TrafficQuest

CENTRE FOR EXPERTISE ON TRAFFIC MANAGEMENT

TrafficQuest report

Artificial Intelligence and Traffic Management

Current and Future Applications

Colophon

Authors	Erwin Walraven (TNO) Dawn Spruijtenburg (TNO) Isabel Wilmink (TNO) Max Schreuder (TNO)
Date	10 December 2021
Version	1.0
Published by	TrafficQuest Postbus 2232 3500 GE UTRECHT
Information	Henk Taale
Telephone	+31 88 798 24 98
Cover picture	Jeroen van den Heuvel

TrafficQuest is een samenwerkingsverband van

TrafficQuest
CENTRE FOR EXPERTISE ON TRAFFIC MANAGEMENT

Artificial Intelligence and Traffic Management

Current and future applications

10 December 2021

Table of contents

Foreword

This year's TrafficQuest Challenge focused on the topic of artificial intelligence and traffic management. The project team worked on this challenge with great pleasure. There are many interesting applications at the intersection of traffic management and artificial intelligence. This challenge gave us the opportunity to delve deeper into them. We had many interesting conversations and discussions about these topics. We would like to thank everyone involved for these discussions, as well as for the valuable information gained from the interviews and the workshop.

1. Introduction

Artificial intelligence (AI) is an interdisciplinary field in which systems are developed that are capable of performing tasks that usually require human intelligence (Russell & Norvig, 2020). AI is currently at the centre of scientific interest with applications being developed across a range of fields, making this an important moment to also investigate what AI can do for traffic management. A lot of research is currently being conducted in the field of pattern recognition, intended to make automatic driving possible. But are there other areas within traffic management that could benefit from the possibilities AI offers? A TrafficQuest challenge was carried out to investigate this.

TrafficQuest was the partnership between Rijkswaterstaat, TNO, and TU Delft in the field of traffic management and traffic information. From 2009 to 2016, this cooperation was active in developing, accumulating, applying, and disseminating knowledge about traffic management and traffic information. For more than seven years, TrafficQuest covered the entire field, from the more fundamental, theoretical knowledge about traffic management and traffic information to ‘operational knowledge’ about its application and effectiveness. At the end of 2016, the decision was made to continue on a smaller scale, concentrating activities on a number of current challenges and on the publication of *Verkeer in Nederland* (‘Traffic in the Netherlands’), a booklet containing TrafficQuest’s annual report. The booklet gives an overview of how traffic is currently being managed in the Netherlands and developments in traffic management.

A challenge is a quick-scan expert analysis with a short lead time. TrafficQuest’s challenges are intended to address and delve deeper into specific topics related to traffic management. At present, challenges have been carried out on the replacement of roadside systems by in-car systems, traffic management and traffic safety, the impact of C-ITS use cases, and 3D printing.

For this challenge, TNO was asked by Rijkswaterstaat to look at the potential applications of artificial intelligence (AI) in traffic management and information in the short term (1-5 years). This includes both current and future applications. For current and future applications there are a number of important questions:

- What applications are there now?
- What is their added value?
- What opportunities will there be in the near future (1-5 years)?
- What is needed to realise these opportunities?

The last question primarily concerns the data required to realise opportunities.

This chapter will clarify the purpose and scope of the challenge. In addition, a general introduction concerning the topics of artificial intelligence and traffic management will be given. This will be followed by a description of the research methodology used during the challenge. Finally, there will be a brief look ahead at the remaining chapters.

1.1. Purpose and scope

This challenge provides an overview of current and future applications of AI in traffic management. It paints a picture of the field and its latest, state-of-the-art innovations, and looks ahead at opportunities for the future. Current applications show which AI technologies are already mature, which data sources are useful, and what lessons can be learned from them. The focus is mainly on applications that have already been implemented or tested as pilots. For future applications, the focus is on those applications that domain experts expect to generate the greatest improvements in terms of traffic flow, safety, and sustainability, while also looking at what the possible pitfalls and challenges could be during implementation. For the selected future applications, experts expect implementation to be possible within one to five years.

1.2. Artificial intelligence

Artificial intelligence (AI) is a collective term for all the theory and developments that enable computer systems to perform tasks that usually require human intelligence. Examples include visual perception (e.g. recognising objects) and recognising and processing spoken language. In addition, AI includes the automatic controlling of other systems through algorithms, and making decisions about such systems, so that they can function optimally. An example of this is using AI to control the traffic system to optimise traffic flow. AI is a multidisciplinary field that combines mathematics, computer science, and control engineering with domain-specific knowledge to develop systems that can perform tasks that require intelligence when executed by humans. The literature gives several definitions of AI. For a full reflection, see the book by Russell and Norvig (2020).

Within the field of AI, there are several sub-areas that focus on specific techniques. In practice, these sub-areas do not stand alone; in many cases, applications are being developed that bring together several sub-areas. The division into sub-areas can be done in different ways and is not uniquely defined in the literature. The following sub-areas give a global overview of the field:

- Machine learning: algorithms that are able to learn from data or from experience (Alpaydin, 2020). This includes generating predictions or recognising patterns in data. Supervised learning includes all methods that learn through example and for which the desired output is known (e.g. whether or not a cat is present in a picture). Unsupervised learning does not require this desired output and is used, among other applications, to recognise patterns or clusters in data. Another commonly used term is deep learning (Goodfellow et al., 2016). This is a machine learning technology in which (deep) neural networks are trained using data.
- Natural language processing: understanding and processing of natural language (Raaijmakers, 2021). Chatbots able to understand what humans are saying are an example of this.
- Speech processing: recognising and optionally translating spoken language (Rabiner & Schafer, 2010). This is used, among other places, in mobile phones able to respond to spoken commands. Speech processing always involves spoken language listened to by a computer, while natural-language processing focuses on the analysis of a text after a computer has converted it into data.

- Computer vision: processing still and moving images, and understanding what happens in them (Szeliski, 2010). Think of recognising of objects in photographs, for example.
- Expert systems: systems that possess specific knowledge provided by human experts to solve problems in specific areas (Gupta & Nagpal, 2020). An example is the analysis of a disease profile using questions asked by the system and answered by a human being.
- Planning, scheduling, and controlling: systems that independently determine which actions must be performed to achieve a goal (Kochenderfer, 2015). An example is the control of traffic lights to maximise traffic flow. In this challenge, reinforcement learning (Sutton & Barto, 2018) also falls within this sub-area. This is a popular machine learning technology in which a system independently decides on actions to take and learns to function more efficiently by analysing the results of its actions. This learning process works using a reward function, whereby the system is rewarded for demonstrating the type of behaviour that achieves the system's end goal (e.g. maximising traffic flow).

AI is currently receiving a lot of attention. This recently led to several breakthroughs. One such breakthrough was the development of self-learning systems that are better than humans at playing the game Go (Silver et al., 2016). Although it may sometimes seem like AI is a completely new technology, it has in fact emerged over time in various fields. The roots of AI originally lie in philosophy, logic, mathematics, psychology, and biology, among other places. In the course of the 20th century, this developed into a large overarching field. Recent successes with AI are a result of the ever-increasing computational power of computers and the availability of data. The field is expected to grow even more rapidly in future, which will lead to a host of more intelligent applications in society.

Systems developed using AI are becoming increasingly complex, making it more difficult for humans to understand how they work and why they are making certain decisions or predictions. An example is the opening or closing of a peak hour lane. At present, this is partly being done manually and partly via a rule-based traffic algorithm. When an AI system indicates that it is better to close the peak hour lane, for example, it is desirable for the humans overseeing the system to know why the AI system has come to that conclusion. A current trend is the development of technologies that make AI-based decisions explainable and more transparent (Adadi & Berrada, 2018), making it easier for humans to discover how an AI system has arrived at a particular prediction or decision.

Although AI is currently receiving a lot of attention and has achieved great successes, there are also limitations that could pose risks in its application. An example is recognising objects (e.g. traffic signs) using camera images. It appears that in some cases it is possible to fool a data-trained AI system by making minimal (invisible to humans) changes to these camera images. A red traffic light might, for example, be recognised as a green traffic light, which could lead to dangerous driving behaviour on behalf of automated vehicles using this information to make their decisions. Another example concerns the recognition of fighter planes on satellite images, where the application of small stickers can already prevent the AI system from recognising the planes (Den Hollander et al., 2020). Vulnerabilities such as these show that the practical application of AI is challenging and must be done carefully.

AI has great potential and is expected to make a major contribution to achieving social goals in the near future. To realise these ambitions, an increasing number of AI initiatives are being set up by companies and the government (e.g. within 'The Netherlands AI Coalition'). AI is also starting to play an increasingly major role in university study programmes. Knowledge institutions, such as TNO, are also contributing with umbrella programmes, such as Appl.AI, in which new scientific results in the field of AI are integrated into applications in cooperation with partners.

1.3. Traffic management

Although the definitions of traffic management differ, put simply, the aim of traffic management is to influence traffic supply and demand in such a way that traffic demands and the capacity supply of the network are better matched, both in time and space. The problems encountered on the road network mainly concern specific bottlenecks and moments (i.e. peak hours, incidents, and events). By spreading traffic demand or dynamically adapting the supply of infrastructure, the existing road network can be better utilised. Typical traffic management measures include ramp metering, dynamic speed limits, peak hour lanes, but also traffic information communicated through panels above the road or other channels. The measures are primarily intended to improve accessibility, but they are also increasingly being used to improve road safety (e.g. through queue tail warning) or quality of life (e.g. by using speed limits) (Hoogendoorn et al., 2011).

Traffic management can fulfil various traffic-related functions:

- <u>Monitoring and detecting</u>: the monitoring and detecting of traffic and incidents.
- <u>Informing</u>: signposting, route information, network status, travel times, lane allocation.
- <u>Advising</u>: advising on lanes, speeds, and alternative routes.
- <u>Warning</u>: queue tail warning, queue tail signalling, dangerous situations, disturbances.
- <u>Management and control</u>: reducing speed limits, changing lane allocations, opening or closing lanes, processing height alerts, stopping traffic, overtaking prohibition, dosing, buffering.

Aspects that do not fall under traffic management are mobility management (e.g. measures to avoid peak hour traffic) and payment according to use (e.g. congestion charging). However, such measures can temporarily or locally reduce traffic flows, making them more manageable. The AI applications in this report are categorised according to their traffic management functions. Informing, advising, and warning have been combined, since these categories all focus on transferring information to road users or traffic control centres.

1.4. Research approach

This research looks at both current and future applications of AI in traffic management. Various sources were consulted to gather information and the following activities were undertaken:

- <u>Interviews</u>: Interviews were held with experts in the field of AI and traffic management. The list of interviewed experts can be found in Annex B. These interviews focused on both current and future applications of AI in traffic management. Interviewees were also asked about possible opportunities and obstacles.

- Quick-scan analysis: In addition to the interviews, relevant and current sources describing applications of AI in traffic management were studied. The interviews were leading to the topics and sources that were considered. The focus was on the added value of applications that have already been implemented. The sources studied can be found in Annex A.
- Selection of most promising applications: Based on the information gathered, a selection of the most promising future applications was made. For this purpose, selection criteria, such as estimated success rate and desirability (Annex D), were drawn up. The full list of future applications for selection can be found in Annex C.
- Elaboration of the most promising applications: The most promising applications were discussed and elaborated on in a working session. The participants of the working session are listed in Annex E. As much detail as possible is provided on how applications work, what their implementation will look like, and what impact can be expected from them.

1.5. Report structure

This report is structured as follows. Chapters 2, 3, and 4 cover applications within three traffic management functions: monitoring and detecting (chapter 2); informing, advising, and warning (chapter 3); and management and control (chapter 4). Within these chapters, current and future applications are discussed, and a brief summary is given of what AI can mean for each traffic management function. The insights from the interviews, literature quick scan, and the work session have been integrated into the descriptions of the applications. Chapter 5 gives an overview of all applications. Chapter 6 contains conclusions and recommendations.

2. Monitoring and detecting

The use of monitoring and detecting allows situations on the road to be analysed. The information gathered can be used for decision-making or to identify connections. Think of applications that inspect peak hour lanes or monitor the number of road users at specified times. Various data sources exist (e.g. video images) for which AI can be used to analyse large quantities of data in an automated way.

2.1. Current applications

There are already several applications where AI is used to monitor traffic better and to detect situations. Applications use AI to lift data to a higher level of detail. In addition, new sensors (mainly cameras) are being used, to measure the traffic flows of different modes of transport, for example.

Data enrichment and gaining more insights from data

High-quality data is at the foundation of various AI applications. There are various initiatives to improve the data available for traffic management and traffic information and, by doing so, gain more insights from the data. Various AI techniques are used to create the high-quality data that can generate these insights. In the past, data from loop detectors was often used. Today, alternative sensors are also being used or the data is being enriched. Data from loop detectors alone may be incomplete or inaccurate. Enrichment minimises inaccuracy and maximises coverage.

The British company Vivacity uses video sensors and processes the data using convolutional neural networks. These are neural networks suitable for analysing visual images. In this way, they obtain information on numbers of vehicles/passengers and the classification of the vehicles. The latter is done in more categories than the loop detectors can distinguish, allowing for micro-modalities (e.g. scooters) to be included. Travel times and near-accidents can also be derived from this data. VivaCity's systems are present in many major UK cities, such as London, Manchester, and Bristol. Applications can be built on top of this data. VivaCity also uses recurrent neural networks to make short-term predictions concerning vehicle/passenger numbers and speeds. These are neural networks capable of processing sequences of temporal data. The resulting predictions are replacing the short-term simulations that are normally performed. Predictions are mainly up to three hours in the future. Various potential applications of these predictions, such as controlling via matrix displays, are possible (Vivacity Labs, n.d.).

American company INRIX also uses convolutional neural networks to obtain higher quality data. The INRIX algorithms use floating car data to obtain real-time counts and speeds, resulting in more data. An added advantage is that the entire network can be measured, since locations without loop detectors, such as smaller roads, are also included. As such, implementing the detection of speed and intensity changes is easy to implement. Moreover, the data is stored and archived per one-minute interval, making it possible for future applications to be trained with this data. INRIX's technology can be used to, for example, develop an application for queue tail warning or other safety warnings about speed changes. The information for this is already available and can be presented to road users using matrix displays or in-vehicle applications (INRIX, n.d.). It is vital that these kinds of applications receive extensive testing, since system errors can have large (fatal) consequences.

Detecting mobile-phone use behind the wheel using smart cameras

Police in the Netherlands have a number of smart cameras at their disposal. These cameras can be placed at different locations and continuously moved around. They take pictures of passing cars and use these pictures to decide whether they suspect the car driver of holding a mobile phone. Reducing the handheld use of mobile phones behind the wheel is expected to increase safety on the road (The Netherlands Police, 2019). The smart cameras use deep learning. Herein, the input from the network is a picture; the output is a number between 0 and 1 that indicates whether a phone is being held. Photos in which a phone is suspected of being held are forwarded for verification and manually checked. When the handheld use of a mobile phone has been verified, the vehicle's licence-plate number is used to send the driver a traffic fine. The added value of AI in this application is the automated and large-scale recognition of drivers on their phone. Analysing all this data would be an incredibly labour-intensive task for humans. One of the challenges regarding this application is that, when constructing the model, photos need to be used that have already been manually checked. This manually labelled data can then be used to train the model.

Inspecting peak hour lanes using smart cameras

Rijkswaterstaat also uses camera images in its traffic control centres to inspect peak hour lanes. They observed that a lot of people in these control centres spend a lot of time doing routine work, such as recognising objects on a peak hour lane. This is frequently work that people are unable to perform well for longer periods of time (e.g. due to fatigue). To automate the recognition of objects and unburden the people previously tasked with object recognition, a system has been developed that automatically recognises objects using video images (AI sub-area: computer vision). By automating this part of their work, people in the traffic control centres can focus their attention on other things. When developing the system, it proved to be a challenge to test how well the system functions; especially, whether it functions well enough to take over tasks from people. In traffic control centres, employees are responsible for their own decisions. When an independently operating computer system makes a wrong decision, it can be difficult to allocate responsibility. Systems must therefore be extensively tested before they can be put into practice.

Counting cyclists and pedestrians

Counting cyclists and pedestrians is often a difficult task. When using detection loops on the road surface, for example, the behaviour of cyclists, who sometimes cycle on the pavement, must be considered. Another problem is that some bicycles, such as cargo bikes, have double wheels. Technolution's FlowCube can count individual cyclists and pedestrians in urban areas, thereby providing density indications. When several sensors are linked, it also becomes possible to determine routes and travel times. The system uses camera images that are analysed using deep-learning techniques. Object characteristics are used by the system to distinguish between individual pedestrians and cyclists. AI plays an important role in this, as the system learns how to distinguish between pedestrians and cyclists using annotated data. This eliminates the need to manually implement a complex set of decision rules. A possible example of this system's application is to use the information from FlowCube to estimate boarding and disembarking times at tram stops so that trams can be given a green light more accurately. Another application is to align urban distribution with pedestrian and cyclist peaks, thus avoiding, for example, the peak of pedestrians on Saturday afternoon in many cities. FlowCube is currently being implemented in Groningen, Rotterdam, and Copenhagen, among other places (Technolution Move, n.d.-a).

Optimising traffic lights with AI

Through 'signal analytics', INRIX is able to gain more insight into traffic flows at traffic lights. At present, the only added value is to better understand traffic flows. This understanding can subsequently be used to make more well-founded decisions. The automatic adjustment of traffic-light settings or suggesting adjustments has yet to be implemented. Detailed information on waiting times, traffic demand per direction, and traffic-light performance is available. The application uses INRIX's AI traffic technologies with convolutional neural networks (INRIX, n.d.).

2.2. Future applications

For future applications in the realm of monitoring and detecting, applications were selected that use new technologies and which build on previously developed applications.

Predicting the duration of an incident

Rijkswaterstaat has developed a system that can predict the probability of an incident. This information is used to identify hotspots (locations where the risk of incidents is higher than at other locations). This information is used to position road inspectors in such a way that it increases their likelihood of being in the vicinity of a potential incident. The next relevant step is to look at predicting the duration of incidents. In addition, it may be possible to predict how the effects of incidents spread across time and space. Doing so can create insight into where people are likely to experience a disturbance, and also, where measures or route advice could be useful. Linear regression models can be used to create a prediction model. More advanced technologies, such as decision trees and neural networks, can also be used.

The resulting system serves to support the road inspector. The system can show the expected duration of an incident and justify its reasoning. In this way, the road inspector can use the information provided by the system to make and support his or her decisions. However, the road inspector will still have to make the decisions and assess the results of the system. In general, these kinds of tools increase our understanding of the factors that influence the duration of an incident. At present, estimates concerning the duration of incidents are being made by road inspectors after they have arrived at the scene of an incident. This system would make it possible to obtain a duration estimation prior to the road inspector's arrival at the scene. Road inspectors should obviously be involved in the development of the model.

Examples of data and information that can feed the prediction model are data on historical incidents and their aftermath, data from loop detectors, image positions, ambulance(s) on site, events in the area, and the weather. Context information about the environment (e.g. concerning events) can also be important. It could, for example, be more important to intervene in an incident after an event in the Amsterdam ArenA than at 4 a.m. on a Sunday morning when things are quiet. Faster response times may ameliorate the worst effects.

Training a model on incidents is a challenge since incidents are infrequent and arise from specific situations. Incidents in historical data can differ greatly from each other. Since such a model is trained on historical data, it is important for it to be able to indicate when it does not recognise a situation. Some types of incidents will be so unusual and without precedent in the historical data that the model is unable to accurately make a prediction. The model must therefore be aware of its own limitations, so that in such cases the estimate can be made by an expert.

A possible challenge when developing the model could also be that the information on incidents is incomplete in the data. There has to be enough data on enough incidents to train the model properly. In addition, it can be difficult to link contextual information to incidents, while such contextual information can have a significant impact on traffic flow during incidents. In general, AI is good at predicting large intensity patterns, such as traffic jams. However, it is more difficult to be very accurate and predicting incidents and their duration can be a challenge.

When evaluating the system, it is important to study the extent to which the system’s predictions are an improvement on the road inspectors' estimates. This can already be done during the development of the system, based on historical data and the corresponding estimations made by road inspectors.

Workshop participants expect that it will be possible to put this application into operation within five years. It is a logical next step to Rijkswaterstaat’s previous work, and its added value is clear.

Monitoring network condition with data from multiple parties

There are various parties that collect data within a city or area. Combining this data could lead to better models and conclusions. However, these parties are not always willing to share their data to create a model using aggregated data. A new trend is the development of federated learning techniques, which can arrive at a single model trained on all data sources without the need to share data.

The idea is that one model circulates between parties and that all parties continue training the model using their own data. A well-known application outside the field of traffic management is in hospitals, where each hospital can separately continue training the model (e.g. to estimate the risk of diabetes using patient data). By circulating the model, it becomes unnecessary to share sensitive patient data. Within the field of traffic management, there are various areas where data sharing is sensitive and federated learning promises a potential solution. These applications are elaborated below as a first outline.

In cities, there are an increasing number of parties offering a form of (micro-)mobility or 'Mobility as a Service' (MaaS) (e.g. providers of shared scooters and/or bicycles). These providers collect information about users and their journeys. Being privacy-sensitive and often important for the business case of these providers, many of these parties are understandably unwilling to share their data. The collection of similar information by multiple parties (possibly at multiple locations) allows these parties to collectively make more connections and extract more lessons from the information. Federated learning could possibly be used to model MaaS usage based on journey information, without explicitly sharing this information with all parties.

Developers of navigation systems and route-planning apps have a lot of information at their disposal. They often already collect information about current speeds or the number of travellers on a route to improve estimations of travel and arrival times. Estimates could be improved by combining information from different parties.

When establishing environmental zones, federated learning could, for example, be used to ensure that a limit on emissions from freight traffic is not exceeded. Data from logistics companies about the trips they have planned within the environmental zone could then be shared with a central party, such as the municipality. Based on the number of expected trips and the distribution of these trips over time, a decision can then be made about the allowed trips within the environmental zone.

In federated learning, it is important that the added value of aggregating the data is clear to all parties. To achieve this, a common goal for which the model is to be used must be established. When the added value is clear to all parties, they will hopefully put aside their objections to sharing their data (even in aggregated form). It is possible that for certain applications, the majority of data resides with a single party. In such cases, it is especially important that the added value is clear to all parties. For some applications, it would also be possible to make the supply of data a condition (e.g. to be permitted inside an environmental zone).

2.3. Summary

In terms of traffic monitoring, AI offers various possibilities for deriving more information from collected data. It also supports humans by simplifying or even automating processes. In many cases, data is collected from photographs or video material. Convolutional neural networks can be used to make statements about such images. These technologies have become increasingly sophisticated in recent years, clearing the way for new applications. An important prerequisite for the construction of prediction models based on photos or video material is that, in most cases, what is observable in the images must first be classified by humans. In addition, it is important to constantly monitor the ability of models to recognise objects in photographs, not just during the development of the system. Finally, AI offers monitoring opportunities regarding privacy-sensitive data collected by different parties without the need for this data to be shared. The use of tracking technology for tracing individuals often meets with resistance, making it important to properly communicate the way in which the data is used and stored. It is also necessary for applications to comply with legislation (such as the GDPR).

3. Informing, advising, and warning

Informing, advising, and warning concerns information transfer to road users or traffic control centres. Communicated information can be informative but communication can also take place by way of advice or a warning. Applications include route information, speed advice, and queue tail warning. AI offers opportunities to better inform, advise, and warn road users using multiple data sources. Combining multiple data sources allows for more accurate predictions, which can subsequently be presented to road users.

3.1. Current applications

Current applications for informing, advising, and warning mainly focus on informing and advising (e.g. with information about expected travel times, congestion, and bridge openings).

Speed advice for better traffic flow

In the Brabant In-Car programme, various experiments were carried out with in-car technologies to influence the behaviour of road users. Road users received personalised information via their smartphone or navigation system while driving on the A67. Smoover is one of the implemented projects from the Brabant In-Car programme. During this project, a navigation application for smartphones was developed that gave road users speed advice with the aim of improving traffic flow on the A67. Smoover's underlying systems used two AI technologies. Deep learning was used to make traffic predictions based on historical data and real-time data from the A67. These predictions made it possible to give proactive speed advice, based on the expected situation on the road. Additionally, reinforcement learning was used during simulations to learn which speed advice has a positive effect on traffic flow. Reinforcement learning in this application also has the potential to make the speed advice adaptive, with the system itself learning to become better over time. Both AI components were integrated into the system and tested with a small group of participants on the A67 in 2014 (Smoover, 2015).

Route advice to avoid and prevent traffic jams

In the European project Socrates 2.0, traffic predictions were used to advise better routes to road users (Huisken, 2020). The project developed systems that provide proactive advice. Their aim is not only to avoid existing traffic jams, but also to prevent expected traffic jams. Herein, real-time data is combined with historical data to calculate the probability of congestion at a specific location in the near future. If the probability is high, alternative routes are suggested to road users. Machine learning is used for the predictions, but further details are lacking at the time of writing. This system was tested in the Amsterdam region between December 2019 and the summer of 2020. However, there were fewer traffic jams than usual during the test due to the outbreak of COVID-19, meaning reliable conclusions regarding the effects on traffic have yet to be drawn.

Informing travellers about the opening of a bridge

The Botlek Bridge and Spijkenisse Bridge regularly open for shipping on the Oude Maas, resulting in longer travel times on the road. Working in collaboration with Technolution, Rijkswaterstaat developed a system to predict the opening of these bridges, so that road users are better informed and alternative routes can be better used. The system uses information transmitted by ships to predict whether one of these bridges will open in the near future. Herein, the speed of ships and their distance to the bridges are used. When the probability of opening increases, the message 'no opening expected' on dynamic route-information panels is replaced by 'opening expected'. In this way, road users are proactively informed about bridge openings, and they can decide to take an alternative route if necessary. Further details about the underlying algorithms were not yet known at the time of writing. The system was tested in practice for six months at the beginning of 2020. Regarding the predictions, the message 'no opening expected' on the information panels proved to be more reliable than the message 'opening expected'. This is due to the fact that the exact routes of ships on the Nieuwe Waterweg are difficult to predict.

Improved estimated time of arrival

Working in cooperation with Google, DeepMind researchers have created an improvement for the estimated time of arrival (ETA) in Google Maps (Lange & Perez, 2020). Traffic predictions play an important role in various applications and can in many cases be seen as underlying data on which decisions are based. A specific location in a city being expected to be crowded in sixty minutes could, for example, be a reason to advise people to take different routes.

Google Maps normally uses real-time traffic data for road sections to estimate how long a journey will take. However, in the past, this did not include expectations regarding traffic flow during the journey. In 97 percent of journeys, existing arrival-time estimates were already accurate. Their accuracy was increased even further by taking into account the expected situation on the road. The expected development of traffic is determined with the help of a graph neural network. Neural networks are a widely used tool used for generating predictions from data. Graph neural networks make it possible to generate better predictions in domains where the data can be structured as a graph. This makes it extremely suitable for application to a traffic network. Historical data is used to divide a road network into 'super segments' that share a large volume of traffic. These super segments find a place in the graph neural network. In combination with current data, this leads to an estimate of the expected time of arrival. These developments also show that integrating domain knowledge (in this case, the structure of the road network) into AI models can lead to better predictions compared to generic models that do not use explicit domain knowledge (Lange & Perez, 2020).

3.2. Future applications

For future applications, applications have been selected that are based on new technologies, such as graph neural networks, and build on existing algorithms. Both applications focus on advising traffic control centres and road users.

Predicting travel times and advising routes

Graph neural networks are relatively new models that use a graph structure as input for a neural network. While training, the neural network learns what the dependencies are between the locations in the graph. It is interesting to see what potential these kinds of models have within the context of traffic management.

They make it possible, for example, to make network-wide predictions rather than per individual location, since dependencies between locations can be taken into account more easily. In addition, contextual information, such as the presence of events, can be included. This information could be used to, for example, gain insight into the influence of a football match in the Amsterdam ArenA on traffic in Amsterdam North thirty minutes later. Better predictions of network conditions are especially beneficial in the event of incidents.

Currently, traffic predictions are prepared by traffic control centres. To make their predictions, information regarding factors that can influence the situation on the road is used. A factor could be that there are often a lot of traffic jams the day after the summer holidays ends or after the first rain in a long period. Ideally, these factors would be included in a graph neural network. Additionally, such a model would be designed in consultation with traffic control centres and would serve as a support for them.

The improved network-wide monitoring can be used for various purposes. Firstly, it can serve as a monitoring tool for traffic control centres. They can be alerted to situations that deviate from normal. If needed, they can then apply measures in the network to improve such situations. The added value of the model is that it can be constantly monitored and updated.

In addition, predictions from (graph) neural networks can be used as input for generating route recommendations to individual road users. This could include advice that is in the public interest (e.g. suggesting routes that circumvent public schools) or advice that aims to avoid vulnerable points in the network (e.g. bridges) at times when the network is already crowded. Herein, finding the balance between what is beneficial to individual travellers and the collective is challenging.

When using graph neural networks, the structure of a graph can be directly taken into account during learning. This allows for a contextual representation of the data to be made. For this, it is important to think carefully about the way the data is projected onto the graph. There are, for example, several ways to divide a road network into segments. The way in which this is done affects the functioning of the graph neural network. The scalability of graph neural networks to larger traffic networks is an open challenge.

Insight into compliance behaviour regarding route recommendations is very important. To influence behaviour, it is good to study how certain stimuli (e.g. advice) impact traffic. To this end, it is necessary to merge several data sets. Levels of compliance behaviour can be measured by keeping track of historical route advice and checking the data to see if people followed the advice they were given.

To realise higher levels of compliance behaviour, it is useful to generate an explanation in addition to the advice itself, thus providing road users with more insight into the underlying reason for the advice they are receiving. A simple example is advising a different route and adding that traffic flows better on this route, making the journey calmer while only extending it by five minutes. Another addition could be that the road user is contributing to safety around a school by following the suggested route. Extracting this kind of information from prediction models is an open challenge.

Shockwave damping

There are algorithms, such as SPECIALIST (Hegyi et al., 2008) and COSCAL (Mahajan et al., 2015), that dampen shockwaves using traffic flow theory. They function by setting speed limits that constrain the flow of traffic upstream, thus homogenising traffic speed and dampen shockwaves. Sometimes additional measures are needed upstream, as this is where higher levels of traffic density arise. AI could be used to create algorithms that resolve shockwaves and also reduce their unwanted effects. One such AI technology is reinforcement learning, which can learn how to 'squeeze' the traffic most effectively. A reward function in reinforcement learning could be expressed in vehicle loss hours. In this way, the system can learn to avoid congestion.

It is expected that AI can contribute to the improvement of algorithms in a number of ways. Firstly, it may be possible to develop better decision rules using reinforcement learning. Consider, for example, rules in which the reduction of speed to sixty kilometres per hour is not necessary, but a smaller reduction suffices. The activation of the system could also be improved. In many cases the system is not activated, while there are shockwaves. The conditions under which the system activates are strict and presently often remain unmet. When correcting this, care should be taken not to set the threshold of activation too low. If the system kicks in at every minor traffic flow disruption, it loses credibility. In addition, the type of intervention is immediately determined upon activation. Doing so more dynamically, while adding the option of making adjustments during the intervention based on the reaction of the traffic, could enhance the system's performance.

However, it is important that the self-learning system does not start without knowledge. The current algorithms can be implemented as basic knowledge that the system can improve upon. Certain safety constraints must also be imposed on the algorithm to prevent the system from engaging in undesirable behaviour (such as constantly changing speed limits). New developments in the field of safe reinforcement learning can offer a solution here. Before such a (renewed) system can be used on the road, it is necessary to show what its added value is in different traffic situations (also compared to existing systems). This can be done in simulations.

The algorithm can be developed to work with data from loop detectors, but more detailed data (such as floating car data) can enable the algorithm to develop better decision rules. AI technologies are expected to be able to contribute to dampen or resolve shockwaves, requiring less space and allowing more dynamic control. As such, they can be implemented in more places and in more different situations. However, the implementation of such systems without AI has proven difficult in the past, making it unlikely that an AI-controlled shockwave-mitigation system will be put into practice in the near future.

3.3. Summary

AI-based applications for informing, advising, and warning usually consist of two phases. In the first phase, data is used to generate insights into or predict traffic. Insights or predictions from the first phase serve as input for the second phase, in which the decision is made which information, advice, or warnings to communicate to road users. AI presents opportunities by making it possible to proactively (rather than reactively) present information to road users. New developments in machine learning, such as graph neural networks, are expected to make it possible to do so in a more network-wide manner in the future.

4. Management and control

Management and control concerns intervening in the way traffic is guided, to optimise the flow of traffic. Applications include reducing the speed limit or opening and closing a lane. Management and control can also be done by models. AI methods can play a role in such applications, especially in making decisions to redirect traffic without direct human intervention.

4.1. Current applications

Controlling traffic lights and optimising intersections are clear applications of management and control. Several examples that have already been applied have been found. But there are also other decisions that can be made in an easier or better way using AI.

Adaptive traffic lights

AI is also used to control with the help of the existing infrastructure. AI algorithms are used, for example, to determine when traffic lights turn green and for how long they remain green in areas with multiple intersections. SURTRAC is a system that plans the control of traffic lights for a specific intersection based on information regarding approaching vehicles (Smith et al., 2013). The information from an intersection is then communicated to intersections in the vicinity, so that it can be used during planning. SURTRAC uses AI-based planning and scheduling. It can be seen as an 'online' controller that constantly optimises itself as new information comes in. SURTRAC is a multi-agent system, in which multiple autonomous agents exhibit smart behaviour and cooperate. This is a field at the intersection between artificial intelligence, optimisation, and operations research. SURTRAC uses cameras to detect approaching vehicles. SURTRAC was tested at nine intersections in Pittsburgh (Pennsylvania, USA) and compared to a reference controller that does not continuously optimise when receiving new information. The results show that travel times can be reduced by an average of 26 percent. Emissions were also reduced by 21 percent during the test. Additionally, road users were able to drive faster and were required to stop less.

Smart traffic lights

VivaCity also uses the data it collects (see section 2.1) for an application with traffic lights. They use the collected data to develop the algorithms using reinforcement learning. The algorithm is trained by running many simulations using historical data. In these simulations, the traffic lights are controlled, from which the algorithm learns. The fact that VivaCity has high-quality data means that the algorithm can also be adapted for different types of road users (e.g. when the active modalities, such as cycling or walking, increase). In this case, the traffic lights react faster to current traffic conditions than traffic lights without this AI addition. The traffic lights were implemented in Manchester (Vivacity Labs, 2020).

Determining the positions of road inspectors

Road inspectors are deployed when incidents are reported. They are responsible for ensuring safety and a good traffic flow. Incidents may concern accidents but could also include, for example, a stone on an emergency lane that needs to be removed or a vehicle that has broken down. Road inspectors are deployed within a specific area and receive reports from the traffic control centre. It is important for road inspectors to be able to arrive at the scene quickly, so that problems can be quicky resolved, resulting in less traffic jams and fewer lost hours by vehicles. Rijkswaterstaat recently started using a new system to optimise the position of road inspectors using incident predictions, with the aim of reducing their incident response time. This process consists of two steps. In the first step, a Bayesian model is used that predicts the probability of an incident. This model is estimated using data on historical incidents from a whole year. In addition to information about the incident itself, this includes contextual information, such as the day of the week, the time of day, and the weather. The predictions make it possible to identify 'hotspots' (areas where the probability of incidents is high compared to other locations). In the second step, the predictions are used to optimise the positions of road inspectors. It is important for road inspectors to be in the vicinity of locations where a relatively high probability of incidents is predicted. Finally, this information is passed on to the road inspectors, allowing them to reposition.

The system is currently active in the central region of the Netherlands and will soon be deployed in the south/west of the Netherlands as well. Before the system's roll-out, the average response time was 18 minutes (measured over 24 hours); after the system's roll-out, it was 14 minutes. At the time of writing, it is not yet clear whether this reduction has been achieved through the optimisation of the system. There is the possibility that a different way of working for road inspectors (who now receive information about the position of their colleagues via tablets) lies at the root of the shorter response times. This will be studied further in the future.

4.2. Future applications

For future applications, an extension of the use of AI for traffic control systems has been chosen. In addition, the introduction of an overtaking prohibition for freight traffic was considered following its effectiveness in the past.

Self-learning traffic light controllers and control systems

Reinforcement-learning methods are capable of making computer systems learn how to optimally perform a task. They can also be used to enable already installed systems to adapt to new situations. The expectation is that continuously learning traffic control or ramp metering systems will be better able to control traffic than static (i.e. not self-learning) systems. This could result in better traffic flow at intersections and improved quality of life.

A configuration that can be finetuned, for example, is the duration of a cycle, which is the time it takes a traffic control plan to give all directions a green light. In addition, it is possible to vary which directions are given a green light at the same time and the order in which they receive green is. The 'green time' (the time that a direction, or set of directions, is given a green light) can also be varied. At the same time, there are limitations to be taken into account. The main functions of a traffic light are safety and the separation of traffic in time and space. This means that certain directions cannot be given a green light at the same time. The maximum or minimum time that a phase takes (think of the time the traffic light is yellow) also constitutes a limitation.

There are already parties involved in using AI for traffic light controllers that optimise the computer programmes that run the installations with the help of AI. A possible improvement could be to make the traffic control plans seasonal or time-dependent, or even to have it respond to a situation. Then routes with heavy traffic (e.g. those connected to a stadium or concert hall) could be given more green time. Additionally, the self-learning aspect allows for continuous optimisation. Concerning the construction of a new housing estate, for example, using the self-learning function of the traffic light controller, the routes new residents are using, which as a result are experiencing increasing traffic, could be given more green time.

Although the envisioned performance of intelligent traffic light controllers is still somewhat lacking, it is still expected that with more information about the position and speed of approaching vehicles, better traffic flow can be achieved. New data sources (such as floating car data and images from video sensors) can help. A possible addition is to distinguish between different types of road users or vehicle types. Schoolchildren on bicycles could, for example, be given more green time. Vehicles transporting heavy goods could also be made to stop less to reduce emissions.

Traffic light controllers can be trained using historical data in combination with simulation-generated data. Reinforcement learning occurs by evaluating random actions; as such, it cannot be directly tested in practice. The development of safe reinforcement techniques, in which random actions are defined within certain frameworks, could change this in the future.

For these types of systems to be scalable in practice, it is important that they are able to configure themselves to specific situations. This prevents humans from having to repeatedly configure systems for new situations.

When regularly adjusting the control plan of a traffic light, it is important to take expectations into account. Travellers who regularly pass through an intersection will at some point notice when they are given a green light. Adjustments that change the order in which green lights appear could lead to accidents due to expectations on behalf of road users.

Predicting the effectiveness of an overtaking prohibition

Prohibiting trucks from overtaking means cars are forced to slow down less. This can make the flow of traffic more stable and less likely to stagnate. Cameras make it possible to monitor which vehicles are driving on which lane, and what the traffic flow is at that moment. In addition, it is possible to monitor the past effectiveness of an overtaking prohibition for freight traffic. By collecting a lot of data, it is possible to train a model to predict for a given situation whether an overtaking prohibition will have a positive effect on traffic flow. This model can then be used to decide whether to impose an overtaking prohibition for trucks.

Data from loop detectors and camera images will be used in the development of the system. The use of cameras in particular can lead to resistance. Additionally, it is usually a complicated process to add more cameras to the existing infrastructure. Moreover, sufficient data must be available from periods when overtaking prohibitions were active, so that the system can learn under which circumstances such restrictions have a positive effect.

Overtaking prohibitions can be communicated via the existing infrastructure, matrix displays, or in-car systems. It may be necessary to install additional matrix displays. Overtaking prohibitions can also be communicated through in-car systems, but this requires the adaptation of regulations. At present, it is not possible to oblige people to follow the information they are given in the car.

A dynamic overtaking prohibition is currently only being tested on the A2 motorway. A factor that makes it difficult to effectively implement an overtaking prohibition for trucks is its low compliance rate. Increased compliance rates may generate greater effects. The future advent of self-driving vehicles will resolve the issue of low compliance rates. However, this is not expected to be in the short term.

The expectation is that a model can be constructed based on (convolutional) neural networks, which, using images and other information, can assess whether overtaking prohibitions can have a positive effect. When training the model, it is necessary to distinguish between regular congestion and disruptions caused by freight traffic overtaking. From this, the model must distinguish whether a freight ban is effective. How this is to be derived from the data is still an open challenge.

The elaboration of this application requires confronting a number of currently unanswered questions. Aside from the distinction between regular congestion and congestion caused by overtaking trucks, it is also important to investigate which types of data contain sufficient information on vehicle types and vehicle positions per lane. Therefore, as a first step, it is important to investigate whether a connection can be made between overtaking trucks and congestion.

4.3. Summary

AI offers opportunities to develop systems that can control traffic independently without the need for direct human intervention. An important added value is the potential to increase the adaptiveness of systems, allowing them to independently adapt to new situations. To realise implementation in practice, it is important to be able to validate how systems arrive at their decisions, something which is not always easy with current AI methods. In addition, when validating algorithms in simulations, it is important to be able to simulate realistically, taking into account all the variables of an environment, such as the weather and pedestrian behaviour. Alongside the development of AI methods themselves, this is a significant challenge that deserves a lot of attention.

5. Overview of applications

Current and future applications of AI in traffic management are mapped out in **Fout! Verwijzingsbron niet gevonden.**. Future applications are indicated in blue. The horizontal axis shows the categories of traffic management, and the vertical axis shows the sub-areas of AI. Applications that use technologies from several sub-areas have been placed in the table several times.

Table 1: Overview of current and future applications

	Monitoring and detecting	**Informing, advising, and warning**	**Management and control**
Machine learning and deep learning	Data enrichment and gaining more insights from data Optimising traffic lights with AI Monitoring network condition with data from multiple parties Predicting the duration of an incident	Route advice to avoid and prevent traffic jams Informing travellers about the opening of a bridge Improved estimated time of arrival Predicting travel times and advising routes	Determining the positions of road inspectors Predicting whether an overtaking prohibition will be effective
Natural language processing			
Speech processing			
Computer vision	Detecting mobile-phone use behind the wheel using smart cameras Counting cyclists and pedestrians Inspecting peak hour lanes using smart cameras		
Expert systems			
Planning, scheduling,		Speed advice for better traffic flow	Adaptive traffic lights

	Monitoring and detecting	**Informing, advising, and warning**	**Management and control**
control, reinforcement learning		Shock wave damping	Smart traffic lights Determining the positions of road inspectors Self-learning traffic light controllers and control systems

Machine learning and computer vision technologies, using floating car data and video images as their source data, are mainly used for counting and classifying vehicles. In practice, there is some overlap between the categories since computer vision technologies sometimes also use machine learning. Machine learning is also used to generate predictions. In many cases, these predictions form the foundation of a larger system. Route recommendations or information on travel times can, for example, be calculated using predictions. Algorithms for planning, scheduling, and control are used when systems need to decide what to do independently (e.g. giving speed advice or controlling a traffic light).

During the quick scan, no existing applications were found, and no future applications were conceived for natural-language processing, speech processing, and expert systems. As the current applications were collected via a quick scan, the table may not give an exhaustive overview of all existing applications.

It is clear from this overview that some combinations of traffic management functions and AI sub-areas are less obvious. This does not mean that it is not possible to develop an application that falls into one of these empty cells. However, filling the cells should not be a goal in itself. It is important that AI technology is selected that matches the requirements that arise from its application.

6. Conclusion and recommendations

Artificial intelligence is currently the centre of attention and is seen as a collection of technologies that enable new applications in various domains, in which people are supported or processes are fully automated. TNO and Rijkswaterstaat conducted a TrafficQuest challenge to map out the potential of these technologies for traffic management.

In **traffic monitoring**, new developments in AI allow more data to be derived from existing and new data sources, such as camera images. In recent years, there have been many developments in convolutional neural networks that can learn from images. This will soon enable new monitoring applications that were previously difficult to achieve, such as automatically recognising objects in images. In addition, there are future opportunities for making more network-wide predictions. Moreover, developments in federated learning have the potential to provide insights derived from privacy-sensitive or business-sensitive mobility data originating from multiple parties.

Regarding **informing, advising, and warning**, AI offers opportunities to proactively present information to road users. Applications in many cases consist of two phases: in the first phase, AI is used to predict; in the second phase, these predictions are used to decide which information to present to road users (e.g. the advice to take another route). It can be relevant for road users to know why specific advice is being given (e.g. to increase compliance rates). It is expected that new technologies in the field of explainable AI will contribute to this in the coming years.

Concerning **controlling traffic**, AI makes it possible to develop systems that can independently decide how to direct traffic (e.g. to improve traffic flow). Current traffic management systems usually cannot adapt automatically to new situations. AI-based systems can be made to be adaptive and continuously learn (e.g. through reinforcement learning). This also offers advantages in the design and development of these systems. For example, it used to be necessary to determine all the decision rules for controlling traffic in advance. AI will make it possible in the future for systems to independently learn to control traffic by making decisions and observing the results of these decisions. These self-learning systems can be used for traffic control, ramp metering, or to decide whether to impose overtaking prohibitions.

The current applications of AI in traffic management in this report show that, in practice, the present focus is on performing one specific task (e.g. counting cyclists) using a limited number of data sources (e.g. camera images). This is a logical process, given the fact that such applications are usually developed incrementally. There is great future potential for AI-based applications that combine multiple data sources or address multiple complex tasks in a combined fashion. This could, for example, lead to new insights about traffic being derived from data; insights that are not readily apparent with existing methods and a single data source.

The development of AI-based methods starts with high-quality data, making it of great importance that the data that enables models to make accurate predictions is available. Currently, there is a strong focus on the development of AI technologies. However, it is important to also focus on the collection, pre-processing, and fusion of detailed traffic data, which can then serve as the fuel for models. It is also becoming increasingly important to take into account privacy aspects and the structured storage of collected data. AI models that are constructed using data can draw conclusions about situations similar to those seen in the data, but at present, they are not advanced enough to draw conclusions about situations not previously observed in the data. Combined with data, AI can make a great contribution to traffic management. It is important, however, that AI is not seen as a technology that always has the solution, especially regarding situations about which little information is known.

To make AI applications successful, it is important to include domain knowledge in the development of models and algorithms. AI is capable of deriving insights from data on its own, but the expectation of the experts interviewed is that integrating domain knowledge will always be necessary, even for the simplest predictions. Another important question is what will happen to traditional traffic models in the future. The conclusion of the challenge is that AI will complement these traditional models in the short term, but it is unlikely to replace them completely. However, AI models may be able to generate new insights regarding traffic, which can then be integrated into existing models. In this way, AI can potentially lead to more knowledge about traffic. In addition, existing knowledge about traffic may be used to improve the explainability of AI models.

New developments in AI in recent years have made models and algorithms increasingly complex, which has made it harder for domain experts to understand AI-based decisions. It is crucial for models to be explainable. In this way, end users can, for example, understand why a model has arrived at a specific prediction or why a vehicle has automatically stopped. It is important to emphasise that this explainability has limits, especially when complex models are used to make accurate predictions or decisions. Applications that demand a high degree of explainability with a lot of detail may be better served by models that are less complex but more explainable. A downside of this is that simpler models may reduce the quality of predictions. There will always be a trade-off between the quality of predictions and the degree of explainability. On top of explainability, there is increasing attention to biases, which can lead to undesirable decisions in practice, in data and models. It is, for example, undesirable for specific groups of road users to never get a green light at an intersection.

During the challenge, several applications were identified through which AI can add value to traffic management. The discussions with experts showed us that, in addition to the technical development of applications, it is highly important to focus on the system's evaluation. Doing so is not only necessary during design, but also after application, to ensure that a system continues to function properly. In addition, there may be implications for privacy regulations. And applications with in-car systems may require additional regulations before full deployment is possible in practice. To date, applications are often tested at a pilot level. For AI to play a more prominent role in traffic management, it is necessary to take additional steps, so that applications tested in pilots can lead to structural application on a larger scale. In addition to evaluation, possible safety risks are another important aspect. Whenever AI systems make high-risk decisions on their own, the question arises as to who exactly is responsible for the outcomes of these decisions. This question should always be included in the development of systems. This applies not only to AI systems, but also to software systems in general.

The challenge also showed us that it is valuable for experts in traffic management to learn more about the possibilities offered by AI. The expectation is that it will be crucial for traffic management experts to actively engage with those in the field of AI. In this way, new developments can land in the domain, and we can be sure that new developments in the field of AI are driven by the questions arising in practice. This can be done within the framework of, for example, PhD research at universities on the cutting edge of AI and traffic management. Another valuable initiative is 'The Netherlands AI Coalition', in which various organisations work together to accelerate AI development in the Netherlands and to connect initiatives in this field.

On a European level, further thought is also being given to how AI applications can be used in various domains. The European High-Level Expert Group on AI (AI HLEG) has established guidelines for the deployment of reliable AI. For example, it lays down the requirements that systems must meet in terms of transparency and explainability, but also privacy and data governance. The applications in this challenge and its associated challenges underline the importance of these aspects and show that AI development should not only focus on the technical aspects of, for example, data and algorithms.

Literature

Adadi, A. & Berrada, M. (2018). Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI). *IEEE Access*, *6*, 52138-52160. https://doi.org/10.1109/ACCESS.2018.2870052

Alpaydin, E. (2020). *Introduction to Machine Learning*. MIT Press

Goodfellow, I., Bengio, Y., Courville, A. & Bengio, Y. (2016). *Deep learning*. MIT press.

Gupta, I. & Nagpal, G. (2020). *Artificial Intelligence and Expert Systems*. Mercury Learning & Information.

Hegyi, A., Hoogendoorn, S. P., Schreuder, M., Stoelhorst, H. & Viti, F. (2008). SPECIALIST: A dynamic speed limit control algorithm based on shock wave theory. *2008 11th International IEEE Conference on Intelligent Transportation Systems*, 827-832. https://doi.org/10.1109/ITSC.2008.4732611

Hollander, R. den, Adhikari, A., Tolios, I., Bekkum, M. van, Bal, A., Hendriks, S., ... & Raaijmakers, S. (2020, September). Adversarial patch camouflage against aerial detection. In *Artificial Intelligence and Machine Learning in Defense Applications II* (Vol. 11543, p. 115430F International Society for Optics and Photonics.

Hoogendoorn, S., Taale, H., Wilmink, I., Katwijk, R. van, Immers, B. & Schuurman, H. (2011). *De toekomst van verkeersmanagement* ('The future of traffic management').

Huisken, G., Pepikj, M., Yperman, I., Feitsma, A., Rodrigues, N. & Vlemmings, T. (2020). Optimising Network Traffic Flow with Cooperative Traffic Management in the Amsterdam Region. *Virtual ITS European Congress.*

INRIX. (n.d.). *INRIX AI Traffic*. INRIX. Retrieved 26 May 2021, from https://inrix.com/products/ai-traffic/

Kochenderfer, M. J. (2015). *Decision making under uncertainty: theory and application*. MIT press.

Lange, O. & Perez, L. (3 September 2020). *Traffic prediction with advanced Graph Neural Networks.* https://deepmind.com/blog/article/traffic-prediction-with-advanced-graph-neural-networks

Mahajan, N., Hegyi, A., Weg, G.S. van de & Hoogendoorn, S. P. (2015). Integrated Variable Speed Limit and Ramp Metering Control Against Jam Waves - A COSCAL v2 Based Approach. *2015 IEEE 18th International Conference on Intelligent Transportation Systems*, 1156-1162. https://doi.org/10.1109/ITSC.2015.191

The Netherlands Police. (30 September 2019). *Deployment of smart cameras against distractions in traffic.* https://www.politie.nl/nieuws/2019/september/30/00-inzet-slimme-cameras-tegen-afleiding-in-het-verkeer.html

Raaijmakers, S. A. (2021). *Deep Learning for Natural Language Processing*. Manning Publications.

Rabiner, L. & Schafer, R. (2010). *Theory and applications of digital speech processing*. Prentice Hall Press.

Rijkswaterstaat. (n.d.). *Oude Maas: proefproject brugopeningsvoorspeller Botlekbrug en Spijkenisserbrug* ('Oude Maas: pilot project bridge-opening predictor Botlek Bridge and Spijkenisser Bridge'). Retrieved 9 September 2021, from https://www.rijkswaterstaat.nl/water/projectenoverzicht/oude-maas-proefproject-brugopeningsvoorspeller-botlekbrug-en-spijkenisserbrug

Russell, S. & Norvig, P. (2020). Artificial intelligence: a modern approach. 4th edition. (2020).

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Driessche, G. van den, ... & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484-489.

Smith, S., Barlow, G., Xie, X.-F. & Rubinstein, Z. (2013). Smart Urban Signal Networks: Initial Application of the SURTRAC Adaptive Traffic Signal Control System. *Proceedings of the International Conference on Automated Planning and Scheduling, 23.* https://ojs.aaai.org/index.php/ICAPS/article/view/13594

Smoover. (2015). *Eindrapport* ('Final report').

Sutton, R. S. & Barto, A. G. (2018). *Reinforcement learning: An introduction.* MIT press.

Szeliski, R. (2010). *Computer vision: algorithms and applications*. Springer Science & Business Media.

Technolution Move. (n.d.-a). *Technolution FlowCube: nauwkeurige data over fiets- en voetgangersverkeer* ('Technolution FlowCube: accurate data on bicycle and pedestrian traffic'). Technolution Move NL. Retrieved 26 May 2021, from https://www.technolution.com/move/nl/flowcube/

Technolution Move. (n.d.-b). *MobiMaestro - Technolution Move*. Technolution Move NL. Retrieved 26 May 2021, from https://www.technolution.com/move/nl/mobimaestro/

Vivacity Labs. (9 November 2020). *TfGM chooses Vivacity Labs to deploy AI-controlled road junctions to optimise traffic networks.* Vivacity Labs. https://vivacitylabs.com/greater-manchester-chooses-vivacity-labs-to-deploy-ai-controlled-road-junctions-to-optimise-traffic-networks/

Vivacity Labs. (n.d.). *Home*. Vivacity Labs. Retrieved 28 May 2021, from https://vivacitylabs.com/

Annexes

A. Consulted sources in quick scan

Abduljabbar, R., Dia, H., Liyanage, S. & Bagloee, S. A. (2019). Applications of Artificial Intelligence in Transport: An Overview. *Sustainability*, *11*(1), 189. https://doi.org/10.3390/su11010189
Hoogendoorn, S., Taale, H., Wilmink, I., Katwijk, R. van, Immers, B. & Schuurman, H. (2011). *De toekomst van verkeersmanagement* ('The future of traffic management').
Howell, S. (2018). Meta-analysis of machine learning approaches to short-term urban traffic prediction. *Proc. Scottish Transp. Appl. Res. Conf.(STAR)*, (pp. 1-15).
Huisken, G., Pepikj, M., Yperman, I., Feitsma, A., Rodrigues, N. & Vlemmings, T. (2020). Optimising Network Traffic Flow with Cooperative Traffic Management in the Amsterdam Region. *Virtual ITS European Congress*.
INRIX. (n.d.). *INRIX AI Traffic*. INRIX. Retrieved 26 May 2021, from https://inrix.com/products/ai-traffic/
Krishnakumari, P., Cats, O. & Lint, H. van (2020). A compact and scalable representation of network traffic dynamics using shapes and its applications. *Transportation Research Part C: Emerging Technologies*, *121*, 102850. https://doi.org/10.1016/j.trc.2020.102850
Lange, O. & Perez, L. (3 September 2020). *Traffic prediction with advanced Graph Neural Networks*. https://deepmind.com/blog/article/traffic-prediction-with-advanced-graph-neural-networks
MeBeSafe. (n.d.). *MeBeSafe - Measures for Behaving Safely in Traffic*. MeBeSafe. Retrieved 26 May 2021, from https://www.mebesafe.eu/
Mediator. (n.d.). *MEDIATOR - MEdiating between Driver and Intelligent Automated Transport systems on Our Roads*. Mediator. Retrieved 26 May 2021, from https://mediatorproject.eu/
Meijer, R., Dam, E. van, Bouma, H., Baan, J. & Hurk, S. E. van der (2019). Measurement method for objective cyclist behaviour parameters. *Traffic Injury Prevention*, *20*(sup3), 16-20. https://doi.org/10.1080/15389588.2020.1745196
Niestadt, M., Debyser, A., Scordamaglia, D., & Pape, M. (2019). *Artificial intelligence in transport: Current and future developments, opportunities and challenges*. European Parliamentary Research Service.
Platform Beter Benutten. (n.d.). *Phantom traffic jams*. Retrieved 26 May 2021, from https://beterbenutten.nl/spookfiles
The Netherlands Police. (30 September 2019). *Inzet slimme camera´s tegen afleiding in het verkeer* ('Deployment of smart cameras against distractions in traffic'). https://www.politie.nl/nieuws/2019/september/30/00-inzet-slimme-cameras-tegen-afleiding-in-het-verkeer.html
Rapid Flow Technologies. (n.d.). *Surtrac: Adaptive Traffic Management System for Traffic Optimization*. Retrieved 26 May 2021, from https://www.rapidflowtech.com/surtrac/how-it-works

Smith, S., Barlow, G., Xie, X.-F. & Rubinstein, Z. (2013). Smart Urban Signal Networks: Initial Application of the SURTRAC Adaptive Traffic Signal Control System. *Proceedings of the International Conference on Automated Planning and Scheduling*, *23.* https://ojs.aaai.org/index.php/ICAPS/article/view/13594
Smoover. (2015). *Eindrapport* ('Final report').
Technolution Move. (n.d.-a). *Technolution FlowCube: nauwkeurige data over fiets- en voetgangersverkeer* ('Technolution FlowCube: accurate data on bicycle and pedestrian traffic'). Technolution Move NL. Retrieved 26 May 2021, from https://www.technolution.com/move/nl/flowcube/
Technolution Move. (n.d.-b). *MobiMaestro - Technolution Move*. Technolution Move NL. Retrieved 26 May 2021, from https://www.technolution.com/move/nl/mobimaestro/
Ommeren, C. van, Willems, F., Paardekooper, J. W., Bakri, T. & Beekelaar, R. (2020). *Artificiële intelligentie in mobiliteit en transport* ('Artificial intelligence in mobility and transport'): *Position paper.* https://repository.tudelft.nl/islandora/object/uuid%3Abe6b2833-4410-4700-8151-f0d7d4f52e41
Vivacity Labs. (9 November 2020). TfGM chooses Vivacity Labs to deploy AI-controlled road junctions to optimise traffic networks. *Vivacity Labs.* https://vivacitylabs.com/greater-manchester-chooses-vivacity-labs-to-deploy-ai-controlled-road-junctions-to-optimise-traffic-networks/
Vlahogianni, E. I., Karlaftis, M. G. & Golias, J. C. (2014). Short-term traffic forecasting: Where we are and where we're going. *Transportation Research Part C: Emerging Technologies*, *43*, 3-19. https://doi.org/10.1016/j.trc.2014.01.005
Waze Mobile. (n.d.). *Driving Directions, Traffic Reports & Carpool Rideshares by Waze*. Retrieved 26 May 2021, from https://www.waze.com/
Zhang, J., Wang, F.-Y., Wang, K., Lin, W.-H., Xu, X. & Chen, C. (2011). Data-Driven Intelligent Transportation Systems: A Survey. *IEEE Transactions on Intelligent Transportation Systems*, *12*(4), 1624-1639. https://doi.org/10.1109/TITS.2011.2158001
ZOOF. (n.d.). *ZOOF: Driving smart together!* Retrieved 26 May 2021, from http://www.zoof.nu/home.html

B. Participants expert interviews

The expert interviews were conducted from April to June 2021. The following persons were interviewed:

- Bram Bakker (founder/managing director at Cygnify BV)
- Panchamy Krishnakumari (assistant professor at TU Delft)
- Paul van Koningsbruggen (director mobility at Technolution)
- Stephan Raaijmakers (senior scientist at TNO; professor at Leiden University)
- Petra Delsing (AI quartermaster at IenW)
- Stijn van Balen (product owner datalab at Rijkswaterstaat)
- Fred van der Zeeuw (senior advisor at Rijkswaterstaat)
- Mark Nicholson (CEO at VivaCity)
- Wolfgang Mühlbauer (sales director at INRIX)

- Jose Carazo (solutions architect INRIX)

C. Full list of future applications

A list of 11 possible future applications was drawn up. Of these 11, six applications were selected for the workshop. (These are **highlighted**.) The full list of future applications is:

- **Self-learning signal controllers and traffic control systems: Using reinforcement learning to make traffic control systems adaptive.**
- Use advice to divide vehicles over lanes: Use camera monitoring to determine and recommend the desired distribution of vehicles across lanes to optimise traffic flow.
- **Predicting the effectiveness of an overtaking prohibition for controlling purposes: Use monitoring to determine whether the introduction of an overtaking prohibition for freight traffic has a positive influence on traffic flow, and if so, introduce such a prohibition.**
- Warn of deviating traffic flows: Use anomaly detection to detect deviations in traffic flow, such as traffic driving closer together than usual, and issue warnings.
- **Monitoring network condition with data from multiple parties: Using federated-learning techniques, traffic management applications can be imagined that allow for the development of models based on data that is sensitive or impossible to share.**
- **Predicting travel times and advising routes: Graph neural networks can be used to develop better route advice to communicate to road users. When communicating advice, it is important that road users also receive the reason(s) for the advice.**
- Ecodriving with AI: Using AI, both general and personal speed advice can be given to smooth traffic out and minimise speed differences, thereby improving traffic flow and reducing emissions.
- **Shockwave damping: AI techniques can be used to improve some of the shortcomings of algorithms such as SPECIALIST AND COSCAL.**
- Recognition of traffic signs: Object recognition can be used to detect road signs that road users encounter, and present the information or issue a warning.
- Flexible lane arrangement: Based on the estimate of the expected traffic flow in both directions, the available lanes of the road can be divided between the two directions.
- **Predicting the duration of an incident: Building on Rijkswaterstaat's incident-probability estimate, the duration of incidents can also be predicted.**

D. Selection criteria

The selection criteria have been used to explain why the choice for six applications is considered the most promising. The criteria serve as guidelines; the applications are not scored based on the criteria. The selection criteria are:

- Implementation possible within 1-5 years
- Chance of success (business case, data collection, technical feasibility)
- Desirability (effects and order of magnitude; preconditions to be met; privacy; legislation on the use of data; explainability)
- Distribution over traffic management functions

- Distribution over purpose of the application (sustainability, safety, traffic flow)

E. Workshop on future applications

During a workshop on Thursday 5 August 2021, the future applications were discussed and further elaborated on. All participants were divided into two groups, and three applications were discussed per group. The workshop ended with a plenary session in which the applications were discussed with the whole group, and the results and recommendations for the future were reflected upon. The following participants attended the workshop:

- Panchamy Krishnakumari (assistant professor at TU Delft)
- Ajaya Adhikari (scientist at TNO)
- Marco Schreuder (advisor at Rijkswaterstaat)
- Rudi Kraaijeveld (traffic expert and advisor for Rijkswaterstaat)
- Henk Taale (consultant at Rijkswaterstaat; assistant professor at TU Delft)
- Max Schreuder (project leader at TNO)
- Isabel Wilmink (scientist at TNO)
- Erwin Walraven (scientist at TNO)
- Dawn Spruijtenburg (scientist at TNO)