

Research Article

Road-Type Detection Based on Traffic Sign and Lane Data

Zoltán Fazekas ¹, Gábor Balázs ², Csaba Gyulai ³, Péter Potyondi ³
and Péter Gáspár ¹

¹Institute for Computer Science and Control (SZTAKI), Eötvös Loránd Research Network (ELKH), Kende u. 13-17, H-1111, Budapest, Hungary

²Zukunft Mobility GmbH, D-85055 Ingolstadt, Marie-Curie-Straße 5/5a, Ingolstadt, Germany

³Robert Bosch Ltd., 104, Gyömrői út, H-1103, Budapest, Hungary

Correspondence should be addressed to Zoltán Fazekas; zoltan.fazekas@sztaki.hu

Received 21 December 2021; Revised 1 April 2022; Accepted 19 April 2022; Published 21 May 2022

Academic Editor: Arkatkar Shriniwas

Copyright © 2022 Zoltán Fazekas et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Establishing the current road type constitutes a significant assistance to car drivers, as, by default, the road type determines the legal speed limit. Although there are GPS- and map-based navigation systems that can retrieve the actual road type and speed limit and some can even access and indicate current traffic volumes, it was our aim to develop and test a software prototype of a road-type detection (RTD) system that relies solely on video and sensor data collected on board. Such a system can still work during GPS signal outages. The study presents a heuristic approach to RTD that is based on type and distance data relating to traffic control devices (TCDs) installed along a road. The road is used by an ego vehicle with an on-board smart camera looking ahead and with a number of vehicular sensors. A complex processing step—not detailed in the study—detects TCDs with reasonable probability and error rate and locates them with respect to a 3D coordinate frame fixed to the ego vehicle. The prototype system takes data describing the detected TCDs as its input. This data are then evaluated in a multiscale manner by computing empirical statistics of occurrences over short, medium, and long patches of road. Such an evaluation is carried out in conjunction with each considered road type, and the resulting values are compared to respective reference values. Heuristics is then used in decision-making to resolve any interscale and interroad-type disaccords. The proposed decision rules take into account the possibility of TCDs having been missed and of faulty detections. Short preprocessed synchronised video and signal sequences recorded in different countries and road environments were used for testing the prototype system. These short sequences were carefully strung together into coherent chains. Distance-based recognition precisions 78.9% and 88.9% were gained for European (continental) and for UK roads, respectively.

1. Introduction

Many car drivers, particularly the ones who use obsolete navigation systems in their cars, know the surreal feeling of being lost in space (i.e., not knowing whether on or off the road) and time (more precisely, estimated arrival time). Such a feeling arises when the driver uses some obsolete navigational system and/or software version along some recently built patches of roads. Clearly, some prior effort exerted in updating the navigational software application could prevent the majority of such navigational hiccups. Nonetheless, even taking such a precaution, situations, in which GPS signals are not receivable, do occur [1].

According to an unpublished survey carried out by Bosch some years back, for more than 5% of the path length along the roads surveyed within Austria, Germany, Hungary, and the UK, the current road type was wrongly given by a then state-of-the-art GPS-based navigation system. Considering this worse than expected GPS coverage, Bosch Hungary initiated a joint research and development (R&D) project on road-type detection (RTD). Obviously, GPS-based routing and detection systems have been getting more and more reliable since then; however, their coverage is still not 100%.

Over time, public road networks tend to grow (i.e., new roads are built, and when ready, these are incorporated into

the road network). However, road networks may also shrink, which is a serious concern for the drivers. Occasionally, such shrinkages occur unexpectedly, e.g., some roads get blocked (e.g., due to an accident) or get damaged (e.g., due to water pipe breakage). In other cases, the shrinkage is linked to roadworks or to temporary road closures. These are the cases when the on-board RTD capability gains real importance.

Normally, these and similar situations can be algorithmically detected based on video, point cloud, and/or sensor data. Certain detection methods rely on data collected on the spot [2], while others consider data gathered from some wider neighbourhood [3]. Also, one could differentiate between static sensors (e.g., fixed traffic surveillance cameras [4] and LiDARs [5] installed over or by the roads) and mobile ones (e.g., cameras installed on board an ego car [6] or cameras and/or GPS receivers carried by a number of probe vehicles floating in the traffic [7–9]) used for the purpose.

In road transport and road traffic-related applications, fixed cameras are used mostly in busy junctions, and their role is to safeguard the continuity and the undisturbed state of the road traffic there. The video stream originating from such cameras reaches and can be inspected and analysed in urban or regional traffic control centres. Fixed cameras are also installed to monitor straight patches of road. Such installations appear mostly along expressways and motorways, where the vehicle speeds are usually rather high. Often, even these cameras—and their respective individual or central traffic surveillance software—cannot track each and every vehicle passing by, for this reason, individual speed warnings are usually not possible. Even so, lane-based speed measurements with speed warnings could be a viable option for controlling speeders. Such warnings could be displayed on varying message signs (VMSs). Authorities responsible for road management and road safety need, however, to consider costs involved in providing and maintaining such electronic solutions against the benefits achieved through their use. Presently, however, such detection and warning facilities do not provide full coverage, not even for major roads, and will not do so in the foreseeable future either.

It is a common practice—at least in our country, but presumably in other countries as well—to install cameras/speed detectors at the entry points of built-up areas. These then view and monitor incoming vehicles and check their speed and warn, if necessary, the drivers to decrease the vehicle speed below the legal limit.

Though installations of fixed cameras—together with some textual or sign-based feedback means—at regular distances over the road network could result in higher security, increased traffic safety, and faster first response—in case of accidents—and even in improved traffic moral, particularly if the roads are monitored continuously in some automated and/or artificial intelligence supported manner, in respect to the present target application (i.e., RTD), there are other competitive static/fixed alternatives. For instance, road-side communication units could broadcast the local—fixed or varying—road type/speed limit using some V2X communication channel. An even simpler and quite driver-friendly approach in providing speed-limit/road-type

information in an on-going manner is repetition. For instance, in the Netherlands, the local speed limit is repeated—using small versions of the traffic signs—on every single line pole along the roads (at least along certain roads), so whenever in doubt, the driver can watch out for the next line pole on the road side and see the applicable speed limit. Unfortunately, this simple solution is not part of the road management practice elsewhere in Europe.

Now, let us turn our attention to mobile camera-based solutions. In the context of road transport, it often means that one or more camera is installed on board an ego car or on an ego vehicle. These terms pop up many times herein and pops up for good reason. Terms “ego car” and “ego vehicle” have come to the automotive field from computer vision (CV), and according to [10], the latter term is defined as follows: “subject connected and/or automated vehicle, the behaviour of which is of primary interest in testing, trialling, or operational scenarios.” The term “ego vehicle” can be used interchangeably with “subject vehicle” and “vehicle under test.” The term is directly related to “ego motion” used in CV, see the definition of this term in [11].

Let us now consider briefly the detection of roadworks—from the above examples for the construction of the road network—using an on-board camera installed on the ego car. Roadwork sites nowadays are carefully railed off and marked garishly with a series of warning traffic signs (TSs) and with a series of caution lights. Such sites can be identified and located, e.g., by an extended, camera-based TS recognition (TSR) system [12] installed on board an ego car. Also, shorter than usual lateral distances of the detected TSs from the current lane—i.e., the lane used by the ego car—could hint at roadworks being conducted there.

The edification of this example is that certain image-based road transport related on-board measurements and data analysis, e.g., for providing some novel or at least uncommon ADAS functions, can be implemented through appropriate extension and/or cooperation of different driver assistance/autonomous driving (AD) subsystems available on board [13, 14]. Various subsystems of the advanced driver assistance systems (ADAS) and of AD systems are often used—directly or indirectly—in smaller scale detection solutions [15, 16]. In such solutions, the ADAS/AD subsystems tend to profoundly rely on and have access to intelligent cameras, though in many cases data originating from other sensors is also drawn upon.

These sensors provide information on the situation within the car [17] and on the environment surrounding the car [18], e.g., in respect of the location, shape, and quality of road [6, 19], on the location, size, type, shape, colour, co-occurrence, etc., of traffic control devices (TCDs), such as road markings [20] and TSs and traffic lights [21], and of various static and dynamic elements of the road environment [22], which are often represented in a local dynamic map [23].

Lim and Bräunl have recently published a methodological review as a pre-print [24] on the topic of visual road detection and recognition for AD applications. They look at approaches, methods, and procedures of road detection and

recognition and evaluate and compare the practical implementations of these.

Though road detection/recognition is not the same as RTD targeted herein, the former has high relevance to the latter, as in fact, vice versa. This is because road detection/recognition and RTD go hand in hand with each other: for RTD, one has to know where the road is within the image or video frame and can then analyse the corresponding image region closely; while if one intends to detect the road or the driveable area of the road, it is good to know (e.g., it saves computing time and/or reduces the need for costly computing resources suitable for automotive applications) what sort of road to look for in an image or in a video sequence.

The first part of the review by Lim and Bräunl focuses on conventional road detection algorithms and approaches, which can distinguish between roads from nonroad regions. In the second part, they survey state-of-the-art machine learning techniques that have already been applied to visual road recognition tasks. The authors are primarily interested in convolutional neural network (CNN)-based techniques that are applied for semantic segmentation. In a dedicated section, the authors overview some relevant implementations coming from the industrial/commercial sector and mention major alliances formed in this sector in regards to R&D tasks pertaining to road detection and recognition.

Though, the review cites more than hundred research papers in its target field and serves, therefore, as a good starting point for anyone who intends to start research in the field, here we comment only on two points from the section dedicated to industrial/commercial implementations and alliances.

Firstly, according to the mentioned section, a start-up company called “comma.ai” specialises in providing assisted and AD systems to the consumer market. Their goal is to achieve full AD with existing road vehicles with after-market devices. Herein, as well as in some of recent publications, we have taken a somewhat similar approach in regards to novel or at least not that common ADAS functions, methods, and subsystems [12–14, 25]. We rely and build upon existing ADAS functions/subsystems, e.g., lane-keep assist (LKA) and TSR that are widely available in modern production cars.

Secondly, also in the mentioned section, Lim and Bräunl refer to AD functions, algorithms, and systems developed by original equipment manufacturers (OEMs) of sensors and computers (e.g., Mobileye, Nvidia, Velodyne, and FLIR), international corporates (e.g., Google and Uber), and automotive manufacturers (e.g., BMW, Volvo, and Daimler). Clearly, the enumeration was not meant to be exhaustive, but still another category, namely, that of the automotive component and subsystem manufacturers, such as Bosch and Continental, could also have been mentioned either as a separate group or as a subgroup of the automotive manufacturers. For this reason, we refer here also to two other reviews [26, 27], as well as to a conference paper [28] and to a technical news [29]. These communications highlight some important activities within the field, which are carried out by this (sub-) group. These activities include extensive R&D projects, as well as

setting up useful publicly accessible AD datasets, and last but not least forming important industrial alliances with road vehicle manufacturers. The activities at Bosch provide a backdrop to our present pilot study.

Intervehicular (i.e., vehicle-to-vehicle, V2V) and vehicle-to-infrastructure (V2I), vehicle-to-network (V2N), and vehicle-to-pedestrian (V2P) communications are fundamental technological driving forces that make AD happen, spread, and proliferate [30, 31]. New solutions, e.g., the side-link communication [32, 33], open new horizons, and time scale in local communications, and thereby make possible to avoid or mitigate critical situations arising in busy road traffic. Considering these technologies in the context of RTD, on the one hand, the use of V2X communication means and technology could increase its reliability; on the other hand, complete reliance on software used by or on the measurements and computations carried out by other vehicles is not necessarily a good idea. It is particularly true in regards to safety critical automotive applications and to applications that can accrue financial consequences (e.g., longer routes and speeding tickets) as in case of navigation on roads and also RTD. For instance, how should one dare to modify a proprietary or even open road database without proper moderation of the incoming data? What about accidents arising due to such modifications/extensions? Who will be responsible for the data that turn out to be simply wrong or even worse corrupted? Clearly, the security and legal responsibility aspect requires further research and reliable solutions for integrating external and/or crowd-sourced data in this context. Furthermore, there are realistic cases when there is no information available from other road traffic participants (e.g., in the sparse road traffic late at night), so reliable camera-based road-type detection for the ego-vehicle is still a primary research and development target.

The notion of road type is used herein as a collective term for various public roads that share some important characteristics. The notion includes motorways, express roads, main roads, and other roads [34]. These road types are associated with different spatial arrangements (e.g., lane layout and number of lanes), geometrical dimensions, and allowable road connections. Though, being car drivers ourselves, it is quite clear for us what these perceptible similarities and differences between these road types are, but as engineers and researcher—developing transport and vehicular systems and applications—we should look also at the motivations and the reasons behind these. In the civil and transport engineering literature, roads are categorized by their functional forms. The choice of the functional form of a road is controlled by standards and design guidelines. As Wolhuter phrases it in the Functional Classification of Roads section of his Geometric Design of Roads Handbook [35], “For geometric design, the most useful form of classification is functional classification, as it defines the spectrum of road usage from pure mobility to pure accessibility. This, in turn, supports the selection of the design speed and the design vehicle. These two parameters, in combination with current and anticipated traffic volumes, define geometric standards of the horizontal and vertical

alignment and intersections or interchanges and definition of the cross section.”

Since the writing of the above cited handbook, road vehicles of increased automation and connectivity have appeared on public roads. A recent review looks into the necessary and/or induced changes in the geometric and other design parameters of the roads and the road infrastructure (e.g., lanes can be narrower for AD vehicles) due to the fast-evolving, highly automated, and connected transport and vehicle technology [36]. These important changes mean that, in the above quotation, the automation and connectivity level of design vehicles should be taken as a crucial design parameter in the road design from now on.

At some later stage of road design process, the type and location of the necessary TCD including TSs and lane markers are specified. On the one hand, the—default and actual—speed limits and other related road characteristics are set in accordance with the road’s functional categorization. On the other, from these speed limits—either indicated by explicit speed limit signs or by traffic signs indicating the actual road type—and from the other road characteristics, it is possible to infer—either statistically or logically—the type of the road currently used. For instance, a speed limit of 90 km/h traffic sign will not occur by a road or street within a built-up area, and a speed limit of 30 km/h traffic sign will not occur on motorways and expressways (except perhaps in case of some on-going roadworks, or in case of certain critical, or extraordinary traffic situations and even then possibly indicated by variable sign boards). These and similar statistical or logical observations are used in inferring the road type of the road currently used.

Research related to our present topic, that is to automatic identification/detection of the current road type, appears in the smart driving and AD literature with far less emphasis and as the topic of significantly smaller number of publications than, say, the intensively researched and widely communicated driver assistance tasks of lane detection and TSR do.

The knowledge of the current road type and therefore also its on-board identification, i.e., the automated RTD ADAS/AD function, are important for the following main reasons. Firstly, at many road locations, it is the road type that determines the legal speed limit, as this is the default case. Secondly, different rules apply for and different TCDs are expected along different road types and the driver should be aware—for obvious traffic safety reasons—of the applicable rules and of the probable TCDs that may appear. Thirdly, the knowledge of the actual road type can be used to optimize fuel consumption and to choose appropriate mode of operation (e.g., speed, acceleration and deceleration ranges, and appropriate gear) for the ego vehicle, not only in case of an autonomous vehicle [37] but also in case of a smart vehicle driven by a human driver [38]. Fourthly, the work load of a human driver can be better estimated, traced, and monitored by the smart ego vehicle if it has on-going access to the current road-type information. Fifthly, if two roads are located very close to each other and run parallel to each other, then the RTD function could help the on-board navigation device to properly place the ego car on its map.

In the next few paragraphs, the history of the RTD research is summarized via sketching the main contributions of five relevant publications written in the last decade.

A paper by Tang and Breckon from 2011 analyses images—taken by a low-cost forward-looking camera mounted on board an ego car—in order to distinguish between on-road and off-road environments [39]. Their method computes and relies on various colour and texture features extracted from multiple regions of interest within each image. A trained classifier was used to resolve this two-class classification problem. The multiclass road environment classification problem relating to the off-road, urban, major/trunk road, and multilane motorway/carriageway environments was also addressed in their paper. A good classification performance—achieved at a near real-time classification rate—was reported for the former problem, while the results in respect of the latter problem were somewhat less impressive.

In 2012, the topic of RTD popped up again in a conference paper by Taylor et al. [38]. They analysed the driving speed and some other signals made available by various intravehicular sensors (e.g., steering wheel angle, gear position, and suspension movement). Thereby, they analysed the manner in which drivers drove on the road, rather than the road itself or its environment. Indeed, the car drivers assisted the RTD rather than the other way round. The authors, referred to, had access to and compared their results to those achieved by an existing, but unpublished RTD system developed by a major car manufacturer. According to the authors, the unpublished system followed a model-based approach, while they used a data mining approach instead. The input data used by Taylor et al. were collected from the controller area network (CAN). The data collection exercise encompassed a single car that was driven by a number of drivers on UK roads. Several data mining and temporal analysis techniques—along with a number of ensemble classifiers—were deployed for the purpose of RTD. According to the findings reported in the paper, the random forest ensemble algorithm with access to summaries of the speed and the steering data for the last few seconds achieved a good classification performance.

In 2014, Huang et al. classified urban roads into three classes based on their conditions, namely, into clean roads without lane markings, simple roads with lane markings and with a small amount of disturbances (e.g., vehicles), and complex roads with lane markings and with a large number of disturbances [40]. They used this classification to improve their lane detection results. Their lane detection method processed bird’s eye view images—gained through appropriate transformation from forward-looking camera images—and implemented class-dependent strategies to complete any missing/hidden lane marking segments.

Also, in 2014, Slavkovikj et al. presented an image-based algorithm for classification of paved and unpaved roads [41]. Their method was programmed to learn discriminative features from training data in an unsupervised manner. For validation purposes, the authors set up a road image dataset that consists of a total of 20,000 sample images. These sample images had been taken at thousands of different paved and the unpaved

road locations. The presented experimental results—in regard to images of the mentioned dataset—indicate that their algorithm can achieve good performance for the aforementioned two-class road classification problem.

In 2016, Seeger et al. presented a number of methods for road-type classification [42] based on fused occupancy grids. These occupancy grids were built from a number of point clouds acquired by various LIDAR and radar sensors mounted on an ego car and from the spatial reconstruction derived from stereo images acquired by an on-board stereo camera. The occupancy grids corresponding to different time instants were analysed separately, thereby information was lost on previous classifications and semantic labelling. The road types considered in the study were freeways, highways, parking lots, and urban roads. The authors compared the performance of an end-to-end convolutional neural network classifier to that of a support vector machine that had been trained on hand-crafted features. They tested the various methods on a dataset that contains 700 local occupancy grids (LOGs)—for training purposes—in regard to each of the mentioned four road categories and 150 LOGs—for testing purposes—again in regard to each road category. Some of their methods achieved test accuracies over 90%.

In 2019, Balado et al. presented results concerning an important subtask of RTD, namely, concerning the road environment semantic segmentation [43]. Their input data consisted of mobile laser scanning point clouds that were automatically divided into sections during the preprocessing stage. The authors presented competitive classification results—achieved via the use of the PointNet deep learning architecture—in respect of the road surface, ditches, embankments, guardrails, borders, and fences, as well as other road objects and vegetation. All these road infrastructure elements and objects serve as important traits in the identification of road types. Table 1 provides a summary of the RTD methods and systems cited herein.

2. Materials and Methods

2.1. The Development of Software Prototype for Road-Type Detection. In the frame of an R&D project, it was our aim to build a software prototype of the RTD ADAS/AD function that can determine the type of a road based solely on video and vehicular data originating from on-board camera and from other on-board sensors, respectively, and on data derived from such data. In other words, GPS information and map data were not to be relied upon.

On the one hand, such a function should be supranational, in the sense that it can determine the type of a road that is located in any country of the world, regardless of the hand of traffic used there, regardless of the varying styles of TCDs across countries, and—theoretically—regardless of the intensity of road traffic. On the other hand, the RTD could be more precise if the country, where this ADAS function is to be used, is known by the RTD subsystem in advance.

Road scenes can be rather complicated and difficult to grasp instantly by a human driver or by an autonomous self-driving system. Therefore, in the given context, some

intelligent filtering mechanism that excludes data concerning TCDs that do not pertain to the ego vehicle is of great significance [44].

Even if well-tested processing modules are used and even if the mentioned intelligent reliable filtering mechanisms are in place, the reliance on prior processing steps and intermediate results should not be absolute. In practice, during the execution of complex processing steps—such as image segmentation, object recognition, and tracking—errors inevitably occur over time, e.g., TSs installed along the road can be missed due to complete or partial occlusion by vehicles, the intelligent object filtering might be impeded for the same reason, or TSs can be misinterpreted by the TSR module due to their fading colour. Against these hopefully infrequent detection, tracking, and classification processing errors, the RTD method must be robust.

In order to confine the R&D work and to make good use of the available software resources, the development of a partial software prototype (PSP) of the RTD ADAS/AD function was agreed upon. This meant that existing proprietary solutions could be utilized for the complex processing steps mentioned above.

2.2. The Development Framework Chosen for the Road-Type Detection. The abovementioned existing solutions were provided as modules to be executed within the Automotive Data and Time-triggered Framework (ADTF) [45]. This automotive software development framework was created and has been marketed by a global supplier of embedded and connected software products and services for the automotive industry. Consecutive versions of ADTF have been serving the needs of developers of automotive measurement, data processing, and data communication solutions for more than a decade now and since then have been used worldwide in numerous projects implementing, testing, and evaluating various ADAS/AD functions, see e.g., [46, 47].

As the existing modules that we could rely on in building the software prototype were developed to be used in ADTF, it was quickly decided that also the RTD PSP should be implemented in and run by the same framework. Furthermore, it was decided that the RTD PSP would take type and distance data relating to certain TCDs installed/painted along the road used by the ego car as its input. To compute these data, existing proprietary software modules were used in respect of real-time TCD recognition and real-time object localisation. The former included the subtasks of TSR, lane-marking, and lane detection, while the latter did the subtasks of lateral localisation of the ego car within its current lane and the TS localisation with respect to the ego car.

The respective intermediate results were stored—also in a synchronised manner—together with the synchronised video and vehicular data in compound data sequences for further processing. We note here that video sequences inherent in these compound sequences can be played real time with ADTF, while the associated actual signal values can be displayed numerically or as registered curves along with the video sequence also in real time.

TABLE 1: Summary of the most important features of the cited research papers on road-type detection methods and systems.

Authors and year	Aim	Sensors	Features	Performance	Remark
Tang and Breckon, 2011	(i) On/off road	(i) Low-cost forward-looking camera	(i) Frame-based processing (ii) Multiple ROIs (iii) Using colour and texture features	Good	(i) Involving an unknown number of vehicles
As above	(i) Several road types	As above	As above	Medium	As above
Taylor et al., 2012	(i) Several road types	(i) Intravehicular sensors	(i) Recording of CAN data (ii) Random forest ensemble classifier	Medium	(i) Involving a single vehicle (ii) Relies heavily on driving speed and style
Huang et al., 2014	(i) "Strange" road types (ii) Improve lane detection results	(i) Forward-looking camera	(i) Transforming images into bird's eye view images (ii) Class-dependent strategies to complete missing or hidden lane markings (iii) Vehicles as disturbances	Specific to the concrete application	(i) Involving an unknown number of vehicles
Slavkovikj et al., 2014	(i) Paved and the unpaved roads	(i) Camera viewing in varied directions (ii) Also Google street view images	(i) Learning discriminative features from training data in an unsupervised manner (ii) A dataset of 20,000 road-surface subimages was created	Good results (in respect of the database)	(i) Involving an unknown number of vehicles (ii) Image sequences taken by cyclists
Seeger et al., 2016	(i) Freeways, high-ways, parking lots, and urban roads (ii) Sensor fusion	(i) Automotive LiDARs (ii) Automotive radars (iii) Forward-looking on-board stereo camera	(i) Based on fused local occupancy grids (LOGs) taken at discrete time instants (ii) No tracking between LOGs	Very good	(i) Involving an unknown number of vehicles (ii) 4 * 700 LOGs for training and 4 * 150 LOGs for testing
Balado et al., 2019	(i) Semantic segmentation of the road environment (ii) Road surface, ditches, embankments, guard rails, etc.	(i) Automotive LiDARs	(i) Automatic sectioning of the point cloud (ii) Using the PointNet deep learning architecture	Very good	(i) Important subtask of road-type detection (ii) Involving a single vehicle

A plug-in module can be used to extract a video sequence and/or various signal values from a stored compound sequence, process the associated data, and display the results synchronised to the video in real time. In line with these, the PSP was implemented as a plug-in module, which was added to the other processing modules. The synchronised data communication between the modules is guaranteed by the ADTF, and the detection results—in this case RTD results—could be shown in real time. The processing carried out within a plug-in module is very similar to the processing carried out by a real-time routine—developed for the same purpose—that is executed within an on-board intelligent automotive camera.

2.3. The Inputs of the Partial Software Prototype. In the frame of the R&D project mentioned above, a few hundred short, i.e., from 10-second to 3-minute long, compound sequences were made available for us. These compound sequences were

annotated with the ground truth road types (i.e., roads within built-up area, country roads, expressways, and motorways). The annotation was carried out by the authors of the present paper based on the inherent video sequences stored within the compound sequences.

The RTD PSP developed in the project relies on a number of input signals that can be extracted from the stored compound sequences. These input signals comprise type, distance, and dimensional data related to TCDs. The method implemented in the RTD PSP considers different pieces of information for the different road types. The input signals relied on are the width of the current lane, types of the detected TSs—including those types that are directly linked to one of the road types—as well as the lateral, longitudinal, and vertical positions of the TSs with respect to ego car and spatial frequencies of the detected TSs.

2.4. The Relationship between the Input Signals and the Road Types. Firstly, the statistical relationships between the input

signals and the road types had to be empirically established and analysed, as such empirical data provide the basis for inferring the current road type from the input signals. If, for example, most of the input signals have values that are characteristic to a particular road type, then the chances are good that the actual patch of road belongs to that road type.

One of the input signals chosen for RTD purpose was the width of the current lane. The lane widths had been computed frame by frame in the preprocessing phase by the lane detection module mentioned in Section 2.2, and this lane-width data had been stored in the compound sequences. These stored lane widths were sampled for the purpose of RTD at equidistant path lengths along the ego-car's trajectory. The sampling process is sketched in Figure 1.

Still as part of the data gathering, the number of occurrences was counted for each input signal and road-type pair in the prerecorded compound sequences. We note here that some of these sequences were used solely for data gathering and training purposes, while different ones were used for testing. The full range of each input signal was partitioned into subranges. Table 2 shows, for example, the partitioning of the full lane-width range into subranges or "bins." The lane-width distribution over these bins is given in the same table for each road type. These empirical distributions are also presented graphically in the bar chart of Figure 2 in different colours.

As it can be seen from the bar chart, the most likely bin both for built-in area roads and for country roads is Bin 3, more precisely, Bin 3 for roads within the built-up areas, and Bin 3 for country roads, respectively. According to the top row of Table 2, Bin 3 corresponds to the sampled lane-widths falling between 3.2 m and 3.6 m. For expressways and motorways, on the contrary, the most likely bin is Bin 4, more precisely Bin 4 for expressways, and Bin 4 for motorways, respectively. According to the top row of Table 2, Bin 4 corresponds to lane widths falling between 3.6 m and 4.0 m.

According to the above observation, expressway and motorway lanes tend to be somewhat wider than lanes of roads within built-up areas and of country roads. This shift in lane width is, of course, intentional and is due to the road design standards and guidelines that serve to ensure the high level of road safety. This shift in lane widths reflects the intended functions of the roads and constitutes an aspect of their functional forms. With respect to RTD, this shift in lane widths serves as a clue for decision-making.

The described partitioning helps to determine the characteristic signal values for each road type, but it does not provide an obvious basis for comparison across road types. Clearly, there exist several mathematically well-founded statistical classification methods with proven optimality and these could have been used for the purpose [48]. However, to allow for intuition and traceability, a unified and simple evaluation scheme was devised to represent the influence of input signal values on the decision to be made about road types and could provide an easy-to-comprehend basis for the comparison of bins across road types. For this scheme, an indicator value—called "score"—was introduced. Each bin within the full range of each input signal was associated with




FIGURE 1: Sampling lane width of the current lane.

a rating called "score" for each road type. See the example of the lane-width scores in Table 3.

The way we assign scores to bins is pragmatic: the scores are derived through simple calculations. The scores aim to characterize the relative likeliness of the considered road types for a given input value. If, for instance, the lane currently used by the ego car is 3 m wide, then, according to Table 3, the ego car is probably being driven in a built-up area that has the highest score, namely, 5, though the location could perhaps be a country road that has the second highest score, namely, 3. On the contrary, the location is probably neither an expressway nor a motorway location as for both road types the associated scores are rather low, namely, -5.

The score associated with a road type (rt) and a bin number (bn)—denoted by $\text{score}_{rt,bn}$ —has been calculated according to the following simple formula:

$$\text{score}_{rt,bn} = \text{minscore} + (\text{maxscore} - \text{minscore}) \cdot \frac{\text{occurrences}_{rt,bn}}{\max_{bn}(\text{occurrences}_{rt,bn})} \quad (1)$$

The occurrences $_{rt,bn}$ can be looked up in Table 2 for road type rt and bin number bn , while the minscore and the maxscore values were chosen experimentally, in a way that a negative score value associated with a certain road type—especially when the negative score summed up with corresponding scores calculated for TS densities and TS types still remains negative—makes the given road type a highly unlikely candidate, while some positive score value indicates that the given road type is a realistic choice.

Having assigned scores to each bin of each input signal for each road type, a look-up table (LUT) can be constructed. Then, the scores for any value of any input signal can be extracted for any of the road type from this LUT.

In our final road-type choice for a given road location, we take into account the scores of each considered input signal for that location, as well as the scores corresponding to a number of road locations that have already been passed by the ego car. These processing steps towards RTD are detailed in the following sections.

TABLE 2: The number of lane-width occurrences—measured and counted at equidistant discrete path lengths—within each lane-width range, or “bin,” computed for the chains of compound sequences used for training.

Lane-width ranges	<2.8 m	2.8–3.2 m	3.2–3.6 m	3.6–4 m	4–4.4 m	4.4–4.8 m	>4.8 m
Built-up area	926	4800	7554	4929	1580	612	996
Country road	2634	12080	22999	7782	2138	1575	678
Expressway	125	454	3422	7133	3065	129	98
Motorway	51	770	9548	17434	3549	590	735




FIGURE 2: Empirical distributions of the sampled lane widths over the lane-width subranges (bins) according to Table 2.

TABLE 3: The “scores”—associated with the lane-width “bins” shown in Table 2—used in the proposed RTD method.

Lane-width ranges	<2.8 m	2.8–3.2 m	3.2–3.6 m	3.6–4 m	4–4.4 m	4.4–4.8 m	>4.8 m
Associated bin number	1	2	3	4	5	6	7
Built-up area	−4	5	12	6	−2	−5	−4
Country road	−4	3	12	0	−4	−5	−5
Expressway	−6	−5	3	12	2	−6	−6
Motorway	−6	−5	4	12	−2	−5	−5

2.5. Cumulative Effect of the Input Signals. As the video sequence—inherent in a compound sequence—runs and the data stored in the compound sequence is processed by the ADTF modules, each considered input signal takes on some value. At the next road location to be considered, the PSP looks up the scores corresponding to the actual input values for each road type from the LUT. The lane-width subarray of this LUT is given in Table 3. The local scores for a road location—derived using the mentioned LUT—appear in the corresponding columns of Table 4. Note that the same lane width—i.e., 3.0 m—of the current lane is assumed here as in the example given in the previous section; this way the first numeric column in Table 4 is the same as the score column for Bin 2 in Table 3.

The extracted scores of the individual input signals are summed within each row of the table, resulting in sum of scores for each road type. The final decision on the perceived road type is based on these summed scores. The sampling and accumulation of the input signals is done in an equidistant manner, more precisely, equidistant over

the path covered by the ego vehicle, and the sum of scores corresponding to these road locations are stored in first-in-first-out (FIFO) queues. Such a score queue is shown in Table 5. The summed scores for the most recent road location are also indicated in the table. These appear at the top of the queues, while the prior summed scores appear below. As the car moves, the oldest row of scores eventually drops out, so after the initial fill up, the length of the queue is kept constant.

2.6. Distance Scales for RTD. The score queues are evaluated at three different distance scales. In these scales, data corresponding to road locations within a short range, within a medium range, and within a long range from the current location are evaluated. So, for example, the short-range (SR) evaluation takes into account the scores of the most recently covered patch of road. The corresponding scores appear in the first few rows of the score queues. In the sample score

TABLE 4: The local scores computed for lane width, TS types, and the spatial frequency of the TSs—measured/taken at a particular road location—in respect to the considered road types. These scores are summed up for each road type in the right most column.

	Scores computed for lane width	Scores computed for TS types	Scores computed for TS spatial frequency	Sum of these scores
Built-up area	5	15	10	30
Country road	3	5	7	15
Expressway	-5	-5	0	-10
Motorway	-5	-10	1	-14

TABLE 5: The score queues for the considered road types.

Most recent road location	Included in short range	Included in medium range	Included in long range	Scores for built-up area	Scores for country road	Scores for expressway	Scores for motorway
✓	✓	✓	✓	30	15	-10	-14
	✓	✓	✓	29	10	-12	-15
	✓	✓	✓	26	14	-12	-16
	✓	✓	✓	27	10	-12	-15
	✓	✓	✓	28	11	-12	-12
		✓	✓	20	11	-12	-12
		✓	✓	26	9	-10	-12
	
		✓	✓	13	9	-5	4
		✓	✓	13	9	-5	5
			✓	5	3	2	10
			✓	6	3	2	12
		
			✓	5	3	6	15
			✓	6	0	8	16

TABLE 6: Aggregated short-range scores computed from the scores appearing in Table 5.

Short-range aggregate score for built-up area	Short-range aggregate score for country road	Short-range aggregate score for expressway	Short-range aggregate score for motorway
140	60	-58	-72

queues presented in Table 5, the first five data rows constitute the SR.

The medium-range (MR) evaluation looks at a somewhat larger number of rows, while the long-range (LR) evaluation considers an even larger number of rows of the score queues. These ranges are indicated with ticks in the data rows. The summed scores are aggregated—again simply summed—for each distance scale. As an example, the aggregated SR scores for the considered road types—calculated from the summed scores given in Table 5—are presented in Table 6.

The motivations for evaluating the score queues over different distance scales are to enable the RTD to react quickly to abrupt road-type changes, on the one hand, and to make robust and reliable RTDs that rely on the continuity of the road characteristics along a given road, on the other.

Admittedly, the multiscale heuristic RTD approach detailed below is a hand-crafted approach for detecting change in stochastic signals in the given context. It should be noted that this topic, i.e., change detection in stochastic signals, has a vast mathematically well-founded literature on its own right, see [49, 50].

In case of a sudden road-type change, the subsequent scores associated with the new road type suddenly start to

grow. MR evaluation of the scores could be used to support or defy the road-type candidates chosen by SR and the LR evaluations, especially, if the two candidate road-types are different. The MR, however, will not be used in the heuristic decision rules proposed herein.

Having derived the aggregated scores for SR, MR, and LR, respectively, a method that determines the road type from these aggregated scores is required. The method should take into account the dynamics of road-type changes in the given country and should compare the aggregated road-type scores within and across the distance scales.

2.7. Categories Used in the Road-Type Selection. A heuristic approach was devised for the abovementioned comparison, which introduces categories that are used in conjunction with each distance scale. Based on the evaluation of the aggregated scores, a road type may be assigned to each category for each distance scale. These categories are as follows.

Very best (i.e., road type with the best aggregated score by far): the aggregated score for a given road type—over a certain distance scale—exceeds each of the other road-types’ aggregated scores—over the same scale—at least by some

predefined value (different predefined threshold values are used in the different distance scales).

Greatest: the aggregated score for a given road type—over a certain distance scale—exceeds each of the other road-types' aggregated scores over the same scale.

Second greatest: the aggregated score for a given road type—over a certain distance scale—exceeds each but one of the other road-types' aggregated scores over the same scale.

Worst-by-far (i.e., the road-type with the lowest aggregated score by far): the aggregated score for a given road type—over a certain distance scale—is less than each of the other road-types' aggregated scores, respectively—over the same scale—at least by some predefined value (different predefined threshold values are used in the different distance scales).

For each distance scale, the aggregated scores are checked against the above criteria. Over each distance scale, each category is associated with at most one road type at a particular road location. To ensure this property—even when there are a number of identical aggregated scores over some distance scale—the following convention is used: if there are two or more equally high aggregated scores over a particular distance scale, then the greatest and the second greatest categories are assigned randomly amongst these high aggregated scores.

In Table 7, the SR categories assigned to the four road-types considered herein are presented for a particular road location as an example. These categories have been calculated from the aggregated scores given in Table 6. The score difference that was required for the “very best” category in SR was 40, while for the “worst-by-far” category in SR, it was 20. Using these thresholds, the “very best” category was assigned to a road type, namely, to the roads within built-up area, while the “worst-by-far” category was not assigned to any.

2.8. Dealing with Incorrect Input Signals. Besides the SR, MR, and LR evaluations, a further somewhat special evaluation is made. It takes into account all the input signals and detected TCDs except for the TSs directly indicating a road type.

The motivation to include this category is to make the detector capable of filtering out the cumulative effect of the road-type-related TSs, such as “start of expressway” or “end of expressway” TSs. The latter TS appears at the road location shown in Figure 3, but the information conveyed by the TS does not apply for the ego car. If this TS is still considered, then it deflects the TS statistics and possibly even the perceived road type.

This fourth evaluation has been implemented only for the long-range computations and abbreviated as LWO.

To illustrate the multiscale approach taken towards RTD, two screenshots taken while running the RTD PSP are included herein. These show two road locations in Hungary and display the corresponding multiscale categories for them. The categories and their associated road types appear at bottom of the screenshots; these are shown in Figures 4 and 5, respectively. The categories and their associated road types are repeated in Tables 8 and 9, respectively, for better

legibility. We use there the denominations applied throughout in the present study (e.g., “medium range” is used instead of “midterm,” “built-up area” instead of “in-town,” and “very best” instead of “Kicker;” the latter was abbreviated as “K” in the screenshot).

2.9. Heuristic Decision Rules Used in the RTD Method. Taking the category values determined for the three different distance ranges (i.e., SR, MR, and LR), as well as for LWO, and their associations with various road types as input, sequential decision rules—i.e., rules to be evaluated from the first one proceeding to consecutive ones until one of the rules evaluates true—are used to determine the perceived road-type. Though, in the present rule set, the MR labels are not used at all, neither are certain categories in SR and LR (e.g., SR second greatest); it should be noted that these could be still useful, if the rule set is to be revised (e.g., other road features are also to be taken into account).

The evaluation of the decision rules is carried out in the following general order. The evaluation starts with that of the SR-related rules. These rules are evaluated firstly to ensure that the RTD software reacts promptly to sudden changes in the road type. The evaluation the LWO is carried out next. The motivation for this special “range” coming next was to ensure that incidental faulty or highly unlikely TS detections are identified as such at an early stage of the decision-making so that their disadvantageous effect can be minimised. Then, the rules concerning the LR categories are evaluated.

The cases which are specifically addressed with rules are given below. We note here that similar, but fewer decision rules, had been used in conjunction with a road environment-type (RET) detection method, which was proposed in [11]. The method presented there relies on TS and crossroad data. A tabular representation of these rules was suggested in [25]. Herein, the rules are represented graphically—in the form of rule-tables—using a slightly modified version of the representation proposed therein. The following colour convention is used in the rule tables; green cell: the perceived road type will be the one associated with that category according to the given rule; pink cell: the perceived road type will not be the one associated with that category; it will be the one appearing in the green cell. Still the given cell is referred to in the given rule. Greyish cell: the road type will be set unknown if the rule is applicable; white cells with grey characters: categories and associated road types are not referred to in the rule. Cells left blank: categories and associated road types not used at all in the rule set. Between cells, solid and dashed lines represent equality and nonequality, respectively. These are placed according to the conditions appearing in the rules. These should be read as “the road type associated with the category given in the first cell is equal/not equal to the one associated with the category given in second cell.”

Rule 1: if there is a road type associated with the SR very best category and it differs from the one associated with LR greatest, then the former road type is chosen as perceived. Explanation of the rule: the situation described in the rule typically occurs when there is a clear-cut, easy-to-detect road-type change implied by a reliable detection of an

TABLE 7: Short-range categories assigned to road types based on the aggregated scores in Table 6. The score difference required for the “very best” category—in this scale—was 40, while for the “worst-by-far” category—also in this scale—it was 20.

Short-range categories for built-up area	Short-range categories for country road	Short-range categories for expressway	Short-range categories for motorway
Very best and greatest	Second greatest	N/A	N/A



FIGURE 3: A directly road-type-related TS—namely, the “end of expressway” TS appearing in the right-hand side of the photo—that does not apply for the ego car moving in the left lane.



FIGURE 4: Screenshot showing a road location on a motorway near Budapest, Hungary. The multiscale category values displayed at the bottom are repeated in Table 8 for readability.

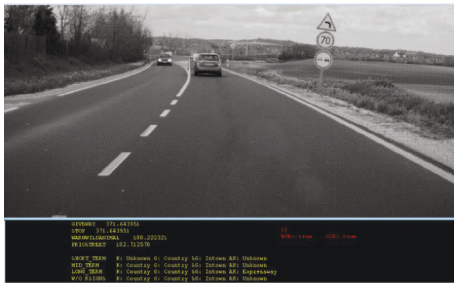


FIGURE 5: Screenshot showing a road location on a country road in Hungary.

explicit road-type-related TS. The tabular representation of the rule is given in Table 10.

Rule 2: if there is a road type associated with the SR worst-by-far and it is the same as the one associated with LR greatest, then the road type associated with the SR greatest is declared as perceived. The tabular representation of the rule is given in Table 11.

Rule 3: if there is a road type associated with the LWO very best and it is the same as the one associated with LR second greatest, then this common road type is chosen as perceived. Explanation of the rule: the situation described typically occurs when an incorrectly identified explicit road-type-related TS diverts the LR greatest; in this case, the LR evaluation of the TSs and of the road data—disregarding

TABLE 8: Multiscale category values derived for the road location shown in Figure 4.

SR	
Very best	Motorway
Greatest	Motorway
Second greatest	Country road
Worst-by-far	N/A
MR	
Very best	N/A
Greatest	Country road
Second greatest	Motorway
Worst-by-far	Expressway
LR	
Very best	N/A
Greatest	Built-up area
Second greatest	Country road
Worst-by-far	Expressway
LWO	
Very best	N/A
Greatest	Built-up area
Second greatest	Country road
Worst-by-far	Motorway

road-type-related TSs—may set the situation straight. The tabular representation of the rule is given in Table 12.

Rule 4: if there is a road type associated with the LWO worst-by-far, and it is the same as the one associated with LR greatest; furthermore, the road type associated with LWO greatest is the same as the one associated with the LR second greatest; then, the latter road type is declared as perceived. Explanation of the rule: this typically happens when an explicit road-type-related TS is detected incorrectly. The tabular representation of the rule is given in Table 13.

Rule 5: if the road type associated with the LR greatest is the same as the one associated with the LWO worst-by-far, then the perceived road type is set to unknown. The tabular representation of the rule is given in Table 14.

Rule 6a: if the road type associated with the LR greatest is the same as the one associated with LWO greatest, then this common road type is chosen as perceived. The tabular representation of the rule is given in Table 15.

TABLE 9: Multiscale category values derived for the road-location shown in Figure 5.

SR	
Very best	N/A
Greatest	Country road
Second greatest	Built-up area
Worst-by-far	N/A
MR	
Very best	Country road
Greatest	Country road
Second greatest	Built-up area
Worst-by-far	N/A
LR	
Very best	Country road
Greatest	Country road
Second greatest	Built-up area
Worst-by-far	Expressway
LWO	
Very best	Country road
Greatest	Country road
Second greatest	Built-up area
Worst-by-far	N/A

TABLE 10: Tabular representation of Rule 1.

Short range	Long-range	Long-range w/o TSs
Very best		Very best
Greatest	Greatest	Greatest
	2nd greatest	2nd greatest
Worst-by-far		Worst-by-far

TABLE 11: Tabular representation of Rule 2.

Short range	Long-range	Long-range w/o TSs
Very best		Very best
Greatest	Greatest	Greatest
	2nd greatest	2nd greatest
Worst-by-far		Worst-by-far

TABLE 12: Tabular representation of Rule 3.

Short-range	Long-range	Long-range w/o TSs
Very best		Very best
Greatest	Greatest	Greatest
	2nd greatest	2nd greatest
Worst-by-far		Worst-by-far

TABLE 13: Tabular representation of Rule 4.

Short-range	Long-range	Long-range w/o TSs
Very best		Very best
Greatest	Greatest	Greatest
	2nd greatest	2nd greatest
Worst-by-far		Worst-by-far

Rule 6b: if the road type associated with the LR greatest is the same as the one associated with LWO second greatest,

TABLE 14: Tabular representation of Rule 5.

Short range	Long-range	Long-range w/o TSs
Very best		Very best
Greatest	Greatest	Greatest
	2nd greatest	2nd greatest
Worst-by-far		Worst-by-far

TABLE 15: Tabular representation of Rule 6a.

Short-range	Long-range	Long-range w/o TSs
Very best		Very best
Greatest	Greatest	Greatest
	2nd greatest	2nd greatest
Worst-by-far		Worst-by-far

TABLE 16: Tabular representation of Rule 6b.

Short-range	Long-range	Long-range w/o TSs
Very best		Very best
Greatest	Greatest	Greatest
	2nd greatest	2nd greatest
Worst-by-far		Worst-by-far

TABLE 17: Tabular representation of Rule 6c.

Short-range	Long-range	Long-range w/o TSs
Very best		Very best
Greatest	Greatest	Greatest
	2nd greatest	2nd greatest
Worst-by-far		Worst-by-far

TABLE 18: Tabular representation of Rule 7.

Short-range	Long-range	Long-range w/o TSs
Very best		Very best
Greatest	Greatest	Greatest
	2nd greatest	2nd greatest
Worst-by-far		Worst-by-far

then this common road type is chosen as perceived. The tabular representation of the rule is given in Table 16.

Rule 6c: if the road type associated with the LR second greatest is the same as the one associated with LWO greatest, then this common road type is chosen as perceived. The tabular representation of the rule is given in Table 17.

Rule 7: if none of the above conditions have been satisfied, then the perceived road type is set to unknown. The tabular representation of the rule is given in Table 18.

The sequential application of the above rules is shown in Figure 6. The rules that have one or more green cell in their tabular representations (i.e., Rules 1, 2, 3, 4, 6a, 6b, and 6c) lead to some concrete road type (e.g., motorway) on their “applicable” branches, while the rules with one or more grey cell in their tabular representations (i.e., Rules 5 and 7) lead to unknown road type. Note that Rule 7 is not represented explicitly in Figure 6; it is

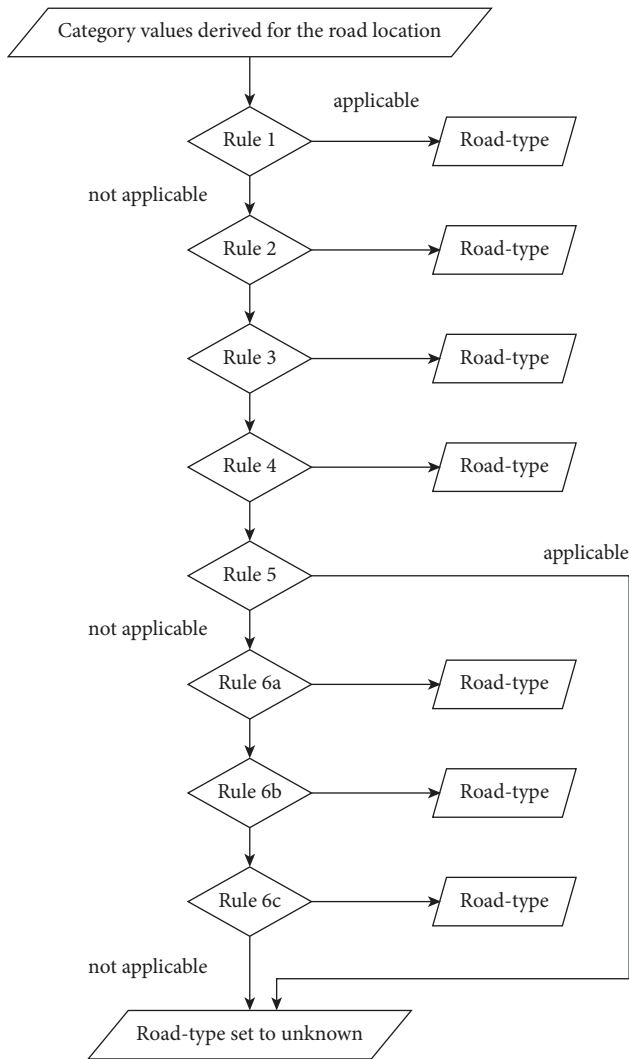


FIGURE 6: The sequential application of Rules 1–5, 6a, 6b, and 6c leading to the decision on the current road type. *Note.* Rule 7 is represented only as the “not applicable” branch of Rule 6c.




FIGURE 7: The steps through which the road-type decision is reached for road location shown in Figure 3.

instead represented by the “not applicable” branch of Rule 6c.

To illustrate the use of the above decision rule set, let us now use the road locations shown in Figures 4 and 5 as examples. The category values derived for these road locations are given in Tables 8 and 9, respectively.

Let us consider first the road location shown in Figure 4. As the road type associated with the SR very best category—i.e., motorway—differs from the one associated with LR greatest—i.e., road in built-up area—in Table 8, the former road type is identified as perceived based on Rule 1. No further rules need

to be evaluated in this case, as it can be seen in Figure 7. The road-type classification reached is correct in this case.

For the road location shown in Figure 5, the first five rules in the heuristic rule set need to be evaluated before Rule 6a finally leads to a positive decision on the road type. The process of decision-making is shown in Figure 8. The road location is perceived as Country road. Again, the road-type classification reached is correct.

2.10. Evaluation of Road Data with the RTD Software Prototype. A car-based data collection had been carried out by Bosch engineers and their support team along roads of Austria, Germany, Hungary, and the UK before the launch of the R&D project mentioned in Section 2.1. They used an unspecified number of cars equipped with intelligent mono cameras. These cameras were developed by Bosch for automotive applications.

The video recordings and the corresponding vehicular data were stored as compound sequences that can be opened, played/displayed, evaluated, and processed using ADTF modules. These compound sequences were later at some stage pre-processed using various low- and medium-level image processing, image understanding, object recognition, object tracking, computer vision and spatial measurement ADTF modules, and the relevant outputs and results were stored back in the sequences.

An assortment of several hundred incoherent, unsorted, and preprocessed short compound sequences were made available for the authors from the collected road data. These sequences were then annotated with the ground truth road types by the authors. A good portion of these annotated sequences were used in the algorithm development—i.e., for training and validation—purposes; others were used for testing purposes.

Having viewed all the road video sequences made available for the project, we believe that the total number of cars that had been involved in recording these sequences was between two and four.

The road-type composition of training data is presented in Table 19, while its hand-of-traffic—or rule-of-road—composition is given in Table 20.

The speed of the ego car(s) was always below 50 km/h for built-up area roads, below 90 km/h for country roads, below 110 km/h for expressways, and below 130 km/h for motorways. However, for most of the time, the driving speed was close to the respective speed limits. The traffic appeared to be relatively light in the short sequences that we received. All the recordings were taken at daylight.

Though different weather conditions occurred during the data collection and were then recorded in short sequences, the effect of the unfavourable weather conditions on the RTD was not systematically analysed. This was the case also for road objects (e.g., roundabouts and flyovers) that were not in direct focus of the present study. Still, such road objects do occur in the recordings, as can be verified for a particular UK chain that is described in Table 21.

We chose to experiment only with reliable TS and lane detections, reliable according to the proprietary ADTF detector modules used in our experiments. These reliability




FIGURE 8: The steps through which the road-type decision is reached for road location shown in Figure 5.

TABLE 19: The road-type composition of the training road data.

Total length of roads within built-up areas	76 (km)
Total length of country roads	91
Total length of expressway	54
Total length of motorways	58
Total road length	279

TABLE 20: Hand-of-traffic composition of training road data.

Total length of roads with right-hand traffic (recorded within Austria, Germany, and Hungary)	203 (km)
Total length of roads with left-hand traffic (recorded within the UK)	91
Total road length	279

values were stored in and were read from the preprocessed compound sequence. The TS and lane detections were considered for further processing, only if their reliability values were over a fixed threshold. So, if the detections were unreliable—e.g., due to some unfavourable weather and light conditions or due to the fading colours on TSs or due to dim lane markers—for a longer period or for a longer patch of road; then, the RTD was suspended and the road type was set to unknown.

The outputs of the mentioned detector modules included, among others, the number of detected lanes, the geometrical characterisation, and the types of detected lane markings, the types of detected TSs, and the distance to TSs measured from the ego car.

In order to ensure that the testing of the RTD software prototype was carried out in a consistent, replicable, and fair manner, furthermore to make up for the lack of longer coherent compound sequences and at the same time utilize the short ones that were available, several chains of such compound sequences

were compiled. These chains were then of more practical lengths: ranging from few minutes to about half an hour.

Table 21 summarizes the road objects and features appearing in the video sequences inherent within a particular chain. The video sequences listed in the table were recorded in the UK, so they show left-hand traffic. This fact can be verified via viewing characteristic snapshots from these videos, which are presented as shown in Figure 9. This particular chain of compound sequences was used in testing the RTD software prototype.

The chains compiled from component sequences feature patches of road of the considered road types, as well as a good selection and variety of road-type changes. The latter trait of the chains ensures that not only the RTD capability but also the road-type change detection capability of the software prototype was put to the test.

It should be noted that, during the manual editing/compilation of each chain, special care was taken to render the component sequences in such a manner and order that there are no harsh differences—in weather, daytime, hand-of-traffic, number of lanes, etc.—between adjacent component sequences. As a consequence, each of the chains used in our experiments covers and shows a more or less realistic and consistent road journey. The compiled chains feature roads in Germany, Austria, Hungary, and the UK. A few of the chains focus on roads near Budapest, Hungary; Figures 4 and 5 show road locations from such a chain.

3. Results and Discussion

The RTD software prototype was tested on a number of compiled chains. In numerous cases, its operation and the generated outputs were closely followed and monitored by a member of the testing team. As part of the testing

TABLE 21: Brief description of the component sequences that make up a particular compiled chain. All its component sequences had been recorded in at daytime on UK roads. The following abbreviations are used: COU=country road, BUA=built-up area, and MWY= motorway.

Seq. no.	Weather conditions	Road type(s)	Road objects encountered	Number of lanes	Number of TSs	Extraneous TSs	Duration (s)
1	Cloudy, wet	COU, MWY	TSs, roundabout, MWY-entry	1, then 3	7	3	126
2	Cloudy	MWY, COU	TSs, flyovers, MWY-exit	3	6	1	123
3	Cloudy, wet	COU	TSs, roundabout	1	13	3	349
4	Clear	COU, BUA	TSs, roundabout	1	15	0	300
5	Clear	BUA	TSs, curving streets	1	0	0	56
6	Clear	BUA	Not applicable	1	0	0	20
7	Clear	BUA	Traffic lights	1	0	0	19
8	Clear	BUA	Roundabout	1	0	1	20
9	Cloudy	BUA, COU	TSs	1	14	1	248
10	Clear	MWY	TSs, flyovers, and MWYs merging	3	16	3	338
Total					71	12	1599



(a)



(b)



(c)



(d)



(e)



(f)



(g)



(h)

FIGURE 9: Continued.



FIGURE 9: Characteristic snapshots from the video sequences inherent within the component sequences of the chain summarized in Table 21. The figures are numbered with the sequence numbers given therein.




FIGURE 10: Diagrams of the RTD software prototype's recognition precision with respect to time (a) and the path length (b), respectively, for the chain described in Table 21. The duration of this chain is about 27 minutes; the total path length covered by the ego vehicle is about 28 km.

procedure, precisions were calculated by comparing the ground-truth road types to the road types detected by the prototype. The precisions were calculated over time, as well as over the distance covered by the ego car. In the time-based calculations, the precisions were computed on a frame-by-frame basis, while for the path-length-based calculations, frames were sampled in an equidistant manner over the path.

In the RTD publications cited in Section 1, the RTD precisions of 80% or higher were mentioned. Some more recent papers could achieve road-type precision over 90%. Most of these results, however, were not derived through real time or real-time capable solutions.

Furthermore, telling that a precision of 80% is good and of 95% is excellent is not that clear cut, it is important to understand the concrete application targeted, or phrased in another way, how and for what purpose the results are used. Is the task at hand safety critical? Does it involve paying costs like fines? The same precision achieved can be acceptable or not, depending on the given context and circumstances of the very concrete application of RTD.

The two precisions for the UK chain described in Table 21 are presented diagrammatically in the two subfigures of Figure 10. The graphs show how the precision values changed with time and with path length, respectively. These graphs are presented in the top and the bottom subfigure, respectively. In

TABLE 22: The achieved recognition precision of the RTD software prototype for different chains of compound sequences.

Chains of compound sequences showing	Time based (%)	Distance based (%)
Roads in Austria, Germany, and Hungary	77.2	78.9
Roads in the UK	89.4	88.9
Roads near Budapest (Hungary)	84.9	85.0

the figures, the coloured vertical lines indicate the changes in the ground truth road type. The light green background indicates that the RTD was correct in the given moment or at the given road location, while the orange background indicates incorrect RTD in that moment or at that road location. Table 22 provides the precisions achieved for the test chains.

4. Conclusions

The software prototype implementing the RTD method described herein was tested thoroughly. It performed adequately in detecting the road types both in case of the European (continental) roads and of the UK roads. The categorizing method used herein brought satisfying results and seems to be applicable in this context.

The future work to enhance the RTD software prototype presented herein includes taking more input signals into account (e.g., road/lane curvature, presence of crossroads, and number of lanes), using the types and along-the-route locations of other TCDs (e.g., traffic lights and pedestrian crossings) in the classification, combining/fusing the various input values, and considering and testing mathematically more well-founded classification and change detection methods that have already performed well in similar applications. The review by Lim and Bräunl [24] has goaded us to consider and suggest a further research direction within RTD, namely, RTD directly from video stream using CNNs and/or recurrent neural networks (RNNs). Having analysed the road detection results referred to and summarized in the review, this research direction could bear fruit in the near future.

The statistical data used for the training of the detector were collected in parallel with the development of the method; therefore, the number of the samples was moderate and should hence be considerably increased. Furthermore, there is regional and country-wise variability of the road features (e.g., with respect to TS density, average lane-width, and TS designs). In order to handle this sort of variability, an international data collection on a large scale should be carried out so that the developers active in the road detection and RTD field could fully understand the international and interregional differences in road conditions and road features and their effect on the classification process and precision.

We intend to collect further empirical data from other European countries and also from countries outside Europe to make the software prototype robust to radically different road environments. Then, with such international road data collected, for handling relatively small (e.g., regional) differences in

the functional forms and in topographical environments, algorithms with learning capabilities could be deployed to adapt thresholds used in our present work.

Data Availability

The synchronised video and vehicular signal sequences—referred to in the text as compound sequences—and the synthetic chains—compiled from such compound sequences—that are used to support the findings of this study have not been made available because of commercial confidentiality.

Conflicts of Interest

The authors declare that there are no conflicts of interest regarding the publication of this article.

Acknowledgments

This research was supported by the Ministry of Innovation and Technology NRDI Office within the framework of the Autonomous Systems National Laboratory Program.

References

- [1] T. G. Reid, B. Chan, A. Goel et al., “Satellite navigation for the age of autonomy,” in *Proceedings of the 2020 IEEE/ION Position, Location and Navigation Symposium*, pp. 342–352, IEEE, Portland, OR, USA, April 2020.
- [2] R. Spinneker, C. Koch, S. Park, and J. J. Yoon, “Fast fog detection for camera based advanced driver assistance systems,” in *Proceedings of the 17th International IEEE Conference on Intelligent Transportation Systems*, pp. 1369–1374, Qingdao, China, October 2014.
- [3] A. Pascale, F. Deflorio, M. Nicoli, B. Dalla Chiara, and M. Pedroli, “Motorway speed pattern identification from floating vehicle data for freight applications,” *Transportation Research Part C: Emerging Technologies*, vol. 51, pp. 104–119, 2015.
- [4] L. Yu, P. Chen, W. Liu, G. Kang, and A. G. Hauptmann, “Training-free monocular 3D event detection system for traffic surveillance,” in *Proceedings of the 2019 IEEE International Conference on Big Data*, pp. 3838–3843, IEEE, Los Angeles, CA, USA, December 2019.
- [5] J. Wu, H. Xu, and J. Zhao, “Automatic lane identification using the roadside LiDAR sensors,” *IEEE Intelligent Transportation Systems Magazine*, vol. 12, no. 1, pp. 25–34, 2020.
- [6] F. Kortmann, P. Fassmeyer, B. Funk, and P. Drews, “Watch out, pothole! Featuring road damage detection in an end-to-end system for autonomous driving,” 2021, https://www.researchgate.net/profile/Felix-Kortmann-2/publication/355486932_Watch_out_pothole_Featuring_Road_Damage_Detection_in_an_End-to-end_System_for_Autonomous_Driving/links/%2061754b27eef53e51e1dd273d/Watch-out-pothole-Featuring-Road-Damage-Detection-in-an.
- [7] W. Ma and S. Qian, “High-resolution Traffic Sensing with Autonomous Vehicles,” 2019, <https://arxiv.org/abs/1910.02376>.
- [8] O. Popescu, S. Sha-Mohammad, H. Abdel-Wahab, D. C. Popescu, and S. El-Tawab, “Automatic incident detection in intelligent transportation systems using aggregation of traffic parameters collected through V2I communications,” *IEEE Intelligent Transportation Systems Magazine*, vol. 9, no. 2, pp. 64–75, 2017.

- [9] M. Yashina, M. A. Belova, and A. I. Mokhov, "Method of video-measurements of traffic flow characteristics at a road junction," *Journal of Information Technology Management*, vol. 12, no. 3, pp. 34–43, 2020.
- [10] BSI, "CAV Vocabulary," 2020, <https://www.bsigroup.com/en-GB/CAV/cav-vocabulary/ego-vehicle/>.
- [11] Wikipedia, "Visual odometry," 2022, https://en.wikipedia.org/wiki/Visual_odometry#Egomotion.
- [12] Z. Fazekas, G. Balázs, L. Gerencsér, and P. Gáspár, "Locating roadworks sites via detecting change in lateral positions of traffic signs measured relative to the ego-car," *Transportation Research Procedia*, vol. 27, pp. 341–348, 2017.
- [13] Z. Fazekas, G. Balázs, and P. Gáspár, "A heuristic approach to road environment-type detection from traffic sign data," in *Proceedings of the 4th Conference on Control and Fault Tolerant Systems*, pp. 288–293, Casablanca, Morocco, September 2019.
- [14] Z. Fazekas, L. Gerencsér, and P. Gáspár, "Detecting change in road environment via analysis of marked point processes associated with traffic signs," *Transportation Research Procedia*, vol. 22, pp. 75–84, 2017.
- [15] R. E. Haas, S. Bhattacharjee, and D. P. F. Möller, "Advanced driver assistance systems," in *Smart Technologies*, K. B. Akhilesh and D. P. F. Möller, Eds., Springer, Singapore, pp. 345–371, 2020.
- [16] C. Badue, R. Guidolini, R. V. Carneiro et al., "Self-driving cars: a survey," *Expert Systems with Applications*, vol. 165, Article ID 113816, 2021.
- [17] M. J. S. Emashharawi, O. O. Khalifa, N. Abdul Malik, and N. F. Abdul Malek, "Computer vision based driver assistance drowsiness detection," Edited by K. Isa, Z. Zain, R. Mohd-Mokhtar et al., Eds., in *Proceedings of the 12th National Technical Seminar on Unmanned System Technology 2020*, vol. 770, Singapore, Springer, 2022.
- [18] M. M. Trivedi, T. Gandhi, and J. McCall, "Looking-in and looking-out of a vehicle: computer-vision-based enhanced vehicle safety," *IEEE Transactions on Intelligent Transportation Systems*, vol. 8, no. 1, pp. 108–120, 2007.
- [19] F. Kortmann, H. Peitzmeier, N. Meier, J. Heger, and P. Drews, "Enabling road condition monitoring with an on-board vehicle sensor setup," in *Proceedings of the 2019 IEEE Sensors*, pp. 1–4, Montreal, QC, Canada, October 2019.
- [20] H. U. Khan, A. R. Ali, A. Hassan, A. Ali, W. Kazmi, and A. Zaheer, "Lane detection using lane boundary marker network with road geometry constraints," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 1834–1843, Snowmass, CO, USA, March 2020.
- [21] Y. Qin, Y. Chen, and K. Lin, "Quantifying the effects of visual road information on drivers' speed choices to promote self-explaining roads," *International Journal of Environmental Research and Public Health*, vol. 17, no. 7, Article ID 2437, 2020.
- [22] F. Damerow, Y. Li, T. Puphal, B. Flade, and J. Eggert, "Intersection warning system for occlusion risks using relational local dynamic maps," *IEEE Intelligent Transportation Systems Magazine*, vol. 10, no. 4, pp. 47–59, 2018.
- [23] J. Eggert, D. A. Salazar, T. Puphal, and B. Flade, "Driving situation analysis with relational local dynamic maps (R-LDM)," in *Proceedings of the Symposium on Future Active Safety Technology*, p. 8, Nara, Japan, 2017.
- [24] K. L. Lim and T. Bräunl, "A Methodological Review of Visual Road Recognition Procedures for Autonomous Driving Applications," 2019, <https://arxiv.org/abs/1905.01635>.
- [25] Z. Fazekas, G. Balázs, A. Boulmakoul, and P. Gáspár, "Visualization of traffic sign related rules used in road environment-type detection," in *Proceedings of the MOSATT*, pp. 53–57, IEEE, Kosice, Slovakia, November 2019.
- [26] E. Marti, M. A. de Miguel, F. Garcia, and J. Perez, "A review of sensor technologies for perception in automated driving," *IEEE Intelligent Transportation Systems Magazine*, vol. 11, no. 4, pp. 94–108, 2019.
- [27] A. S. Mohammed, A. Amamou, F. K. Ayevide, S. Kelouwani, K. Agbossou, and N. Zioui, "The perception system of intelligent ground vehicles in all weather conditions: a systematic literature review," *Sensors*, vol. 20, no. 22, p. 6532, 2020.
- [28] T. R. Kaushek Kumar, S. Thiruvikkraman, R. Gokul, A. Nirmal, and R. Karthika, "Evaluating the scalability of a multi-object detector trained with multiple datasets," in *Proceedings of the 5th International Conference on Intelligent Computing and Control Systems (ICICCS)*, pp. 1359–1366, IEEE, Madurai, India, May 2021.
- [29] J. M. Amend, "Bosch AI-camera sharper than human eye," 2019, <https://www.wardsauto.com/technology/bosch-ai-camera-sharper-human-eye>.
- [30] A. Zanella, A. Bazzi, G. Pasolini, and B. M. Masini, "On the impact of routing strategies on the interference of ad hoc wireless networks," *IEEE Transactions on Communications*, vol. 61, no. 10, pp. 4322–4333, 2013.
- [31] C. K. Toh, J. C. Cano, C. Fernandez-Laguia, P. Manzoni, and C. T. Calafate, "Wireless digital traffic signs of the future," *IET Networks*, vol. 8, no. 1, pp. 74–78, 2019.
- [32] B. Vejlggaard, "Sidelink: unlocking the full potential of device communication with 5G," 2020, <https://www.bell-labs.com/institute/blog/sidelink-unlocking-full-potential-device-communication-5g/>.
- [33] A. Bazzi, A. O. Berthet, C. Campolo, B. M. Masini, A. Molinaro, and A. Zanella, "On the design of sidelink for cellular V2X: a literature review and outlook for future," *IEEE Access*, vol. 9, Article ID 97953, 2021.
- [34] European Road Safety Observatory, "Motorways 2018," p. 41, 2018, https://ec.europa.eu/transport/road_safety/sites/default/files/pdf/ersosynthesis2018-motorways.pdf.
- [35] K. Wolhuter, *Geometric Design of Roads Handbook*, p. 626, CRC Press, Boca Raton, FL, USA, 2015.
- [36] M. M. Rana and K. Hossain, "Connected and autonomous vehicles and infrastructures: a literature review," *International Journal of Pavement Research and Technology*, pp. 1–21, 2021.
- [37] A. Chehri and H. T. Mouftah, "Autonomous vehicles in the sustainable cities, the beginning of a green adventure," *Sustainable Cities and Society*, vol. 51, Article ID 101751, 2019.
- [38] P. Taylor, S. S. Anand, N. Griffiths, F. Adamu-Fika, A. Dunoyer, and T. Popham, "Road type classification through data mining," in *Proceedings of the Fourth International Conference Automotive User Interfaces and Interactive Vehicular Applications*, pp. 233–240, Portsmouth, NH, USA, October 2012.
- [39] I. Tang and T. P. Breckon, "Automatic road environment classification," *IEEE Transactions on Intelligent Transportation Systems*, vol. 12, no. 2, pp. 476–484, 2011.
- [40] J. Huang, H. Liang, Z. Wang, Y. Song, and Y. Deng, "Lane marking detection based on adaptive threshold segmentation and road classification," in *Proceedings of the IEEE International Conference on Robotics and Biomimetics*, pp. 291–296, Bali, Indonesia, December 2014.
- [41] V. Slavkovikj, S. Verstockt, W. D. Neve, S. V. Hoecke, and R. V. Walle, "Image-based road type classification," in

- Proceedings of the Twenty Second International Conference on Pattern Recognition*, pp. 2359–2364, Stockholm, Sweden, August 2014.
- [42] C. Seeger, A. Müller, L. Schwarz, and M. Manz, “Towards road type classification with occupancy grids,” in *Proceedings of the IEEE Intelligent Vehicles Symposium*, vol. 2, Gothenburg, Sweden, June 2016.
- [43] J. Balado, J. Martínez-Sánchez, P. Arias, and A. Novo, “Road environment semantic segmentation with deep learning from MLS point cloud data,” *Sensors*, vol. 19, no. 16, p. 3466, 2019.
- [44] A. Bódis-Szomorú, T. Dabóczy, and Z. Fazekas, “A lane detection algorithm based on wide-baseline stereo vision for advanced driver assistance,” in *Proceedings of the Seventh Conference of Hungarian Association for Image Processing and Pattern Recognition*, D. Chetverikov and T. Szirányi, Eds., pp. 142–152pp. 142–, Budapest, Hungary, 2009.
- [45] Elektrobit, “ADAS development environment EB Assist ADTF (automotive data and time-triggered framework),” 2021, <https://www.elektrobit.com/products/eb-assist/adtf/>.
- [46] M. Kuttila, P. Pyykönen, P. van Koningsbruggen, N. Pallaro, and J. Pérez-Rastelli, “The DESERVE project: towards future ADAS functions,” in *Proceedings of the International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation*, pp. 308–313, IEEE, Agios Konstantinos, Greece, July 2014.
- [47] P. Nickel and T. Schöpfner, “Sensors replace senses - ADAS data acquisition for functional development,” *ATZelectronics worldwide*, vol. 16, no. 3, pp. 48–51, 2021.
- [48] L. Tang, X. Yang, Z. Kan, and Q. Li, “Lane-level road information mining from vehicle GPS trajectories based on naïve Bayesian classification,” *ISPRS International Journal of Geo-Information*, vol. 4, no. 4, pp. 2660–2680, 2015.
- [49] E. S. Page, “Continuous inspection schemes,” *Biometrika*, vol. 41, no. 1-2, pp. 100–115, 1954.
- [50] G. Lorden, “Procedures for reacting to a change in distribution,” *The Annals of Mathematical Statistics*, vol. 42, no. 6, pp. 1897–1908, 1971.