
22.1 Introduction

With autonomous driving, a technological system will replace humans in driving automobiles. The car industry, universities, and large IT companies, are currently working on implementing functions permitting a technological system to take on vehicle operation. Their focus is on the tasks that are also done by humans: of perception, cognition, deciding how to act (planning) and carrying out this behavior (acting). In addition, humans possess further capabilities not directly connected to driving a vehicle. For example, learning directly changes people's capacity to tackle tasks. In the driver-vehicle-environment system, this human capability raises a question: Will the technological system that is to replace humans also exhibit a capacity to learn? In the most diverse fields, primarily IT-driven ones, there are learning and learned systems of the most varied kinds, which rival conventional analytical systems in their performance. What marks out vehicle automation, though, is firstly its relevance to safety; secondly, how cars as a product additionally differ from other IT-industry goods in their system life cycles. Both of these particularities, with their challenges and attempts at solutions, are the subject of this chapter. Attention is also given to collective learning in the context of autonomous driving, as directly exchanging with and copying from the learned is one of the particular advantages machine learning has over the human version.

Replicating human learning in machine learning occupies a whole area of research. It is expected that examining the processes of human learning will both provide a deeper

W. Wachenfeld (✉)

Institute of Automotive Engineering, Technische Universität Darmstadt,
64287 Darmstadt, Germany

H. Winner

Institute of Automotive Engineering, Technische Universität Darmstadt, FZD, 64287
Darmstadt, Germany

understanding of this process and improve applied machine-learning methods. With reference to the currently understood difference between both forms of learning, machine learning is understood in this chapter as an algorithm generated by a human. The running of the software follows these algorithms, just as with all other software. Our objective is not to compare a learning human with a learning robot. Rather, it is to discuss why, whether, and with which challenges and approaches machine learning is possible in its current form in autonomous driving. This chapter portrays the view of vehicle technology in particular on this question, and is based on experience from the literature for the area of machine learning.

22.2 Vehicle, Environment and Learning Drivers

When a driver drives a vehicle in any environment, it constitutes target-oriented behavior. According to Rasmussen [1, 2], people display behavior here that can be divided into three areas. They behave according to their skills, rules, and knowledge (see Fig. 21.6 in Chap. 21). Skill-based behavior is described by Rasmussen [1] as stimulus-response automatism, which people overcome in routine everyday situations without placing intensive demands on cognitive capacities.

Rule-based behavior proves to be more cognitively demanding. It requires associative classification in addition to perception and motor actions. People in this case match the recognized situation to a known rule and, based on this, select from a repertoire of behavioral rules. People have learned these rules purposely, or noticed (“saved”) them in past situations and behavior. This gives people the ability to identify similar situations and transfer learned rules to them.

Should situations arise that are new for people, and for which they have no trained behavior, then their response will be knowledge-based. People try, based on their trained knowledge, to generate and evaluate the alternative ways to act that are available to them. The subjectively optimal alternative is selected and carried out.

Rasmussen’s [1] understanding of target-oriented behavior makes clear what is meant by a learning driver. At the beginning of their “career,” a driver builds up their basic knowledge of road traffic in theory courses at driving school. This is then tested. In the process, this knowledge builds on what has already been learned by living in society. In addition, rule-based behavior is trained via theory and practice lessons. Upon obtaining their driver’s license, people can then drive on public roads without further supervision (exception: driving when accompanied with a license at 17 in Germany). However, drivers at this point in time have neither learned all the rules nor processed the knowledge needed for their future lives on the road. With each new piece of experience they gain, their behavior transforms from knowledge-based to rule-based, and from rule-based to skill-based. Training thus enables an increase in efficiency in human behavior [1].

If we consult the figures for car drivers involved in accidents per million passenger kilometers, then according to Oswald and Williams [3, 4] the risk falls with increasing

human age, until it starts to rise again from 40 to 50. According to Burgard [5], experience gained with age is responsible for this, alongside character skills (personality) and mental and physical prerequisites. If we view the assimilation of experience as a learning process, then the capacity to learn contributes to improved driving skills [6].

If road traffic followed clear and known rules, people would not need to behave in the ways described above. Road traffic, however, is an open system consisting of static and dynamic objects, and a multitude of environmental factors such as light levels and rain. Even if only to a limited extent compared to beginner drivers, experienced road users continue to come across unknown situations that need dealing with. It is because people show just this knowledge-, rule- and skill-based behavior that today's road, with its efficiency, accuracy, and safety, is possible.

Further, this behavior leads to individually varied driving behavior. In one and the same situation, drivers act differently and have different preferences in selecting distances, speeds and acceleration.

It is precisely these capabilities enabled by skill-, rule-, and knowledge-based behavior that will be taken out of driving upon automation, and replaced with corresponding capabilities in driving robots.

22.3 Learning Technical Systems

The term machine learning stands for a research area dealing with methods for designing algorithms. A particular feature of these algorithms is the automatic improvement of technical systems based on experience. In this, the automatic improvement follows rules and measures previously defined by the human developer. The hotly debated ideas of completely free and creatively acting machines, were they to exist at all, are not our focus here. In employing machine learning, a clearly defined task is needed, with accompanying assessment metrics and (training) data.

An oft-used quote for the definition of machine learning is from Mitchell [7]:

A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P , if its performance at tasks in T , as measured by P , improves with experience E .

According to Mitchell [7], machine learning has proved itself in service particularly in the following circumstances:

- (a) There is a large quantity of data in a database that, by implication, potentially contains valuable information that can be extracted automatically.
- (b) People only have limited understanding of a certain area and thus lack the knowledge for effective algorithms.
- (c) In tasks demanding dynamic adjustment to changing conditions.

How these definitions and their usage areas fit to autonomous driving is discussed in Sect. 22.4. First, however, we shall take a look at the processes of machine learning. With the aid of examples from various fields, we will explore the range of usage possibilities before examining the task of autonomous driving in particular.

22.3.1 Overview of the Various Machine-Learning Processes

Breiman [8] noted the following in 2001:

In the past fifteen years, the growth in algorithmic modeling applications and methodology has been rapid. It has occurred largely outside statistics in a new community – often called machine learning.

In [9], the first machine-learning operations are traced back to McCulloch and Pitts in 1948, among others. The multiplicity of learning processes thus makes a detailed description of all of them impossible. Therefore, the following will describe categories of learning problems into which we may at least group the processes [10–12].

Supervised learning

Supervised learning, or learning with a teacher, is characterized by pre-assessed training data (labeled data). This training data—the experience that machine learning builds on—contains the input and output parameters of a learning problem. In the classification of whether airbags are deployed in an accident or not, for example, training data would consist of acceleration values and the corresponding assessments (deploy or not). These assessments must be made by an expert/teacher, for example, or by observation over time. The learning process then uses these empirical values to determine the output for newly observed input values. In the transfer of empirical values onto new input values, we can distinguish [10] between *lazy* (memory-based) learning and *eager* (model-based) learning.

In lazy learning, the training data are saved during the learning and a similarity measure is defined. This similarity measure can vary in its complexity, and ranges from the simple Euclidian distance to complex distances in case-based reasoning. When output values are sought for new input values, the most similar training data to the new case are determined and output values derived from them. This procedure corresponds to transductive inference [13]. In eager learning, in contrast, a global model based on training data is constructed during the training phase (induction). Output values for new cases are obtained by deduction from the model.

Unsupervised learning

For unsupervised learning, or learning without a teacher, training data are also on hand, although in this case with no assessment or output values (unlabeled data). In this learning process, the aim is to find a structure in the data, and classify the data according to the structure. The training data are here used to reveal these structures and, based on this, to classify newly observed input values.

Reinforcement learning

Reinforcement learning differs from both of the previous processes in that few or no training data are available to begin with. The training data required for a striven-for improvement are collected by an agent¹ itself by carrying out the task to be optimized according to a fixed scheme. An assessment of the execution of the task feeds back into the learning process, forming a training dataset of input and output values that are used for further optimization steps. The approaches of reinforcement learning are exposed to the so-called innovation dilemma, for “exploration” and “exploitation” contradict each other. March [15] describes it thus:

Exploration includes things captured by terms such as search, variation, risk taking, experimentation, play, flexibility, discovery, innovation. Exploitation includes such things as refinement, choice, production, efficiency, selection, implementation, execution.

Corresponding to the learning problem, a balance for both has to be found, as, on the one hand, an optimal execution is sought in a partly unknown search area, and, on the other, this search is limited by external conditions such as costs, safety and time.

Alongside the question of whether, and in what form, training data are available, the learning problem can also be distinguished based on the use of training data.

Batch learning

In batch, or offline learning, a set of training data is applied at a time point to make use of the learning methods. If the learning method produces a model, for example, this model is not updated by further experience gained during its use.

Online learning

Online learning is characterized by an iterative process, in which new experiences are incorporated into the learning process. The aim is to continually optimize how the task is tackled, in the process incorporating experience from the operation. This results in system behavior changing from experience and thus over time.

These different types of learning problems require the use of various methods of machine learning [12]. These range from decision trees, artificial neural networks, and genetic algorithms to support vector methods, instance-based learning, hidden Markov models, value iteration and Q-learning, etc. These methods have in common that how well they deal with a learning problem depends on three fundamental properties. Firstly, the methods are only in a position to optimally solve the learning problem if data (experience) relevant and representative to the operation are used in sufficient quantity for learning. Secondly, the same applies for the quality of training data, so that the handling of noisy, inaccurate or partial data is especially necessary for actual measured variables. Alongside the test data, the assessment of performance (P) represents a further challenge. The

¹Russell [14] defines it thus: “An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through effectors.”.

methods will also only correctly deal with the learning problem if the assessment is valid with real data for the whole operational area.

22.3.2 Examples

The following examples serve as a small sample for areas that can be addressed with the processes of machine learning.

Airbag deployment [16, 17]

The learning problem consists of classifying sensor values to order the deployment of a vehicle airbag. A classifier is learned for this, which puts accidents either in the class of “deploy” or that of “do not deploy.” In this example, the dataset of an accident has 30 dimensions, and consists, for instance, of acceleration, pressure and impact-sound sensors at various positions in the vehicle. For 40 training datasets, in which the sensor values are recorded for representative accidents, it is annotated (labeled) whether an airbag deployment would be necessary or not.

Vehicle-path control system with artificial neural network [18]

The ALVINN (Autonomous Land Vehicle In a Neural Network) project had the aim of positioning a vehicle on its optimal path within its lane. The training data consisted of the input parameters of the individual pixels of a camera image, and the associate output parameters of the steering angle. These are recorded during the journey of a human driver. An artificial neural network whose 960 input nodes receive the individual values of the 30×32 pixels of the camera image is learned. These input nodes are connected via 4 hidden nodes with the 30 output nodes, which each stand for a different curvature.

22.4 Automation that Replaces the Learning Driver

For the car industry, machine learning and artificial intelligence are of interest not only for the automation of driving, but also in other fields such as design, production and after-sales management [19]. These areas, and that of infotainment, are not our concern, however. Our focus is on vehicle automation, which, as presently understood, consists of the following components in moving a vehicle on public roads:

1. Perception of the environmental and vehicle state variables
2. Cognition of these state variables to arrive at a representation of the world
3. Behavior planning based on this representation
4. Execution of selected behavior

These components give the properties that have, according to Mitchell (see Sect. 22.3), already led to successful machine-learning operations. Increasing perception of environmental and vehicle state variables gives an enormous amount of machine-readable

information. Firstly, this is down to continued increases in sensor performance and signal processing power, so that a detailed picture of the world is available for machine processing. Secondly, the numbers of sensors and vehicles equipped with them are rising as we approach full automation. It follows from this that the quality and quantity of training data for machine learning is increasing. The second property especially applies for certain areas of cognition and behavior decisions, since work on the human processes that are to be replaced is largely theoretical at this stage (Sect. 22.2). Based on the EU-funded Human Brain Project and associated studies, it is evident that there are still many questions that remain unsolved, and that the knowledge needed for effective algorithms is lacking. The third property is generated by road traffic. As already described in Sect. 22.2, the world in which vehicles move requires adjustments to be made to changing environmental conditions. This is why machine learning in automation suggests itself to be implemented in such a way that it will be able to adjust to these changes.

Alongside these three motivators for employing machine learning, however, there are particular challenges opposed to it. What the four components needed for autonomous driving (perception, cognition, behavior planning, behavior execution) particularly have in common, as opposed to other machine-learning applications, is the intervention in the actual behavior of the vehicle. This means that, regardless of where undesired behavior appears in this chain, it may develop into a breakdown or an accident. In order to debate the deployment of machine-learning processes in vehicle automation, we will now classify safety-related systems and allocate vehicle automation to one of these classes.

22.4.1 Safety Systems

According to ISO26262, safety is the “absence of unreasonable risk.” The extent to which a system in a car (or also generally) impacts this safety can be determined in the following ways ([16] extended):

1. Not safety-relevant

Errors in these systems do not lead to any dangers for persons or the environment. Language recognition, where machine-learning processes are often employed on [20], is not safety-relevant, for example, when used in infotainment. Such systems are therefore already in use; although errors do regularly crop up to some extent, they do not negatively impact safety.

2. Safety-relevant

A system is deemed relevant to safety if an error in it may result in danger to persons or the environment.

- a) Decision-making support systems

Here, the decision maker has the choice of whether to act on the system’s suggestions or not. An anesthetist, for example, receives suggestions for dosages based on information on the patient, operation, and previous experience. A system error would endanger the patient, though only when the anesthetist follows the suggestion [21].

b) Systems for monitoring and diagnosis

A system error leads to a warning failing to appear and can, if the error is not otherwise spotted, become a danger to people and the environment. If a diagnostic system in industrial machinery fails, the absence of this diagnosis may be dangerous [22].

3. Safety-critical

Systems where an error directly leads to persons or the environment being endangered.

(a) Supervised/correctable automation

If an action is carried out automatically without additional conformation, then a failure of an automatic system directly leads to a hazard for people or the environment. If the system is additionally supervised by a human, and if the possibility of correcting it is provided, then this hazard may be avoided. The person supervising brings the system back under control, so that human involvement creates fault-tolerance. It is to be kept in mind here that, particularly with increasing automation of processes and the further relieving of human tasks, people's capacities to supervise an automated process will decrease [23]. The so-called congestion assistant, as partially automated vehicle control, represents just such a safety-critical system, as faulty system behavior would pose a direct danger. This danger, however, is addressed by the driver's supervision, because the congestion assistant is developed in such a way that humans, by intervening, correct the faulty system behavior. The system is designed for controllability.

(b) Unsupervised/non-correctable automation

The most critical form of automation in terms of safety is that with no possibility of correction. Without supervision, a failure of the system leads to danger and, depending on the situation, damage to people or the environment. Fully automated driving falls in this category, as, by definition, the vehicle's occupants are no longer supervising it. Undesired behavior, or a failure or malfunction not addressed by the system, thus leads directly to people and the environment being put in danger and possibly harmed.

This categorization and the classification of autonomous driving in the *unsupervised automation* category shows why machine-learning processes as they currently stand cannot simply be carried over from the two levels less critical to safety (non-critical and relevant). Not without reason are the safety-critical examples from Sect. 22.3.2 reports on studies without reference to non-supervised application.

The application of machine-learning processes in unsupervised or non-correctable automation requires further differentiation, as different challenges arise depending on the point in the vehicle's system life cycle.

22.4.2 Challenges and Problem-Solving Approaches in the Various Phases of the System Life Cycle

In this account, vehicle life cycles are broken up into five phases—research, development, operation, service and change of user/end of vehicle life—as various challenges await machine learning in each phase.

22.4.2.1 Research

When machine learning is applied in the research phase, the aim is mainly to establish what its processes can do. The examples range from online, offline, supervised and unsupervised, to reinforcement learning. Exemplary training datasets are drawn on as databases. The same applies when assessing process performance and robustness. This is carried out under controlled and/or supervised conditions, based on exemplary test data or test runs. In particular, controlled conditions and/or deploying trained test drivers (Category 3a Sect. 22.4.1) make faults tolerable, so that there are a multitude of existing examples. Accordingly, proving its safety does not belong to the challenges of using machine learning. The challenge here lies in accessing representative data for its later area of application. The question thus arises of whether the research findings can be transferred to the development and operational phases of the system life cycle.

22.4.2.2 Development

Learning during the development phase can be compared with offline learning. As many application-relevant training data as possible are selectively gathered, for example to learn a model during development.

The fulfilling of safety demands in the results is, as in all other safety-relevant vehicle components, verified and validated, so that the vehicle can be released for production and operation. Afterwards, neither the learned model nor the classification change further. The learning process does not work online and is not adaptive, and thus does not use further data gained in operation as training data to update the model or classification. During use, it thus constitutes a time-invariant system, where the known methods for verification and validation remain valid. But it should be borne in mind that the results of the various machine-learning processes can vary in their interpretability. For example, learned decision trees of limited extent, or a manageable set of learned rules, are easy to interpret [24] and thus permit the use of white-box test procedures [25]. Other methods such as random forests or sub-symbolic neural networks, on the other hand, are difficult for testers to interpret and thus represent a black-box.

For such complex components, proving safety poses a large challenge compared to analytical models. Because the brute force test,² as for most analytical models, is not

²The brute force approach stands for testing all possible combinations of the parameters influencing the object under test behavior.

suitable [26] for systems with high input dimensionality, the following four counter-measures from Otte [24] are often used as a result:

- (a) Breaking down problems of many dimensions into submodels of fewer dimensions, so that the submodels can be interpreted and validated by experts
- (b) Employing reference solutions which permit the safety of the learned components to be analyzed
- (c) Limiting the input, output and state variables to specified value ranges, such as those of the training data, for example. These limits may be static, but also dependent on other variables
- (d) Limiting the dynamics of the input, output and state variables to minimal, maximum, positive or negative changes per time unit

Each of these measures restricts the potential of machine learning in order to allow testing of the learned model.

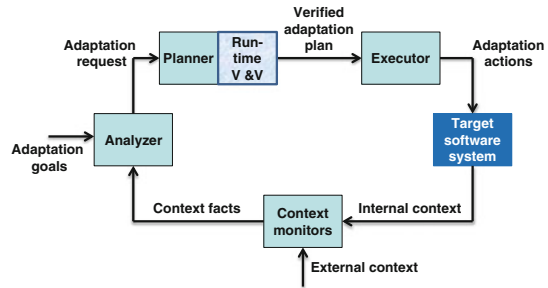
22.4.2.3 Operation

When the fully developed and manufactured vehicle is put into operation, data accrue on the real usage area, the static environment, other users and their behavior, and the vehicle user and occupants. In addition, the vehicle has data on its machine behavior over time. This directly available new information, which was previously inaccessible, encourages the use of online learning processes, and thus also adaptive systems. The vehicle thus becomes a time- or experience-variant system. This further degree of freedom in a changing system with no additional supervision results in a particular challenge in testing and safety validation, one that has not yet been solved for time-invariant autonomous systems, see Chap. 21. Basically, there are two options for how a system that changes during operation could be made safe. One option is to limit its adaptivity to a clearly defined enveloping area, such as the strategies of Adaptive Transmission Control [27]. Here, input, output and state spaces are restricted to a few parameters [28], so that, from where we stand, validation and verification during the development phase appears to be possible. Should this restriction contradict the purpose of machine learning during operation, then an online check of the changing, time-variant and complex system becomes necessary [29]. The following two approaches can in turn be applied here [29].

Runtime verification and validation

In contrast to conventional procedures of verification and validation during the development process by the developer, the system applies verification and validation methods during operation [26, 29, 30]. In principle, the adaptation process is seen as a feedback loop. Figure 22.1 shows the four stages of observation, analysis, planning and execution, which, according to Tamura et al. [30], are necessary for a structured examination of the adaptation process. This picture can be directly transferred to online learning processes. If a system adaptation is detected in this process, this then needs checking with runtime verification and validation processes before being applied to the software running. These

Fig. 22.1 Adaptation process as feedback loop with runtime verification and validation as found in [30]



checks include whether, due to the applied changes to the system, it stays within the viability zone. For autonomous driving, this means checking that any changes to the driving robot comply with safety regulations before implementing them.

The procedures applied here, such as model checking or theorem proving, reach their limits at this point, as described in Chap. 21. Using a software-in-the-loop procedure may also be possible, although it is questionable if there is sufficient processing power for such processes in production vehicles.


Validation and verification via monitoring and fault tolerance

If a system safety check prior to updating the software is not possible, a fault may arise and lead to failure and thus also danger. To avoid this, the system is to be designed as fault tolerant. Such a fault-tolerant system fundamentally requires two components [31]. Firstly, monitoring of system states and behavior to decide, based on an analysis, whether there is a fault. Secondly, redundancy is required that can be transferred to in case of a fault. Figure 22.2 shows this set-up schematically. This principle corresponds to a human supervisor who takes over a part-automated vehicle control system when a fault occurs.

What this shows is that machine learning during operation poses a challenge, especially for validating safe behavior. Both procedures—runtime verification and validation and verification and validation through monitoring—need a way to measure safe driving. Approaches for such measurements are outlined in Sect. 22.4.3.

Another approach is supplied here by a comparison with human learning. Road users accept that humans, with no further screening, learn based on the measures they take while driving and adjust their behavior accordingly. Checking whether adjusted behavior complies with traffic regulations does not transpire directly but, rather, sporadically via

Fig. 22.2 Dynamic redundancy according to [31]



police speeding and traffic checks. In addition, other road users report highly erratic behavior, such as accidents, so that human drivers receive feedback in their learning process. Transferred to the technical system, this would mean that machine learning leads directly to an adjustment of behavior. The test or check of whether this adjustment was legitimate would follow in retrospect from other road users, the police or a special supervisory authority. This approach reduces the demands placed on online validation and verification, as the skills of other road users are incorporated. If this test does not take place until after the functions have been modified, and if there is also no other possibility of intervention (in the control system), operating the updated but untested functions would pose higher risks.

22.4.2.4 Service

Alongside machine learning during development and operation, there is a further phase in the system life cycle where learning can take place. As part of servicing, training data gathered by the vehicle can be downloaded and the vehicle functions updated. This does not necessarily require the vehicle's physical presence [32]. This procedure opens up the feedback loop in Fig. 22.1. Corresponding to a further development stage, training data and planned adaptations can be tested, so that the software system can be updated later after safety has been assured. As these methods of machine learning can involve personal data leaving the vehicle, security as well as safety needs to be borne in mind. For a more in-depth look at this, see Chap. 24.

22.4.2.5 Change of User/End of Vehicle Life

Should, as is hoped, a vehicle be able to personalize how it drives for a user, or optimize this for an area of operation, then, with a change of user or at the end of its life, these learned capabilities and knowledge should stay with the user and not the vehicle. This capacity then becomes of particular interest when, for example, ownership patterns change as in Vehicle on Demand (see Chap. 2) where the user does not purchase a car but only its mobility service. In principle, it is not difficult for a technical system to transfer knowledge. It is actually a strength of artificial systems that they can transfer this information without the long-winded learning process. This is examined in more depth in Sect. 22.5.

22.4.3 Measures of Safe Driving

As described in the previous chapters, an evaluation of the vehicle's driving in safety terms is needed for both the machine-learning process in general and verification and validation during operation. In the first instance, it can of course be seen retrospectively if an accident has taken place, with what impact speed and energy, and how it came about. This measurement has the drawback that the accident should not, if possible, have taken

place in the first place. It follows from this that, for a safe system, accidents will be an extremely rare event and thus hardly suitable for learning from.

What is needed is an evaluation that classifies a journey as unsafe before it exceeds physical driving limits. To this end, we shall differentiate between deterministic and stochastic procedures in hazard assessment, as found in [33, 34].

22.4.3.1 Deterministic Procedures in Hazard Assessment

Following [35], we shall in addition distinguish between identifiers from driving dynamics and identifiers from distances. The easiest values to determine, though with limited informative value, are the limits for lateral and longitudinal acceleration and yaw rate. In the 100 Cars Study [33], longitudinal accelerations of greater than 0.7 g are used as triggers for detecting unsafe situations. Identifying critical situations with variables from the ego-vehicle is not adequate, as other road users can bring even stationary vehicles into dangerous situations. It is easy to see that a vehicle's safety is also influenced by other traffic in its immediate surroundings.

If the environment is at the beginning first reduced to parallel traffic on one lane in each direction, time-to-collision (*TTC*) or its reciprocal value give an indication of how safe a situation is. For example, in ISO 22839 (Forward vehicle collision mitigation systems),

$$TTC = \frac{x_c}{v_r}$$

defines the time that elapses until the ego-vehicle collides with an object at distance x_c , assuming the relative speed $v_r = (v_{ego} - v_{obj})$ remains constant. This measure can be applied to both the preceding and the following vehicle. From the *TTC*, Chan [34] defines a criticality index for which it is assumed that the severity of an accident is proportional to the square of the speed:

$$\text{Criticality Index} = \frac{v^2}{TTC}$$

If the speed between the vehicles does not remain constant, the enhanced *TTC* is used. In ISO 22839, this is derived thus:

$$ETTC = \frac{v_r - \sqrt{v_r^2 - 2 \cdot a_r \cdot x_c}}{a_r}.$$

This equation only applies so long as the relative acceleration $a_r = (a_{obj} - a_{ego})$ remains constant. If a vehicle comes to a stop, the equation needs adjusting, see Winner [35] on this. These values give an assessment of a state, equivalent to a situation analysis. To enable an assessment of driving safety for all situations, the following procedures from [36] are recommended:

- Number of Conflicts (NOC)
- Time Exposed TTC (TET)
- Time Integrated TTC (TIT)

The names of the procedures speak for themselves, so that we shall not dwell on them here, but rather refer the reader to [36, 37]. Further simple measure include time headway

$$t_h = \frac{x_c}{v_{\text{ego}}}$$

and the necessary delay a_{req} to avoid a rear-end collision. Further approaches based on the foregoing can be found in [37–39].

When vehicles, objects and accelerations moving in a lateral direction are included, for example, when studying situations at junctions, the metrics for a safety assessment have to be expanded. Easily the simplest approach is Post-Encroachment Time (PET). This temporal variable is defined by [40] as “time between the moment that the first road user leaves the path of the second and the moment that the second road user reaches the path of the first.”


When observing the lateral direction, there are more objects than just the two object-vehicles in relation to the ego-vehicle. These can influence the safety assessment statically, but also in principle dynamically from any direction on the vehicle plane. Tamke’s [41] approach is, for example, to identify the distance, and the distance’s time derivation, to all objects in the vicinity of the ego-vehicle using Euclidian norms. The TTC is determined based on these variables by ascertaining the time until contact with the vehicle’s bodywork using predictions of behavior. This approach is still classified as a deterministic procedure at this point, as the behavior of the ego-vehicle and the dynamic objects follow the “constant turn and constant acceleration approach.” The approach per se, however, also allows predictions to be made non-deterministically.

Deterministic procedures are suitable for ex-post safety assessments of a carried-out maneuver or situation. If an online assessment of situations as they happen and a prediction on a time horizon greater than 1 s is desired, however, then its application per se will be flawed as long as the situation contains uncertainties. Uncertainties here imply that the development of a situation over time not only has a potential deterministic procedure, as is often assumed in “constant turn and constant acceleration” models of simplicity, but rather that many different situations can occur. For example, in calculating the TTC, the car in front could, instead of slowing down, accelerate to alleviate a critical situation. In order to incorporate these uncertainties into the hazard assessment, we will now take a brief look at stochastic procedures.

22.4.3.2 Stochastic Methods in Hazard Assessment

A hazard assessment of two vehicles simply passing each other, as shown in Fig. 22.3, clearly encourages the use of stochastic procedures. If deterministic behavior and an extrapolation of speed are assumed, passing like this would pose no danger, as the

Fig. 22.3 Comparison of predictions in a passing situation: deterministic (*above*) and stochastic (*below*)



trajectories and areas occupied by the vehicles do not cross. This looks different when uncertainties are taken into consideration. If a human-driven vehicle (Fig. 22.3 lower lane, black) does not stay in its lane, this can be hazardous for the autonomous vehicle (Fig. 22.3 upper lane, orange) and the immediate surroundings.

The associated risk results from the probability of an accident and its potential severity. Both values involve unknowns that need to be estimated as accurately as possible.

The principal approach here [42] for determining probability is to predict the trajectories of the ego-vehicle and objects in its vicinity, based on measured state variables. Due to the uncertainties involving human drivers, sensors and actuators mentioned above, as well as the interaction between the objects, there is not only one trajectory, but rather a probability distribution of the states of all objects over time. If the potential areas occupied by the objects overlap, an accident is probable. The possible states of the objects result from the dynamic models used and the defined limits for the dynamic variables. In [43], a single-track model for vehicles is used here to find a compromise between prediction accuracy and computing time. In addition, dynamic variables such as acceleration and steering rate are limited, according to the analysis of the author [43], to non-critical and typical values. In [43, 44] an overview of alternative methods can be found. All methods have in common that the dynamic simulation of a vehicle cannot proceed analytically [41], so that numerical methods with accompanying discretization and simplification are to be used.

Including severity in the hazard assessment is approximated in [45] by a relative assessment of each accident based on its inelastic collision. This corresponds to the approach based on the Potential Collision Energy (PCE) [36]. Further, however, no approach has been found that determines both severity and probability in combination. One reason for this is imprecise analytical regression methods [46], another is computing-intensive Finite Element Methods (FEM), which are currently used in other fields for determining accident severity [47]. Meier et al. [46] provides a new approach

here based on symbolic regression. Using a crash-situation database (produced from FEM calculations), regression functions are learned which predict accident severity in a situation to a few milliseconds. This uses pre-crash information such as vehicle mass, speeds, collision point and collision angle. The downside of this approach is its limited interpretability, as the regression model does not include any physical variables. If this approach gave a valid prediction for severity, it would be possible to extend an assessment of risk into one of severity.

The above-mentioned procedures are not yet sufficient to assess adaptive automated driving's safety, for the following reasons. Firstly, all the approaches are based on a series of simplifications, such as leaving out weather conditions, simplifying the driving dynamics, or not including sensor uncertainties. Secondly, the procedures as they stand do not provide any validated and combined ascertainment of the probabilities and severity of accidents. A general definition and assessment of vehicle-driving safety therefore does not currently exist.

However, the ongoing progress in vehicle automation does provide enabling factors. These are, firstly, removing the uncertainty of the driver (ego-vehicle). Although human behavior continues to be present in the form of other road users, the trajectory of the ego-vehicle is known within its control performance. Moreover, current vehicles' sensor performance is rising, reducing uncertainties in the state of the object. Further, additional information on the surroundings is exchanged via V2X communication, thus improving the points of reference for hazard assessment in both quality and quantity.

22.5 Automation as Part of a Learning Collective Group

In this chapter, our analysis of learning vehicles has so far restricted itself to the system life cycle of a particular vehicle. What vehicle automation implicitly brings with it, however, is the duplication of hardware and software across a whole group of vehicles in the aimed-for mass production. Vehicles operating in road traffic would consequently have the same capabilities. On the one hand, this has the drawback that driving errors, breakdowns and accident types would affect not only one vehicle but the whole group. On the other, it gives adaptive systems a further degree of freedom. The option of exchanging data makes collective learning during operation possible. Two approaches may in principle be distinguished here [48]: agent-based machine learning and machine learning with agents. In the former (also often called agent-based swarm intelligence), the learning collective system is made up of a multitude of networked agents with limited cognitive capacity. The behavior of animals such as ants or bees is often drawn on in this connection [49]. In contrast, though, machine learning with agents also consists of several agents,

these employ the procedures described in Sect. 22.3. These two approaches differ in the fundamental stages of machine learning in terms of:

- generation of collective experience
- collective performance evaluation
- derivation of learned models and knowledge

Essentially, the approaches can be applied as soon as data from real situations has been gathered. This raises questions: At which position of a group is relevant information gathered from the data and assessed in terms of the learning problem and performance measurements? Can learning methods be used that are based on this? Not least for reasons of limited data-transfer bandwidth, it will be necessary to transfer pre-processed data or even learned models and knowledge rather than raw sensor data. So long as the participating agents/vehicles belong to a series and software release, there will be challenges for the integrity of the transferred information, although not in compatibility and trustworthiness between agents. So-called homogeneous teams exist in which how and what is being learned is known among the group. If the aim, however, is to expand the database which learning methods are applied to, vehicles with different software or even of different manufacturers could also be networked to each other. This amounts to a collective of heterogeneous vehicles in which potentially different machine-learning procedures are used, and the knowledge representation is likewise heterogeneous. Reference [50] gives examples from other areas of how such collectives could in principle be handled. Essentially, as well as vehicle robots, there are other agents, such as smartphones or also, in future, service robots, which likewise gather data covering completely different areas to those accessible to the driving robot. The currently highest conceivable level of connectivity, and by some distance the largest database, is the internet. The autonomous vehicle as a web-enabled device permits a great number of applications and functions, for better or worse. Not least, IBM's Watson Project shows that part of the knowledge archived in the internet is also comprehensible to machines. Information made accessible on the internet by any authority would therefore not need to be learned through experience first to influence an autonomous vehicle's behavior. On the other hand, access to any (unauthorized or anonymous) source threatens to create problems in terms of both road traffic safety and data security (see Chap. 24).

22.6 Conclusion

Machine learning is of great interest to current research, as the quality and quantity of available data is constantly increasing and, in addition, vehicle automation is throwing up questions that can only be partly solved with conventional analytical approaches. In carrying research findings over to the development of functions in the driving of autonomous vehicle, proving their safe use in unsupervised safety-critical systems with no possibility of correction poses the greatest challenge. Therefore, according to the authors'

knowledge, learned models that do not change after testing and approval can already be found in current mass-production vehicles. Systems such as Adaptive Transmission Control are not the subject of our examination due to their negligible adaptive parameters in clearly defined limited value ranges. However, learning and adaptation while in operation is giving automation an extra degree of freedom.

Exploiting this degree of freedom is motivated by the possibility of optimizing autonomous journeys, compensating for the loss of people's capacity to adjust and learn, and individualizing vehicle-driving. This chapter has highlighted that applying machine learning in vehicle automation during operation requires greater attention in terms of both road safety and data security. There is currently no valid measure for assessing traffic safety in respect of risk. Therefore, it seems clear that the use of adaptive machine-learning processes while in operation first requires a fault-tolerant set-up with redundant conventional systems, with the latter serving to assess road safety. It is thus to be expected that machine learning while in operation (adaptive systems) will initially only optimize vehicle automation within the framework of the conventional system.

Due to the challenges highlighted with respect to proving the safe behavior of time-invariant adaptive systems, it appears necessary to do intensive further research on runtime verification and validation. The same holds for the above-mentioned applied non-adaptive systems and their demonstration of safety. Although there is copious literature on examples that have already been successfully introduced, these mostly come from other fields that do not place comparable demands on a product.

In this chapter, we have examined the question of data security only peripherally—for a more in-depth discussion, see Chap. 24. Learning approaches require data and hence information on occupants, the vehicle and the environment. Data protection is thus of the same relevance as road safety. It should be pointed out that the environment also contains people, which implies protecting their data. The quality and quantity of the sensors needed for vehicle automation is a boost for machine learning on the one hand, but on the other they are regarded with suspicion from a data protection point of view. One special property of vehicle sensors is, in addition, that they currently are to a great extent not physically covered when the vehicle or function is not active. With the application of machine learning, data security should also accordingly be addressed directly.

Nevertheless, the application of networked agents and learning systems has its advantages, whose effects are not to be underestimated. Collectively learning agents need not acquire available knowledge slowly. Instead, this information can be copied and pasted to the next vehicle or software generation. This, in combination with access to the large quantity of electronically recorded information, has the potential to change the driving of vehicles, road traffic and thus people's overall (mobility) behavior. This makes the findings gained for driving equally interesting to research into medical robots, and also

domestic robots with human contact. Due to the similar conditions, the reverse also applies, which speaks for close cooperation between vehicle technology and robotics.

Open Access This chapter is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, duplication, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, a link is provided to the Creative Commons license and any changes made are indicated.

The images or other third party material in this chapter are included in the work's Creative Commons license, unless indicated otherwise in the credit line; if such material is not included in the work's Creative Commons license and the respective action is not permitted by statutory regulation, users will need to obtain permission from the license holder to duplicate, adapt or reproduce the material.

References

1. Rasmussen, J.: Skills, Rules, and Knowledge; Signals, Signs, and Symbols, and Other Distinctions in Human Performance Models. *IEEE Transactions On Systems, Man, and Cybernetics* **SMC-13**(3), 257–266 (1983)
2. Donges, E.: Fahrerverhaltensmodelle. In: Winner, Hakuli, Wolf (Hrsg.) *Handbuch Fahrerassistenzsysteme*, pp. 15–23 (2011)
3. Oswald, W.D.: Automobilismus und die “gefährlichen Alten”. In: G. Schmidt (Hrsg.) *Technik und Gesellschaft. Automobil und Automobilismus*, vol. 10, pp. 183–195 (1999)
4. Williams, A.F.: Teenage drivers: patterns of risk. *Journal of safety research* **34**(1), 5–15 (2003)
5. Burgard, E.: Fahrkompetenz im Alter: Die Aussagekraft diagnostischer Instrumente bei Senioren und neurologischen Patienten. Dissertation, LMU (2005)
6. Funk, W., Grüninger, M., Dittrich, L., Goßler, J., Hornung, C., Kreßner, I., Libal, I., Limberger, S., Riedel, C., Schaller, S.: Begleitetes Fahren ab 17-Prozessevaluation des bundesweiten Modellversuchs. *Berichte der Bundesanstalt für Straßenwesen. Unterreihe Mensch und Sicherheit*.(213) (2010)
7. Mitchell, T.M.: *Machine Learning*. McGraw-Hill series in computer science. McGraw-Hill, New York (1997)
8. Breiman, L.: Statistical Modeling: The Two Cultures (with comments and a rejoinder by the author). *Statist. Sci.*, 199–231 (2001)
9. Carbonell, J., Michalski, R., Mitchell, T.: An Overview of Machine Learning. In: Michalski, R., Carbonell, J., Mitchell, T. (Hrsg.) *Machine Learning. Symbolic Computation*, pp. 3–23. Springer Berlin Heidelberg (1983)
10. Ertel, W.: *Grundkurs Künstliche Intelligenz. Eine praxisorientierte Einführung*, 3rd edn. Lehrbuch. Springer Fachmedien, Wiesbaden (2013)
11. Sewell, M.: *Machine Learning*. <http://machine-learning.martinsewell.com/> (zuletzt geprüft 15.07.2014) (2009)
12. Sammut, C. (ed.): *Encyclopedia of machine learning*. 78 tables. springer reference. Springer, New York (2011)
13. Shaoning Pang, Nikola Kasabov: Inductive vs transductive inference, global vs local models: SVM, TSVM, and SVMT for gene expression classification problems - Neural Networks. *Proceedings. 2004 IEEE International Joint Conference on* (2004)
14. Russell, S., Norvig, P., Intelligence, A.: *A modern approach. Artificial Intelligence*. Prentice-Hall, Englewood Cliffs **25** (1995)

15. March, J.G.: Exploration and exploitation in organizational learning. *Organization science* **2**(1), 71–87 (1991)
16. Nusser, S.: Robust Learning in Safety-Related Domains. *Machine Learning Methods for Solving Safety-Related Application Problems*, Otto-von-Guericke-Universität Magdeburg (2009)
17. Nusser, S., Otte, C., Hauptmann, W., Leirich, O., Krätschmer, M., Kruse, R.: Maschinelles Lernen von validierbaren Klassifikatoren zur autonomen Steuerung sicherheitsrelevanter Systeme. *at-Automatisierungstechnik Methoden und Anwendungen der Steuerungs-, Regelungs- und Informationstechnik* **57**(3), 138–145 (2009)
18. Pomerleau, D.: Neural Network Vision for Robot Driving. In: Hebert, M., Thorpe, C., Stentz, A. (Hrsg.) *Intelligent Unmanned Ground Vehicles*, vol. 388. The Springer International Series in Engineering and Computer Science, pp. 53–72. Springer US (1997)
19. Gusikhin, O., Rychtyckyj, N., Filev, D.: Intelligent systems in the automotive industry: applications and trends. *Knowl Inf Syst* **12**(2), 147–168 (2007). doi: [10.1007/s10115-006-0063-1](https://doi.org/10.1007/s10115-006-0063-1)
20. Deng, L., Li, X.: Machine Learning Paradigms for Speech Recognition: An Overview. *IEEE Trans. Audio Speech Lang. Process.* **21**(5), 1060–1089 (2013). doi: [10.1109/TASL.2013.2244083](https://doi.org/10.1109/TASL.2013.2244083)
21. Olivier Caelen, Gianluca Bontempi, and Luc Barvais: LNAI 4594 - Machine Learning Techniques for Decision Support in Anesthesia
22. Widodo, A., Yang, B.-S.: Support vector machine in machine condition monitoring and fault diagnosis. *Mechanical Systems and Signal Processing* **21**(6), 2560–2574 (2007). doi: [10.1016/j.ymssp.2006.12.007](https://doi.org/10.1016/j.ymssp.2006.12.007)
23. Bainbridge, L.: Ironies of automation. *Automatica* **19**(6), 775–779 (1983). doi: [10.1016/0005-1098\(83\)90046-8](https://doi.org/10.1016/0005-1098(83)90046-8)
24. Otte, C.: SCI 445 - Safe and Interpretable Machine Learning: A Methodological Review. In: Moewes, C., Nürnberger, A. (Hrsg.) *Computational intelligence in intelligent data analysis. Studies in computational intelligence*, vol. 445. Springer, Berlin, New York (2013)
25. Burgdorf, F.: Eine kunden- und lebenszyklusorientierte Produktfamilienabsicherung für die Automobilindustrie, KIT Scientific Publishing; Karlsruher Institut für Technologie (2010)
26. Taylor, B.J.: *Methods and procedures for the verification and validation of artificial neural networks*. Springer (2006)
27. Nelles, O.: Lernfähige Fuzzy-basierte Fahrstrategie für automatische Getriebe. In: Isermann, R. (Hrsg.) *Modellgestützte Steuerung, Regelung und Diagnose von Verbrennungsmotoren*, pp. 233–250. Springer Berlin Heidelberg (2003)
28. Cao, C.T., Kronenberg, K., Poljansek, M.: Adaptive transmission control. Google Patents. <http://www.google.com/patents/US5954777> (1999)
29. Dahm, W.: Perspectives on Verification and Validation in Complex Adaptive Systems, Notre Dame University. Workshop on Verification and Validation in Computational Science (2011). Accessed 22 July 2014
30. Tamura, G., Villegas, N., Müller, H., Sousa, J., Becker, B., Karsai, G., Mankovskii, S., Pezzè, M., Schäfer, W., Tahvildari, L., Wong, K.: Towards Practical Runtime Verification and Validation of Self-Adaptive Software Systems. In: Lemos, R. de, Giese, H., Müller, H., Shaw, M. (Hrsg.) *Software Engineering for Self-Adaptive Systems II*, vol. 7475. Lecture Notes in Computer Science, pp. 108–132. Springer Berlin Heidelberg (2013)
31. Isermann, R.: *Fault-diagnosis systems. An introduction from fault detection to fault tolerance*. Springer, Berlin, New York (2006)
32. Stokar, R. von: Software-Updates Effiziente Nutzung von Connected Cars. *ATZ Elektron* **9**(1), 46–51 (2014). doi: [10.1365/s35658-014-0387-7](https://doi.org/10.1365/s35658-014-0387-7)

33. Guo, F., Klauer, S., Hankey, J., Dingus, T.: Near Crashes as Crash Surrogate for Naturalistic Driving Studies. *Transportation Research Record: Journal of the Transportation Research Board* **2147**(-1), 66–74 (2010). doi: [10.3141/2147-09](https://doi.org/10.3141/2147-09)
34. Ching-Yao Chan (ed.): *Defining Safety Performance Measures of Driver-Assistance Systems for Intersection Left-Turn Conflicts*. Intelligent Vehicles Symposium, 2006 IEEE. Intelligent Vehicles Symposium, 2006 IEEE (2006)
35. Winner, H., Geyer, S., Sefati, M.: Maße für den Sicherheitsgewinn von Fahrerassistenzsystemen. In: Winner, H., Bruder, R. (Hrsg.) *Maßstäbe des sicheren Fahrens*. 6. Darmstädter Kolloquium Mensch + Fahrzeug. Ergonomia Verlag, Stuttgart (2013)
36. Dijkstra, A., Drolenga, H.: *Safety effects of route choice in a road network. Simulation of changing route choice*. SWOV Institute for Road Safety Research, Leidschendam, Netherlands (2008)
37. Yang, H.: *Simulation-based evaluation of traffic safety performance using surrogate safety measures* (2012)
38. Zhang, Y., Antonsson, E.K., Grote, K.: A new threat assessment measure for collision avoidance systems. *Intelligent Transportation Systems Conference, 2006. ITSC'06*. IEEE
39. Jansson, J.: *Collision avoidance theory with application to automotive collision mitigation*. Linköping studies in science and technology. Dissertations, vol. 950. Dept. of Electrical Engineering, Univ., Linköping (2005)
40. Horst, A. R. A. van der: *A time-based analysis of road user behaviour in normal and critical encounters*. Institute for Perception TNO, Soesterberg, Netherlands (1990)
41. Tamke, A., Dang, T., Breuel, G.: A flexible method for criticality assessment in driver assistance systems. *Intelligent Vehicles Symposium (IV)*, 2011 IEEE
42. Althoff, M., Stursberg, O., Buss, M.: Model-Based Probabilistic Collision Detection in Autonomous Driving. *Intelligent Transportation Systems, IEEE Transactions on* **10**(2), 299–310 (2009). doi: [10.1109/TITS.2009.2018966](https://doi.org/10.1109/TITS.2009.2018966)
43. Eugen Käfer: *Situationsklassifikation und Bewegungsprognose in Verkehrssituationen mit mehreren Fahrzeugen*, Universität Bielefeld (2013). Accessed 7 July 2014
44. Althoff, D., Wollherr, D., Buss, M.: Safety assessment of trajectories for navigation in uncertain and dynamic environments. *Robotics and Automation (ICRA)*, 2011 IEEE International Conference on
45. Althoff, D., Kuffner, J., Wollherr, D., Buss, M.: Safety assessment of robot trajectories for navigation in uncertain and dynamic environments. *Auton Robot* **32**(3), 285–302 (2012). doi: [10.1007/s10514-011-9257-9](https://doi.org/10.1007/s10514-011-9257-9)
46. Meier, A., Gonter, M., Kruse, R.: Symbolic Regression for Precrash Accident Severity Prediction. In: Polycarpou, M., Carvalho, A. de, Pan, J.-S., Woźniak, M., Quintian, H., Corchado, E. (Hrsg.) *Hybrid Artificial Intelligence Systems*, vol. 8480. *Lecture Notes in Computer Science*, pp. 133–144. Springer International Publishing (2014)
47. Mukherjee, S., Chawla, A., Mohan, D., Singh, M., Dey, R.: Effect of vehicle design on head injury severity and throw distance variations in bicycle crashes. *Proceedings From 20th International Technical Conference on the Enhanced Safety of Vehicles*. Lyon (2007)
48. Czarnowski, I., Jędrzejowicz, P.: Machine Learning and Multiagent Systems as Interrelated Technologies. In: Czarnowski, I., Jędrzejowicz, P., Kacprzyk, J. (Hrsg.) *Agent-Based Optimization*, vol. 456. *Studies in computational intelligence*, pp. 1–28. Springer Berlin Heidelberg (2013)
49. Miller, P.: *Die Intelligenz des Schwarms. Was wir von Tieren für unser Leben in einer komplexen Welt lernen können*. Campus-Verl., Frankfurt am Main [u.a.] (2010)
50. Gifford, C.M.: *Collective Machine Learning: Team Learning and Classification in Multi-agent Systems*. Dissertation, University of Kansas (2009)