UP-Drive

Automated Urban Parking and Driving

Research and Innovation Action - Horizon 2020
Grant Agreement Number 688652
Start date of the project: 01-01-2016, Duration: 48 months

Deliverable D 10.2
Periodic Technical Report 2
Part B
Period covered by the report: from 01.01.2017 to 30.06.2018 Status: Revision 1.0

Lead contractor for this deliverable: VW
Due date of deliverable: 30.06.2018
Coordinator: VW

Project co-funded by the European Commission within HORIZON 2020		
Dissemination Level		
PU	Public	X
PP	Restricted to other programme participants (including the Commission Services)	
SEAM	Restricted to partners of the SEAM Cluster (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Summary for publication

The project UP-Drive aims at pushing the state of the art of automated driving and focuses on urban environments. Automotive experts expect the automated driving technology to radically improve the safety, comfort and affordability of individual mobility. Whilst many driver assistance systems - addressing mainly highway-like environments - are being introduced into the market and successively propagated from the premium car segment down to economy vehicles, automation of urban traffic remains an open research topic.

The main goal of UP-Drive is therefore to push forward the perception, localization and reasoning abilities of autonomous vehicles. In the course of the project, we will build a prototype car system capable of driverless operation in complex urban environments. Our focus is placed on residential areas and speeds up to 30 km/h.

During the second period, the project has focused on building up the second vehicle platform and bring closure to the first development and integration cycle in terms of environment perception, life-long localization and mapping, scene understanding and decision-making and navigation. Integration and testing of all relevant components have been successfully completed at a system level.

Figure 1: Test vehicles with integrated sensors: second vehicle platform (left) and first vehicle platform (right).

Contributing Partners

	Company/Organisation	Name
Document Manager	VW	Wojciech Derendarz
Partner 1	VW	Jens Spehr
Partner 1	VW	Dominik Vock
Partner 1	VW	Heiko Kurz
Partner 2	IBM	Martin Rufli
Partner 3	UTC	Ion Giosan
Partner 4	ETH	Mathias Buerki
Partner 5	ČVUT	Václav Hlaváč

Contents

1 Explanation of the work carried out by the beneficiaries and overview of the progress

1.1 Objectives

The main objective of the project is to build a prototype car system able to perform automated operation in complex urban environments with speeds up to 30km/h. To achieve this goal, substantial progress needs to be made in the key technologies of perception, localization and reasoning.

The project has chosen a spiral approach to system development as depicted in Figure 2. This essentially means that the UP-Drive system will be built in two consecutive iterations. The main purpose of the first iteration is to have an integrated system early and thus to be able to identify the key issues early enough so that they can be thoroughly addressed within the development phase of the second iteration.

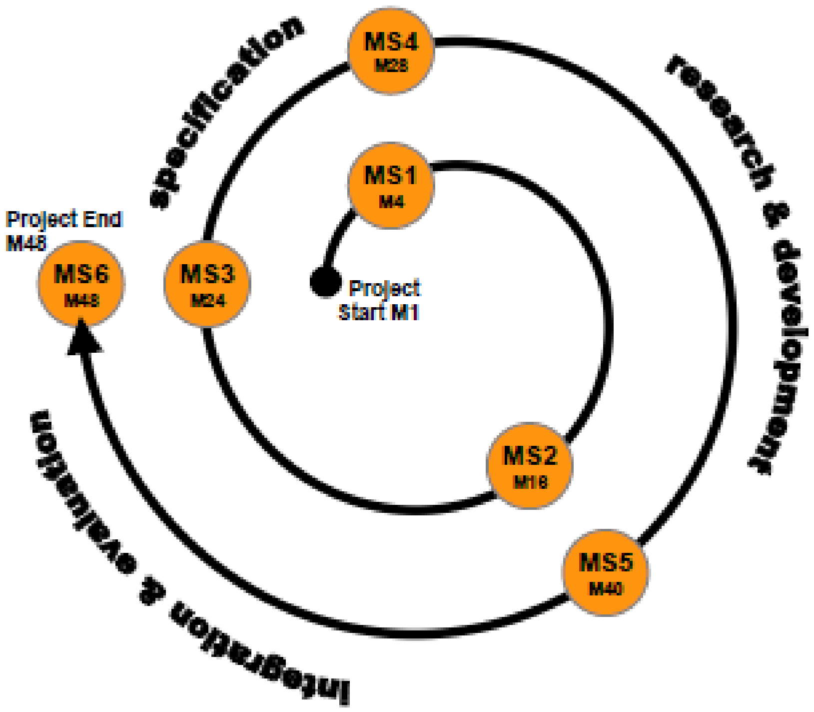

Figure 2: UP-Drive project spiral life cycle model

Based on that approach, 6 milestones have been defined for the project. These are summarized in the following table.

Table 1: UP-Drive project milestones

Number	Milestone title	Due Month
1	Work has started, first specification completed	4
2	First development phase completed	16
3	First integration and testing phase completed	24
4	Revised specifications completed	28
5	All developments completed	40
6	Project completed	48

During the second period (Months 13-30) the project has focused on building up the second vehicle platform and bring closure to the first development and integration cycle in terms of environment perception, long-life localization and mapping, scene understanding and decision-making and navigation. Integration and testing of all relevant components have been successfully completed at a system level.

According to the Description of Action - and expressed in more formal terms - the objectives for the second period of the project were to complete Milestone 2, 3 and 4. Descriptions of the milestones are quoted below and accompanied by report on current status.

MS2 - First development phase completed (Month 16)

Requirement / Expectation	*Status*
Initial versions of components have been developed and individually tested	True for all key contributions from all Work Packages and Partners
First scientific results published (at least 5 papers have been submitted to relevant conferences)	At the time of writing this report, the project partners have published around 30 publications
Submitted deliverables: D2.1, D2.2, D3.2, D4.1, D6.1, D7.1, D8.1, D8.2, D8.3, D9.3	All submitted

The milestone was succesfully achieved. However, it needs to be noted, that the development phase of the first system iteration has taken longer than initially planned. This is especially true in case of Work Package 4. The reasons for that as well as the minor consequences for the project are discussed in the conclusion of Section 1.2.4.

MS3 - First integration and testing phase completed (Month 24)

Requirement / Expectation	*Status*
The first cycle of the spiral development model has been completed	True for all key contributions from all Work Packages and Partners
First integration week has been carried out	Integration week 1 in M24
Integration and testing of initial version of components has been completed at a system level	True for all key contributions from all Work Packages and Partners
Mapping front-end is able to provide useful maps. Map completeness and scale is yet limited.	True
Car is able to perform basic automated driving in urban environments. Functionality and performance is limited.	First automated drives in M23. Since then improved stability and extended functionality.
Submitted deliverables: D2.3, D4.2, D4.3, D5.2, D6.2, D7.2, D8.4, D9.4, D9.5, D9.6	D2.3, D4.2, D4.3, D5.2, D6.2, D7.2, D9.4 submitted. D9.5 and D9.6 are expected in M33; D8.4 in M34

The most important element of the milestone is the mid-term demonstration, which has been postponed from M24 to M33. Thus the milestone cannot be considered completed. However, as the first automated drives had been performed already in M23 and since then the system stability has been continously improved and the functionality extended to include all key components from all Work Packages and Partners, the project believes that the delay will have no negative impact on the final project outcome. Currently the system functionality is being tested using both vehicle platforms.

MS4 - Revised specifications completed (Month 28)

Requirement / Expectation	*Status*
Second integration week has been carried out	Integration week 2 in M29
At least 12 papers have been submitted to relevant conferences or journals	About 30 papers submitted
Submitted deliverables: D1.2, D2.4, D4.4	D1.2, D4.4 submitted. D2.4 is under preparation and expected in M34

The most important element of the milestone is the updated version of requirements and specification for the vehicle platform, such that the second test vehicle can be built according to these specifications. This has been successfully achieved. The second vehicle platform has received a redesigned sensor setup and computer cluster and is fully functional. The work on documenting all the vehicle functionality is delayed and expected to complete in M34. As this is only a minor issue, we consider the milestone to be succesfully completed.

It should be noted that the number of submitted papers is much larger than expected for this milestone - and is actually already higher than expected for the whole project.

Summary

Project has achieved excellent scientific results and managed to integrate key results into one system. Thus we consider the project to be on track to success. The following subsections give more details on the achieved results.

1.2 Explanation of the work carried per WP

Work in UP-Drive project has been partitioned into 10 packages. The work package structure is depicted in Figure 2, which is self-explanatory. It should be mentioned, however, that the main effort is placed on the 4 work packages (numbered 4-7 and marked green) which are devoted to what the project believes to be the key challenges of automated driving.

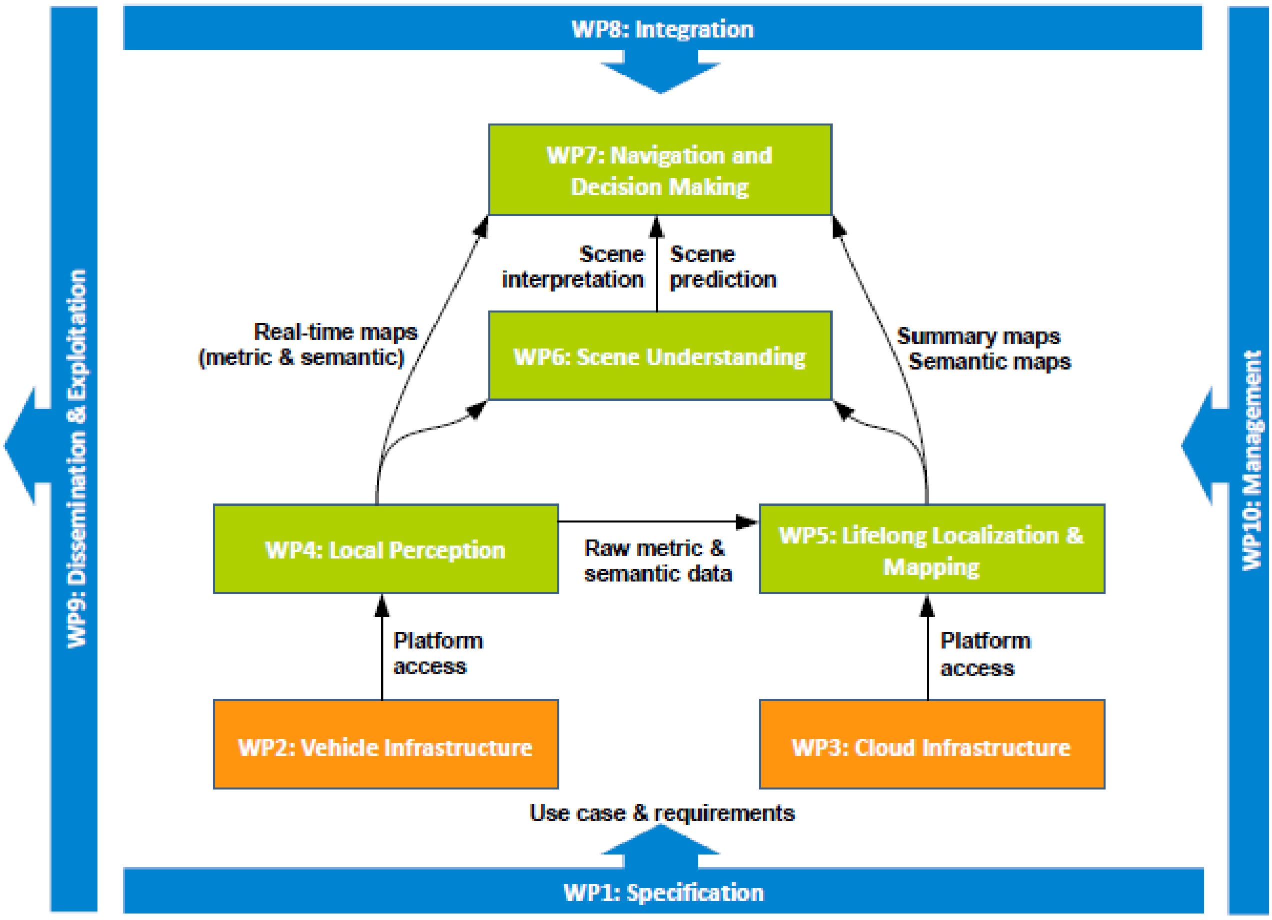

Figure 3: Work package structure

1.2.1 WP1 Requirements, System & Components Specification and Architecture

In WP1 end user needs and requirements are analyzed with respect to the proposed applications which in turn lead to a concretized application specification. Based on these requirements, the structural and functional architectures of the system and components are specified. During the second period of the project (M13-M30) the objective was to update the specification carried out in the first period (M1-M12).

Task 1.1 End-user needs and requirements analysis

The end-user needs and the derived requirements were kept unchanged (see D1.1).

Task 1.2 Application analysis and requirement specification

The requirements for perception derived from the use-case analysis were updated and improved in accordance with the experience acquired in the past two and half years of the project. These updates were based on the carried out experiments and implementations. The result of this analysis is synthesized in the chapter 2.2 from D1.2.

The work has been carried out by UTC.

Task 1.3 Definition and specification of system architecture

The definition and specification of the system architecture was kept unchanged (see D1.1).

Task 1.4 System hardware and software components specification

The experiments demonstrated the need of better sensory system. The depth and accuracy provided by the 16 layer LIDAR sensors were not enough to fulfil the perception goals. A new solution based on 32 layer LIDARs was introduced.

Similarly, in order to improve the semantic segmentation capabilities for the frontal area, two color cameras with different horizontal field of view were mounted at the rear-view mirror level. The advantages come from the more rich color information provided at two different resolutions, and from the closer position of the cameras with respect to the frontal 32 layer LIDARs.

The closer position of the two types of sensors to the front LIDARs increase the overlapping of the image and 3D point cloud views and diminish the problems generated by occlusions - false semantic associations. The updated sensory system specification is presented in chapter 3.3.1 from D1.2.

New calibration and cross-calibration procedures and goals for the new sensory setup were defined in chapter 3.3.3.A1 from D1.2.

The processing power foreseen for the perception task was not enough for the multiple redundant 360 degree perception. An improvement of the processing power was proposed and experimented. A new solution based on ZOTAC EN 1080K is under implementation.

The 3D terrain perception task received a new denomination and specification - DIFF map and non-static object segmentation in the scene. DIFF Map is a low-level 3D geometric map of the road surface and its close vicinity in the entire city. It encompasses road surface, curbs, posts etc. It consists of a probabilistic elevation map and simple repeatable detectable structures (planes, 3D lines). It is envisioned as an "average" map of permanent parts of the environment. Based on the difference between this "average" map and the current instantaneous map the dynamic entities are detected and labeled easier (see chapter 3.3.3.B2 from D1.2).

The work has been carried out by UTC, VW and CVUT.

Outlook

The objectives for this Work Package have been achieved. The Deliverable D1.2 was compiled. The work package is on track, all the project requirements and the component specifications have been set.

1.2.2 WP2 Vehicle Infrastructure

The main objective of Work Package 2 is to provide test platforms for the project. During the second project period, WP2 focused its efforts on building the 2nd vehicle platform (planned for month 24 - January 2018) and to make it fully operational (planned for month 28 - May 2018). The objectives have been successfully achieved, yet with slight delay: in months 30 and 32 respectively. Details are given in the following sections.

Task 2.1 Vehicle platform setup

This task involves the hardware buildup of the vehicle, including the mechanical and electrical installation of the computer and sensor systems and their proper configuration. The second test vehicle platform "Superwolle" is based on the update of the fully electric VW e-Golf. The vehicle is a mirror version of the hardware and software of the first vehicle platform with upgrades of the LiDAR (32-layer LiDAR), Tri-focal system and computer cluster. The buildup phase is completed, thus the sensor system, computer system, communication system (in-vehicle and between vehicles) as well as HMI and safety elements are fully integrated.

In Figure 4 both vehicle platforms are shown for comparison.

Figure 4: Second vehicle platform Superwolle (left) and first vehicle platform (right).

The second vehicle platform is fully operational and thus used for the development work in all relevant WPs. A thorough analysis of the sensor setup (calibration and data integrity validation of the sensors) and report on communication capabilities along with a summary of the high-level processing framework and overview of the safety elements and policies is under preparation and will be provided in D2.4.

Task 2.2 Drive by wire functionality

T2.2 enables automated operation of the test vehicles through actuation of all relevant vehicle systems: gas, brakes, steering wheel, parking brake and signaling lights. During the second period, the efforts were concentrated on implementing identical drive-by-wire functionality as in the first vehicle platform. This has been achieved by the modification of the topology of the vehicle's Controller-Area-Network (CAN), installation of vehicle-gateways and programming thereof. More details on that have been given in Deliverable D2.1. In addition to that, work on integrating a new, more powerful, custom-made interface for the steering system has been performed. We expect this work to be finished in 2018 and to enable more natural behavior when driving curves. Finally, wa are working on enabling backward driving to allow for performing parking manouvers or 3 point turns.

Task 2.3 Low-level data acquisition, processing & communication framework

Task 2.3 is considered completed. The results have been described in detail in Periodic report 1 and D2.1.

Task 2.4 High level system debugging & maintenance framework

The main objective of this task is to provide the necessary tools for system integration, maintenance and debugging in order to analyze the low-level vehicle system issues directly in the car. The following tools have been implemented in both test vehicle platforms.

- data recorder that can easily be activated by the driver
- system health monitoring
- 3D visualisation of messages from sensor data to trajectory planning
- a central cockpit display providing all relevant system components information
- software deployment tool, easing the boot-up of software releases in the distributed system
- configuration and distribution of vehicle test parameters
- mission and black box recording system running in the background

T2.3 work is ongoing and on track, thus ensuring integration of status information and live visualization for all modules.

Task 2.5 Calibration & data integrity validation of the sensor system

The second vehicle platform has been upgraded with 32-layer LiDARS. This required modification and further extension of the calibration chamber. The lower number of rays of the new 32- layered LiDAR hit the ground in close distance to the vehicle thus making the correct calibration of the sensors difficult. Because of this, only the surfaces of the 3 vertical walls were usable for the calibration, significantly limiting the calibration quality for the 32-layered LiDAR. To overcome this limitation, the calibration chamber was extended by an additional wall segment (Figure 5, right), which is mounted with a 20° pitch angle against the vertical walls.

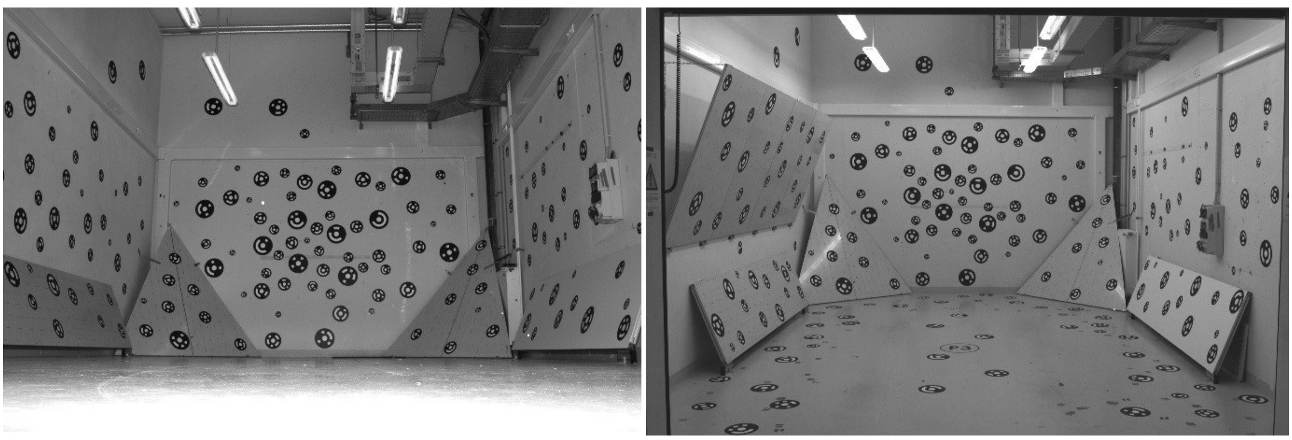

Figure 5: (a) Comparison of calibration chamber (before: left; extended setup: right).

This modification allowed proper calibration of the sensor system.

Task 2.6 Reference sensor integration

During the second period, identical reference system for ego-motion estimation as in the 1st vehicle platform was implemented and commissioned in the 2nd vehicle platform. This task is considered completed.

Task 2.7 Vehicle maintenance and update

Task 2.7 is a continuous activity throughout the project. Regular maintains of the vehicle test platforms were performed, thus ensuring fully operational conditions and usability for development and testing. The most significant update of 1st vehicle within Period 2 was the replacement of the 3 out of 5 16-layered LIDARs by their 32-layered versions.

Conclusion & Outlook

The work in this work package has been performed mainly by VW. The delivery of the second vehicle platform was slightly delayed, but had no negative impact on the overall project progress. Moreover, the updates to the sensor system and computer cluster defined in WP1 have been fully implemented in the 2nd vehicle. The first 6 Tasks are considered completed. In the upcoming months we expect to update the steering interface in both cars, update the computer cluster of 1st vehicle and add yet more GPU power to both cars.

1.2.3 WP3 Cloud Infrastructure

Work package 3 has the aim of providing a project-wide foundational layer for the storage, processing, maintenance and sharing of data across devices, as well as the (automated) compilation, testing, deployment and operation of software modules within the project.

The first development and implementation cycle ended in M12 and has been reported in the previous periodic report, and in more detail in D3.1 and and D3.2. Starting with M25, an assessment of the various deployed services has been completed. This has resulted in the following tasks performed in an incremental second development and integration cycle:

- Setup of hardware stack: extension of physical storage to currently 90TB to accommodate growing project data acquisitions (further extensible based on project needs).
- Design and implementation of the development and deployment tool chain: implementation of a Gitlab, Travis CI, Docker and Kubernetes based development and deployment toolchain. Application of this workflow to a growing set of deployed cloud services.
- Setup and implementation of the bulk data storage service: migration to a more extensible cloud middleware based on Kubernetes.

For the remaining two tasks (Task 3.1, and Task 3.4) no need for an adaptation was identified. They hence remain in the state described in D3.2.

WP3 is on track of successfully completing all tasks by M32. Activities performed during the second implementation and integration cycle will be documented in more detail in the upcoming report D3.3. Below we provide a high level description.

Task 3.2 Specification, setup and maintenance of hardware stack

The UP-Drive hardware stack consists of compute nodes, GPU accelerators, memory, storage and communication infrastructure, and has been scoped based on the functional and application requirements derived in WP1. The scoping is described in detail in D3.2.

As storage needs have been growing considerably over the past year (this is largely due to the near-daily acquisition of test datasets), an extension of physical storage to currently 90TB has been completed. Further extensions based on project needs will be facilitated by the completed migration of the bulk data storage service to Kubernetes (see Task 3.5).

Task 3.3 Design and implementation of the development and deployment framework

This task involves the specification and implementation of a development and deployment toolchain consisting of a source code management system, a configuration management system and a continuous integration system. Within the first development cycle the Gitlab source code management system has been deployed for all partners' use.

As part of the second integration cycle the configuration management system allows to declaratively specify and automatically build and deploy a selected software stack. It has been implemented based on Travis CI and the Portus Docker registry. Deployment of services has been implemented based on the Kubernetes container orchestration service containerisation and/or virtualisation as key implementation techniques. The automated build, test and deployment toolchain has been applied to a growing set of the project's cloud services. Deliverable 3.3 will provide additional detail on the toolchain and applications.

Task 3.5 Design and implementation of the bulk data storage service

In the first development cycle the bulk data storage service has been scoped, designed and implemented in the form of an object store. To this end, OpenStack Swift has been selected, which is ideal for storing unstructured data that can grow without bound at a central server.

During the second development and implementation cycle, the OpenStack Swift storage service has been migrated to run on the project's Kubernetes Container orchestration service. At the same time the maximal storage space has been extended from 10TB (3 times replicated) to presently 30TB (3 times replicated). This setup is more easily configurable, provides better means to ensure availability (via e.g., load balancing) and can be seamlessly extended as storage needs of the project grow further. Besides the raw storage service that comes with a command line interface, we have implemented a WebDAV proxy service that enables a range of GUI clients to connect to the object store.

Outlook

WP3 is on track of successfully completing all tasks by M32. Deliverable 3.3 will detail the development and integration tasks conducted during the second project cycle in detail.

1.2.4 WP4 Perception

This work package provides the vehicle-side online sensing functionality required for automated driving in low-speed urban environments. The work package covers the following primary aspects:

- Specification and design of on-board sensing (covered by Task 4.1 in M1-M8 and M25-M28) [UTC, CVUT, VW]. An on-board 360-degree multi-sensorial perception solution was designed based on a set of well selected sensors including laser scanners, radars, large field of view area view cameras, and trifocal cameras disposed in a configuration offering good coverage, redundancy and good measurement accuracy. In the same time this approach offers access to the powerful perception functions of the individual sensors.
- Spatio-temporal and appearance based low-level representation (covered by Task 4.2 in M5-M32) [UTC, CVUT]. The redundancy and complementarity of the individual sensors are exploited throughout complex calibration, motion correction, temporal and spatial alignment and fusion processes to provide a unified, enriched and more accurate spatio-temporal and appearance based low-level representation. This spatio-temporal and appearance based 360-degree representation integrates at pixel level, both in the 3D point cloud and in the 2D image space, the 3D position, the 3D motion, the 2D intensity or color, the image segment and the image class information.
- Refinement of detection, tracking, and classification capabilities and environment sensorial representation (covered by Tasks 4.3-4.7 in M5-M42) [UTC, CVUT, VW]. We utilize the richness of the sensor setup and of the unified low-level representation as a building block not only for detection of traffic participants but also for the detection of terrain and road infrastructure. Starting from the selected sensors and the associated perception solutions and using the new spatio-temporal and appearance based low level representation new solutions are investigated for objects or object parts better detection, tracking and classification.

The main goals of WP4 in the second period (M13-M30) of the project were to provide final specification and design of on-board sensing perception system architecture and then to develop an initial version of the low-level and high-level perception functions. Because the sensor setup got much more rich than assumed at the time of writing the proposal and changed even in the second period of the project, our plan from the proposal has evolved:

- Task 4.1 Specification and design of on-board sensing
- Task 4.1bis Calibration, cross-calibration, synchronization (newly introduced)
- Task 4.2 Spatio-temporal and appearance based low level representation
- Task 4.3 Perception adaptation to adverse visibility conditions
- Task 4.4 Road infrastructure perception
- Task 4.5 Real-time 3D terrain perception received a new denomination: DIFF map and non-static object segmentation in the scene
- Task 4.6 Road users and signaling perception

- Task 4.7 Sensor fusion based perception refinement

Several modules were identified and specified in terms of interface (input and output data structures), functionality (their main tasks) and performance (evaluation of the output). Their interconnection and mapping on the above mentioned tasks is further presented. Some possible solutions were studied, evaluated and implemented.

Task 4.1 Specification and design of on-board sensing

An initial design, implementation and evaluation of the low-level perception functions were accomplished and described in D4.2. The main objective was to accurately build the spatio-temporal and appearance based low level representation (STAR) which is the key element in achieving high accuracy for the higher level perception functions like: object detection from point cloud and image data, parking spot detection, signaling perception, road users detection, tracking and classification.

An initial design and implementation of the high-level perception functions were accomplished and described in D4.3. The main objective was to build high-level environment perception functions that uses the low-level data coming from all available sensors. In order to make the perception system robust, functions should exist that use only data coming from a specific type of sensors (e.g. either cameras or LIDARs) such that the environment may be perceived even in the case when some sensors cannot provide reliable data (e.g. night, adverse weather conditions, failures). In the case when data is available from all sensors, the spatio-temporal and appearance based low level representation (STAR) is built and used. An initial design and implementation of high-level perception functions based on STAR representation is proposed.

The specification and design of on-board sensing initially described in D4.1 were updated and improved in accordance with the project needs. These updates were based on the experiments and implementations carried out in the past two and half years of the project. They are presented in more detail in D4.4.

During the project the 360 degree technology advanced. The sensory system was upgraded with LIDARs with 32 channels and a variable vertical resolution. The new sensors provide better (isotropic) resolution in the central part of the field-of-view. The other characteristics of the sensor are similar. All the updates in the sensor model were described in chapter 2.1 of D4.4. The offline calibration procedures for all sensors including LIDARs, area-view cameras, trifocal cameras and RADARs were revised and their design is updated in chapter 2.2 of D4.4.

The design of the 3D point cloud segmentation and classification was updated and presented in chapter 3.2 of D4.4. Object detection can be performed based on information only from the LIDAR point clouds. This offers reliable depth information but it is sparse and semantic information can be more difficult to extract. Classification of objects is deferred onto a later stage when semantic information from images is ready.

An update to the specification and design of DIFF map and non-static object segmentation in the scene was carried out and presented in chapter 3.3 of D4.4. It refers to the functionalities like off-line DIFF Map update, on-line DIFF Map-aided current data-point labeling to permanent/transient, grouping of such data-points to geometrically consistent objects, possibly classification of such objects to semantic classes.

The design of the area-view images segmentation module was updated and presented in chapter 4.1 of D4.4. In order to achieve a 360 degree surround semantic perception, a semantic segmentation procedure is applied on all four (front, left, rear, right) undistorted area view images using a fully convolutional neural network. It has an encoder-decoder architecture: the encoder network computes features at different scales and the decoder network combines the features to obtain a higher resolution representation.

A design of an object detection and classification module was proposed and presented in chapter 4.2 of D.4.4. For visual detection the unwarped area view images are considered. A convolutional neural network that achieves object detection based on Faster R-CNN is used for detecting relevant objects and road users from the environment. The output consists in a list of 2D bounding boxes in images along with their class and instance ID.

The design of the enhanced detection, classification and localization based on low-level spatio-temporal and appearance based representation (STAR) was updated and presented in chapter 5.2 of D4.4. Two approaches were designed. The first is the image semantic segmentation based detection, classification and localization where STAR is especially useful for static objects which are small but are easily segmented in the image domain: poles, traffic signs, curbs, lane markings. The second is point cloud based detection, classification and localization refinement using semantic segmentation where the objects detected from the point clouds can be classified using the STAR representation by determining the highest frequency of semantic class labels corresponding to the contained 3D voxels.

The work has been carried out by UTC, VW and CVUT.

Task 4.1bis Calibration, cross-calibration, synchronization (newly introduced)

A sensor calibration module was implemented and presented in chapter 2 of D4.2. It includes an efficient calibration procedure with high accuracy results for the intrinsic and extrinsic calibration parameters of all available sensors: area-view cameras, trifocal cameras, 360 degree LIDARs, 4 layer LIDARs and RADARs. The objective was to obtain a very precise calibration which is crucial for building the STAR.

The work has been carried out by CVUT.

Task 4.2 Spatio-temporal and appearance based low-level representation

A 3D points correction module was implemented and presented in chapter 4 of D4.2. It consists in the motion correction approach for the 360 degree and 4 layer LIDARs point clouds, trifocal camera and RADAR 3D objects. It also contains an approach for laser based depth-correction of trifocal camera 3D objects. These aspects are very important for having a correct data alignment in the STAR building process, in the situation when the ego-vehicle is moving.

An optical flow computation module was developed and presented in chapter 5 of D4.2. It defines a solution for optical flow computation from area-view images. The relative motion determined by the optical flow vector is further used in the STAR.

The STAR representation was developed and presented in chapter 6 of D4.2. The representation is obtained by low level fusion of the information collected from the following sensors: area-view cameras, 360 degree LIDARs, 4 layer LIDARs, trifocal cameras and RADARs. Additional to the sensorial information, optical flow computation and odometry based 6DoF ego-motion were considered.

The work has been carried out by UTC.

Task 4.3 Perception adaptation to adverse visibility conditions

A perception adaptation module was described in in chapter 3 of D4.2. It contains an approach for image enhancement in fog conditions which is important for the subsequent modules that process images with reasonable quality even in adverse weather conditions.

The work has been carried out by UTC.

Task 4.4 Road infrastructure perception and Task 4.6 Road users and signaling perception

A road surface detection and 3D object segmentation module based on 3D voxels using only the LIDAR point cloud data was implemented and described in chapter 2.1 of D4.3. The proposed approach for ground and obstacle segmentation consists of two main steps. The first step is the detection of the road surface, which is done initially on each individual 360-degree LIDAR data, and then the results are fused in order to get a global road surface. The second step is the detection of 3D obstacles by a proposed voxel representation that allows the fusion of all the LIDAR outputs. Given the five 360-degree LIDARs (and the additional data from 4-layer LIDARs), the LIDARs measurements are first fused into the same representation (the voxel space), and then the 3D obstacles are detected by grouping.

A semantic segmentation of images and an instance based object segmentation module from area-view image data were implemented and described in chapter 3.1 of D4.3. In the segmentation process, the most relevant 23 visual classes for the traffic scene were chosen. Classes were selected based on their frequency and compatibility with other public datasets. A common deep CNN based architecture was proposed to achieve both semantic segmentation and instance segmentation over area view images. The implementation is built on a custom C++/ CUDA framework based on Nvidia CuDNN. It runs in real time for semantic segmentation of images. Instance based object segmentation has high computational costs and struggle to perform at over four frames/second on high-end GPUs.

A parking spot detection module from LIDAR point cloud data was implemented and described in chapter 2.2 of D4.3. The proposed approach consists of two steps: continuous aggregation of 3D point data into moving occupancy grid, and query-based classification of parking spots. A probabilistic volumetric approach for representing the local map in the vicinity of the ego-vehicle was used and the LIDAR sensor readings were aggregated into a 3D occupancy grid. The occupancy grid is being continuously kept up-to-date. The classification is based only on a number of cells that are observed to be occupied (with some probability) and a number of cells that are observed to be free (with some probability).

Regarding the signaling perception of road users from image data, a module that is capable of detecting the vehicle taillight was developed and described in chapter 3.2 of D4.3. Two methods have been explored to extract candidate taillight regions. Extracted candidate regions are paired by comparing their sizes, positions and colors. Each taillight of a detected pair is then tracked using Kalman filtering in order to deal with false negatives.

An enhanced road users perception module by associating semantic labels to 3D objects was implemented and described in chapter 4.1 of D4.3. During the information fusion between the LIDAR point cloud and the segmentation results, semantic labels can be transferred to 3D points erroneously. This happens when the LIDAR sensors observes an object that is occluded in the camera view. The LIDAR points that originate from the occluded object are projected onto the occluding object in the camera view. Any semantic label transferred to the occluded object in this way is a mistake. This problem was addressed and solved by depth maps generation and analysis. After the cloud of LIDAR measurements is enhanced with semantic information, the semantic class for each 3D obstacle is established. This is done statistically, in order to compensate for various errors in the semantic class of the LIDAR points (on the borders of dynamic obstacles, or due to the different view point of the cameras versus LIDARs).

A module that detects road users starting from the low-level STAR representation was developed and presented in chapter 4.2 of D4.3. A grouping of entities based on the image view and the semantic segmentation is available in the form of an object id. This representation has the advantage over the raw 3D points that additional information is available before grouping of the points is performed. Based on this representation road users can be detected by grouping 3D points from the birds eye view together. The first stage entails grouping together points that have the same semantic class and same object id. The second stage involves validation and correction of the objects. Several constraints are applied to cuboids and the 3D points.

A module for road users tracking and object position correction was developed and presented in chapter 4.3 of D4.3. The objects detected each frame from the 3D point cloud are provided in the tracking process in the form of oriented cuboids that can move in the lateral and longitudinal directions. The Unscented Kalman Filter is used for tracking. A data association procedure based on object's 3D dimensions, orientation and location of the bounding box is employed for associating the tracks to the measurements.

The work has been carried out by UTC and CVUT.

Task 4.5 Real-time 3D terrain perception

A module for real-time 3D terrain perception named DIFF map was developed. DIFF Map is a low-level 3D geometric map of the road surface and its close vicinity in the entire city. It encompasses road surface, curbs, posts etc. It consists of a probabilistic elevation map and simple repeatable detectable structures (planes, 3D lines).

It is envisioned as an "average" map of permanent parts of the environment. DIFF Map data for a city area come from repeated drives through the area. The data acquisition process is not be pre-planned, DIFF Map update would happen on an irregular basis as a side-effect of normal driving.

Comparison of the DIFF Map with the current on-line low-level LiDAR frame is planned as a means of attention-free fast non-specific object detection (or "spotting") via the principle of background subtraction. The structured part of DIFF Map is planned to aid accurate vehicle self-localization.

The work has been carried out by CVUT.

Task 4.7 Sensor fusion based perception refinement

The main objective of this task is the high level fusion and representation of the environment. This includes the occupancy grid estimation, object estimation and the module called Road Graph.

Whilst occupancy grid has been designed and implemented to represent static obstacles in the scene, the object fusion keeps track of (potentially) moving objects. Both modules can be understood as low-pass filtering mechanism for the noisy sensor data. Until very recently, both modules have been used to combine high-level information provided by different sensors: laser scanners, radars and cameras. This is a parallel approach to the preffered STAR approach from Task 4.2. However, as the development of the STAR approach has taken longer than initially planned, the high-level fusion enabled providing Work Packages 6 and 7 with usable data early on in the project.

Now that obstacle data based on the low level sensor fusion approach is being provided, the role of the object fusion and occupancy grid computation is expected to get reduced to just temporal fusion of the data from that virtual sensor.

Finally, we mention the module called Road Graph - designed and implemented to represent the local road topology and geometry around the ego-vehicle. In fact RoadGraph is the result of work on intersection of many Work Packages as it deals with description of road geometry (WP4 and WP5), road topology (WP5) and - based on that road topology - interpretation of traffic rules or computation of conflict areas (WP6).

All the above modules have been described in more detail in deliverables D4.1 and especially D4.3. They have been successfully tested in the vehicle and provide data to the prediction and motion planning steps.

The work in this Task has been carried out by VW.

Outlook

The objectives of this Work Package have been achieved. The development of the object detection based on the STAR approach, developed by UTC has taken longer than initially planned - the first automated drive based on this approach was performed in M32, rather than in M24 as planned at the time of writing the proposal. It needs to be noted however, that early on, the decision was made by the project to build this kind of novel approach from scratch and to accommodate for the rich sensor setup and thus some substantial delay was anticipated. Moreover, a classical approach to object detection was followed in parallel by VW, which enabled Work Packages dependent on perception data (WP6 & WP7) to proceed without any delay.

Summarizing it can be said that - although with some delay - the objectives for this Work Package have been achieved. The deliverables D4.2, D4.3 and D4.4 were compiled. All the perception tasks have already provided relevant scientific results and the key modules are being regularly tested in the car in real-life driving scenarios.

1.2.5 WP5 Lifelong Localization & Mapping

The goal of WP5 is to provide a compact, customized, and metrically accurate map representation with which the vehicles can localize reliably in long-term and large-scale operations in urban environments, and where semantic data can be associated, stored and queried. The obtained localization and map representations are furnished to other WPs, which take higher level decisions for the self-driving task.

Overview - Progress With Respect to Milestones

Milestones M1 and M2 have been completed and reported in D10.1.

MS3 *First integration and testing phase completed*:

- The visual mapping framework is fully operational. Map data can be collected over indefinite time spans and a visual map can be created, optimized, and updated for long-term use. Two scientific publications have been published to international conferences in relation to this topic.
- The visual metric localization pipeline has been fully integrated and tested both in the Wolle and in the Kermit vehicle and is fully operational (see video link on homepage). A first evaluation report has been compiled and submitted in deliverable D5.2. Further, more extensive tests and evaluations are currently ongoing and will be reported in deliverable D5.3.
- A first version of the reference frame alignment has been developed and integrated on the vehicles. Testing is currently still ongoing.

Task-Related Progress

Description	**Related Tasks**
Data Integrity Verification	*T5.1*
Software tools have been developed that can verify the data integrity of recorded map data, report on any inconsistencies and integrity violations that would hinder the use of the respective data for mapping purposes. In particular, availability, timestamp synchronization, delay, and missing samples are checked for all sensor streams. Appropriate parameter configuration allows for tolerating a certain amount of missing data in a flexible way.	

Description	**Related Tasks**
Map Management and Cloud-Based Map Optimization and Summarization	*T5.6, T5.2, T5.3*
Significant progress has been made towards integration of the map backend (map optimization e.g., Bundle Adjustment, map summarization) on the IBM cloud infrastructure. This greatly improves the scalability of the long-term map maintenance, as nearly arbitrary amounts of RAM and CPU power is made available. A docker image with the map backend application has been prepared and is ready to be deployed on the IBM cloud infrastructure over respective web services (running on Kubernetes; deployed from a Portus Docker registry). It has been tested and the full set of map backend functionality is now available on the IBM cloud. The challenges related to the long-term management of visual maps for localization, with a special focus on scalability and data transmission efficiency in a shared-map scenario as in line with the UP-Drive scenario, have been researched and the respective findings published at IV 2018 [1]. The map backend, including storage format, serialization&deserialization, Bundle-Adjustment and map summarization, has been released open-source and presented at ICRA 2018 [2].	

Description	**Related Tasks**
Metric Localization	*T5.4*
The metric visual localization pipeline has been fully integrated on both vehicles (Wolle, Kermit) and is operational. In the recent months, all efforts have been focused on improving the system robustness and stability on long lasting sorties. In addition to that, specific challenges of the visual localization system such as global localization/bootstrapping have been further researched in the context of a Master's Thesis [3]	

Description	**Related Tasks**
LiDAR Integration	*T5.1, T5.2, T5.3, T5.4*
In alignment with the reviewer's comments after the first review meeting, we aim at extending our visual localization with the use of LiDAR data in order to increase the robustness with respect to long-term appearance change and in areas of small visual feature density. In this regard, a first iteration of requirements specification and integration into the map backend has been carried out. In particular, LiDAR data recorded by the Wolle vehicle can be parsed and stored in the map map format. Furthermore, the topic of LiDAR-based global localization has been researched in the context of a Master's Thesis [4], and we aim at integrating the respective findings and algorithms in the UP-Drive project in the upcoming months.	

Description	**Related Tasks**
Semantic Data Analysis	*T5.5*
In support of scene understanding (WP6) recorded data needs to be acquired, condensed, stored and made queryable / searchable. An initial implementation capable of extracting (from recorded mission files), storing and querying parking spot occupancy information has been completed. The storage and querying backend consists of an Elasticsearch service operating on the UP-Drive Kubernetes cluster. Implemented querying capabilities include (i) queries for last known occupancy status, or (ii) expected occupancy status of a set of specified parking spots. The implementation of functionality for (i) ingesting live data as the car completes a mission, and (ii) querying for available parking spots during a mission are currently in progress. An extension of the service to support additional semantic categories of interest to WP6 (such as car densities, pedestrian flows) is planned for the second implementation and integration cycle of the project.	

[2] Map Management for Efficient Long-Term Visual Localization in Outdoor Environments, Bürki et al., 2018
[2] maplab: An Open Framework for Research in Visual-inertial Mapping and Localization, Schneider et al., 2018
[3] Visual Global Outdoor Localization for an Autonomous Car, David Vogt, 2017
[4] Place Recognition with Data-Driven Descriptors Using 3D Point Clouds, Lukas Schaupp, 2018

Description	**Related Tasks**
Reference Frame Alignment	*T5.2*
Reference frame alignment describes the process of aligning and deforming a visual map to associated (RTK-)GPS tracks and/or to HD navigation maps. The aim is to derive a continuous transformation between global (GPS, HD navigation map) and visual map coordinate frames. A first version of the alignment functionality – consisting of an offline cloud-processing part and an online in-vehicle lookup part – has been developed and integrated. The offline part has been implemented as a robust nonlinear least squares optimization. Initial tests have been successfully carried out on the Wolle vehicle, verifying the functionality of the alignment pipeline and all involved interface to/from the visual map. A more in-depth evaluation of the quality of the alignment is currently ongoing.	

Outlook In conclusion, significant progress has been achieved on all tasks in WP5. A first version of the visual localization system is available and has been fully integrated. More in-depth testing and evaluation of all components is ongoing, and further improvements with respect to robustness are planned in the remaining project months. In addition to that, the integration of LiDAR data into the online localization pipeline will be a second key focus until the end of the project.

Work package 5 is on track.

1.2.6 WP6 Scenario Understanding

The WP6 work reported in this periodic report builds on the Deliverable D6.2 named First development and integration cycle of scene understanding, which was written in 2017.

The main purpose of *UP-Drive* scene/scenario understanding module is to explore sensory percepts delivered by WP4 (Perception), fuse it with the contextual information about the local driving environment, other traffic participants, which is relevant to a particular traffic situation. This information is used for the traffic scenario interpretation, prediction of its participants behavior/intention in a short horizon of one to typically three seconds.

The team benefited from having access to the data from a real *UP-Drive* experimental car in period two. Used methods rely on machine learning techniques which need a huge amount of relevant data. For now, the experimental car is being driven autonomously only in the compound of VW factory in Wolfsburg. This implies that the collection of tested scenarios and induced events is limited. However, experimental cars driven by humans can collect datasets also in general traffic conditions. The data-sets available suffice for developing modules and their integration to the experimental car. The overall performance is naturally hindered by the lack of huge amounts of relevant annotated datasets. This is not unique to this project. All other competing teams face the same challenge. On the other hand, we have shown in *UP-Drive* that the chosen concept works. If more data labeled date is available, the performance can be improved.

Task 6.1 Long-term semantic scene understanding

As it was described in D10.1, this task is located at the intersection of WP5 and WP6. Percepts provided by WP4 constitute the input. Knowledge is accumulated and maintained by the perception and lifelong mapping modules over extended periods.

Semantics is provided by particular object detectors. The percepts origin is either in the sensor itself, e.g. in the multifocal camera used in *UP-Drive* experimental cars, or in more sophisticated methods implemented in WP4 as, e.g. the supersensor developed by UTC partner. Experiments with the multifocal camera showed that *UP-Drive* cannot rely on them as a solely source of correctly semantically labeled objects. This experience stresses the importance of the supersensor. Object detectors from optical images and/or range information have been relying on machine learning.

One developed example in the year 2 of *UP-Drive* is the interplay between WP4 and WP5 performed in WP6, Task 6.1. The development runs in parallel. The WP4 functionality enables to detect free parking spots in a large parking lot while driving along the row of perpendicularly parked cars. Our solution to that is that so far WP6 has been based on the object data generated by the classical fusion approach provided by VW partner (described in WP4, Task 4.7). and that future work involves shifting to data generated by the superposition approach (STAR)

The interplay with the cloud infrastructure is crucial for the task. The parking spot aggregation over time via cloud performed by IBM partner is one example of this activity illustrating cooperation between tasks T6.1 and T5.5. The future work will use similar technology for estimating a car density, pedestrian flows, etc.

Task 6.2 Scenario based scene understanding

Semantics is undoubtedly an important prerequisite of the scene/scenario understanding. The complexity of traffic scenes and its dynamics is rather challenging for the scene understanding. *UP-Drive* project uses the following two paradigms:

1. **Detailed map**. *UP-Drive* project as well as most of competing industrial approaches relies on a rather detailed map of the environment. This map used in *UP-Drive* is named *RoadGraph* and serves as the core representation providing an bidirectional interface between WP6 and other WPs. The detected objects treated by WP4, WP5 and WP6 are used by a rather sophisticated and fast trajectory planner, which implements most of the WP7 (Decision-making and Navigation) functionalities. This is the approach chosen in *UP-Drive* project.
2. **Scenarios**. A potentially very complex traffic scenes are decomposed of into a finite set of scenarios. This approach allows detecting a particular scenario and start a prepared module which treats it. The most developed scenario in WP6 so far has been the prediction of pedestrians' behavior.

The typical navigation of a self-driving ego-car relies on a following servo-loop provided a detailed road map is available:

- The road map represents also the road infrastructure with the precision up to a driving lane.
- The map contains information about crossroads, pedestrian crossings, traffic lights, etc.
- The ego-car localization is provided by the visual localization in *UP-Drive*. The ground truth is generated by GPS.
- The path to be driven is provided by standard navigation systems known from cars or mobile devices.
- The ego-car runs a servo-loop, which it attempts to stay close to the prescribed path.

There is an opportunity to improve understanding of a dynamic traffic scene by employing qualitative reasoning.

There might be a natural desire to avoid a strong dependence on the detailed road map. It would need a local scene interpretation from the sensory information, continuous incremental development of a ego-car own local map, its fusion with the road map if available. A local scene and scenario representation has to be maintained. The deliberation tools need to cope with uncertainty in real sensory measurements. The approach enables to use high-level reasoning which could for instance implement traffic rules naturally.

We had a working workshop discussing such an approach in Wolfsburg on April 27, 2018. V. Hlaváč from ČVUT team explained the potential, which the high-level qualitative reasoning could bring to *UP-Drive*. It was decided by the *UP-Drive* coordinator after a discussion taking into account other self-driving activities on the coordinator side (VW) that *UP-Drive* approach in WP6 will be modified. *UP-Drive* integration efforts will stay within the detailed map and scenarios framework. ČVUT will provide and integrate the pedestrian prediction and potentially the prediction for the uncontrolled 4-way intersection, which is described in more detail in T6.3 future work part.

The job-split among *UP-Drive* partners VW and ČVUT will be changed, too. ČVUT will study and develop speculative methods based on qualitative, high-level logic reasoning-based approaches. The approach might benefit from having *UP-Drive* data but is not limited to it. Moreover, in order to be able to focus the resources on the scientific aspects, ČVUT will not integrate the developed methods to *UP-Drive* test vehicles, as that would incur high integration efforts. Instead, VW will provide and integrate an rule/heuristics-based functionality for all traffic scene scenarios - except for pedestrians and potentially for the uncontrolled 4-way intersection - provided by ČVUT.

Task 6.3 Scene prediction

For scenario understanding, the correct estimation of other traffic participants' intentions and the prediction of their motions is crucial. Pedestrians are very vulnerable road users. Therefore, we are focusing to their motion prediction and their intention estimation. The correct motion prediction and the induced reaction of the autonomous car can significantly affect the safety of the pedestrians.

Motion prediction and intention estimation

We considered and tested several approaches to motion prediction and intention estimation: simple physic-based motion prediction, Kalman filters, Bayesian networks (BNs), and heuristics.

The *Pedestrian_Intention_Predictor* module developed by ČVUT is responsible for processing pedestrian data from the object sensor, predicting their short-term trajectory and estimating the intention of pedestrians to cross roads on crosswalks. The implemented solution processes data from pedestrian movement history and locations of nearby roads and crosswalks and uses learned Bayesian Network for the estimation of pedestrian intention near a crossing based on distance measurements between the pedestrian, the road and the crossing in two consecutive time points and Kalman filter for predicting pedestrian trajectories. The delay between the time points of measurement was chosen to be one second (representing the current situation and the situation one second ago).

The proposed Bayesian network for the intention estimation of pedestrians near the zebra crossing works with measured/predicted input data, which are captured by sensors used in *UP-Drive* project car:

- pedestrian perpendicular distance to the zebra crossing;
- pedestrian perpendicular distance to the road;
- absolute distance to the zebra crossing;
- relative angle between the pedestrian‘s heading and the direction towards the zebra crossing.

The predictions acquired by the proposed Bayesian network are shown in the visualization (Figure 6), where the pedestrians are represented by colored squares. The movement direction and the probability of crossing the street for each pedestrian is displayed too. The color of each pedestrian is changing depending on the probability that the pedestrian crosses the street. The blue color represents zero probability (e.g. a pedestrian is walking away from the road). Color is changing from blue through purple to red, which denotes the highest probability.

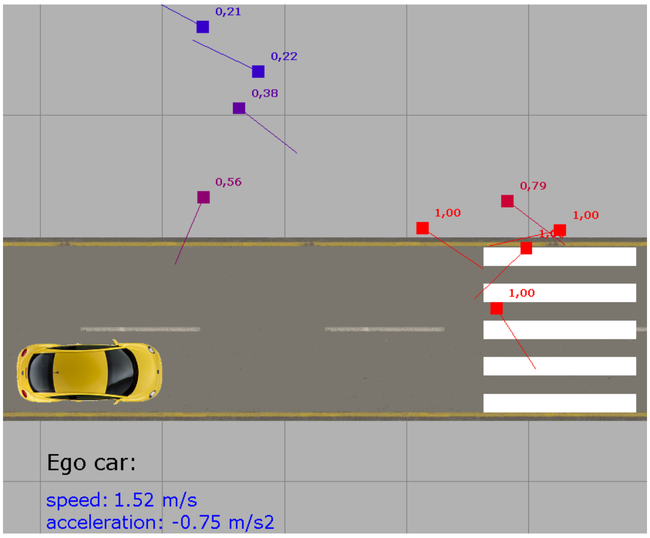

Figure 6: Bayesian network prediction visualization.

In some situations, however, the ego-car spots the pedestrians in the vicinity of a pedestrian crossing relatively late with respect to the car movement, resulting in late predictions and sudden braking of the car in front of the crossing. This behaviour is not desirable. Also, the inference by the Bayesian network is a relatively expensive operation, which might cause problems in situations where a large group of pedestrians is crossing the street. Therefore, additional intention estimation step (early crossing detection) was implemented that performs the estimation using a simple heuristics, which is run in situations where the pedestrian is very near a crossing. This heuristics is fast and reasonably defensive. It handles all the cases in which the pedestrian is about to enter the vicinity of the crossing zone in a few seconds. The rest of the cases is handled by the Bayesian network.

The trajectory adjustment is performed by a simple particle simulation. The zone around the crossing is divided to several subzones (Figure 7) and the pedestrian is represented by a particle with direction vector and velocity. The crossing entry and exit are visualized by a thick blue and green lines.

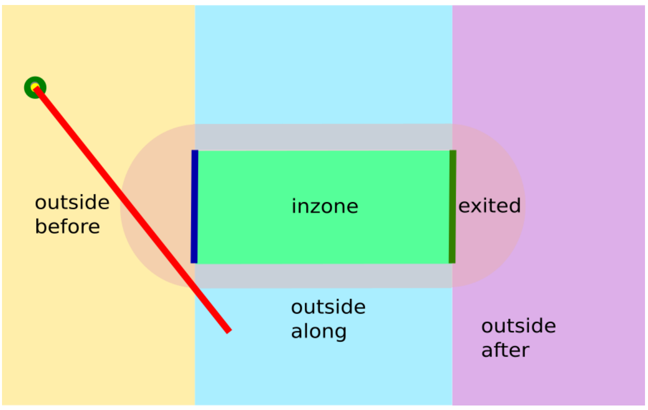

Figure 7: Zones around the crossing affecting the particle movement.

The particle current goal and movement is governed by the zone where it resides at a particular instant. The goal then induces a distinct set of forces applied to the particle in each prediction time step (usually 0.1 seconds). The particle usually starts in the *Outside before* zone. Based on the current zone, the next behaviour is determined roughly as follows:

- If the current zone is *Outside before*, the next goal will be the entry line;

- If the current zone is *Inzone*, the next goal will be exit line;
- If the current zone is *Outside along*, the goal will be *Inzone*;
- Otherwise (the current zone is *Exited* or *Outside after*), the particle will keep the current heading.

This approach proved to be reasonably fast and robust. The heuristics is being extensively tested in the car and new updates of our software module, focusing mainly on reducing the occurrence of false positives, are being implemented.

We have been focusing lately to the integration of this system into the scene/scenario understanding software framework, which is running in the *UP-Drive* research vehicle. We have also worked on improving and testing the proposed system.

In the future, we plan focusing Task 6.3 to more robust motion prediction methods for pedestrians. Task 6.3 will also develop other scenarios than the current one devoted to pedestrians close to the pedestrian crossing.

The next scenario, the development of which ČVUT has started within T6.3 lately, deals with the uncontrolled 4-way intersection. An incoming car and the ego-car drive straight initially. The ego-car has to predict from the incoming car behavior what is the probability that the incoming car turns left (in countries that drive on the right-hand side) and consequently crosses the ego-car trajectory.

AnTool

The extensive use of machine learning methods in many research areas requires large amount of training data. In order to train and evaluate these methods, the data needs to be annotated. This can be performed manually, semi-automatically, or automatically. The manual annotations can be resource consuming, therefore development of tools for automatic annotation of data is beneficial.

We have created *AnTool* – Object intentions annotation tool, which focuses on annotations of traffic–related data. Based on the input data, which are map and detected object data (such as object classification, position, speed, etc.), the intentions of traffic participants can be annotated. These annotations are assigned to the traffic participants based on user-defined rules. The annotations can be further used for development of systems for intention estimation of traffic participants for scenario understanding in autonomous cars.

The uniqueness of the presented tool is in the type of data it is intended to work with. That is, the combination of complex map data describing the layout of the static entities in the environment and dynamic object data detected by various sensors of the ego-car. We are using AnTool in *UP-Drive* project to annotate future actions (intentions) of traffic participants. We are provided by the coordinator the dat-files recording particular test drives mostly in the VW factory in Wolfsburg. The ground-truth, the behaviour of all involved traffic participants with respect to a selected current instant, are available from the recorded dat-file. *AnTool* allows to deal with this data effectively.

Task 6.4 Self-assessment

Data from Work Packages 4-6 need to be analyzed to identify potential data integrity issues or infer context-dependent system limitations.

We see a separate object verification and stabilization step as an important aspect of scene understanding. The core of the verification and stabilization step lies not only in the fusion of the inputs from sensors or examining their temporal consistency but in verification of the consistency with the context data, that for example can be extracted from digital maps. Some substantial work on this aspect has been performed by VW and will be reported in Deliverable 6.3. In the same document, VW's ongoing work on detecting and handling occlusions will be presented.

V. Hlaváč from ČVUT in Prague visited UPC team in Cluj-Napoca on July 26, 2018 also to discuss the cooperation in WP6. UPC team stated that they concentrate their efforts on the supersensor in WP4. It is expected, that work on improvement of perception algorithms will have high positive impact on data quality and will thus render verification and stabilization functionality at least partially obsolete.

Conclusion and Outlook

The work in this package has been lead by ČVUT, with contributions from IBM in Task 6.1 and VW in Tasks 6.2-6.4. This resulted in a number of modules that have been integrated into the vehicle system as well as scientific contributions. In order to maximize the scientific output and to optimize the car performance, partners have agreed to a new task distributions, with ČVUT concentrating on exploration of speculative approaches to object motion prediction, VW pursuing the heuristics-based approaches on same topic and UTC devoting their attention to perception topics of WP4. Role of IBM remains unchanged. The work package is on track.

1.2.7 WP7 Decision-making and Navigation

In Period 2 substantial progress has been achieved in Work Package 7. The topics that have witnessed the most developments include Tactical and Trajectory Planning. Most importantly, a trajectory planner based on dynamic programming has been implemented from scratch and shown to cope well with urban scenarios. This development led to some changes in the architecture within WP7, which have been described in detail in Deliverable D7.2. Work performed in WP7 has been verified both by using a simulation environment, as well as through numerous test drives in the UP-Drive vehicles.

Task 7.1 Route Planning

This module has already provided the necessary functionality in Period 1. In Period 2 only slight changes have been performed to accommodate for the changes in architecture of WP7.

Task 7.2 Tactical Planning

The task of the tactical planning has been revisited whilst working on the new trajectory planning paradigm. One important change is that as we moved from two separate trajectory planners (for road following and parking) to just one, the need for orchestrating them disappeared. The other important change relates to the way the tactical decisions are being made. The relevant data - such as information about lane and road boundaries, zebra crossings, intersections or obstacles - are encoded into so called *scene masks* which are represented as 2D grids. Based on that information, costs for any of the possible maneuvers are computed directly in the trajectory planning stage.

Task 7.3 Trajectory Planning

The initial trajectory planning modules used at the beginning of Period 1 were originating from previous projects - dealing with parking and highway driving. The concepts behind those modules turned out ill-suited to the complex task of urban driving. Thus - building on top of previous experiences - a new trajectory planner was designed and implemented from scratch.

The main concepts behind the resulting trajectory planner are:

- Vehicle motion is forward simulated using motion primitives (acceleration and steering commands).
- Exploration of vehicle motion in a 2D grid.
- Encoding of various costs (collisions, crossing a solid line, lane boundary, etc...) related to visiting of different cells of the 2D grid.
- Optimizing for lowest cost using dynamic programming.
- Usage of GPUs to parallelize computations and achieve real-time performance.

The Deliverable D7.2 gives more details on the concept and its implementation.

Task 7.4 Trajectory Control

Work in this task has initially focused on adapting the controllers developed in other projects to the UP-Drive vehicle platform. This work lead to first automated drives with UP-Drive test vehicle early in Period 2 (Month 15) and allowed for full-system testing. Currently we are working on improving and extending our control interface. Two areas are of interest.

First, a much more precise controller - offering more trajectory fidelity is being analyzed and integrated. This is accompanied by work on enabling a custom - more powerful - steering interface. We expect much more path / trajectory fidelity as well as more natural dynamics of the vehicle once this work is completed. Second, changes to the vehicle interface and controller, which will enable backward driving and parking manouvers, are underway.

Task 7.5 Mission Executive

The mission executive module has been adapted to the changes in the architecture. This includes the simplification due to the fact that only one trajectory planner is being used.

Outlook

The work performed in this work package has lead to 5 out of 11 use-cases from D1.1 already being well addressed:

- Keeping safety distance to leading vehicle
- Constantly adapt the lateral distance to static objects within the own lane
- Compute collision free trajectories to handle pedestrians crossing the street
- Stop at zebra crossings to let pedestrians pass by
- Do lane changes when necessary

The following three use-cases seem well manageable from the perspective of WP7, but pose some challenges to perception and prediction modules:

- Plan collision free trajectories to overtake pedestrians walking along the lane
- Decide whether to wait behind or to pass a vehicle standing in the own lane
- Compute a safe strategy to overtake and pass obstacles in presence of oncoming traffic

Improvements of the vehicle controller will enable handling of additional two use-cases:

- Smooth driving within own lane while keeping the track error very low
- Search for free parking spots and park in and out

Finally, the one use-case that already has been addressed, but will require further attention, is:

- Follow traffic rules while passing intersections

The work on the Tasks 7.1 - 7.5 has been performed by VW, is on track and will continue as outlined above. The report on the algorithmic updates as well as performance of WP7 software will be given in Deliverable D7.3.

1.2.8 WP8 System Integration and Evaluation

Work Package 8 deals with the topics of system integration and evaluation. Period 2 has witnessed a lot of activities in this area - involving all partners and technical work packages. Those activities and their results are described in the following sections.

Task 8.1 Integration plan

The integration plan that has been composed in Period 1 as Deliverable D8.1 has been continuously revised and adapted by the team. The project has adopted a strategy in which partial results have been integrated step by step to the baseline system provided by VW. This turned out crucial for resolving the dependencies - as different parts of the system are in different stages of development and thus synchronous or sequential integration of all results into one system would not be efficient or even feasible.

Task 8.2 System-wide data acquisition

In a distributed project like UP-Drive shared data-sets are key to collaboration. The Deliverable D8.2 from the early phase of P2 lists the data-sets, that have been acquired by Month 14. Since then, the data-set collection has grown considerably - exceeding 10TB in storage space. In order to simplify detection of issues, automated scripts for checking data-set integrity have been introduced and will be continuously extended and improved.

Task 8.3 Integration and test tools and processes

In order to simplify the integration process, a number of tools has been prepared: UP-Drive SDK, SW for playback of data, simulation environment. One important result from Period 2 is the extension of the functionality of the simulation environment, resulting in more realistic behavior of simulated agents and seamless integration into the real system.

Another important outcome is establishing of an orderly test procedure starting from downloading updates of the different system components, including performing tests in the car in a real-world scenario and providing feedback on the observed performance or spotted issues.

More details on the tools and processes can be found in deliverable D8.3.

Task 8.4 System integration

System integration is an activity that is being performed in a number of different ways. First of all, all partners work with the data-sets originating from the project. This greatly helps to understand and prepare for the challenges to be expected in the car. Second, integration weeks and days have been organized, which boost direct collaboration and are very efficient means to tackle integration issues. Within Period 2 a total of over 30 integration days have been organized, which amounts to approximately 9% of working days for that period. Finally, numerous tests of UP-Drive software stack have been performed by the VW team and the results shared with the rest of the project team.

Task 8.5 System evaluation and validation

So far the overall system evaluation has been somewhat limited: a number of real-life traffic tests have been performed and the module performance has been judged - but in purely qualitative terms. Now that the first version of the system has been fully integrated, we expect that the evaluation task will gain on importance.

Conclusion & Outlook

The most relevant outcome of WP8 is an integrated system running in the car - or in case of the mapping pipeline - in the cloud. Its first version has been initially planned for M24. Whilst the project had performed automated drives with partial functionality already in M23, the overall milestone is estimated to be delayed by 9 months - the midterm demonstration has been postponed and aligned with the EC Review in M33. We see the prolonged development phase - especially in the WP4 - as the main reason for that delay. However, compared to other projects of similar or even lower complexity (e.g. V-Charge) we see UP-Drive system to be converging at similar - sufficiently fast - rate.

Apart from the aforementioned delay, the integration work is proceeding as expected and desired: all work packages have provided significant results and all partners have integrated their contributions into one system. The integration work is being lead by VW.

In the Period 3 we expect to continuously extend the feature set of UP-Drive software stack in all work packages and thus improve the overall driving competence and performance.

1.2.9 WP9 Dissemination, Exploitation & Knowledge Management

Work package 9 has the aim of grouping together all dissemination, exploitation and knowledge management tasks within UP-Drive. The objectives of this work package are:

- To transfer knowledge both within the consortium and to the outside world.
- To set up and maintain a comprehensive set of dissemination tools and mechanisms.
- To promote the results developed within the project.

Within the second reporting period, the following deliverables were due:

- D9.3 Initial exploitation plan (M18)
- D9.4 Initial dissemination report (M21)
- D9.5 Mid-term demonstration (M24)
- D9.6 Press video (M27)

D9.3 and D9.4 represent initial versions of the exploitation plan and the dissemination report, respectively. They are living documents that are regularly updated and will be re-submitted by project end. D9.5 and D9.6 have undergone delays: the mid-term demonstration has been postponed to M33 to coincide with the second review meeting. Press video filming has been postponed to M32. This was partly to weather reasons but also allowed for using of the new vehicle platform. The film is being produced and is expected in M33.

Task 9.1 Management of knowledge

Internal knowledge management is subsumed by WP3, and the Gitlab service (wiki, source code management) and Swift cloud object store (mission recordings) in particular. Their implementation is reported in D3.1 – D3.3. Internal documentation and usage of the source code management facility is substantial and continuously growing. The focal point for public knowledge management is assumed by the external project web-page (D9.1). The web-page offers an overview on the project and it maintains continued updated records to documentation, scientific reports and project-related papers. The press video (D9.6) is intended to serve as a prominent, succinct introduction to the project on the web-page's landing page.

Task 9.2 Dissemination

Dissemination activities up to M21 have been reported in D9.4. Preparations for the press video and the mid-term demonstration are underway and will be reported in D9.8: final dissemination report.

Task 9.3 Exploitation

Exploitation activities up to M18 have been reported in D9.3. Ongoing activities will be added to the document and reported in D9.7: final exploitation plan.

Task 9.4 Interaction with the public

Interaction with the public has taken the following parts: leaflet, webpage, presentations / invited talks. They are listed in D9.4. Similarly, the press video will be hosted on the webpage and its innovative play will hopefully garner substantial interest.

Outlook

During the final project year, the final exploitation plan (D9.7, M40) and the final dissemination report (D9.8, M45) will be due. They will be extended and finalised versions of the respective initial deliverables. Most importantly, by the end of the project a public demonstration with participation of press representatives will be performed.

1.2.10 WP10 Project management

Work package 10 spans the scientific, administrative and financial management and monitoring of the project. The monitoring of project progress as well as project-wide decision making takes place in the Steering Committee meetings. Those are held monthly per telephone conference and twice a year as a physical meeting. Regular SC phone conferences were organized, minutes written and actions were followed up. The daily project management and administration were assured and the status of all internal reports and deliverables has been monitored continuously. Additionally, regular Workshops and integration weeks were organized to ensure smooth overall system integration and knowledge transfer.

Table 2: Workshops and integration weeks

Workshop / Integration week	**Place and date**
Calibration workshop	Wolfsburg, February 2017
WP6 workshop	Wolfsburg, February 2017
Localisation workshop	Wolfsburg, February 2017
Perception workshop	Wolfsburg, February 2017
Localisation workshop	Wolfsburg, March 2017
EC Review	Wolfsburg, March 2017
Perception workshop	Wolfsburg, September 2017
Localisation workshop	Wolfsburg, September 2017
Integration week	Wolfsburg, December 2017
Localisation workshop	Wolfsburg, May 2018
Integration week	Wolfsburg, May 2018

The meetings in person have been complemented by very active online collaboration with number of wiki pages and reported issues both approaching more than 200. In summary, the adopted management structure functions properly. The project has a clear vision of its product and is working toward this goal and is thus on track.

2 Conclusions

Overall the project has progressed as planned in Period 2 and delivered very promising results in the challenging field of urban automated driving. Two vehicles with up-to-date systems are fully build-up and operational. The integration and testing of all relevant components has been successfully completed at a system level. The results achieved are encouraging and open up new opportunities for further development. The project partners published around 30 scientific publications during the two and a half years, thus clearly showing the excellence of the results.

The collaboration between the industry and academia partners are of mutual benefit. The academic partners provide fresh innovative ideas that are regularly picked up by the industrial partners. On the other hand, the academic partners have the unique opportunity to test their software and algorithms under real-life conditions from a perspective of a full system. We believe that this form of cooperation is having positive impact on exploitation of results and the quality of scientific output.

Given the complexity of the project - the high interdependency between Partners and Work Packages, it is not surprising that the project consortium experienced some challenges as well. The development phase turned out to be longer than assumed in the initial - and very ambitious - plan. This is especially true for the new low-level sensor fusion approach developed within WP4. Due to the prolonged development phase, the midterm demonstration has been postponed and aligned with the EC Review in M33. However, due to the applied risk mitigation strategies, the delays in one part of the project had only minor effects on the other parts of the project. Moreover, compared to other projects of similar or even lower complexity (e.g. V-Charge) we see UP-Drive system to be converging at similar - sufficiently fast - rate.

Another challenge - which unfortunately still partially remains - is the submission of all Deliverables on time. It results on the one hand from the high number of public deliverables planned in the project and on the other hand from consortium strong focus on scientific publications and integration work.