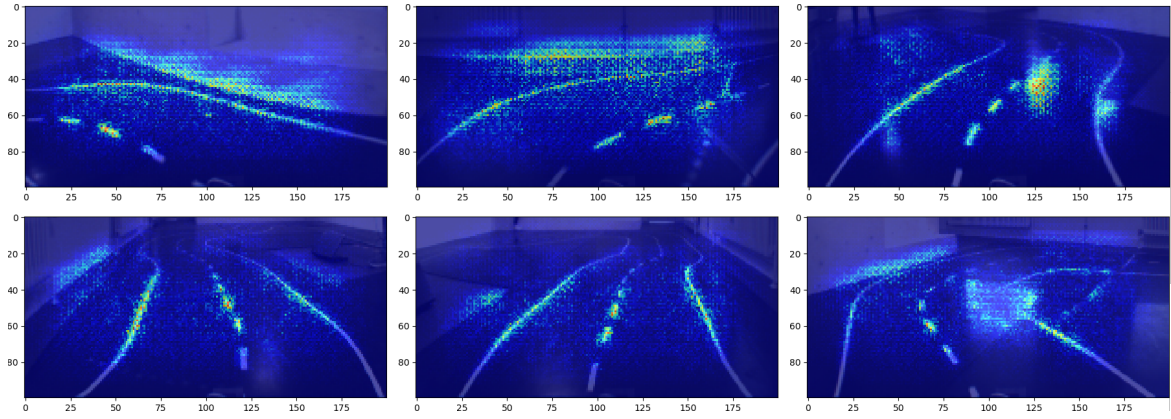


Figure 22: Saliency maps of the trajectory prediction model.

We decided to use the trajectory prediction approach for the final race since we managed to make the car much faster on both the office and the final track compared to the end-to-end approach. Two days before the race, our best lap time was about 8.5 seconds. To get a sense of that speed, there was no competitor who was able to drive a lap under 7 seconds. The best lap-time performed by a human driver was about 6.9 seconds.

Despite being reasonably fast, the car did leave the track quite frequently. However, it was capable of returning back to the track surprisingly often. We think that the camera's large overview and the specially recorded situations to return to the ideal line definitely helped in this regard. The overall driving behavior was much as we expected. It was following the ideal line most of the times and did manage to accelerate on straight lines and was braking in front of curves.

Our strategy was to sacrifice stability for speed. Even though the end-to-end model which was trained on the center line dataset would have probably been making fewer mistakes, there was no chance of winning the competition had it been used. The rules gave room to drive riskfully since we only had to perform a single valid round. Our best valid lap-time during the final race was 10.0 seconds which we achieved the third place with. The winner's lap-time on this day was 8.93 seconds.

It turned out that all cars including ours had severe problems with different lighting conditions. There was a short preparation time at 4 p.m. where the teams were allowed to perform the last runs on the track. Nobody ever trained their car at that time of day and lighting. The result was that the vehicles were barely able to drive a single lap without leaving the track. By asking around, we figured out that every competitor was using deep learning approaches to process the camera frames. This raises questions about the overall generalization ability of neural networks.

5.4. Limits

There are several reasons why we think that our trajectory approach hasn't reached its full potential. The decision to lowering the speed in curves didn't stem from the fact that the wheels started to slip. The vehicle rather seemed to not react aggressively enough. However, we do not think that there is an issue regarding the computational speed since we measured that the system processes 20 frames per second on average. In our opinion, this should be more than enough visual information between two consecutive frames to react even to sharp turns just in time. Also, we believe that after replacing the servo, the steering speed wasn't an issue as well.

What we do believe is that the actual problem source was rooted in the non-systematic error of the odometry which was produced by the weird anomaly of the wheel encoder output. Even after calibration, the odometry comprised too large deviations from the actual driven path. Since the odometry constructed the trajectories which were used as the labels in our training data, the non-systematic error propagated into the predictions of the neural network. Furthermore, we noticed that the steering servo produced a rather non-linear translation between the input angle and the actual angle of the wheels which is another source of inaccuracy.

5.5. Possible Improvements

There has been lots of research on how to implement localization methods for robots. More advanced techniques such as *simultaneous localization and mapping* (SLAM) may be employed in order to establish a more accurate position estimate. In order to use SLAM, it is necessary to provide odometry data as well as periodic scans of the environment which are often acquired by *LIDAR* sensors. Besides dead-reckoning with wheel encoders, there exist several other methods on how to retrieve odometry. For example, *visual odometry* [27] determines the travelled distance using camera images. Zhang and Singh propose a method to extract odometry data from LIDAR [43]. Furthermore, one can combine information from different sources using *bayes filters* such as the *Kalman filter* or the *particle filter* in order to further improve the state estimate [38]. For instance, we could fuse the wheel encoder odometry with data obtained from a relatively cheap *inertial measurement unit* (IMU) which measures linear acceleration and rotational velocity [10].

There could be more work into finding a network architecture that may generalize better. For example, one may use *recurrent neural networks* which incorporate information about the past. This could especially help in tight curves where the camera cannot see much of the upcoming section of the track so it's not possible to make reliable trajectory predictions. Using some type of memory, the network would eventually memorize the shape of the track before entering the turn.

Since it's quite demanding to compute the output of the neural network on the odroid, it may be advantageous to add additional hardware to increase the frame rate. For

instance, the *Intel Movidius Neural Compute Stick* [1] is a very compact computing unit which is optimized for deep learning applications. It's connected via USB and specified to consume only 1 ampere of current. Therefore, it would be suitable for our use case.

We found out that splitting the dataset into training and validation loss by random sampling seemed to pose problems in our case. Because of large similarities between the images while recording with high frame rates, the validation data becomes very similar to the training data. Consequently, this destroys the meaning behind the validation loss since it should actually reflect how well the model performs on new data. To provide evidence to this hypothesis, instead of random sampling, we used the first 20% of our samples as validation data and the remaining 80% as training data. Figure 23 shows a comparison between both methods by training the same network with the same hyperparameters over 50 epochs. We can see that using random sampling, the training and validation loss is almost exactly the same. However, with non-random sampling, the validation loss is significantly lower than the training loss. Thus the neural network performs considerably worse on non-training data. This example demonstrates that we should carefully choose the training and validation data in order to retain a metric of how well the car will perform on the racing track under untrained conditions. If we then still observe an increase in validation loss, we'll clearly know that the model overfits so we may apply techniques discussed in Section 3.3.4.

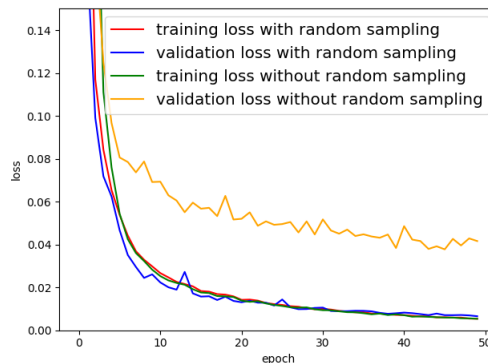


Figure 23: Comparison between training and validation loss using random sampling and non-random sampling to create the training and validation datasets.

In order to unleash the full potential of the trajectory prediction approach, the driving model has to be adapted to take the vehicle's physical properties such as tire friction into account.

6. Conclusion

In this work, we presented two methods for autonomous racing with deep learning. Of course, this problem could have also been solved using traditional image processing techniques including contour and edge detection. However, this requires lots of hand-written rules and tweaking to figure out where the lanes are, how to filter noise such as reflections and how to steer accordingly. It's far more elegant to create a system that learns this processing by itself.

Unfortunately, the end-to-end approach turned out to be suboptimal when applied to racing. Since the neural network tries to mimic the manual driving behavior, the system is limited by human capabilities and does not scale well to higher velocities. Our trajectory prediction approach overcomes this problem by using deep learning only as a subcomponent to predict the vehicle's optimal future path. This provides the flexibility to incorporate an arbitrarily complex driving model to transform a trajectory into steering commands. In addition, collecting training data becomes much easier as it's just a process of guiding the vehicle along the desired path on the track, regardless of the speed.

Due to strange problems with our wheel encoders, we were not able to show the full potential of this approach. Even the best neural network architecture is worthless without good training data. Better sensors and more advanced algorithms are required in order to create quality localization data to train the model with. We believe that if this had been the case, we would have performed considerably better at the Deep Berlin Robocars challenge.

Still there remain some open questions regarding deep learning aspects. Despite the fact that our adaption of *Nvidia's* PilotNet architecture worked in principle, it is unclear whether there exist better architectures and hyperparameter configurations for our particular application. For example, we don't know what would have happened if we had used fewer convolutional layers. Overly complex models are known to generalize worse.

Further investigations will be needed to answer the question why almost every car performed poorly on lighting conditions it hasn't been trained on. It's not clear whether this was caused by problems in the network architecture or in the training data. Therefore, future work has to concentrate on the generalization capabilities of neural networks. Nevertheless, visualization techniques such as saliency maps already help to detect the existence of such problems.

The results achieved by applying deep learning to image classification and object detection make it a promising technique to design future autonomous systems. However, this thesis showed that neural networks are not able to provide complete end-to-end solutions for all applications yet. Until now, we should rather consider them a powerful tool to help solving complex problems as part of a processing pipeline.

References

- [1] Intel neural compute stick webpage. URL <https://software.intel.com/en-us/neural-compute-stick>.
- [2] The imaging source dfm 22buc03-ml webpage. URL <https://www.theimagingsource.com/products/board-cameras/usb-2.0-color/dfm22buc03ml/>.
- [3] Donkey car homepage. URL <http://donkeycar.com>.
- [4] *Semesterprojekt Hochautomatisiertes Fahren Dokumentation*, 2018. URL https://www2.informatik.hu-berlin.de/~hs/Lehre/2017-WS_SP-HAF/20180524_SPHAF-Dokumentation.pdf.
- [5] S. Albarqouni, C. Baur, F. Achilles, V. Belagiannis, S. Demirci, and N. Navab. Aggnet: Deep learning from crowds for mitosis detection in breast cancer histology images. *IEEE Transactions on Medical Imaging*, 35(5):1313–1321, May 2016. ISSN 0278-0062. doi: 10.1109/TMI.2016.2528120.
- [6] K. H. Ang, G. Chong, and Y. Li. Pid control system analysis, design, and technology. *IEEE Transactions on Control Systems Technology*, 13:559–576, 2005.
- [7] M. Bojarski, P. Yeres, A. Choromanska, K. Choromanski, B. Firner, L. D. Jackel, and U. Muller. Explaining how a deep neural network trained with end-to-end learning steers a car. *CoRR*, abs/1704.07911, 2017. URL <http://arxiv.org/abs/1704.07911>.
- [8] D. M. Bradley. *Odometry : Calibration and error modeling*. 2005.
- [9] M. Burgess. Roborace is building a 300kph ai supercar – no driver required, March 2018. URL <https://www.wired.co.uk/article/roborace-car-formyla-e-robocar-uk-autonomous-race-denis-sverdlov>.
- [10] B.-S. Cho, W.-s. Moon, W.-J. Seo, and K. Baek. A dead reckoning localization system for mobile robots using inertial sensors and wheel revolution encoding. *Journal of Mechanical Science and Technology*, 25, 11 2011. doi: 10.1007/s12206-011-0805-1.
- [11] R. Collobert and S. Bengio. Links between perceptrons, mlps and svms. In *Proceedings of the Twenty-first International Conference on Machine Learning*, ICML '04, pages 23–, New York, NY, USA, 2004. ACM. ISBN 1-58113-838-5. doi: 10.1145/1015330.1015415. URL <http://doi.acm.org/10.1145/1015330.1015415>.
- [12] R. Collobert and J. Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th International Conference on Machine Learning*, ICML '08, pages 160–167, New York,

- NY, USA, 2008. ACM. ISBN 978-1-60558-205-4. doi: 10.1145/1390156.1390177. URL <http://doi.acm.org/10.1145/1390156.1390177>.
- [13] M. de la Iglesia Valls, H. F. C. Hendrikx, V. Reijgwart, F. V. Meier, I. Sa, R. Dubé, A. R. Gawel, M. Bürki, and R. Siegwart. Design of an autonomous racecar: Perception, state estimation and system integration. *CoRR*, abs/1804.03252, 2018. URL <http://arxiv.org/abs/1804.03252>.
- [14] emo berlin.de. Abschlussrennen der deep berlin robocars challenge, September 2018. URL <https://www.emo-berlin.de/de/newsarchiv/news/abschlussrennen-der-deep-berlin-robocars-challenge/>.
- [15] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [16] J. B. H. II and A. H. Waibel. A novel objective function for improved phoneme recognition using time-delay neural networks. *IEEE Trans. Neural Networks*, 1(2):216–228, 1990. doi: 10.1109/72.80233. URL <https://doi.org/10.1109/72.80233>.
- [17] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *CoRR*, abs/1502.03167, 2015. URL <http://arxiv.org/abs/1502.03167>.
- [18] S. Jain. An overview of regularization techniques in deep learning (with python code), April 2018. URL <https://www.analyticsvidhya.com/blog/2018/04/fundamentals-deep-learning-regularization-techniques/>.
- [19] L. F. Johann Borenstein. Umbmark: a benchmark test for measuring odometry errors in mobile robots, 1995. URL <https://doi.org/10.1117/12.228968>.
- [20] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. URL <http://arxiv.org/abs/1412.6980>.
- [21] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’12, pages 1097–1105, USA, 2012. Curran Associates Inc. URL <http://dl.acm.org/citation.cfm?id=2999134.2999257>.
- [22] F.-F. Li, J. Johnson, and S. Young. Lecture 3: Loss functions and optimization. April 2017. URL http://cs231n.stanford.edu/slides/2017/cs231n_2017_lecture3.pdf.
- [23] F.-F. Li, J. Johnson, and S. Young. Lecture 6: Training neural networks, part 1. April 2017. URL http://cs231n.stanford.edu/slides/2017/cs231n_2017_lecture6.pdf.

- [24] F.-F. Li, J. Johnson, and S. Young. Lecture 7: Training neural networks, part 2. April 2017. URL http://cs231n.stanford.edu/slides/2017/cs231n_2017_lecture7.pdf.
- [25] A. Liniger, A. Domahidi, and M. Morari. Optimization-based autonomous racing of 1:43 scale rc cars. *Optimal Control Applications and Methods*, 36:628–647, 09 2015. doi: 10.1002/oca.2123.
- [26] K. Majek. Self-racing cars 2017 – udacity team soulless, 2017. URL <https://karolmajek.pl/self-racing-cars-2017-udacity-team-soulless/>.
- [27] D. Nistér, O. Naroditsky, and J. R. Bergen. Visual odometry. *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, 1:I–I, 2004.
- [28] B. Póczos and R. Tibshirani. Convex optimization cmu-10725 quasi newton methods. URL <http://www.stat.cmu.edu/~ryantibs/convexopt-F13/lectures/11-QuasiNewton.pdf>.
- [29] M. Quigley, K. Conley, B. P Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y Ng. Ros: an open-source robot operating system, 01 2009.
- [30] U. Rosolia, A. Carvalho, and F. Borrelli. Autonomous racing using learning model predictive control. In *2017 American Control Conference (ACC)*, pages 5115–5120, May 2017. doi: 10.23919/ACC.2017.7963748.
- [31] A. E. Sallab, M. Abdou, E. Perot, and S. Yogamani. Deep reinforcement learning framework for autonomous driving. *CoRR*, abs/1704.02532, 2017.
- [32] D. Shapiro. Go, autonomous speed racer, go! nvidia drive px 2 to power world’s first robotic motorsports competition, April 2016. URL <https://blogs.nvidia.com/blog/2016/04/05/roborace/>.
- [33] D. Shapiro. Self-racing cars kick off first autonomous vehicle track day, June 2016. URL <https://blogs.nvidia.com/blog/2016/06/03/autonomous-vehicles/>.
- [34] S. Sharma. Activation functions: Neural networks, Sep 2017. URL <https://towardsdatascience.com/activation-functions-neural-networks-1cbd9f8d91d6>.
- [35] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *CoRR*, abs/1312.6034, 2013. URL <http://arxiv.org/abs/1312.6034>.
- [36] A. Simpson. Self-driving car steering angle prediction based on image recognition. 2017.

- [37] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. A. Riedmiller. Striving for simplicity: The all convolutional net. *CoRR*, abs/1412.6806, 2014. URL <http://arxiv.org/abs/1412.6806>.
- [38] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*.
- [39] A. S. Walia. Activation functions and it's types-which is better?, May 2017. URL <https://towardsdatascience.com/activation-functions-and-its-types-which-is-better-a9a5310cc8f>.
- [40] T. Wang, D. J. Wu, A. Coates, and A. Y. Ng. End-to-end text recognition with convolutional neural networks. In *Proceedings of the 21st International Conference on Pattern Recognition (ICPR2012)*, pages 3304–3308, Nov 2012.
- [41] P. Werbos and P. J. (Paul John. Beyond regression : new tools for prediction and analysis in the behavioral sciences /. 01 1974.
- [42] Y.-x. Yuan. Step-sizes for the gradient method. 1999.
- [43] J. Zhang and S. Singh. Loam: Lidar odometry and mapping in real-time. In *Robotics: Science and Systems Conference*, July 2014.
- [44] J. Zhong. Pid controller tuning: A short tutorial. 2006.

A. Code Snippets

A.1. End-to-End Model Implementation using Keras

```
1 def build_end_to_end_model():
2     img_in = Input(shape=(100, 200, 3)), name="img_in")
3     x = img_in
4     x = BatchNormalization(axis=3)(x)
5     x = Convolution2D(filters=24, kernel_size=(5, 5),
6                       strides=(2, 2), activation="relu")(x)
7     x = Convolution2D(filters=36, kernel_size=(5, 5),
8                       strides=(2, 2), activation="relu")(x)
9     x = Convolution2D(filters=48, kernel_size=(5, 5),
10                      strides=(2, 2), activation="relu")(x)
11    x = Convolution2D(filters=64, kernel_size=(3, 3),
12                      strides=(1, 1), activation="relu")(x)
13    x = Convolution2D(filters=64, kernel_size=(3, 3),
14                      strides=(1, 1), activation="relu")(x)
15    x = Flatten()(x)
16    x = Dense(units=100, activation="relu")(x)
17    x = Dense(units=50, activation="relu")(x)
18    x = Dense(units=10, activation="relu")(x)
19
20    angle_out = Dense(units=1, name="angle_out",
21                      activation="linear")(x)
22    throttle_out = Dense(units=1, name="throttle_out",
23                          activation="linear")(x)
24
25    model = Model(inputs=[img_in], outputs=[angle_out, throttle_out])
26    model.compile(optimizer="adam",
27                  loss={"angle_out": "mean_squared_error",
28                        "throttle_out": "mean_squared_error"})
29    return model
```

A.2. Trajectory Prediction Model Implementation using Keras

```
1 def build_trajectory_prediction_model(args):
2     img_in = Input(shape=(100, 200, 3)), name="img_in")
3     x = img_in
4     x = BatchNormalization(axis=3)(x)
5     x = Convolution2D(filters=24, kernel_size=(5, 5),
6                       strides=(2, 2), activation="relu")(x)
7     x = Convolution2D(filters=36, kernel_size=(5, 5),
8                       strides=(2, 2), activation="relu")(x)
9     x = Convolution2D(filters=48, kernel_size=(5, 5),
10                      strides=(2, 2), activation="relu")(x)
11    x = Convolution2D(filters=64, kernel_size=(3, 3),
12                      strides=(1, 1), activation="relu")(x)
13    x = Convolution2D(filters=64, kernel_size=(3, 3),
14                      strides=(1, 1), activation="relu")(x)
15    x = Flatten()(x)
16    x = Dense(units=100, activation="relu")(x)
17    x = Dense(units=50, activation="relu")(x)
18    x = Dense(units=10, activation="relu")(x)
19
20    loss = {}
21    outputs = []
22    for i in range(3):
23        name_x = "coord_x_out_" + str(i)
24        name_y = "coord_y_out_" + str(i)
25        coord_x_out = Dense(1, name=name_x, activation="linear")(x)
26        coord_y_out = Dense(1, name=name_y, activation="linear")(x)
27        loss[name_x] = "mean_squared_error"
28        loss[name_y] = "mean_squared_error"
29        outputs.append(coord_x_out)
30        outputs.append(coord_y_out)
31
32    model = Model(inputs=[img_in], outputs=outputs)
33    model.compile(optimizer="adam", loss=loss)
34    return model
```

B. Trajectory Prediction Approach Parameters

parameter	value
c_l	3.1697932657906436
c_r	3.189793265790643
T	305
\mathcal{D}	(0.6, 1.2, 1.8)
b	0.32
ν_1	2.5
ν_2	1.5
χ	0.3
K_p	$5 \cdot 10^{-3}$
K_i	0
K_d	$1 \cdot 10^{-3}$

C. Source Code

C.1. Odroid

C.1.1. C++ Source Code

AsyncCapture.h

```
1 //
2 // Created by philipp on 17.08.18.
3 //
4
5 #ifndef CAR_ASYNC_CAPTURE_H
6 #define CAR_ASYNC_CAPTURE_H
7
8 #include <opencv2/opencv.hpp>
9 #include <opencv2/core/mat.hpp>
10
11 #include <gst/gst.h>
12 #include <gst/app/gstappsink.h>
13
14 namespace car
15 {
16
17 class AsyncCapture
18 {
19 public:
20     using OnFrameCapturedCallback = std::function<void(cv::Mat & frame)>;
21
22     explicit AsyncCapture(int frameWidth = 744, int frameHeight = 480, int frameRate
23         ↪ = 60);
24
25     ~AsyncCapture();
26
27     AsyncCapture(const AsyncCapture &) = delete;
28
29     AsyncCapture(AsyncCapture && other) noexcept;
30
31     AsyncCapture & operator=(const AsyncCapture &) = delete;
32
33     AsyncCapture & operator=(AsyncCapture && other) noexcept;
34
35     void start(const OnFrameCapturedCallback & callback);
36
37     void stop();
38
39 private:
40     int frameWidth, frameHeight, frameRate;
41
42     OnFrameCapturedCallback onFrameCapturedCallback;
43
44     GstElement * pipeline = nullptr;
45     GstElement * source = nullptr;
46     GstElement * inputCaps = nullptr;
47     GstElement * queue = nullptr;
48     GstElement * bayer2rgb = nullptr;
49     GstElement * convert = nullptr;
50     GstElement * appsink = nullptr;
51
52     static GstFlowReturn appsinkFrameCallback(GstAppSink * appsink, gpointer data);
53
54     void createPipeline();
55 }
```

```

55     void ensureReadyState();
56 };
57
58 }
59
60 #endif //CAR_ASYNC_CAPTURE_H

```

AsyncCapture.cpp

```

1 //
2 // Created by philipp on 17.08.18.
3 //
4
5 #include "camera/AsyncCapture.h"
6
7 namespace car
8 {
9
10 AsyncCapture::AsyncCapture(int frameWidth, int frameHeight, int frameRate)
11     : frameWidth(frameWidth), frameHeight(frameHeight), frameRate(frameRate)
12 {
13     if (!gst_is_initialized())
14         gst_init(nullptr, nullptr);
15
16     createPipeline();
17     ensureReadyState();
18 }
19
20 AsyncCapture::~AsyncCapture()
21 {
22     stop();
23     gst_object_unref(pipeline);
24 }
25
26 AsyncCapture::AsyncCapture(AsyncCapture && other) noexcept
27     : pipeline(other.pipeline)
28 {
29     other.pipeline = nullptr;
30 }
31
32 AsyncCapture & AsyncCapture::operator=(AsyncCapture && other) noexcept
33 {
34     gst_object_unref(pipeline);
35     pipeline = other.pipeline;
36     other.pipeline = nullptr;
37     return *this;
38 }
39
40 void AsyncCapture::start(const AsyncCapture::OnFrameCapturedCallback & callback)
41 {
42     onFrameCapturedCallback = callback;
43     GstAppSinkCallbacks callbacks = {nullptr, nullptr, appsinkFrameCallback};
44     gst_app_sink_set_callbacks(GST_APP_SINK(appsink), &callbacks, this, nullptr);
45
46     gst_element_set_state(pipeline, GST_STATE_PLAYING);
47     gst_element_get_state(pipeline, nullptr, nullptr, GST_CLOCK_TIME_NONE);
48 }
49
50 void AsyncCapture::stop()
51 {
52     gst_element_set_state(pipeline, GST_STATE_NULL);
53     gst_element_get_state(pipeline, nullptr, nullptr, GST_CLOCK_TIME_NONE);
54 }
55
56 GstFlowReturn AsyncCapture::appsinkFrameCallback(GstAppSink * appsink, gpointer data)

```

```

57 {
58     auto this_ = static_cast<AsyncCapture *>(data);
59     auto sample = gst_app_sink_pull_sample(appsink);
60     bool success{false};
61     if (sample)
62     {
63         auto buffer = gst_sample_get_buffer(sample);
64         auto caps = gst_sample_get_caps(sample);
65         auto structure = gst_caps_get_structure(caps, 0);
66         int width, height;
67         gst_structure_get_int(structure, "width", &width);
68         gst_structure_get_int(structure, "height", &height);
69         GstMapInfo info;
70         if (gst_buffer_map(buffer, &info, GST_MAP_READ))
71         {
72             const auto frameData = info.data;
73             // 'data' now contains a pointer to readable image data
74             cv::Mat frame{height, width, CV_8UC3, frameData};
75             this->onFrameCapturedCallback(frame);
76             // unmap after use
77             gst_buffer_unmap(buffer, &info);
78             success = true;
79         }
80     }
81     gst_sample_unref(sample);
82     if (!success)
83     {
84         cv::Mat emptyFrame{};
85         this->onFrameCapturedCallback(emptyFrame);
86     }
87     return GST_FLOW_OK;
88 }
89
90 void AsyncCapture::createPipeline()
91 {
92     pipeline = gst_pipeline_new("pipeline");
93
94     source = gst_element_factory_make("tcambin", nullptr);
95     if (!source)
96         throw std::runtime_error("'tcambin' could not be initialized! Check tiscamera
97             ↪ installation");
98
99     inputCaps = gst_element_factory_make("capsfilter", nullptr);
100     std::ostringstream oss;
101     oss << "video/x-bayer, width=" << frameWidth
102         << ", height=" << frameHeight
103         << ", framerate=" << frameRate << "/1";
104     g_object_set(inputCaps, "caps", gst_caps_from_string(oss.str().c_str()), nullptr);
105
106     queue = gst_element_factory_make("queue", nullptr);
107     g_object_set(queue, "leaky", true, "max-size-buffers", 2, nullptr);
108
109     bayer2rgb = gst_element_factory_make("bayer2rgb", nullptr);
110
111     convert = gst_element_factory_make("videoconvert", nullptr);
112
113     appsink = gst_element_factory_make("appsink", nullptr);
114     g_object_set(appsink, "caps", gst_caps_from_string("video/x-raw, format=BGR"),
115         ↪ nullptr);
116
117     assert(pipeline && source && inputCaps && queue && bayer2rgb && convert && appsink)
118         ↪ ;
119
120     gst_bin_add_many(
121         GST_BIN(pipeline),
122         source, inputCaps, queue, bayer2rgb, convert, appsink, nullptr);

```

```

120     auto result = gst_element_link_many(
121         source, inputCaps, queue, bayer2rgb, convert, appsink, nullptr);
122     assert(result);
123 }
124
125 void AsyncCapture::ensureReadyState()
126 {
127     GstState state;
128     if ((gst_element_get_state(source, &state, nullptr, GST_CLOCK_TIME_NONE) ==
129         ↪ GST_STATE_CHANGE_SUCCESS) &&
130         state == GST_STATE_NULL)
131     {
132         gst_element_set_state(source, GST_STATE_READY);
133         gst_element_get_state(source, nullptr, nullptr, GST_CLOCK_TIME_NONE);
134     }
135 }
136 }

```

Config.h

```

1 //
2 // Created by philipp on 09.07.18.
3 //
4
5 #ifndef CAR_CONFIG_H
6 #define CAR_CONFIG_H
7
8 #include "Json.h"
9
10 namespace car
11 {
12     namespace config
13     {
14
15         std::string directory();
16
17         nlohmann::json open();
18
19     }
20 }
21
22 #endif //CAR_CONFIG_H

```

Config.cpp

```

1 //
2 // Created by philipp on 09.07.18.
3 //
4
5 #include "utils/Config.h"
6 #include <fstream>
7
8 namespace car
9 {
10     namespace config
11     {
12
13         std::string directory()
14         {
15             auto userHome = std::string{std::getenv("HOME")};
16             return userHome + "/.car";
17         }
18     }
19 }

```



```

18
19 nlohmann::json open()
20 {
21     auto configPath = directory() + "/config.json";
22     std::ifstream file{configPath};
23     nlohmann::json j;
24     file >> j;
25     return j;
26 }
27
28 }
29 }

```

Recorder.h

```

1 //
2 // Created by philipp on 20.06.18.
3 //
4
5 #ifndef CAR_RECORDER_H
6 #define CAR_RECORDER_H
7
8 #include <nodelet/nodelet.h>
9 #include <ros/ros.h>
10 #include <logging/MessageOutputStream.h>
11 #include <atomic>
12 #include <opencv2/opencv.hpp>
13 #include <opencv2/core/mat.hpp>
14 #include <NetworkingLib/Time.h>
15 #include <camera/AsyncCapture.h>
16 #include <mutex>
17
18 #include "car/WheelTicks.h"
19 #include "car/Odometry.h"
20 #include "car/SetThrottle.h"
21 #include "car/SetAngle.h"
22 #include "car/UssDistance.h"
23 #include "car/StopRecording.h"
24 #include "car/RemoteState.h"
25
26 namespace car
27 {
28
29 class Recorder : public nodelet::Nodelet
30 {
31 public:
32     void onInit() override;
33
34     Recorder();
35
36 private:
37     ros::NodeHandle nh;
38
39     ros::Subscriber wheelTicksSubscriber;
40     ros::Subscriber odometrySubscriber;
41     ros::Subscriber setThrottleSubscriber;
42     ros::Subscriber setAngleSubscriber;
43     ros::Subscriber ussDistanceSubscriber;
44     ros::Subscriber stopRecordingSubscriber;
45     ros::Subscriber remoteStateSubscriber;
46
47     MessageOutputStream messageOutputStream;
48
49     float ticksRL{0};
50     float ticksRR{0};

```

```

51     std::mutex ticksMutex;
52
53     std::atomic<float> x{0};
54     std::atomic<float> y{0};
55
56     std::atomic<float> setThrottle{0.0f};
57     std::atomic<float> setAngle{0.0f};
58     std::atomic<float> distance{0.0f};
59
60     std::atomic<bool> recording{true};
61     std::atomic<bool> recordingSequence{true};
62
63     AsyncCapture capture;
64     cv::Mat currFrame;
65     std::mutex frameCopyMutex;
66     std::atomic<bool> receivedFrame{false};
67
68     std::string recordingsDirectory;
69     networking::time::Duration frameTime;
70
71     void wheelTicksCallback(const WheelTicks::ConstPtr msg);
72
73     void odometryCallback(const Odometry::ConstPtr msg);
74
75     void setThrottleCallback(const SetThrottle::ConstPtr msg);
76
77     void setAngleCallback(const SetAngle::ConstPtr msg);
78
79     void ussDistanceCallback(const UssDistance::ConstPtr msg);
80
81     void stopRecordingCallback(const StopRecording::ConstPtr msg);
82
83     void remoteStateCallback(const RemoteState::ConstPtr msg);
84
85     void record();
86
87     std::string createRecordingDirectory();
88
89     void writeCsvRecord(std::ofstream & dataFile, const std::string & imgFilename);
90
91     void readConfig();
92 };
93
94 }
95
96 #endif //CAR_RECORDER_H

```

Recorder.cpp

```

1 //
2 // Created by philipp on 20.06.18.
3 //
4
5 #include <pluginlib/class_list_macros.h>
6 #include <thread>
7 #include "recorder/Recorder.h"
8 #include <boost/filesystem.hpp>
9 #include <utils/Config.h>
10
11 PLUGINLIB_EXPORT_CLASS(car::Recorder, nodelet::Nodelet);
12
13 namespace car
14 {
15
16 Recorder::Recorder()

```

```

17 : messageOStream(nh, "recorder")
18   , capture(744, 480, 60)
19 {}
20
21 void Recorder::onInit()
22 {
23   messageOStream.write("onInit", "START");
24
25   wheelTicksSubscriber = nh.subscribe("WheelTicks", 1000, &Recorder::
    ↪ wheelTicksCallback, this);
26   odometrySubscriber = nh.subscribe("Odometry", 1, &Recorder::odometryCallback, this)
    ↪ ;
27   setThrottleSubscriber = nh.subscribe("SetThrottle", 1, &Recorder::
    ↪ setThrottleCallback, this);
28   setAngleSubscriber = nh.subscribe("SetAngle", 1, &Recorder::setAngleCallback, this)
    ↪ ;
29   ussDistanceSubscriber = nh.subscribe("UssDistance", 1, &Recorder::
    ↪ ussDistanceCallback, this);
30   stopRecordingSubscriber = nh.subscribe("StopRecording", 1, &Recorder::
    ↪ stopRecordingCallback, this);
31   remoteStateSubscriber = nh.subscribe("RemoteState", 1, &Recorder::
    ↪ remoteStateCallback, this);
32
33   readConfig();
34
35   capture.start(
36     [&](auto & frame)
37     {
38       std::lock_guard<std::mutex> lock{frameCopyMutex};
39       currFrame = frame.clone();
40       receivedFrame = true;
41     });
42
43   std::thread recorderThread{[&] { this->record(); }};
44   recorderThread.detach();
45
46   messageOStream.write("onInit", "END");
47 }
48
49 void Recorder::wheelTicksCallback(const WheelTicks::ConstPtr msg)
50 {
51   std::lock_guard<std::mutex> lock{ticksMutex};
52   ticksRL += msg->rearLeftTicks;
53   ticksRR += msg->rearRightTicks;
54 }
55
56 void Recorder::odometryCallback(const Odometry::ConstPtr msg)
57 {
58   x = msg->xDistance;
59   y = msg->yDistance;
60 }
61
62 void Recorder::setThrottleCallback(const SetThrottle::ConstPtr msg)
63 {
64   setThrottle = msg->throttle;
65 }
66
67 void Recorder::setAngleCallback(const SetAngle::ConstPtr msg)
68 {
69   setAngle = msg->angle;
70 }
71
72 void Recorder::ussDistanceCallback(const UssDistance::ConstPtr msg)
73 {
74   distance = msg->distance;
75 }

```

```

76
77 void Recorder::stopRecordingCallback(const StopRecording::ConstPtr msg)
78 {
79     recording = false;
80 }
81
82 void Recorder::remoteStateCallback(const RemoteState::ConstPtr msg)
83 {
84     if (msg->id == 0)
85         recordingSequence = msg->value != 0 ? true : false;
86 }
87
88 void Recorder::record()
89 {
90     auto dataDirName = createRecordingDirectory();
91
92     std::ostringstream oss;
93     oss << dataDirName << "/data.csv";
94     std::ofstream dataFile{oss.str()};
95
96     dataFile << "ticksRL,ticksRR,x,y,setThrottle,setAngle,distance,recording,img\n";
97
98     auto last = networking::time::now();
99
100    for (std::size_t counter = 0; recording; counter++)
101    {
102        // keep constant frame rate
103        auto now = networking::time::now();
104        auto lastFrameTime = now - last;
105        if (lastFrameTime < frameTime)
106            std::this_thread::sleep_for(frameTime - lastFrameTime);
107        last = networking::time::now();
108
109        while (!receivedFrame);
110        cv::Mat frame{};
111        {
112            std::lock_guard<std::mutex> lock{frameCopyMutex};
113            frame = currFrame;
114        }
115        if (frame.empty())
116            continue;
117
118        // save image
119        oss = std::ostringstream{};
120        oss << "img-" << counter << ".jpg";
121        auto imgFilename = oss.str();
122        oss = std::ostringstream{};
123        oss << dataDirName << "/" << imgFilename;
124        auto imgPath = oss.str();
125        cv::imwrite(imgPath, frame);
126
127        writeCsvRecord(dataFile, imgFilename);
128    }
129
130    messageOStream.write("MSG", "stopped recording");
131 }
132
133 std::string Recorder::createRecordingDirectory()
134 {
135     std::ostringstream oss;
136     auto time = std::time(nullptr);
137     oss << recordingsDirectory << "/data-" << std::put_time(std::localtime(&time), "%Y
138         ↪ -%m-%d-%H-%M-%S");
139     auto dirName = oss.str();
140     boost::filesystem::path dir{dirName};
141     boost::filesystem::create_directory(dir);

```

```

141     return dirName;
142 }
143
144 void Recorder::writeCsvRecord(std::ofstream & dataFile, const std::string &
    ↪ imgFilename)
145 {
146     float ticksRLCopy{0}, ticksRRCopy{0};
147     {
148         std::lock_guard<std::mutex> lock{ticksMutex};
149         ticksRLCopy = ticksRL;
150         ticksRRCopy = ticksRR;
151         ticksRL = 0;
152         ticksRR = 0;
153     }
154
155     dataFile << ticksRLCopy << ","
156             << ticksRRCopy << ","
157             << x.load() << ","
158             << y.load() << ","
159             << setThrottle.load() << ","
160             << setAngle.load() << ","
161             << distance.load() << ","
162             << recordingSequence.load() << ","
163             << imgFilename << "\n";
164 }
165
166 void Recorder::readConfig()
167 {
168     auto j = config::open();
169
170     if (j.count("recordingsDirectory") > 0)
171         recordingsDirectory = j.at("recordingsDirectory").get<std::string>();
172     else
173         recordingsDirectory = config::directory() + "/recordings";
174
175     int fps = 60;
176     if (j.count("recordingFPS") > 0)
177         fps = j.at("recordingFPS").get<int>();
178     frameTime = std::chrono::nanoseconds(1000000000 / fps);
179 }
180
181 }

```

RemoteControl.h

```

1 //
2 // Created by philipp on 20.06.18.
3 //
4
5 #ifndef CAR_REMOTECONTROL_H
6 #define CAR_REMOTECONTROL_H
7
8 #include <nodelet/nodelet.h>
9 #include <ros/ros.h>
10 #include <logging/MessageOStream.h>
11
12 #include "car/RemoteAngle.h"
13 #include "car/RemoteThrottle.h"
14
15 namespace car
16 {
17
18 class RemoteControl : public nodelet::Nodelet
19 {
20 public:

```

```

21     void onInit() override;
22
23     RemoteControl();
24
25 private:
26     ros::NodeHandle nh;
27
28     ros::Subscriber remoteThrottleSubscriber;
29     ros::Subscriber remoteAngleSubscriber;
30
31     ros::Publisher setThrottlePublisher;
32     ros::Publisher setAnglePublisher;
33
34     MessageOutputStream messageOutputStream;
35
36     bool controlThrottle{false};
37     bool controlSteeringAngle{false};
38
39     void remoteThrottleCallback(const RemoteThrottle::ConstPtr & msg);
40
41     void remoteAngleCallback(const RemoteAngle::ConstPtr & msg);
42
43     void readConfig();
44 };
45
46 }
47
48 #endif //CAR_REMOTECONTROL_H

```

RemoteControl.cpp

```

1 //
2 // Created by philipp on 20.06.18.
3 //
4
5 #include <pluginlib/class_list_macros.h>
6 #include <utils/Config.h>
7 #include "remoteControl/RemoteControl.h"
8 #include "car/SetThrottle.h"
9 #include "car/SetAngle.h"
10
11 PLUGINLIB_EXPORT_CLASS(car::RemoteControl, nodelet::Nodelet);
12
13 namespace car
14 {
15
16 RemoteControl::RemoteControl()
17     : messageOutputStream(nh, "remoteControl")
18 {}
19
20 void RemoteControl::onInit()
21 {
22     messageOutputStream.write("onInit", "START");
23
24     readConfig();
25
26     if (controlThrottle)
27     {
28         remoteThrottleSubscriber = nh.subscribe("RemoteThrottle", 1, &RemoteControl::
                ↪ remoteThrottleCallback, this);
29         setThrottlePublisher = nh.advertise<SetThrottle>("SetThrottle", 1);
30     }
31
32     if (controlSteeringAngle)
33     {

```

```

34     remoteAngleSubscriber = nh.subscribe("RemoteAngle", 1, &RemoteControl::
      ↪ remoteAngleCallback, this);
35     setAnglePublisher = nh.advertise<SetAngle>("SetAngle", 1);
36 }
37
38 messageOStream.write("onInit", "END");
39 }
40
41 void RemoteControl::remoteThrottleCallback(const RemoteThrottle::ConstPtr & msg)
42 {
43     SetThrottle throttleMsg;
44     throttleMsg.throttle = msg->throttle;
45     setThrottlePublisher.publish(throttleMsg);
46 }
47
48 void RemoteControl::remoteAngleCallback(const RemoteAngle::ConstPtr & msg)
49 {
50     SetAngle angleMsg;
51     angleMsg.angle = msg->angle;
52     setAnglePublisher.publish(angleMsg);
53 }
54
55 void RemoteControl::readConfig()
56 {
57     auto j = config::open();
58     if (j.count("enableThrottleRemoteControl") > 0)
59         controlThrottle = j.at("enableThrottleRemoteControl").get<bool>();
60     if (j.count("enableSteeringAngleRemoteControl") > 0)
61         controlSteeringAngle = j.at("enableSteeringAngleRemoteControl").get<bool>();
62 }
63
64 }

```

RemoteControlMessage.h

```

1 //
2 // Created by philipp on 21.06.18.
3 //
4
5 #ifndef CAR_REMOTECONTROLMESSAGE_H
6 #define CAR_REMOTECONTROLMESSAGE_H
7
8 #include "NetworkingLib/Message.h"
9 #include <boost/algorithm/string.hpp>
10 #include <utility>
11
12 namespace car
13 {
14
15     struct RemoteControlMessage
16     {
17         RemoteControlMessage() = default;
18
19         RemoteControlMessage(std::string key, std::string value)
20             : key(std::move(key)), value(std::move(value))
21         {}
22
23         std::string key;
24         std::string value;
25     };
26
27 }
28
29 namespace networking
30 {

```

```

31 namespace message
32 {
33
34 template<>
35 struct Decoder<car::RemoteControlMessage>
36 {
37     void operator()(car::RemoteControlMessage & message, const std::string & data)
38         ↪ const
39     {
40         std::vector<std::string> split;
41         boost::split(split, data, [](char c) {return c == ',';});
42         message.key = split.at(0);
43         message.value = split.at(1);
44     };
45 };
46 }
47 }
48
49 #endif //CAR_REMOTECONTROLMESSAGE_H

```

RemoteControlMessageReceiver.h

```

1 //
2 // Created by philipp on 23.08.18.
3 //
4
5 #ifndef CAR_REMOTECONTROLMESSAGERECEIVER_H
6 #define CAR_REMOTECONTROLMESSAGERECEIVER_H
7
8 #include <nodelet/nodelet.h>
9 #include <ros/ros.h>
10 #include <logging/MessageOStream.h>
11
12 #include "NetworkingLib/DatagramReceiver.h"
13
14 #include "RemoteControlMessage.h"
15
16 namespace car
17 {
18
19 class RemoteControlMessageReceiver : public nodelet::Nodelet
20 {
21 public:
22     void onInit() override;
23
24     RemoteControlMessageReceiver();
25
26 private:
27     ros::NodeHandle nh;
28
29     ros::Publisher remoteThrottlePublisher;
30     ros::Publisher remoteAnglePublisher;
31     ros::Publisher remoteStatePublisher;
32
33     MessageOStream messageOStream;
34
35     networking::Networking net;
36
37     networking::message::DatagramReceiver<RemoteControlMessage>::Ptr messageReceiver;
38
39     void receiveMessages();
40
41     void handleMessage(const RemoteControlMessage & message);
42 };

```



```

43 }
44 }
45
46 #endif //CAR_REMOTECONTROLMESSAGE RECEIVER_H

```

RemoteControlMessageReceiver.cpp

```

1 //
2 // Created by philipp on 23.08.18.
3 //
4
5 #include <pluginlib/class_list_macros.h>
6 #include "remoteControl/RemoteControlMessageReceiver.h"
7 #include <boost/algorithm/string.hpp>
8
9 #include "car/RemoteThrottle.h"
10 #include "car/RemoteAngle.h"
11 #include "car/RemoteState.h"
12
13 PLUGINLIB_EXPORT_CLASS(car::RemoteControlMessageReceiver, nodelet::Nodelet);
14
15 namespace car
16 {
17
18 RemoteControlMessageReceiver::RemoteControlMessageReceiver()
19 : messageOStream(nh, "remoteControlMessageReceiver")
20 {}
21
22 void RemoteControlMessageReceiver::onInit()
23 {
24     messageOStream.write("onInit", "START");
25
26     remoteThrottlePublisher = nh.advertise<RemoteThrottle>("RemoteThrottle", 1);
27     remoteAnglePublisher = nh.advertise<RemoteAngle>("RemoteAngle", 1);
28     remoteStatePublisher = nh.advertise<RemoteState>("RemoteState", 1);
29
30     messageReceiver = networking::message::DatagramReceiver<RemoteControlMessage>::
31         ↪ create(net, 10288);
32     receiveMessages();
33
34     messageOStream.write("onInit", "START");
35 }
36
37 void RemoteControlMessageReceiver::receiveMessages()
38 {
39     messageReceiver->asyncReceive(
40         networking::time::Duration::max(),
41         [&](const auto & error, const auto & message, const auto & host, auto port)
42         {
43             if (error)
44                 return;
45             this->handleMessage(message);
46             this->receiveMessages();
47         });
48 }
49
50 void RemoteControlMessageReceiver::handleMessage(const RemoteControlMessage & message
51         ↪ )
52 {
53     if (message.key == "angle")
54     {
55         RemoteAngle angleMsg;
56         angleMsg.angle = std::stof(message.value);
57         remoteAnglePublisher.publish(angleMsg);
58     }
59 }

```

```

57     else if (message.key == "speed")
58     {
59         RemoteThrottle throttleMsg;
60         throttleMsg.throttle = std::stof(message.value);
61         remoteThrottlePublisher.publish(throttleMsg);
62     }
63     else if (message.key == "state")
64     {
65         RemoteState remoteState;
66         std::vector<std::string> split;
67         boost::split(split, message.value, [](char c) { return c == ','; });
68         remoteState.id = std::stoi(split[0]);
69         remoteState.value = std::stoi(split[1]);
70         remoteStatePublisher.publish(remoteState);
71     }
72 }
73
74 }

```

Stm.h

```

1  #ifndef ENVIRONMENT_H
2  #define ENVIRONMENT_H
3
4  #include <nodelet/nodelet.h>
5  #include <ros/ros.h>
6  #include <logging/MessageOStream.h>
7
8  #include "NetworkingLib/Networking.h"
9  #include "VeloxProtocolLib/Connection.h"
10
11 #include "car/SetAngle.h"
12 #include "car/SetThrottle.h"
13
14 namespace car
15 {
16
17 class Stm : public nodelet::Nodelet
18 {
19 public:
20     void onInit() override;
21
22     Stm();
23
24 private:
25     ros::NodeHandle nh;
26
27     ros::Publisher odometryPublisher;
28     ros::Publisher wheelTicksPublisher;
29
30     ros::Subscriber setThrottleSubscriber;
31     ros::Subscriber setAngleSubscriber;
32
33     MessageOStream messageOStream;
34
35     networking::Networking net;
36     veloxProtocol::Connection::Ptr veloxConnection;
37
38     float throttleGain{0.0};
39
40     void readConfig();
41
42     void onStmDataReceived();
43
44     void setThrottleCallback(const SetThrottle::ConstPtr & msg);

```

```

45
46     void setAngleCallback(const SetAngle::ConstPtr & msg);
47 };
48
49 }
50 #endif

```

Stm.cpp

```

1 #include <pluginlib/class_list_macros.h>
2 #include <ros/ros.h>
3 #include <utils/Config.h>
4
5 #include "stm/Stm.h"
6 #include "car/Odometry.h"
7 #include "car/WheelTicks.h"
8
9 PLUGINLIB_EXPORT_CLASS(car::Stm, nodelet::Nodelet);
10
11 namespace car
12 {
13
14 Stm::Stm()
15     : messageOStream(nh, "Stm")
16 {}
17
18 void Stm::onInit()
19 {
20     using namespace std::chrono_literals;
21
22     messageOStream.write("onInit", "START");
23
24     readConfig();
25
26     odometryPublisher = nh.advertise<Odometry>("Odometry", 1);
27     wheelTicksPublisher = nh.advertise<WheelTicks>("WheelTicks", 1000);
28
29     setThrottleSubscriber = nh.subscribe("SetThrottle", 1, &Stm::setThrottleCallback
    ↪     , this);
30     setAngleSubscriber = nh.subscribe("SetAngle", 1, &Stm::setAngleCallback, this);
31
32     veloxConnection = veloxProtocol::Connection::create(net);
33     veloxConnection->open(
34         "/dev/ttySAC0",
35         [this]
36         { onStmDataReceived(); },
37         [this]
38         { messageOStream.write("ERROR", "UART was closed"); });
39
40     messageOStream.write("onInit", "END");
41 }
42
43 void Stm::onStmDataReceived()
44 {
45     auto stmOdometry = veloxConnection->getOdometry().get();
46     Odometry odometryMsg;
47     odometryMsg.speed = stmOdometry.speed;
48     odometryMsg.xDistance = stmOdometry.xDistance;
49     odometryMsg.yDistance = stmOdometry.yDistance;
50     odometryMsg.yawAngle = stmOdometry.yawAngle;
51     odometryMsg.steeringAngle = stmOdometry.steeringAngle;
52     odometryPublisher.publish(odometryMsg);
53
54     WheelTicks wheelTicksMsg;
55     wheelTicksMsg.rearLeftTicks = (uint32_t) stmOdometry.yawAngle;

```

```

56     wheelTicksMsg.rearRightTicks = (uint32_t) stmOdometry.steeringAngle;
57     // we treat the 'steeringAngleTimestamp' value as the delta time in us between the
    ↪ last ticks update
58     wheelTicksMsg.deltaTime = (double) stmOdometry.steeringAngleTimestamp * 1e-6;
59     wheelTicksPublisher.publish(wheelTicksMsg);
60 }
61
62 void Stm::setThrottleCallback(const SetThrottle::ConstPtr & msg)
63 {
64     // float throttle = std::max(std::min(msg->throttle * throttleGain, 1.0f), -1.0f);
65     float throttle = msg->throttle * throttleGain;
66     veloxConnection->setSpeed(throttle);
67 }
68
69 void Stm::setAngleCallback(const SetAngle::ConstPtr & msg)
70 {
71     veloxConnection->setSteeringAngle(msg->angle);
72 }
73
74 void Stm::readConfig()
75 {
76     auto j = config::open();
77     throttleGain = j.at("throttleGain");
78 }
79
80 }

```

C.1.2. Python Source Code

autonom_drive.py

```

1  #!/usr/bin/env python
2  import rospy
3  from utils import *
4  from keras.models import load_model
5  from keras import backend as K
6  from car.msg import RemoteThrottle, WheelTicks, SetAngle, SetThrottle
7  from threading import Thread, Lock
8  import cv2
9  import numpy as np
10 import signal
11 import datetime
12 from image_stream import ImageStream
13 import motion
14 import odometry
15 from parameters import *
16
17
18 class AutonomDrive:
19     def __init__(self):
20         self.running = True
21         self.contrast = 1.0
22         self.remote_throttle = 0.0
23         self.speed = 0.0
24         self.curr_angle_estimator = motion.CurrentAngleEstimator(ANGLE_CHANGE_RATE,
    ↪ confidence_duration=ANGLE_ESTIMATOR_CONFIDENCE_DURATION)
25         self.remote_throttle_sub = rospy.Subscriber("RemoteThrottle", RemoteThrottle
    ↪ , self.remote_throttle_callback, queue_size=1)
26         self.wheel_ticks_sub = rospy.Subscriber("WheelTicks", WheelTicks, self.
    ↪ wheel_ticks_callback, queue_size=1)
27         self.set_angle_pub = rospy.Publisher("SetAngle", SetAngle, queue_size=1)
28         self.set_throttle_pub = rospy.Publisher("SetThrottle", SetThrottle,
    ↪ queue_size=1)

```

```

29     self.model_path = None
30     self.read_config()
31     self.image_stream = ImageStream(self.contrast)
32     self.model = load_model(self.model_path)
33     rospy.init_node("autonom_drive", log_level=rospy.INFO)
34     signal.signal(signal.SIGINT, self.signal_handler)
35
36     def read_config(self):
37         config = open_config()
38
39         if "model" not in config:
40             raise AttributeError("attribute 'model' is missing in config.json")
41         model_filename = str(config["model"])
42         models_dir = config["modelsDirectory"] if "modelsDirectory" in config else os
43             ↪ .path.join(get_config_dir(), "models")
44         self.model_path = os.path.join(models_dir, model_filename)
45
46         if "numTensorflowThreads" in config:
47             num_threads = int(config["numTensorflowThreads"])
48             K.set_session(K.tf.Session(config=K.tf.ConfigProto(
49                 ↪ intra_op_parallelism_threads=num_threads,
50                 ↪ inter_op_parallelism_threads=num_threads)))
51
52         if "contrast" in config:
53             self.contrast = float(config["contrast"])
54
55     def signal_handler(self, sig, frame):
56         self.running = False
57         self.image_stream.stop()
58
59     def remote_throttle_callback(self, msg):
60         self.remote_throttle = msg.throttle
61
62     def wheel_ticks_callback(self, msg):
63         self.speed = odometry.speed(msg.rearLeftTicks, msg.rearRightTicks, msg.
64             ↪ deltaTime, DEFAULT_ODOMETRY_PARAMS)
65
66     def publish(self, angle, throttle):
67         set_angle_msg = SetAngle()
68         set_angle_msg.angle = angle
69         self.set_angle_pub.publish(set_angle_msg)
70
71         set_throttle_msg = SetThrottle()
72         set_throttle_msg.throttle = throttle
73         self.set_throttle_pub.publish(set_throttle_msg)
74
75     def run(self):
76         self.image_stream.start()
77         while self.running:
78             if not self.image_stream.running:
79                 break
80
81             image = self.image_stream.read()
82             image = np.expand_dims(image, axis=0)
83
84             prediction_start_time = datetime.datetime.now()
85
86             prediction = self.model.predict(image, batch_size=1)
87             coords = np.array([(float(prediction[2 * i]), float(prediction[2 * i
88                 ↪ + 1])) for i in range(int(len(prediction)/ 2))])
89
90             angle = motion.circular_regression(coords[0:2], ANGLES, WHEEL_BASE)
91             speed = 1.5 if abs(coords[-1][0]) > 0.3 else 2.5
92
93             throttle = speed * self.remote_throttle
94             self.publish(np.rad2deg(angle), throttle)

```

```

90
91         now = datetime.datetime.now()
92         prediction_duration = now - prediction_start_time
93         rospy.loginfo("\nprediction duration: %f\nframe processing duration: %f\n
94             ↪ \n", \
95                 prediction_duration.total_seconds(), self.image_stream.
96                 ↪ last_frametime)
97
98 if __name__ == "__main__":
99     try:
100         drive = AutonomDrive()
101         drive.run()
102     except rospy.ROSInterruptException:
103         pass

```

image_stream.py

```

1  import cv2
2  from threading import Thread
3  import gi
4  import numpy as np
5  import numexpr as ne
6  import datetime
7  from utils import *
8  import sys
9
10 gi.require_version("Tcam", "0.1")
11 gi.require_version("Gst", "1.0")
12
13 from gi.repository import Tcam, Gst, GLib
14
15
16 class ImageStream:
17     def __init__(self, contrast, crop_top=CROP_TOP, scale_width=IMAGE_WIDTH,
18         ↪ scale_height=IMAGE_HEIGHT, frame_rate=60):
19         Gst.init(sys.argv) # init gstreamer
20         self.contrast = contrast
21         self.crop_top = crop_top
22         self.scale_width = scale_width
23         self.scale_height = scale_height
24         self.frame_rate = frame_rate
25         self.frame = None
26         self.running = False
27         self.init_pipeline()
28         self.last_timestamp = datetime.datetime.now()
29         self.last_frametime = 0
30
31     def init_pipeline(self):
32         source = Gst.ElementFactory.make("tcambin")
33         input_caps = Gst.ElementFactory.make("capsfilter")
34         input_caps.set_property("caps", Gst.Caps.from_string("video/x-bayer,
35             ↪ framerate=" + str(self.frame_rate) + "/1"))
36         queue = Gst.ElementFactory.make("queue")
37         queue.set_property("leaky", True)
38         queue.set_property("max-size-buffers", 2)
39         bayer2rgb = Gst.ElementFactory.make("bayer2rgb")
40         convert = Gst.ElementFactory.make("videoconvert")
41         crop = Gst.ElementFactory.make("videocrop")
42         crop.set_property("top", self.crop_top)
43         balance = Gst.ElementFactory.make("videobalance")
44         balance.set_property("contrast", self.contrast)
45         scale = Gst.ElementFactory.make("videoscale")
46         output = Gst.ElementFactory.make("appsink")

```

```

45     output.set_property("caps", Gst.Caps.from_string("video/x-raw, format=RGB,
    ↪ width=" + str(self.scale_width) + ", height=" + str(self.scale_height))
    ↪ )
46     output.set_property("emit-signals", True)
47     output.connect("new-sample", self.on_new_sample)
48
49     self.pipeline = Gst.Pipeline.new()
50     self.pipeline.add(source)
51     self.pipeline.add(input_caps)
52     self.pipeline.add(queue)
53     self.pipeline.add(bayer2rgb)
54     self.pipeline.add(convert)
55     self.pipeline.add(crop)
56     self.pipeline.add(balance)
57     self.pipeline.add(scale)
58     self.pipeline.add(output)
59
60     source.link(input_caps)
61     input_caps.link(queue)
62     queue.link(bayer2rgb)
63     bayer2rgb.link(convert)
64     convert.link(crop)
65     crop.link(balance)
66     balance.link(scale)
67     scale.link(output)
68
69     self.pipeline.set_state(Gst.State.READY)
70     if self.pipeline.get_state(10 * Gst.SECOND)[0] != Gst.StateChangeReturn.
    ↪ SUCCESS:
71         raise RuntimeError("Failed to start video stream.")
72
73     def on_new_sample(self, sink):
74         sample = sink.emit("pull-sample")
75         if sample:
76             buf = sample.get_buffer()
77
78             caps = sample.get_caps()
79             width = caps.get_structure(0).get_value("width")
80             height = caps.get_structure(0).get_value("height")
81
82             try:
83                 res, mapinfo = buf.map(Gst.MapFlags.READ)
84                 img_array = np.asarray(bytearray(mapinfo.data), dtype=np.uint8)
85
86                 # Performance-critical section here!
87                 # Keep this in mind if there's any change in the preprocessing!
88                 frame = img_array.reshape((height, width, 3))
89                 frame = cv2.cvtColor(frame, cv2.COLOR_RGB2YUV)
90                 self.frame = ne.evaluate("frame / 127.5 - 1.0")
91
92                 now = datetime.datetime.now()
93                 self.last_frametime = (now - self.last_timestamp).total_seconds()
94                 self.last_timestamp = now
95
96             finally:
97                 buf.unmap(mapinfo)
98
99             return Gst.FlowReturn.OK
100
101     def start(self):
102         self.pipeline.set_state(Gst.State.PLAYING)
103         while self.frame is None:
104             pass
105         self.running = True
106
107     def stop(self):

```

```

108     self.pipeline.set_state(Gst.State.PAUSED)
109     self.pipeline.set_state(Gst.State.READY)
110     self.pipeline.set_state(Gst.State.NULL)
111     self.frame = None
112     self.running = False
113
114     def read(self):
115         return self.frame

```

motion.py

```

1  # -*- coding: utf-8 -*-
2  import numpy as np
3  import odometry
4  from scipy.spatial.distance import euclidean
5  from scipy.optimize import minimize
6  import datetime
7  from parameters import *
8
9
10 def angle_to_turn_radius(angle, wheelbase):
11     return wheelbase / np.tan(angle) if angle != 0.0 else float("inf")
12
13
14 def turn_radius_to_angle(turn_radius, wheelbase):
15     return np.arctan(wheelbase / turn_radius) if turn_radius != 0.0 else 0.0
16
17
18 def circular_regression(coords, angles, wheelbase):
19     min_err = float("inf")
20     final_angle = 0
21     for angle in angles:
22         err = 0
23         if angle == 0.0:
24             for c in coords:
25                 err += c[0] ** 2
26         else:
27             r = angle_to_turn_radius(angle, wheelbase)
28             for c in coords:
29                 err += (np.sqrt((c[0] - r) ** 2 + c[1] ** 2) - abs(r)) ** 2
30         if err < min_err:
31             min_err = err
32             final_angle = angle
33     return final_angle
34
35
36 def circular_regression_numerically(coords, angles, wheelbase):
37     def err(angle, coords):
38         if angle == 0:
39             return np.mean([c[0]**2 for c in coords])
40         r = angle_to_turn_radius(angle, WHEEL_BASE)
41         return np.mean([(np.sqrt((c[0] - r) ** 2 + c[1] ** 2) - abs(r)) ** 2 for c in
42             ↪ coords])
43
44     min_angle = np.deg2rad(-30)
45     max_angle = np.deg2rad(30)
46
47     result = minimize(err, 0, (coords), bounds=[(min_angle, max_angle)], method="L-
48         ↪ BFGS-B", options={"maxiter": 10})
49     return result["x"]
50
51 def compute_turns(coords, wheel_base):
52     coords1 = coords[0:3]
53     angle1 = circular_regression(coords1, ANGLES, wheel_base)

```



```

53
54     direction2 = odometry.direction(coords, 1)
55     odom2 = coords[1:4]
56     distance2 = odometry.distance(odom2)
57     coords2 = odometry.trajectory_coords(odom2, 0, direction2, 2, distance2
    ↪ / 2 - 0.00001)
58     angle2 = circular_regression(coords2, ANGLES, wheel_base)
59
60     return angle1, angle2
61
62
63 def compute_future_angle(curr_angle, set_angle, angle_change_rate, time):
64     angle_direction = 1 if set_angle >= curr_angle else -1
65     future_angle = curr_angle + (angle_direction * angle_change_rate * time)
66     return min(future_angle, set_angle) if angle_direction == 1 else \
67         max(future_angle, set_angle)
68
69
70 class CurrentAngleEstimator:
71     def __init__(self, turn_speed, initial_curr_angle=0.0, initial_set_angle=0.0,
    ↪ confidence_duration=0.5):
72         """
73         confidence_duration: If the angle hasn't changed much for '
    ↪ confidence_duration' seconds,
74             we can assume that we're driving the set angle.
75         """
76         self.turn_speed = turn_speed
77         self.curr_angle = initial_curr_angle
78         self.last_set_angle = initial_set_angle
79         self.last_set_timestamp = datetime.datetime.now()
80         self.last_hard_turn_timestamp = datetime.datetime.now()
81         self.confidence_duration = confidence_duration # s
82
83     def update(self, set_angle):
84         now = datetime.datetime.now()
85         last_set_angle = self.last_set_angle
86         last_set_timestamp = self.last_set_timestamp
87         self.last_set_angle = set_angle
88         self.last_set_timestamp = now
89
90         if abs(set_angle - self.curr_angle) > DELTA_ANGLE_TOLERANCE:
91             self.last_hard_turn_timestamp = now
92
93         if (now - self.last_hard_turn_timestamp).total_seconds() >= self.
    ↪ confidence_duration:
94             self.curr_angle = set_angle
95             return
96
97         self.curr_angle = compute_future_angle(self.curr_angle, last_set_angle, self.
    ↪ turn_speed, (now - last_set_timestamp).total_seconds())

```

odometry.py

```

1  # -*- coding: utf-8 -*-
2  import numpy as np
3  from scipy.spatial.distance import euclidean
4  from scipy.optimize import minimize
5  import pandas as pd
6  from matplotlib import pyplot as plt
7  from matplotlib import cm
8  from keras.models import load_model
9  import cv2
10 import utils
11 import motion
12 import datetime

```

```

13 from parameters import *
14
15
16 def normalize_vector(x):
17     n = np.linalg.norm(x)
18     return 0 if n == 0 else x / n
19
20
21 def follow(odom, idx, dist, backwards=False):
22     inc = -1 if backwards else 1
23     i0 = idx
24     i1 = i0 + inc
25     if i1 < 0 or i1 >= len(odom):
26         raise ValueError("Not enough coordinates!")
27     p0 = odom[i0]
28     p1 = odom[i1]
29     acc_dist = 0
30     next_dist = euclidean(p0, p1)
31
32     while acc_dist + next_dist <= dist + FLOAT_TOLERANCE:
33         acc_dist += euclidean(p0, p1)
34         i0 = i1
35         i1 = i1 + inc
36         if i1 < 0 or i1 >= len(odom):
37             raise ValueError("Not enough coordinates!")
38         p0 = odom[i0]
39         p1 = odom[i1]
40         next_dist = euclidean(p0, p1)
41
42     if acc_dist == dist:
43         return p0, i0
44     direction = normalize_vector(p1 - p0)
45     rest_dist = dist - acc_dist
46     v = rest_dist * direction
47     return v + p0, i0
48
49
50 def distance(odom):
51     if len(odom) < 2:
52         raise ValueError("Cannot measure the distance of odom of size < 2.")
53     dist = 0
54     for i in range(len(odom) - 1):
55         dist += euclidean(odom[i], odom[i + 1])
56     return dist
57
58
59 def plot_odometry(odom, ax=None, start_idx=0, end_idx=-1, color="gray", strength=1,
60                 ↪ elements=None):
61     x = [coord[0] for coord in odom]
62     y = [coord[1] for coord in odom]
63     if elements is None:
64         return ax.plot(x[start_idx:end_idx], y[start_idx:end_idx], "o", markersize=
65             ↪ strength, color=color)[0]
66     elements.set_data(x[start_idx:end_idx], y[start_idx:end_idx])
67
68 def plot_future(odom, idx, odom_ax=None, future_ax=None, elements=None):
69     out_elements = {}
70     p = odom[idx]
71
72     if elements is None:
73         out_elements["current_pos"] = odom_ax.plot(p[0], p[1], "o", markersize=5,
74             ↪ color="b")[0]
75         out_elements["future"] = []
76     else:
77         elements["current_pos"].set_data(p[0], p[1])

```

```

76
77 for i in range(utils.NUM_COORDS_TO_PREDICT):
78     q, _ = follow(odom, idx, (i + 1) * COORD_SPACING)
79     if elements is None:
80         out_elements["future"].append(odom_ax.plot(q[0], q[1], "o", markersize
            ↪ =5, color="r")[0])
81     else:
82         elements["future"][i].set_data(q[0], q[1])
83
84     # find the index where we start to move
85     start_of_motion = idx
86     while np.array_equal(direction(odom, start_of_motion), (0, 0)): start_of_motion
            ↪ += 1
87
88     coords = trajectory_coords(odom, idx, direction(odom, start_of_motion),
            ↪ NUM_COORDS_TO_PREDICT, COORD_SPACING)
89     x = [coord[0] for coord in coords]
90     y = [coord[1] for coord in coords]
91     angle = motion.circular_regression(coords[0:2], ANGLES, WHEEL_BASE)
92     r = motion.angle_to_turn_radius(angle, WHEEL_BASE)
93     label = "ground truth angle: {}".format(np.rad2deg(angle))
94     if elements is None:
95         out_elements["coords"] = future_ax.plot(x, y, "o", markersize=5, color="r")
            ↪ [0]
96         out_elements["zero"] = future_ax.plot(0, 0, "o", markersize=5, color="b")[0]
97         out_elements["circle"] = plot_circle(r, label, future_ax, color="r", pos
            ↪ =(0.05, 0.15))
98     return out_elements
99 else:
100     elements["coords"].set_data(x, y)
101     elements["zero"].set_data(0, 0)
102     plot_circle(r, label, elements=elements["circle"])
103
104
105 def plot_prediction(idx, model, data_dir, odom_ax=None, future_ax=None, image_ax=None
            ↪ , elements=None):
106     out_elements = {}
107
108     image = utils.load_image(data_dir, "img-" + str(idx) + ".jpg")
109     if elements is None:
110         out_elements["image"] = image_ax.imshow(mark_image(image)) if image_ax is not
            ↪ None else None
111     else:
112         if elements["image"] is not None:
113             elements["image"].set_data(mark_image(image))
114         image = utils.preprocess_image(image)
115         image = np.expand_dims(image, axis=0)
116         prediction = model.predict(image, batch_size=1)
117         coords = np.array([(float(prediction[2 * i]), float(prediction[2 * i + 1])) for i
            ↪ in range(int(len(prediction) / 2))])
118     x = [coord[0] for coord in coords]
119     y = [coord[1] for coord in coords]
120     angle = motion.circular_regression(coords[0:2], ANGLES, WHEEL_BASE)
121     r = motion.angle_to_turn_radius(angle, WHEEL_BASE)
122     label = "predicted angle: {}".format(np.rad2deg(angle))
123
124     if elements is None:
125         out_elements["future"] = future_ax.plot(x, y, "o", markersize=5, color="g")
            ↪ [0]
126         out_elements["circle"] = plot_circle(r, label, future_ax, color="g")
127     return out_elements
128 else:
129     elements["future"].set_data(x, y)
130     plot_circle(r, label, elements=elements["circle"])
131
132

```

```

133 def plot_circle(r, label, ax=None, color="b", pos=(0.05, 0.05), elements=None):
134     out_elements = {}
135
136     if elements is None:
137         circle = plt.Circle((r, 0), r, color=color, fill=False)
138         ax.add_artist(circle)
139         out_elements["circle"] = circle
140         text_props = dict(boxstyle='round', facecolor='wheat', alpha=0.5)
141         out_elements["text"] = ax.text(pos[0], pos[1], label,
142             transform=ax.transAxes, fontsize=10, verticalalignment='bottom', bbox=
                ↪ text_props)
143
144     else:
145         elements["circle"].set_radius(r)
146         elements["circle"].center = (r, 0)
147         elements["text"].set_text(label)
148
149     if elements is None:
150         return out_elements
151
152 def mark_image(image):
153     image = image.copy()
154     image[279 + 1][:] = [255, 0, 0]
155     image[279 + 0][:] = [255, 0, 0]
156     image[279 + -1][:] = [255, 0, 0]
157
158     image[306 + 1][:] = [255, 0, 0]
159     image[306 + 0][:] = [255, 0, 0]
160     image[306 + -1][:] = [255, 0, 0]
161
162     image[364 + 1][:] = [255, 0, 0]
163     image[364 + 0][:] = [255, 0, 0]
164     image[364 + -1][:] = [255, 0, 0]
165
166     # 0.5m
167     cv2.line(image, (493, 240), (744, 444), (0, 255, 0), 2)
168     cv2.line(image, (251, 240), (0, 444), (0, 255, 0), 2)
169     # 1.0m
170     cv2.line(image, (744, 326), (532, 240), (0, 255, 0), 2)
171     cv2.line(image, (0, 326), (212, 240), (0, 255, 0), 2)
172
173     cv2.line(image, (372, 240), (372, 480), (0, 255, 0), 2)
174     return image
175
176
177 def rotation_matrix(angle):
178     """ angle in radians """
179     c, s = np.cos(angle), np.sin(angle)
180     return np.array((c, -s), (s, c))
181
182
183 def direction(odom, idx):
184     if idx == 0:
185         return np.array((0, 0))
186     c = odom[idx]
187     i = idx - 1
188     d = c - odom[i]
189     while i > 0 and np.array_equal(d, (0, 0)):
190         i -= 1
191         d = c - odom[i]
192     return d
193
194
195 def trajectory_coords(odom, idx, direction, num_coords, coord_spacing):
196     if np.array_equal(direction, (0, 0)):
197         ValueError("'direction' must not be (0, 0).")

```

```

198     # handle divide by zero case
199     if direction[0] == 0.0:
200         angle = 0.0 if direction[1] >= 0 else np.pi
201     else:
202         angle = np.arctan(direction[1] / direction[0])
203         angle += 0.5 * np.pi if direction[0] < 0 else 1.5 * np.pi
204     rotation = rotation_matrix(-angle)
205     coords = np.empty([num_coords, 2])
206     for i in range(num_coords):
207         coord, _ = follow(odom, idx, coord_spacing * (i + 1))
208         coord -= odom[idx]
209         coord = np.dot(rotation, coord)
210         coords[i] = coord
211     return coords
212
213
214 class OdometryBuilder:
215     def __init__(self, resolution):
216         self.resolution = resolution
217         self.odom = np.array([(0, 0)])
218         self.p = np.array((0, 0))
219
220     def left(self, s):
221         coords = np.empty([int(s / self.resolution), 2])
222         for i in range(len(coords)):
223             coords[i][0] = self.p[0] - (i + 1) * self.resolution
224             coords[i][1] = self.p[1]
225         self.p = self.p - (s, 0)
226         if len(coords) == 0 or coords[-1][0] != self.p[0]:
227             coords = np.append(coords, [self.p], axis=0)
228         self.odom = np.append(self.odom, coords, axis=0)
229         return self
230
231     def right(self, s):
232         coords = np.empty([int(s / self.resolution), 2])
233         for i in range(len(coords)):
234             coords[i][0] = self.p[0] + (i + 1) * self.resolution
235             coords[i][1] = self.p[1]
236         self.p = self.p + (s, 0)
237         if len(coords) == 0 or coords[-1][0] != self.p[0]:
238             coords = np.append(coords, [self.p], axis=0)
239         self.odom = np.append(self.odom, coords, axis=0)
240         return self
241
242     def up(self, s):
243         coords = np.empty([int(s / self.resolution), 2])
244         for i in range(len(coords)):
245             coords[i][0] = self.p[0]
246             coords[i][1] = self.p[1] + (i + 1) * self.resolution
247         self.p = self.p + (0, s)
248         if len(coords) == 0 or coords[-1][1] != self.p[1]:
249             coords = np.append(coords, [self.p], axis=0)
250         self.odom = np.append(self.odom, coords, axis=0)
251         return self
252
253     def down(self, s):
254         coords = np.empty([int(s / self.resolution), 2])
255         for i in range(len(coords)):
256             coords[i][0] = self.p[0]
257             coords[i][1] = self.p[1] - (i + 1) * self.resolution
258         self.p = self.p - (0, s)
259         if len(coords) == 0 or coords[-1][1] != self.p[1]:
260             coords = np.append(coords, [self.p], axis=0)
261         self.odom = np.append(self.odom, coords, axis=0)
262         return self
263

```

```

264
265 def reference_odometry_cw(resolution=0.1):
266     builder = OdometryBuilder(resolution)
267     builder.right(3).down(3).left(3).up(3)
268     return builder.odom
269
270
271 def reference_odometry_ccw(resolution=0.1):
272     builder = OdometryBuilder(resolution)
273     builder.right(3).up(3).left(3).down(3)
274     return builder.odom
275
276
277 def load_data(csv):
278     data_df = pd.read_csv(csv, usecols=["x", "y", "ticksRL", "ticksRR"])
279     odom = data_df[["x", "y"]].values.astype(np.float64)
280     ticks = data_df[["ticksRL", "ticksRR"]].values.astype(np.float64)
281     return odom, ticks
282
283
284 def compress_ticks(ticks, steps):
285     n = int(len(ticks) / steps)
286     compressed_ticks = np.empty([n, 2])
287     for i in range(n):
288         dts_left = 0
289         dts_right = 0
290         for j in range(steps):
291             dts_left += ticks[i * steps + j][0]
292             dts_right += ticks[i * steps + j][1]
293         compressed_ticks[i][0] = dts_left
294         compressed_ticks[i][1] = dts_right
295     return compressed_ticks
296
297
298 def ticks_to_odometry(ticks, params):
299     odom = np.empty([len(ticks), 2])
300     x = 0
301     y = 0
302     yaw = 0
303     for i, item in enumerate(ticks):
304         ticks_rl, ticks_rr = item
305         dts_left = ticks_rl * params.dts_tick_left
306         dts_right = ticks_rr * params.dts_tick_right
307         dts = (dts_left + dts_right) / 2.0
308         dt_yaw = (dts_right - dts_left) / params.track_width
309         x += np.cos(yaw) * dts * 0.001
310         y += np.sin(yaw) * dts * 0.001
311         yaw += dt_yaw
312         odom[i][0] = x
313         odom[i][1] = y
314     return odom
315
316
317 def speed(ticks_left, ticks_right, delta_time, params):
318     dts_left = ticks_left * params.dts_tick_left
319     dts_right = ticks_right * params.dts_tick_right
320     dts = (dts_left + dts_right) / 2.0
321     return (dts * 0.001) / delta_time
322
323
324 def mean_squared_distance_error(odom1, odom2, resolution, ax=None):
325     dist1 = distance(odom1)
326     dist2 = distance(odom2)
327     dist = min(dist1, dist2)
328     n = int(dist / resolution) - 1
329     #err = (dist1 - dist2) ** 2

```

```

330     err = 0
331     for i in range(1, n):
332         p1, _ = follow(odom1, 0, i * resolution)
333         p2, _ = follow(odom2, 0, i * resolution)
334         if ax is not None:
335             ax.plot([p1[0], p2[0]], [p1[1], p2[1]], color="black")
336         err += euclidean(p1, p2) ** 2
337     return err / n
338
339
340 def reference_error(x, reference_odom_to_ticks_list, resolution):
341     """
342     reference_odom_to_ticks_list:
343     [(reference_odom_1, [ticks_1, ticks_2, ticks_3]),
344      (reference_odom_2, [ticks_1, ticks_2, ticks_3]),
345      ...
346      (reference_odom_n, [ticks_1, ticks_2, ticks_3])]
347     """
348     params = OdometryParams(x[0], x[1], x[2])
349     err = 0
350     n = 0
351     for item in reference_odom_to_ticks_list:
352         reference_odom = item[0]
353         ticks_list = item[1]
354         n += len(ticks_list)
355         for ticks in ticks_list:
356             odom_reconstructed = ticks_to_odometry(ticks, params)
357             err += mean_squared_distance_error(reference_odom, odom_reconstructed,
358                                               ↪ resolution)
359
360     return err / n
361
362 def mean_center_error(x, cw, ccw):
363     params = OdometryParams(x[0], x[1], x[2])
364     odoms_cw = [ticks_to_odometry(t, params) for t in cw]
365     odoms_ccw = [ticks_to_odometry(t, params) for t in ccw]
366     center_err_cw = np.linalg.norm(mean_center_of_gravity(odoms_cw))
367     center_err_ccw = np.linalg.norm(mean_center_of_gravity(odoms_ccw))
368     distance_err = np.mean([abs(12 - distance(odom)) for odom in np.append(odoms_cw,
369     ↪ odoms_ccw, axis=0)])
370     print("center_err: " + str(max(center_err_cw, center_err_ccw)))
371     print("distance_err: " + str(distance_err))
372     return max(center_err_cw, center_err_ccw) + distance_err
373
374 def mean_offset_error(x, cw, ccw):
375     err = 0
376     params = OdometryParams(x[0], x[1], x[2])
377     odoms_cw = [ticks_to_odometry(t, params) for t in cw]
378     odoms_ccw = [ticks_to_odometry(t, params) for t in ccw]
379
380     err += np.mean([(odom[int(0.175 * len(odom))][1] - 0.0)**2 for odom in odoms_cw])
381     err += np.mean([(odom[int(0.175 * len(odom))][1] - 0.0)**2 for odom in odoms_ccw
382     ↪ ])
383
384     err += np.mean([(odom[int(0.375 * len(odom))][0] - 3.0)**2 for odom in odoms_cw])
385     err += np.mean([(odom[int(0.375 * len(odom))][0] - 3.0)**2 for odom in odoms_ccw
386     ↪ ])
387
388     err += np.mean([(odom[int(0.55 * len(odom))][1] - (-3.0))**2 for odom in odoms_cw
389     ↪ ])
390     err += np.mean([(odom[int(0.55 * len(odom))][1] - 3.0)**2 for odom in odoms_ccw])
391
392     err += np.mean([(odom[int(0.75 * len(odom))][0] - 0.0)**2 for odom in odoms_cw])
393     err += np.mean([(odom[int(0.75 * len(odom))][0] - 0.0)**2 for odom in odoms_ccw])

```

```

391
392     err += np.mean([np.linalg.norm(odom[-1])**2 for odom in odoms_cw])
393     err += np.mean([np.linalg.norm(odom[-1])**2 for odom in odoms_ccw])
394
395     return err / 10.0
396
397
398 def mean_center_of_gravity(odoms):
399     cog = np.array([odom[-1] for odom in odoms])
400     x = np.mean(cog[:,0])
401     y = np.mean(cog[:,1])
402     return np.array((x, y))
403
404
405 def manual_calibration(fig, ax, ticks_list, params):
406     isarray = isinstance(ax, (list, tuple, np.ndarray))
407     lines_list = []
408     mean_distance = 0
409     for i, ticks in enumerate(ticks_list):
410         odom = ticks_to_odometry(ticks, params)
411         mean_distance += distance(odom)
412         axis = ax[i] if isarray else ax
413         lines, = axis.plot([o[0] for o in odom], [o[1] for o in odom], "ro",
414                             ↪ markersize=1)
414         lines_list.append(lines)
415     mean_distance /= len(ticks_list)
416
417     text_props = dict(boxstyle='round', facecolor='wheat', alpha=0.5)
418     label = "dts left: {} \ndts right: {} \ntrack width: {} \nmean distance: {}".format(
419         params.dts_tick_left, params.dts_tick_right, params.track_width,
420         ↪ mean_distance)
420     axis = ax[0] if isarray else ax
421     text = axis.text(0.05, 0.95, label, transform=axis.transAxes, fontsize=14,
422                     verticalalignment='top', bbox=text_props)
423
424     def event_handler(event):
425         dts_step = 0.005
426         track_step = 0.5
427         if event.key == "t":
428             params.dts_tick_left += dts_step
429         elif event.key == "g":
430             params.dts_tick_left -= dts_step
431         elif event.key == "i":
432             params.dts_tick_right += dts_step
433         elif event.key == "k":
434             params.dts_tick_right -= dts_step
435         elif event.key == "ü":
436             params.track_width += track_step
437         elif event.key == "ä":
438             params.track_width -= track_step
439         mean_distance = 0
440         for i, ticks in enumerate(ticks_list):
441             odom = ticks_to_odometry(ticks, params)
442             mean_distance += distance(odom)
443             lines_list[i].set_data([o[0] for o in odom], [o[1] for o in odom])
444         mean_distance /= len(ticks_list)
445         text.set_text("dts left: {} \ndts right: {} \ntrack width: {} \nmean distance
446             ↪ : {}".format(
446             ↪ params.dts_tick_left, params.dts_tick_right, params.track_width,
447             ↪ mean_distance))
447
448     fig.canvas.mpl_connect("key_press_event", event_handler)
449     plt.ion()
450     while True:
451         plt.pause(0.1)
452

```



```

453
454 def UMBmark_optimize(b, D_L, D_R, dts_L, dts_R, cog_cw, cog_ccw, L):
455     """
456     b: wheel base: distance between the two wheels
457     D_L: diameter of left wheel
458     D_R: diameter of right wheel
459     dts_L: left wheel travel per encoder pulse
460     dts_R: right wheel travel per encoder pulse
461     cog_cw: center of gravity in clockwise direction
462     cog_ccw: center of gravity in counter-clockwise direction
463     L: length of the square the robot has driven
464     """
465     alpha = (cog_cw[0] + cog_ccw[0]) / (-4 * L)
466     beta = (cog_cw[0] - cog_ccw[0]) / (-4 * L)
467
468     R = (L / 2) / (np.sin(beta / 2))
469     Ed = (R + (b / 2)) / (R - (b / 2))
470     D_a = (D_R + D_L) / 2
471     D_L = (2 / (Ed + 1)) * D_a
472     D_R = (2 / ((1 / Ed) + 1)) * D_a
473     c_L = 2 / (Ed + 1)
474     c_R = 2 / ((1 / Ed) + 1)
475     dts_L *= c_L
476     dts_R *= c_R
477
478     E_b = (np.pi / 2) / ((np.pi / 2) - alpha)
479     b *= E_b
480
481     return b, D_L, D_R, dts_L, dts_R
482
483
484 def load_square_ticks(compression_steps=5):
485     cw = [compress_ticks(load_data("/media/philipp/Transcend/odometryCalibration/cw/"
486     ↪ + str(i) + ".csv")[1], compression_steps) for i in range(5)]
487     ccw = [compress_ticks(load_data("/media/philipp/Transcend/odometryCalibration/ccw
488     ↪ /" + str(i) + ".csv")[1], compression_steps) for i in range(5)]
489
490     cw = [cw[i] for i in [2, 3]]
491     ccw = [ccw[i] for i in [0, 1, 2, 4]]
492
493     return cw, ccw

```

parameters.py

```

1  # -*- coding: utf-8 -*-
2  import numpy as np
3
4
5  class OdometryParams:
6      def __init__(self, dts_tick_left, dts_tick_right, track_width):
7          """ all units are given in millimeters """
8          self.dts_tick_left = dts_tick_left #
9          self.dts_tick_right = dts_tick_right
10         self.track_width = track_width
11
12         def __str__(self):
13             return "dts_tick_left: {} mm\ndts_tick_right: {} mm\ntrack_width: {} mm".
14             ↪ format(
15                 self.dts_tick_left, self.dts_tick_right, self.track_width)
16
17     ANGLES = [np.deg2rad(deg) for deg in list(range(-30, 31))]
18     LOOK_AHEAD_DISTANCE = 2 # m
19     GRAVITY_ACC = 9.81 # m / s^2
20     WHEEL_BASE = 0.32 # m

```

```

21 #WHEEL_BASE = 0.2304363
22 # the wheels turn from full left (-30) to full right (30) in approx. 0.5 seconds
23 ANGLE_CHANGE_RATE = np.deg2rad(20) # rad / s
24 MAX_ANGLE_DELAY_DISTANCE = 0.2 # m
25 # defines a delta angle tolerance for comparing two angles:
26 # angle_1 ~ angle_2 if abs(angle_1 - angle_2) <= DELTA_ANGLE_TOLERANCE
27 DELTA_ANGLE_TOLERANCE = np.deg2rad(1) # rad
28 STRAIGHT_ANGLE_TOLERANCE = np.deg2rad(10)
29 ANGLE_ESTIMATOR_CONFIDENCE_DURATION = 0.5 # s
30 STATIC_FRICTION_COEF = 1.0
31
32 IMAGE_HEIGHT, IMAGE_WIDTH, IMAGE_CHANNELS = 100, 200, 3
33 INPUT_SHAPE = (IMAGE_HEIGHT, IMAGE_WIDTH, IMAGE_CHANNELS)
34 COORD_SPACING = 0.6 # m
35 NUM_COORDS_TO_PREDICT = int(LOOK_AHEAD_DISTANCE / COORD_SPACING)
36 CROP_TOP = 200
37
38 FLOAT_TOLERANCE = 10e-6
39 DISC_RESOLUTION = 120
40 WHEEL_DIAMETER = 0.11 # m
41 # original
42 # DTS_TICK_LEFT = (np.pi * WHEEL_DIAMETER / DISC_RESOLUTION) * 1000.0 # mm
43 # DTS_TICK_RIGHT = (np.pi * WHEEL_DIAMETER / DISC_RESOLUTION) * 1000.0 # mm
44 # TRACK_WIDTH = 275 # mm
45 # end up using
46 DTS_TICK_LEFT = (np.pi * WHEEL_DIAMETER / DISC_RESOLUTION) * 1000.0 - 0.01 + 0.3 # mm
47 DTS_TICK_RIGHT = (np.pi * WHEEL_DIAMETER / DISC_RESOLUTION) * 1000.0 + 0.01 + 0.3 #
    ↪ mm
48 TRACK_WIDTH = 305 # mm
49
50 ERROR_RESOLUTION = 1 # m
51 SHOULD_OPTIMIZE = True
52 OPTIMIZATION_METHOD = "BFGS"
53 MAX_ITERATIONS = 100
54 SQUARE_LENGTH = 3 # m
55 SQUARE_DISTANCE = 4 * SQUARE_LENGTH # m
56 SQUARE_DISTANCE_TOLERANCE = 1
57 DEFAULT_ODOMETRY_PARAMS = OdometryParams(DTS_TICK_LEFT, DTS_TICK_RIGHT, TRACK_WIDTH)

```

saliency.py

```

1 from vis.visualization import visualize_saliency, visualize_cam, overlay
2 from keras import activations
3 from utils import *
4 from vis.utils import utils
5 from keras.models import load_model
6 from matplotlib import pyplot as plt
7 import numpy as np
8 import cv2
9 from keras.models import Model
10 from matplotlib import cm
11
12 # ----
13 # hella images
14 # ---
15 images = [
16     load_image("I:/recordings/data-2018-08-31-12-12-14", "img-300.jpg"), # left,
17     ↪ ideal line
18     load_image("I:/recordings/data-2018-08-31-12-12-14", "img-6175.jpg"), # right,
19     ↪ ideal line
20     load_image("I:/recordings/data-2018-08-31-12-12-14", "img-351.jpg"), # right,
21     ↪ ideal line
22     load_image("I:/recordings/data-2018-08-31-10-42-27", "img-10937.jpg"), # right,
23     ↪ center line

```

```

21     load_image("I:/recordings/data-2018-08-31-10-42-27", "img-625.jpg"), # straight,
        ↪ center line
22     load_image("I:/recordings/data-2018-08-31-10-42-27", "img-9880.jpg"), # right,
        ↪ center line
23
24     # afternoon
25     load_image("I:/recordings/data-2018-09-07-15-59-32", "img-180.jpg"),
26     load_image("I:/recordings/data-2018-09-07-15-59-32", "img-820.jpg"),
27     load_image("I:/recordings/data-2018-09-07-15-59-32", "img-1110.jpg"),
28
29     load_image("I:/recordings/data-2018-09-07-15-59-32", "img-10450.jpg"),
30     load_image("I:/recordings/data-2018-09-07-15-59-32", "img-11300.jpg"),
31     load_image("I:/recordings/data-2018-09-07-15-42-32", "img-390.jpg"),
32 ]
33
34 # ----
35 # flur images
36 # ---
37 # images = [
38 #     load_image("I:/recordings/data-2018-09-08-17-19-41", "img-540.jpg"),
39 #     load_image("I:/recordings/data-2018-09-08-17-19-41", "img-610.jpg"),
40 #     load_image("I:/recordings/data-2018-09-08-17-19-41", "img-660.jpg"),
41 #     load_image("I:/recordings/data-2018-09-08-17-19-41", "img-1460.jpg"),
42 #     load_image("I:/recordings/data-2018-09-08-17-19-41", "img-1515.jpg"),
43 #     load_image("I:/recordings/data-2018-09-08-17-19-41", "img-2090.jpg"),
44 # ]
45
46 model = load_model("G:/car/models/model-040.h5")
47
48 fig, axes = plt.subplots(int(len(images) / 3), 3)
49
50 layer_idx = 6
51 model.layers[layer_idx].activation = activations.linear
52 model = utils.apply_modifications(model)
53 for i, image in enumerate(images):
54     image = preprocess_image(image)
55     grads = visualize_saliency(model, layer_idx, filter_indices=None, seed_input=
        ↪ image, backprop_modifier="guided")
56     jet_grads = (np.delete(cm.jet(grads), 3, 2) * 255.0).astype(np.uint8)
57
58     image = (image + 1.0) * 127.5
59     image = image.astype(np.uint8)
60     image = cv2.cvtColor(image, cv2.COLOR_YUV2RGB)
61
62     axes[int(i/3), int(i%3)].imshow(overlay(image, jet_grads, alpha=0.3))
63
64 plt.show()

```

train.py

```

1  import numpy as np
2  np.random.seed(0)
3  import pandas as pd
4  from sklearn.cross_validation import train_test_split
5  from keras.models import Sequential, Model
6  import keras.regularizers
7  from keras.optimizers import Adam
8  from keras.callbacks import ModelCheckpoint, TensorBoard, CSVLogger
9  from keras.layers import Lambda, Convolution2D, MaxPooling2D, Dropout, Dense, Flatten
        ↪ , Input, BatchNormalization
10 from keras import regularizers
11 from utils import *
12 import argparse
13 import os
14 from time import time

```

```

15 import odometry
16 import json
17 import merge
18
19
20 def load_data(args):
21     X = y = None
22
23     with open(args.training_config, "r") as f:
24         config = json.load(f)
25         recordings = config["recordings"]
26         for recording in recordings:
27             directory = recording["directory"]
28             csv = os.path.join(directory, "data.csv")
29             data_df = pd.read_csv(csv)
30             contains_subsequences = "recording" in data_df
31
32             image_paths_recording = [os.path.join(directory, img.strip()) for img in
33                                     ↪ data_df["img"].values]
34
35             odom_params = read_odometry_params(
36                 recording["odometryParameters"]) if "odometryParameters" in recording
37                 ↪ \
38                                     else DEFAULT_ODOMETRY_PARAMS
39
40             ticks = data_df[["ticksRL", "ticksRR"]].values.astype(np.float64)
41             odom_recording = odometry.ticks_to_odometry(ticks, odom_params)
42
43             for sequence in merge.sequences_from_json(recording["sequences"]):
44                 sequence_start = sequence[0]
45                 sequence_end = sequence[1]
46
47                 subsequences = get_subsequences(data_df, sequence_start, sequence_end
48                 ↪ ) \
49                                     if contains_subsequences \
50                                     else [(sequence_start, sequence_end)]
51
52                 for subsequence in subsequences:
53                     subsequence_start = subsequence[0]
54                     subsequence_end = subsequence[1]
55                     subsequence_odom = np.array([odom_recording[i] for i in range(
56                 ↪ subsequence_start, subsequence_end + 1)])
57                     try: subsequence_coords_list, first_idx, last_idx =
58                 ↪ get_coords_list(subsequence_odom)
59                     except ValueError: continue
60                     subsequence_image_paths = np.array([image_paths_recording[i]
61                 ↪ for i in range(subsequence_start +
62                 ↪ first_idx, subsequence_start +
63                 ↪ last_idx + 1)])
64
65                     X = np.append(X, subsequence_image_paths, axis=0) if X is not
66                 ↪ None else np.array(subsequence_image_paths)
67                     y = np.append(y, subsequence_coords_list, axis=0) if y is not
68                 ↪ None else np.array(subsequence_coords_list)
69
70             X_train, X_valid, y_train, y_valid = train_test_split(X, y, test_size=args.
71                 ↪ test_size, random_state=1)
72             return X_train, X_valid, y_train, y_valid
73
74
75 def get_coords_list(odom):
76     # from every coordinate in the training set we must be able
77     # to travel a distance of (NUM_COORDS_TO_PREDICT * COORD_SPACING) into the future
78     _, last_idx = odometry.follow(odom, len(odom) - 1, NUM_COORDS_TO_PREDICT *
79     ↪ COORD_SPACING, backwards=True)
80     last_idx -= 1

```

```

70     if last_idx < 0:
71         raise ValueError("Not enough coordinates!")
72     # find the index where we start to move + 1
73     # which is where direction() does not return (0, 0)
74     first_idx = 0
75     while np.array_equal(odometry.direction(odom, first_idx), (0, 0)): first_idx += 1
76     # get list of coordinates to predict
77     coords_list = [
78         odometry.trajectory_coords(odom, i, odometry.direction(odom, i),
79             ↪ NUM_COORDS_TO_PREDICT, COORD_SPACING)
80         for i in range(first_idx, last_idx + 1)]
81     return coords_list, first_idx, last_idx
82
83 def get_subsequences(data_df, section_start, section_end):
84     subsequences = []
85     subsequence_start = None
86     recording = False
87     for i in range(section_start, section_end + 1):
88         recording = bool(data_df["recording"][i])
89         if recording:
90             if subsequence_start is None:
91                 subsequence_start = i
92             elif subsequence_start is not None:
93                 subsequences.append((subsequence_start, i - 1))
94                 subsequence_start = None
95     if subsequence_start is not None and recording:
96         subsequences.append((subsequence_start, i))
97     return subsequences
98
99
100 def read_odometry_params(jparams):
101     dts_tick_left_mm = float(jparams["dts_tick_left_mm"]) if "dts_tick_left_mm" in
102         ↪ jparams else odometry.DEFAULT_ODOMETRY_PARAMS.dts_tick_left_mm
103     dts_tick_right_mm = float(jparams["dts_tick_right_mm"]) if "dts_tick_right_mm" in
104         ↪ jparams else odometry.DEFAULT_ODOMETRY_PARAMS.dts_tick_right_mm
105     track_width_mm = float(jparams["track_width_mm"]) if "track_width_mm" in jparams
106         ↪ else odometry.DEFAULT_ODOMETRY_PARAMS.track_width_mm
107     return odometry.OdometryParams(dts_tick_left_mm, dts_tick_right_mm,
108         ↪ track_width_mm)
109
110
111 def build_model_nvidia(args):
112     img_in = Input(shape=INPUT_SHAPE, name='img_in')
113
114     x = img_in
115     x = BatchNormalization(axis=3)(x)
116     x = Convolution2D(24, (5, 5), strides=(2, 2), activation='relu')(x)
117     x = Convolution2D(36, (5, 5), strides=(2, 2), activation='relu')(x)
118     x = Convolution2D(48, (5, 5), strides=(2, 2), activation='relu')(x)
119     x = Convolution2D(64, (3, 3), strides=(1, 1), activation='relu')(x)
120     x = Convolution2D(64, (3, 3), strides=(1, 1), activation='relu')(x)
121     x = Flatten()(x)
122     x = Dense(100, activation="relu")(x)
123     x = Dense(50, activation="relu")(x)
124     x = Dense(10, activation="relu")(x)
125
126     loss = {}
127     outputs = []
128     for i in range(NUM_COORDS_TO_PREDICT):
129         name_x = "coord_x_out_" + str(i)
130         name_y = "coord_y_out_" + str(i)
131         coord_x_out = Dense(1, name=name_x, activation="linear")(x)
132         coord_y_out = Dense(1, name=name_y, activation="linear")(x)
133         loss[name_x] = "mean_squared_error"
134         loss[name_y] = "mean_squared_error"

```

```

131         outputs.append(coord_x_out)
132         outputs.append(coord_y_out)
133
134     model = Model(inputs=[img_in], outputs=outputs)
135     model.compile(optimizer="adam", loss=loss)
136     return model
137
138
139 def train_model(model, args, X_train, X_valid, y_train, y_valid):
140     checkpoint = ModelCheckpoint(get_config_dir() + '/models/model-{epoch:03d}.h5',
141                                 monitor='val_loss',
142                                 verbose=0,
143                                 save_best_only=args.save_best_only,
144                                 mode='auto')
145
146     tensorboard = TensorBoard(log_dir=get_config_dir() + "/logs/tensorboard/{_}_{_}".
147                               ↪ format(args.log_label, time()))
148     csvLogger = CSVLogger(get_config_dir() + "/logs/csv/{_}_{_}.csv".format(args.
149                               ↪ log_label, time()), separator=',', append=False)
150
151     model.compile(loss='mean_squared_error', optimizer=Adam(lr=args.learning_rate))
152
153     cache_capacity = min(int(args.image_memory_size / (IMAGE_HEIGHT * IMAGE_WIDTH *
154                               ↪ IMAGE_CHANNELS * 8)), len(X_train) + len(X_valid))
155     print("Cache capacity: {} images".format(cache_capacity))
156     image_cache = ImageCache(cache_capacity)
157     print("Buffering images...")
158     image_cache.load_images(np.append(X_train, X_valid, axis=0))
159     print("Completed!")
160
161     train_sequence = BatchSequence(args.data_dir, X_train, y_train, args.batch_size,
162                               ↪ True, image_cache)
163     valid_sequence = BatchSequence(args.data_dir, X_valid, y_valid, args.batch_size,
164                               ↪ False, image_cache)
165
166     model.fit_generator(train_sequence,
167                        None,
168                        args.nb_epoch,
169                        max_queue_size=args.max_queue_size,
170                        validation_data=valid_sequence,
171                        validation_steps=None,
172                        callbacks=[checkpoint, tensorboard, csvLogger],
173                        verbose=1,
174                        workers=args.workers,
175                        shuffle=False)
176
177 def s2b(s):
178     s = s.lower()
179     return s == 'true' or s == 'yes' or s == 'y' or s == '1'
180
181 def main():
182     default_training_dir = get_config_dir() + "/trainingData"
183     parser = argparse.ArgumentParser(description='Behavioral Cloning Training Program
184                               ↪ ')
185     parser.add_argument('training_config', help='training config file', type=str)
186     parser.add_argument('-d', help='data directory', dest='data_dir',
187                               ↪ type=str, default=default_training_dir)
188     parser.add_argument('-t', help='test size fraction', dest='test_size',
189                               ↪ type=float, default=0.2)
190     parser.add_argument('-k', help='drop out probability', dest='keep_prob',
191                               ↪ type=float, default=0.5)
192     parser.add_argument('-n', help='number of epochs', dest='nb_epoch',
193                               ↪ type=int, default=50)
194     parser.add_argument('-b', help='batch size', dest='batch_size',
195                               ↪ type=int, default=40)

```

```

186     parser.add_argument('-o', help='save best models only', dest='save_best_only',
187         ↪     type=s2b, default='false')
187     parser.add_argument('-l', help='learning rate', dest='learning_rate',
188         ↪     type=float, default=1.0e-4)
188     parser.add_argument('-x', help='tensorboard log label', dest='log_label',
189         ↪     type=str, default='log')
189     parser.add_argument('-w', help='number of workers', dest='workers',
190         ↪     type=int, default=8)
190     parser.add_argument('-q', help='max queue size ', dest='max_queue_size',
191         ↪     type=int, default=1)
191     parser.add_argument('-m', help='image memory size', dest='image_memory_size',
192         ↪     , type=int, default=12e9)
192     args = parser.parse_args()
193
194     print('-' * 30)
195     print('Parameters')
196     print('-' * 30)
197     for key, value in vars(args).items():
198         print('{:<20} := {}'.format(key, value))
199     print('-' * 30)
200
201     print("Loading data...")
202     data = load_data(args)
203     print("Completed!")
204
205     model = build_model_nvidia(args)
206     train_model(model, args, *data)
207
208
209 if __name__ == '__main__':
210     main()

```

utils.py

```

1  import cv2, os
2  import numpy as np
3  import numexpr as ne
4  import matplotlib.image as mpimg
5  from os.path import expanduser
6  import json
7  import keras
8  from threading import Thread, Condition, Lock
9  import sys
10 is_py2 = sys.version[0] == '2'
11 if is_py2:
12     from Queue import Queue, Empty
13 else:
14     from queue import Queue, Empty
15 from parameters import *
16
17
18 def get_config_dir():
19     if os.name == "nt":
20         return "G:/car"
21     home = expanduser("~")
22     return home + "/.car"
23
24
25 def open_config():
26     f = open(get_config_dir() + "/config.json", "r")
27     return json.load(f)
28
29
30 def load_image_from_path(path):
31     return cv2.imread(path)

```

```

32
33
34 def load_image(data_dir, image_file):
35     return load_image_from_path(os.path.join(data_dir, image_file.strip()))
36
37
38 def crop(image):
39     return image[CROP_TOP:, :, :]
40
41
42 def resize(image):
43     return cv2.resize(image, (IMAGE_WIDTH, IMAGE_HEIGHT), cv2.INTER_AREA)
44
45
46 def normalize(image):
47     return ne.evaluate("image / 127.5 - 1.0")
48
49
50 def convert(image):
51     image = image.astype(np.uint8)
52     return cv2.cvtColor(image, cv2.COLOR_BGR2YUV)
53
54
55 def bgr2yuv(image):
56     return cv2.cvtColor(image, cv2.COLOR_BGR2YUV)
57
58
59 def random_flip(image, coords):
60     if np.random.rand() < 0.5:
61         image = cv2.flip(image, 1)
62         for i, coord in enumerate(coords):
63             coords[i][0] = -coord[0]
64     return image, coords
65
66
67 def apply_contrast(image, ratio):
68     image = image.astype(np.float64)
69     image = np.minimum(image * ratio, 255)
70     return image.astype(np.uint8)
71
72
73 def apply_brightness(image, ratio):
74     image = image.astype(np.uint8)
75     hsv = cv2.cvtColor(image, cv2.COLOR_BGR2HSV)
76     hsv = hsv.astype(np.float64)
77     hsv[:, :, 2] = np.minimum(hsv[:, :, 2] * ratio, 255)
78     hsv = hsv.astype(np.uint8)
79     return cv2.cvtColor(hsv, cv2.COLOR_HSV2BGR)
80
81
82 def random_brightness(image):
83     # HSV (Hue, Saturation, Value) is also called HSB ('B' for Brightness).
84     ratio = np.random.uniform(0.5, 2.0, 1)[0]
85     return apply_brightness(image, ratio)
86
87
88 def augument(image, coords):
89     image, coords = random_flip(image, coords)
90     image = random_brightness(image)
91     return image, coords
92
93
94 def preprocess_image(image, do_crop=True, do_resize=True, do_convert=True,
95     ↪ do_normalize=True):
96     if do_crop:
97         image = crop(image)

```



```

97     if do_resize:
98         image = resize(image)
99     if do_convert:
100        image = convert(image)
101     if do_normalize:
102        image = normalize(image)
103     return image
104
105
106 def preprocess_angle(angle):
107     return angle / 30
108
109
110 def preprocess_throttle(throttle):
111     return throttle
112
113
114 def postprocess_angle(angle):
115     return angle * 30
116
117
118 def postprocess_throttle(throttle):
119     return throttle
120
121
122 class BatchSequence(keras.utils.Sequence):
123     def __init__(self, data_dir, X, y, batch_size, is_training, image_cache):
124         self.data_dir = data_dir
125         self.X = X
126         self.y = y
127         self.batch_size = batch_size
128         self.is_training = is_training
129         self.sequence = self.next_sequence()
130         self.image_cache = image_cache
131
132     def __len__(self):
133         return int(np.floor(len(self.X) / float(self.batch_size)))
134
135     def __getitem__(self, idx):
136         inputs = np.empty([self.batch_size, IMAGE_HEIGHT, IMAGE_WIDTH, IMAGE_CHANNELS
137             ↪ ])
138
139         outputs = [np.empty(self.batch_size) for i in range(2 * NUM_COORDS_TO_PREDICT
140             ↪ )]
141
142         for batch_idx in range(self.batch_size):
143             seq_idx = (idx * self.batch_size + batch_idx) % len(self.sequence)
144             sample_idx = self.sequence[seq_idx]
145             image_path = self.X[sample_idx]
146             image = self.image_cache.get_image(image_path)
147             coords = self.y[sample_idx].copy()
148
149             # augmentation
150             if self.is_training and np.random.rand() < 0.6:
151                 image, coords = augument(image, coords)
152
153             inputs[batch_idx] = normalize(convert(image))
154             for i, coord in enumerate(coords):
155                 outputs[2 * i][batch_idx] = coord[0]
156                 outputs[2 * i + 1][batch_idx] = coord[1]
157
158         return inputs, outputs
159
160     def on_epoch_end(self):
161         self.sequence = self.next_sequence()
162
163     def next_sequence(self):

```

```

161         return np.random.permutation(len(self.X))
162
163
164 class ImageCache:
165     def __init__(self, capacity):
166         self.images = np.empty([capacity, IMAGE_HEIGHT, IMAGE_WIDTH, IMAGE_CHANNELS])
167         self.path_to_idx = {}
168
169     def load_images(self, image_paths):
170         for i in range(min(len(self.images), len(image_paths))):
171             image_path = image_paths[i]
172             self.images[i] = self.load_image(image_path)
173             self.path_to_idx[image_path] = i
174
175     def load_image(self, image_path):
176         image = load_image_from_path(image_path)
177         return preprocess_image(image, do_convert=False, do_normalize=False)
178
179     def get_image(self, image_path):
180         if image_path in self.path_to_idx:
181             return self.images[self.path_to_idx[image_path]]
182         return self.load_image(image_path)

```

C.2. STM32

C.2.1. C Source Code

Note that the following code was originally delivered by Assystem GmbH and includes modifications from us.

mavlink.h

```

1
2 #ifndef MAVLINK_H_
3 #define MAVLINK_H_
4
5 ***** Includes
6     ↪ *****
7 #include <stdlib.h>
8 #include <stdbool.h>
9 #include "per/irq.h"
10 #include "sys/util/io.h"
11 #include "per/uart.h"
12 #include "sys/sigp/sigp.h"
13 #include "sys/systime/systime.h"
14 #include "protocol/velox/mavlink.h"
15 #include "sys/errorhandler/errorHandler.h"
16
17 ***** Public typedefs
18     ↪ *****
19 ***** Macros and constants
20     ↪ *****
21 #define MAVLINK_BAUDRATE 115200
22 #define MAX_BUFFER_LEN 1000
23 ***** Public function prototypes
24     ↪ *****
25 /**

```

```

25  * @brief The init function needs to be called during the initialization of the
      ↪ system.
26  * It sets all local variables to their initial values.
27  */
28  int8_t mavlink_init(void);
29
30  /**
31  * @brief read messages from uart buffer and interpret them
32  */
33  void mavlink_getMessages(void);
34
35  /**
36  * @brief pack and send all messages that are marked as active
37  */
38  void mavlink_sendMessages(void);
39
40  /**
41  * @brief activate outgoing messages that are linked to the specified state
42  */
43  void mavlink_activateStateMessages(uint8_t state);
44
45  /**
46  * @brief deactivate outgoing messages that are linked to the specified state
47  */
48  void mavlink_deactivateStateMessages(uint8_t state);
49
50  #endif /* MAVLINK_H_ */

```

mavlink.c

```

1  /*****
      ↪
2  * Unit in charge:
3  * @file odom.c
4  * @author mbrieske
5  * $Date:: 2017-05-22 17:17:17 #$
6  *
7  * -----
8  * SVN information of last commit:
9  * $Rev:: 1538 $
10 * $Author:: mbrieske $
11 *
12 * -----
13 *
14 * @brief MAVLink see mavlink.h
15 *
16 *****/
      ↪ */
17
18 /***** Includes
      ↪ *****/
19 #include "../mavlink.h"
20 #include "per/eict.h"
21 #include "app/leddebug.h"
22
23 /***** Private typedefs
      ↪ *****/
24
25 #define NUM_MSGS (sizeof mavlinkMessageInfo / sizeof mavlinkMessageInfo[0])
26
27 typedef struct msgmeta {
28     uint8_t msgid; /* mavlink message id */
29     bool active; /* will send message if set to TRUE */
30     uint8_t period; /* signal will be send every [period] * 10 ms */

```

```

31     bool        customRequest; /* true if message was requested via mavlink message
    ↪ */
32 } msgmeta_t;
33
34 ***** Private signals and variables
    ↪ *****
35
36 static const mavlink_message_info_t mavlinkMessageInfo[] = MAVLINK_MESSAGE_INFO;
37 mavlink_system_t mavlinkSystem;
38 msgmeta_t msgMeta[NUM_MSGS];
39 static bool prevTimeout;
40
41 static uint32_t lastTimestamp;
42
43 /* UART */
44 uart_t uart;
45 static uint8_t byteBuffer;
46 static int8_t receiveBuffer[MAX_BUFFER_LEN], transmitBuffer[MAX_BUFFER_LEN];
47 static uint32_t receiveBufferSize = MAX_BUFFER_LEN, transmitBufferSize =
    ↪ MAX_BUFFER_LEN;
48 static uint32_t receiveBufferTimestamp[MAX_BUFFER_LEN], transmitBufferTimestamp[
    ↪ MAX_BUFFER_LEN];
49
50 static uint16_t transmitcounter;
51
52 /* Signals */
53 uint32_t *pSyncOffset_mavlink;
54 heartbeat_t *pHeartbeat_mavlink;
55 carcontrol_t *pCarControl_mavlink;
56 trajectory_t *pTrajectory_mavlink;
57 uint8_t *pState_mavlink;
58 uint8_t *pRequestedState_mavlink;
59 uint64_t *pErrorRegister_mavlink;
60
61 uint32_t *pOdomTimestamp_mavlink;
62 uint32_t *pStAngTimestamp_mavlink;
63 float32_t *pVehSpdOdom_mavlink;
64 float32_t *pVehXDistOdom_mavlink;
65 float32_t *pVehYDistOdom_mavlink;
66 float32_t *pVehYawAngOdom_mavlink;
67 float32_t *pStAngAct_mavlink;
68
69 ***** Private functions
    ↪ *****
70
71 bool configureOutgoingMessage(uint8_t msgid, bool active, uint8_t period, bool
    ↪ customRequest)
72 {
73     if (msgid == MAVLINK_MSG_ID_HEARTBEAT && customRequest) return false; // do not
    ↪ allow heartbeat to be reconfigured by custom requests
74     for (uint8_t i = 0; i < NUM_MSGS; i++)
75     {
76         if (msgMeta[i].msgid == msgid)
77         {
78             if (msgMeta[i].customRequest == false || customRequest == true)
79             {
80                 msgMeta[i].active = active;
81                 msgMeta[i].period = period;
82                 msgMeta[i].customRequest = customRequest;
83             }
84             return true;
85         }
86     }
87     return false;
88 }
89

```

```

90 void configureStateMessages(uint8_t state, bool active)
91 {
92     switch (state)
93     {
94     case SYSTEM_STATE_INITIALIZING:
95         configureOutgoingMessage(MAVLINK_MSG_ID_CMD_REQUEST_CLOCKSYNC, active, 50, false)
96         ↵ ;
97         break;
98     case SYSTEM_STATE_IDLE:
99         configureOutgoingMessage(MAVLINK_MSG_ID_CMD_REQUEST_CLOCKSYNC, active, 10, false)
100        ↵ ;
101        break;
102     case SYSTEM_STATE_RUNNING_EXT:
103         configureOutgoingMessage(MAVLINK_MSG_ID_CMD_REQUEST_CLOCKSYNC, active, 10, false)
104        ↵ ;
105        configureOutgoingMessage(MAVLINK_MSG_ID_ODOMETRY, active, 1, false);
106        break;
107     case SYSTEM_STATE_RUNNING_RC:
108         configureOutgoingMessage(MAVLINK_MSG_ID_CMD_REQUEST_CLOCKSYNC, active, 10, false)
109        ↵ ;
110        break;
111     case SYSTEM_STATE_EMERGENCY:
112         configureOutgoingMessage(MAVLINK_MSG_ID_ERROR, active, 10, false);
113         break;
114     default:
115         io_printf("MAVLINK: switched to unknown state");
116     }
117 }
118
119 bool sendMessage(uint8_t msgid)
120 {
121     mavlink_message_t msg;
122     uint8_t buf[MAVLINK_MAX_PACKET_LEN];
123
124     uint16_t ticksRL;
125     uint16_t ticksRR;
126     if (msgid == MAVLINK_MSG_ID_ODOMETRY)
127         eict_getCounter2(&ticksRL, &ticksRR);
128
129     uint32_t now;
130     uint32_t deltaTime;
131
132     irq_disable();
133     uint32_t syncOffset = *pSyncOffset_mavlink;
134     switch (msgid)
135     {
136     case MAVLINK_MSG_ID_HEARTBEAT:
137         mavlink_msg_heartbeat_pack(mavlinkSystem.sysid, mavlinkSystem.compid, &msg, *
138         ↵ pState_mavlink);
139         break;
140     case MAVLINK_MSG_ID_ERROR:
141         mavlink_msg_error_pack(mavlinkSystem.sysid, mavlinkSystem.compid, &msg, *
142         ↵ pErrorRegister_mavlink);
143         break;
144     case MAVLINK_MSG_ID_ODOMETRY:
145         // use this for real driving:
146         //mavlink_msg_odometry_pack(mavlinkSystem.sysid, mavlinkSystem.compid, &msg, *
147         ↵ pOdomTimestamp_mavlink - syncOffset, *pStAngTimestamp_mavlink - syncOffset
148         ↵ , *pVehSpdOdom_mavlink, *pVehXDistOdom_mavlink, *pVehYDistOdom_mavlink, *
149         ↵ pVehYawAngOdom_mavlink, *pStAngAct_mavlink);
150
151         // use this for calibration:
152         now = systime_getSysTime();
153         deltaTime = now - lastTimestamp;
154         lastTimestamp = now;

```

```

146     mavlink_msg_odometry_pack(mavlinkSystem.sysid, mavlinkSystem.compid, &msg, *
        ↪ pOdomTimestamp_mavlink - syncOffset, deltaTime, *pVehSpdOdom_mavlink, *
        ↪ pVehXDistOdom_mavlink, *pVehYDistOdom_mavlink, (float32_t) ticksRL, (
        ↪ float32_t) ticksRR);
147
148     break;
149 case MAVLINK_MSG_ID_CMD_REQUEST_CLOCKSYNC:
150     mavlink_msg_cmd_request_clocksync_pack(mavlinkSystem.sysid, mavlinkSystem.compid
        ↪ , &msg, 0);
151     break;
152 default:
153     io_printf("MAVLink: not implemented or unknown msgid\n");
154     irq_enable();
155     return false;
156 }
157 irq_enable();
158
159 uint16_t len = mavlink_msg_to_send_buffer(buf, &msg);
160 if (uart.sendring.size - uart.sendring.count > len) {
161     uart_writeString(&uart, (int8_t*) buf, len);
162 }
163 else {
164     io_printf("MAVLink: UART TX buffer full\n");
165     return false;
166 }
167
168 return true;
169 }
170
171 void handleMessage(mavlink_message_t *msg)
172 {
173     uint32_t recTimestamp = systime_getSysTime();
174     irq_disable();
175     uint32_t syncOffset = *pSyncOffset_mavlink;
176     switch(msg->msgid)
177     {
178     case MAVLINK_MSG_ID_HEARTBEAT:
179         pHeartbeat_mavlink->lastReceiveTimestamp = recTimestamp;
180         pHeartbeat_mavlink->mavlink_version = mavlink_msg_heartbeat_get_mavlink_version(
        ↪ msg);
181         pHeartbeat_mavlink->state = mavlink_msg_heartbeat_get_state(msg);
182         break;
183     case MAVLINK_MSG_ID_CARCONTROL:
184         pCarControl_mavlink->lastReceiveTimestamp = recTimestamp;
185         pCarControl_mavlink->vehSpdTgtExt = mavlink_msg_carcontrol_get_vehspd(msg);
186         pCarControl_mavlink->stAngTgtExt = mavlink_msg_carcontrol_get_stang(msg) - 4;
187         break;
188     case MAVLINK_MSG_ID_TRAJECTORY:
189         pTrajectory_mavlink->lastReceiveTimestamp = recTimestamp;
190         pTrajectory_mavlink->timestamp = mavlink_msg_trajectory_get_timestamp(msg) +
        ↪ syncOffset;
191         mavlink_msg_trajectory_get_x(msg, pTrajectory_mavlink->x);
192         mavlink_msg_trajectory_get_y(msg, pTrajectory_mavlink->y);
193         mavlink_msg_trajectory_get_theta(msg, pTrajectory_mavlink->theta);
194         mavlink_msg_trajectory_get_kappa(msg, pTrajectory_mavlink->kappa);
195         mavlink_msg_trajectory_get_v(msg, pTrajectory_mavlink->v);
196         break;
197     case MAVLINK_MSG_ID_CMD_REQUEST_MSG: ; /* semicolon is needed here ... */
198         uint8_t msgid = mavlink_msg_cmd_request_msg_get_msgid(msg);
199         bool_t active = (bool_t) mavlink_msg_cmd_request_msg_get_active(msg);
200         uint8_t period = mavlink_msg_cmd_request_msg_get_period(msg);
201         configureOutgoingMessage(msgid, active, period, true);
202         break;
203     case MAVLINK_MSG_ID_CMD_REQUEST_STATECHANGE:
204         *pRequestedState_mavlink = mavlink_msg_cmd_request_statechange_get_state(msg);
205         break;

```

```

206     default:
207         io_printf("MAVLINK: Don't know how to handle message id %u\n", msg->msgid);
208     }
209     irq_enable();
210 }
211
212 /****** Public functions
213     ↪ *****
214 int8_t mavlink_init(void)
215 {
216     io_printf("MAVLINK: initializing\n");
217     uart.baudrate = MAVLINK_BAUDRATE;
218     uart.channel = UART_CH3;
219     if (uart_init(&uart, receiveBuffer, receiveBufferSize, transmitBuffer,
220         ↪ transmitBufferSize, transmitBufferTimestamp, receiveBufferTimestamp)) {
221         return -1;
222     }
223     mavlinkSystem.sysid = 0;
224     mavlinkSystem.compid = MAV_COMP_ID_STM;
225
226     for (uint8_t i = 0; i < NUM_MSGS; i++)
227     {
228         msgMeta[i].msgid = mavlinkMessageInfo[i].msgid;
229         msgMeta[i].active = false;
230         msgMeta[i].customRequest = false;
231     }
232
233     /* heartbeat is always active, might as well activate it here (and never deactivate
234     ↪ ) */
235     configureOutgoingMessage(MAVLINK_MSG_ID_HEARTBEAT, true, 10, false);
236
237     prevTimeout = true;
238
239     return 0;
240 }
241 void mavlink_getMessages(void)
242 {
243     mavlink_message_t msg;
244     mavlink_status_t status;
245     uint8_t chan = 0;
246
247     while(uart_readByte(&uart, (int8_t*) &byteBuffer))
248     {
249         if (mavlink_parse_char(chan, byteBuffer, &msg, &status))
250         {
251             handleMessage(&msg);
252         }
253     }
254
255     bool timeout = (system_getSysTime() - pHeartbeat_mavlink->lastReceiveTimestamp)
256     ↪ > 2.5e5; // heartbeat timeout after 250ms
257     if (timeout && !prevTimeout)
258     {
259         errorHandler_setError(ERROR_MAVLINK_Timeout);
260         prevTimeout = true;
261     }
262     else if (!timeout && prevTimeout) {
263         errorHandler_clearError(ERROR_MAVLINK_Timeout);
264         prevTimeout = false;
265     }
266 }
267 void mavlink_sendMessages(void)

```

```

268 {
269     transmitcounter++;
270     for (uint8_t i = 0; i < NUM_MSGS; i++)
271     {
272         if (msgMeta[i].active && (transmitcounter % msgMeta[i].period) == 0)
273         {
274             sendMessage(msgMeta[i].msgid);
275         }
276     }
277 }
278
279 void mavlink_activateStateMessages(uint8_t state)
280 {
281     configureStateMessages(state, true);
282 }
283
284 void mavlink_deactivateStateMessages(uint8_t state)
285 {
286     configureStateMessages(state, false);
287 }

```

stangproc.h

```

1  /** *****
    ↪
2  * Unit in charge:
3  * @file stangproc.h
4  * @author hinze
5  * $Date:: 2017-09-06 12:43:27 # $
6  *
7  * -----
8  * SVN information of last commit:
9  * $Rev:: 2158 $
10 * $Author:: mbrieske $
11 *
12 * -----
13 *
14 * @brief The torque processing transforms the valid torque target value into the
    ↪ expected format and range of the actuating module.
15 *
16 *****
    ↪ */
17
18 #ifndef STANGPROC_H_
19 #define STANGPROC_H_
20
21 /***** Includes
    ↪ *****/
22 #include "brd/startup/stm32f4xx.h"
23 #include "dev/servo/servo.h"
24 #include "per/irq.h"
25 #include "sys/systeme/systeme.h"
26
27 /***** Public typedefs
    ↪ *****/
28
29 /***** Macros and constants
    ↪ *****/
30
31 /***** Public function prototypes
    ↪ *****/
32 /**
33 * @brief Initializes the steering angle processing
34 *
35 * @param [in,out] void

```



```

36  *
37  * @return 0
38  */
39  int8_t stangproc_init(void);
40
41  /**
42  * @brief This function process the steering angle provided by application
43  * steering angle target selector in a angle that can be read by the servo.
44  * The process consists in an linear transformation
45  *
46  * @param [in,out] void
47  *
48  * @return void
49  */
50  void stangproc_step(void);
51
52  #endif /* STANGPROC_H_ */

```

stangproc.c

```

1  /*****
   ↪
2  * Unit in charge:
3  * @file stangproc.c
4  * @author hinze
5  * $Date:: 2017-09-06 12:43:27 #$
6  *
7  * -----
8  * SVN information of last commit:
9  * $Rev:: 2158 $
10 * $Author:: neumerkel $
11 *
12 * -----
13 *
14 * @brief The steering processing consist in the linear transformation of the
   ↪ steering angle
15 * in an angle which can be read by the servo device.
16 *
17 *
18 *****/
   ↪ */
19
20 /***** Includes
   ↪ *****/
21 #include "../stangproc.h"
22 #include "sys/util/io.h"
23
24 /***** Private typedefs
   ↪ *****/
25
26 /***** Macros and constants
   ↪ *****/
27
28 /***** Global variables
   ↪ *****/
29 /* Signal pointers for input signals */
30 float32_t* pStAngTgtSel_stangproc;
31
32 /* Signal pointers for output signals */
33 uint32_t* pStAngTimestamp_stangproc;
34 float32_t* pStAngTgt_stangproc;
35 float32_t* pStAngAct_stangproc;
36 float32_t* pSrvAngAct_stangproc; /**< signal from servo */
37
38 /* Parameter pointers */

```

```

39 float32_t* pStAngMax_stangproc;
40 float32_t* pStAngMin_stangproc;
41 float32_t* pStAngTgtFailSafe_stangproc;
42 float32_t* pMinAbsAngSrv1_stangproc;
43 float32_t* pMaxAbsAngSrv1_stangproc;
44
45 /****** Private static variables
46 ↳ *****
47 static float32_t stAngMax = 0;
48 static float32_t stAngMin = 0;
49 static float32_t stAngTgtFailSafe = 0;
50 static float32_t stAngTgt = 0;
51 static float32_t srvAngMax = 0;
52 static float32_t srvAngMin = 0;
53
54 /****** Private function prototypes
55 ↳ *****
56 static void updateParameters(void);
57 static float32_t transformAng(float32_t inAng);
58 static float32_t transformSvrAngToStAng(float32_t servoAngAct);
59
60 /****** Public functions
61 ↳ *****
62 int8_t stangproc_init(void)
63 {
64     stAngMax = 0;
65     stAngMin = 0;
66     stAngTgtFailSafe = 0;
67     stAngTgt = 0;
68     srvAngMax = 0;
69     srvAngMin = 0;
70
71     /* fetch updated parameters */
72     updateParameters();
73     /* set output signals to init values */
74     *pStAngTimestamp_stangproc = 0;
75     *pStAngTgt_stangproc = stAngTgtFailSafe;
76
77     return 0;
78 }
79
80 void stangproc_step(void)
81 {
82     float32_t sigSrvAngAct = 0.0;
83     /* fetch updated parameters */
84     irq_disable();
85     updateParameters();
86     sigSrvAngAct = *pSrvAngAct_stangproc;
87     /* store target value as intermediate taret value */
88     stAngTgt = *pStAngTgtSel_stangproc;
89     /* transform steer angle target value into servo steer angle target value */
90     *pStAngTgt_stangproc = transformAng(stAngTgt);
91     /* transform the actual servo angle value into the actual steer angle */
92     *pStAngAct_stangproc = transformSvrAngToStAng(sigSrvAngAct);
93     /* publish timestamp for steering angle */
94     *pStAngTimestamp_stangproc = systime_getSysTime();
95     irq_enable();
96
97     /* call servo device driver to update to the new target value */
98     #ifndef USE_RX24F
99     servo_updateAngleSrv1();
100     #endif
101 }
102
103 /****** Private functions
104 ↳ *****

```

```

101 static void updateParameters(void)
102 {
103     /* update all parameters to current values from parameter system */
104     stAngMax = *pStAngMax_stangproc;
105     stAngMin = *pStAngMin_stangproc;
106     srvAngMax = *pMaxAbsAngSrv1_stangproc;
107     srvAngMin = *pMinAbsAngSrv1_stangproc;
108     stAngTgtFailSafe = transformAng(*pStAngTgtFailSafe_stangproc);
109 }
110
111 static float32_t transformAng(float32_t inAng)
112 {
113     float32_t ret = 0;
114     float32_t m = 0;
115
116     /* linear transformation of steer angle input into servo steer angle */
117     m = (srvAngMax - srvAngMin) / (stAngMax - stAngMin);
118     ret = m * (inAng - stAngMin) + srvAngMin;
119
120     return ret;
121 }
122
123 /* transform the actual servo angle value into the actual steer angle */
124 static float32_t transformSvrAngToStAng(float32_t servoAngAct)
125 {
126     float32_t retStAngAct = 0.0;
127     float32_t res = 0.0;
128
129     /* linear transformation of servo steer angle input into steer angle */
130     res = (servoAngAct - srvAngMin) / (srvAngMax - srvAngMin);
131     retStAngAct = res * (stAngMax - stAngMin) + stAngMin;
132
133     return retStAngAct;
134 }

```

vehspdctrl.h

```

1
2 /** *****
3     ↪
4     * Unit in charge:
5     * @file vehspdctrl.h
6     * @author Brieske
7     * $Date:: 2017-09-12 08:12:11 #$
8     * -----
9     * SVN information of last commit:
10    * $Rev:: 2022 $
11    * $Author:: Brieske $
12    * -----
13    *
14    * @brief TODO
15    *
16    *
17    * ↪ */
18
19 #ifndef VEHSPEEDCTRL_H_
20 #define VEHSPEEDCTRL_H_
21
22 /***** Includes
23     ↪ *****/
24 #include "brd/startup/stm32f4xx.h"
25 #include "per/irq.h"
26 #include "sys/algo/pid.h"

```

```

26 #include "sys/algo/pt1.h"
27 /***** Public typedefs
    ↳ *****/
28
29 /***** Macros and constants
    ↳ *****/
30
31 /***** Public function prototypes
    ↳ *****/
32 int8_t vehspdctrl_init(void);
33 /**
34  * @brief Initialize the parameters and signals used by the application vehspdctrl
35  *
36  * @return int8_t success
37  */
38 void vehspdctrl_step(void);
39 /**
40  * @brief TODO
41  */
42
43 #endif /* VEHSPDPROC_H_ */

```

vehspdctrl.c

```

1  /** *****
    ↳
2  * Unit in charge:
3  * @file vehspdctrl.c
4  * @author Brieske
5  * $Date:: 2017-09-12 08:12:11 #$
6  *
7  * -----
8  * SVN information of last commit:
9  * $Rev:: 2022 $
10 * $Author:: Brieske $
11 *
12 * -----
13 *
14 * @brief TODO
15 *
16 *****
    ↳ */
17
18 /***** Includes
    ↳ *****/
19 #include "../vehspdctrl.h"
20 #include "app/stm/stm.h"
21 #include "dev/rcrec/rcrec.h"
22 #include "dev/servo/servo.h"
23 #include "app/leddebug.h"
24
25 /***** Global variables
    ↳ *****/
26 /* input signals */
27 float *pVehSpd0dom_vehspdctrl;
28 uint8_t *pState_vehspdctrl;
29 carcontrol_t *pCarControl_vehspdctrl;
30 //float *pVehSpdTgtTraj_vehspdctrl;
31 uint16_t *pRcChannel1_vehspdctrl;
32
33 /* output signals */
34
35 /* parameters */
36

```

```

37 /***** Private typedefs
    ↳ *****/
38
39 /***** Macros and constants
    ↳ *****/
40 #define TDELTA 0.02
41
42 #define PID_K_P 0.005
43 // No I value required since we do not need to correct a constant error.
44 #define PID_K_I 0
45 // D value to avoid overshoot.
46 #define PID_K_D 0.001
47
48 #define PID_GAIN 1
49
50 #define PT1_T 0.08
51 #define PT1_T_OPT (1. / (PT1_T / TDELTA) + 1)
52
53 #define MOT_MAXVAL 10.
54 // #define U_DEADBAND_POSITIVE 0.0583
55 // #define U_DEADBAND_NEGATIVE -0.075
56 #define U_DEADBAND_POSITIVE 0.02
57 #define U_DEADBAND_NEGATIVE -U_DEADBAND_POSITIVE * 4
58
59 #define MOTOR_LIMIT 0.4f
60 #define MOTOR_FULLSTOP -10
61 #define MOTOR_AFTER_FULLSTOP 0
62
63 #define MAX_SPEED 3.0f
64 #define FLOAT_TOLERANCE 0.0001f
65
66 #define MIN(a, b) (a < b ? a : b)
67 #define MAX(a, b) (a > b ? a : b)
68 #define ABS(x) (x < 0 ? -x : x)
69
70 /***** Private static variables
    ↳ *****/
71 static pidCtrl_t pid;
72 static pt1_t pt1;
73
74 static float vehSpdOdom;
75 static uint8_t state;
76
77 static float motTrqTgtSrv;
78
79 // tells whether there was a fullstop on the last step
80 static bool lastFullstop;
81
82 /***** Private function prototypes
    ↳ *****/
83
84 static float spdCtrlAssystem(float w);
85 static void spdCtrl(float destSpd);
86 static void setThrottle(float throttle);
87
88 /***** Public functions
    ↳ *****/
89 int8_t vehspdctrl_init(void)
90 {
91     pCarControl_vehspdctrl->vehSpdTgtExt = 0.0f;
92
93     pt1.K = 1;
94     pt1.T_opt = PT1_T_OPT;
95     pt1.y_old = 0;
96
97     pid.kp = PID_K_P;

```

```

98     pid.ki = PID_K_I;
99     pid.kd = PID_K_D;
100
101     // pid.ki = 0.F;
102     // pid.kd = 0.F;
103
104     pid.gain = PID_GAIN;
105     pid.ta = TDELTA;
106     pid.min = -MOT_MAXVAL;
107     pid.max = MOT_MAXVAL;
108
109     pid_initialize(&pid);
110
111     return 0;
112 }
113
114 void vehspdctrl_step(void)
115 {
116     irq_disable();
117     state = *pState_vehspdctrl;
118     irq_enable();
119
120     switch (state) {
121     case SYSTEM_STATE_INITIALIZING:
122         // motor is disabled, nothing to do here
123         return;
124     case SYSTEM_STATE_IDLE:
125         motTrqTgtSrv = 0;
126         lastFullstop = false;
127     case SYSTEM_STATE_EMERGENCY:
128         spdCtrl(0.F);
129         break;
130     case SYSTEM_STATE_RUNNING_RC:
131         irq_disable();
132         float rcvalue = (float) *pRcChannel1_vehspdctrl;
133         irq_enable();
134         // if (rcvalue < 1400) {
135         //     motTrqTgtSrv = 0;
136         // }
137         // else {
138         motTrqTgtSrv = (rcvalue - 1500.0F) / (float) (MAX_PERIOD_CHANNEL -
139             ↪ MIN_PERIOD_CHANNEL);
140         motTrqTgtSrv *= 0.8;
141         //     motTrqTgtSrv *= 3;
142         //     motTrqTgtSrv = spdCtrl(motTrqTgtSrv);
143         // }
144         break;
145     case SYSTEM_STATE_RUNNING_EXT:
146         irq_disable();
147         float spdTgt = pCarControl_vehspdctrl->vehSpdTgtExt;
148         irq_enable();
149
150         // debugging
151         if (spdTgt != 0.0f)
152             debugLedOn();
153         else
154             debugLedOff();
155
156         spdCtrl(spdTgt);
157         break;
158     // case SYSTEM_STATE_RUNNING_TRAJ:
159     //     TODO system state not implemented yet
160     //     irq_disable();
161     //     float spdTgt = *pVehSpdTgtTraj_vehspdctrl;
162     //     irq_enable();
163     //     motTrqTgtSrv = spdCtrl(spdTgt);

```

```

163 //      break
164     default:
165         // cannot go here
166         break;
167     }
168
169     servo_updateMotorTrq(motTrqTgtSrv);
170 }
171
172 /***** Private functions
173     ↳ *****/
174
175 static float spdCtrlAssystem(float w)
176 {
177     float u = 0; // manipulated value
178     float u_f = 0; // feed forward value
179     float y = 0; // controlled value
180
181     irq_disable();
182     y = *pVehSpdOdom_vehspdctrl;
183     irq_enable();
184
185     y = pt1_calculate(&pt1, y);
186     pid_controller(&pid, y, w, &u, -MOT_MAXVAL, MOT_MAXVAL);
187
188     if (u > 0) {
189         u_f = U_DEADBAND_POSITIVE;
190     }
191     else if (u < 0) {
192         u_f = U_DEADBAND_NEGATIVE;
193     }
194
195     u += u_f;
196
197     return u;
198 }
199
200 static void spdCtrl(float destSpeed)
201 {
202     destSpeed = MIN(destSpeed, MAX_SPEED);
203
204     // Here we first check if speed 0 was received.
205     // In this case we perform a fullstop with maximum brake power.
206     // This might have been implemented differently but there wasn't enough time at
207     // ↳ this point.
208     // However this functionality is definitely required since when left out the PID
209     // ↳ tries to smoothly reduce speed to 0
210     // which is not what we desire when trying to avoid a collision.
211     bool fullstop;
212     if (destSpeed <= FLOAT_TOLERANCE)
213         fullstop = true;
214     else
215         fullstop = false;
216
217     if (!fullstop && lastFullstop)
218     {
219         motTrqTgtSrv = MOTOR_AFTER_FULLSTOP;
220         pid_initialize(&pid);
221     }
222
223     lastFullstop = fullstop;
224
225     if (fullstop)
226     {
227         motTrqTgtSrv = MOTOR_FULLSTOP;
228         return;
229     }

```

```

226     }
227
228     float change = 0; // value calculated by PID
229     float currSpeed = 0; // controlled value
230
231     irq_disable();
232     currSpeed = ABS(*pVehSpdOdom_vehspdctrl);
233     irq_enable();
234
235     //currSpeed = pt1_calculate(&pt1, currSpeed);
236     pid_controller(&pid, currSpeed, destSpeed, &change, -MOT_MAXVAL, MOT_MAXVAL);
237
238     motTrqTgtSrv += change;
239
240     if (motTrqTgtSrv > 0 && motTrqTgtSrv < U_DEADBAND_POSITIVE)
241         motTrqTgtSrv = U_DEADBAND_POSITIVE;
242
243     motTrqTgtSrv = MIN(motTrqTgtSrv, MOTOR_LIMIT);
244 }
245
246 static void setThrottle(float throttle)
247 {
248     throttle = MIN(throttle, 1.0f);
249     throttle = MAX(throttle, 0.0f);
250
251     if (throttle == 0.0f)
252     {
253         motTrqTgtSrv = MOTOR_FULLSTOP;
254         return;
255     }
256
257     motTrqTgtSrv = throttle;
258 }

```

eict.h

```

1
2  /* *****
3     ↪
4     * Unit in charge:
5     * @file eict.h
6     * @author Yoga
7     * $Date:: 2015-05-05 13:31:20 #
8     * -----
9     * SVN information of last commit:
10    * $Rev::          $
11    * $Author::       $
12    * -----
13    * -----
14    * @brief This unit configures an input capture timer triggering on the edges of the
15    ↪ wheel encoder.
16    * -----
17    *****
18    ↪ */
19
20 #ifndef EICT_H_
21 #define EICT_H_
22
23 ***** Includes
24 ↪ *****
25 #include "per/hwallocation.h"

```



```

26 /***** Public typedefs
    ↳ *****/
27
28 /***** Privat typedefs
    ↳ *****/
29
30 /***** Macros and constants
    ↳ *****/
31 /***** Public function prototypes
    ↳ *****/
32 void eict_init(void);
33
34 void eict_getCounter(uint16_t *frontLeftEnc, uint16_t *frontRightEnc, uint16_t *
    ↳ rearLeftEnc, uint16_t *rearRightEnc);
35
36 void eict_getCounter2(uint16_t *rearLeftEnc, uint16_t *rearRightEnc);
37
38 #endif /* EICT_H_ */

```

eict.c

```

1 /** *****/
    ↳
2  * Unit in charge:
3  * @file eict.h
4  * @author Yoga
5  * $Date::2015-05-05 13:31:20 #$
6  *
7  * -----
8  * SVN information of last commit:
9  * $Rev::          $
10 * $Author::       $
11 *
12 * -----
13 *
14 * @brief This unit configures an input capture timer triggering on the both edges of
    ↳ the wheel encoder signal.
15 *
16 *****/
    ↳ */
17
18 /***** Includes
    ↳ *****/
19 #include "brd/startup/stm32f4xx.h"
20 #include "app/config.h"
21 #include "per/eict.h"
22 #include "per/irq.h"
23
24
25 /***** Private typedefs
    ↳ *****/
26 /***** Macros and constants
    ↳ *****/
27 #define SR_INC_MAX      100
28
29 /***** Globale Variable
    ↳ *****/
30
31 /***** Private static variables
    ↳ *****/
32 static uint16_t frontLeftCapture = 0;
33 static uint16_t frontRightCapture = 0;
34 static uint16_t rearLeftCapture = 0;
35 static uint16_t rearRightCapture = 0;
36

```

```

37 static uint16_t rearLeftCapture2 = 0;
38 static uint16_t rearRightCapture2 = 0;
39
40 ***** Private function prototypes
41 ↪ *****
42 static void gpio_config(void);
43
44 static void nvic_config(void);
45
46 ***** Public functions
47 ↪ *****
48 void eict_init(void)
49 {
50 #ifdef WENC_V3
51     uint16_t polarity = TIM_ICPolarity_Falling;
52 #else
53     uint16_t polarity = TIM_ICPolarity_Rising;
54 #endif
55     TIM_ICInitTypeDef TIM_ICStructure;
56
57     /* GPIO configuration */
58     gpio_config();
59
60     /* nvic configuration */
61     nvic_config();
62
63 #ifdef FOUR_WHEEL_ODOM
64     /* TIMER3 channel 1(PB4) Configuration: Input Capture mode */
65     TIM_ICStructure.TIM_Channel = TIM_Channel_1;
66     TIM_ICStructure.TIM_ICFilter = 0x0;
67     TIM_ICStructure.TIM_ICPolarity = polarity; /* rising and falling edges are used
68     ↪ as active edge */
69     TIM_ICStructure.TIM_ICPrescaler = TIM_ICPSC_DIV1; /* triggering on each edge
70     ↪ */
71     TIM_ICStructure.TIM_ICSelection = TIM_ICSelection_DirectTI;
72     TIM_ICInit(WENC_FRONT_TIMER, &TIM_ICStructure);
73
74     /* TIMER3 channel 2(PB5) Configuration: Input Capture mode */
75     TIM_ICStructure.TIM_Channel = TIM_Channel_2;
76     TIM_ICStructure.TIM_ICFilter = 0x0;
77     TIM_ICStructure.TIM_ICPolarity = polarity; /* rising and falling edges are used
78     ↪ as active edge */
79     TIM_ICStructure.TIM_ICPrescaler = TIM_ICPSC_DIV1; /* triggering on each edge
80     ↪ */
81     TIM_ICStructure.TIM_ICSelection = TIM_ICSelection_DirectTI;
82     TIM_ICInit(WENC_FRONT_TIMER, &TIM_ICStructure);
83 #endif
84
85     /* TIMER2 channel 4(PB11) Configuration: Input Capture mode */
86     TIM_ICStructure.TIM_Channel = TIM_Channel_4;
87     TIM_ICStructure.TIM_ICFilter = 0x0;
88     TIM_ICStructure.TIM_ICPolarity = polarity; /* rising and falling edges are used
89     ↪ as active edge */
90     TIM_ICStructure.TIM_ICPrescaler = TIM_ICPSC_DIV1; /* triggering on each edge
91     ↪ */
92     TIM_ICStructure.TIM_ICSelection = TIM_ICSelection_DirectTI;
93     TIM_ICInit(WENC_REAR_TIMER, &TIM_ICStructure);
94
95     /* TIMER2 channel 2(PB3) Configuration: Input Capture mode */
96     TIM_ICStructure.TIM_Channel = TIM_Channel_2;
97     TIM_ICStructure.TIM_ICFilter = 0x0;
98     TIM_ICStructure.TIM_ICPolarity = polarity; /* rising and falling edges are used
99     ↪ as active edge */

```

```

93     TIM_ICStructure.TIM_ICPrescaler = TIM_ICPSC_DIV1;          /* triggering on each edge
    ↪ */
94     TIM_ICStructure.TIM_ICSelection = TIM_ICSelection_DirectTI;
95     TIM_ICInit(WENC_REAR_TIMER, &TIM_ICStructure);
96 #ifdef FOUR_WHEEL_ODOM
97     /* TIMER3 enable counter */
98     TIM_Cmd (WENC_FRONT_TIMER, ENABLE);
99
100    /* Timer 3 enable the CC Interrupt request */
101    TIM_ITConfig(WENC_FRONT_TIMER, TIM_IT_CC1 | TIM_IT_CC2, ENABLE);
102
103    /* Timer 3 clear CC Flag */
104    TIM_ClearFlag(WENC_FRONT_TIMER, TIM_IT_CC1 | TIM_IT_CC2);
105 #endif
106
107    /* TIMER2 enable counter */
108    TIM_Cmd (WENC_REAR_TIMER, ENABLE);
109
110    /* Timer 2 enable the CC Interrupt request */
111    TIM_ITConfig(WENC_REAR_TIMER, TIM_IT_CC4 | TIM_IT_CC2, ENABLE);
112
113    /* Timer 2 clear CC Flag */
114    TIM_ClearFlag(WENC_REAR_TIMER, TIM_IT_CC4 | TIM_IT_CC2);
115 }
116
117 void eict_getCounter(uint16_t *frontLeftEnc, uint16_t *frontRightEnc, uint16_t *
    ↪ rearLeftEnc, uint16_t *rearRightEnc)
118 {
119     irq_disable();
120
121     *frontLeftEnc = frontLeftCapture;
122     *frontRightEnc = frontRightCapture;
123     *rearLeftEnc = rearLeftCapture;
124     *rearRightEnc = rearRightCapture;
125
126     /* clear capture Counter */
127     frontLeftCapture = 0;
128     frontRightCapture = 0;
129     rearLeftCapture = 0;
130     rearRightCapture = 0;
131
132     irq_enable();
133 }
134
135 void eict_getCounter2(uint16_t *rearLeftEnc, uint16_t *rearRightEnc)
136 {
137     irq_disable();
138
139     *rearLeftEnc = rearLeftCapture2;
140     *rearRightEnc = rearRightCapture2;
141
142     /* clear capture Counter */
143     rearLeftCapture2 = 0;
144     rearRightCapture2 = 0;
145
146     irq_enable();
147 }
148
149 #ifdef FOUR_WHEEL_ODOM
150 void TIM3_IRQHandler()
151 {
152     if(TIM_GetITStatus(WENC_FRONT_TIMER, TIM_IT_CC1)) {
153         TIM_ClearITPendingBit(WENC_FRONT_TIMER, TIM_IT_CC1);
154         frontLeftCapture++;
155     }
156 }

```

```

157     if(TIM_GetITStatus(WENC_FRONT_TIMER, TIM_IT_CC2)) {
158         TIM_ClearITPendingBit(WENC_FRONT_TIMER, TIM_IT_CC2);
159         frontRightCapture++;
160     }
161 }
162 #endif
163
164 void TIM2_IRQHandler()
165 {
166     if(TIM_GetITStatus(WENC_REAR_TIMER, TIM_IT_CC4)) {
167         TIM_ClearITPendingBit(WENC_REAR_TIMER, TIM_IT_CC4);
168         rearLeftCapture++;
169         rearLeftCapture2++;
170     }
171     if(TIM_GetITStatus(WENC_REAR_TIMER, TIM_IT_CC2)) {
172         TIM_ClearITPendingBit(WENC_REAR_TIMER, TIM_IT_CC2);
173         rearRightCapture++;
174         rearRightCapture2++;
175     }
176 }
177
178 /****** Private function *****/
179 static void gpio_config(void)
180 {
181
182     GPIO_InitTypeDef GPIO_InitStructure;
183
184     /* GPIOB clock enable */
185     RCC_AHB1PeriphClockCmd(WENC_PIN_CLOCK, ENABLE);
186
187     /* Pin configuration for Timer 2 and 3 */
188     GPIO_InitStructure.GPIO_PuPd = GPIO_PuPd_UP;
189     GPIO_InitStructure.GPIO_Speed = GPIO_Speed_100MHz;
190     GPIO_InitStructure.GPIO_OType = GPIO_OType_PP;
191     GPIO_InitStructure.GPIO_Mode = GPIO_Mode_AF;
192     GPIO_InitStructure.GPIO_Pin = WENC_RR_CHAN_2_PIN | WENC_FL_CHAN_1_PIN |
193     ↪ WENC_FR_CHAN_2_PIN | WENC_RL_CHAN_4_PIN;
194     GPIO_Init(WENC_PORT, &GPIO_InitStructure);
195
196     /* Configure the the GPIO pin as alternate function */
197     GPIO_PinAFConfig(WENC_PORT, WENC_FL_CHAN_1_AF_PIN, WENC_FL_CHAN_1_AF);
198     GPIO_PinAFConfig(WENC_PORT, WENC_FR_CHAN_2_AF_PIN, WENC_FR_CHAN_2_AF);
199     GPIO_PinAFConfig(WENC_PORT, WENC_RL_CHAN_4_AF_PIN, WENC_RL_CHAN_4_AF);
200     GPIO_PinAFConfig(WENC_PORT, WENC_RR_CHAN_2_AF_PIN, WENC_RR_CHAN_2_AF);
201 }
202
203 static void nvic_config(void)
204 {
205     NVIC_InitTypeDef NVIC_InitStructure;
206
207     /* Timer 2 peripheral clock enable */
208     RCC_APB1PeriphClockCmd (WENC_REAR_TIM_CLOCK, ENABLE);
209
210     /* Timer 3 peripheral clock enable */
211     RCC_APB1PeriphClockCmd (WENC_FRONT_TIM_CLOCK, ENABLE);
212
213 #ifdef FOUR_WHEEL_ODOM
214     /* enable the TIM3 global Interrupt */
215     NVIC_InitStructure.NVIC_IRQChannel = TIM3_IRQn;
216     NVIC_InitStructure.NVIC_IRQChannelCmd = ENABLE;
217     NVIC_InitStructure.NVIC_IRQChannelPreemptionPriority = 0x00;
218     NVIC_InitStructure.NVIC_IRQChannelSubPriority = 0x01;
219     NVIC_Init(&NVIC_InitStructure);
220 #endif
221
222     /* enable the TIM2 global Interrupt */

```

```
222     NVIC_InitStructure.NVIC_IRQChannel = TIM2_IRQn;
223     NVIC_InitStructure.NVIC_IRQChannelCmd = ENABLE;
224     NVIC_InitStructure.NVIC_IRQChannelPreemptionPriority = 0x00;
225     NVIC_InitStructure.NVIC_IRQChannelSubPriority = 0x00;
226     NVIC_Init(&NVIC_InitStructure);
227 }
```

Selbständigkeitserklärung

Ich erkläre hiermit, dass ich die vorliegende Arbeit selbständig verfasst und noch nicht für andere Prüfungen eingereicht habe. Sämtliche Quellen einschließlich Internetquellen, die unverändert oder abgewandelt wiedergegeben werden, insbesondere Quellen für Texte, Grafiken, Tabellen und Bilder, sind als solche kenntlich gemacht. Mir ist bekannt, dass bei Verstößen gegen diese Grundsätze ein Verfahren wegen Täuschungsversuchs bzw. Täuschung eingeleitet wird.

Berlin, den December 18, 2018

.....