International Transport Forum

Artificial Intelligence in Proactive Road Infrastructure Safety Management

Summary and Conclusions

187

Roundtable

Artificial Intelligence in Proactive Road Infrastructure Safety Management

Summary and Conclusions

187

Roundtable

The International Transport Forum

The International Transport Forum is an intergovernmental organisation with 63 member countries. It acts as a think tank for transport policy and organises the Annual Summit of transport ministers. ITF is the only global body that covers all transport modes. The ITF is politically autonomous and administratively integrated with the OECD.

The ITF works for transport policies that improve peoples' lives. Our mission is to foster a deeper understanding of the role of transport in economic growth, environmental sustainability and social inclusion and to raise the public profile of transport policy.

The ITF organises global dialogue for better transport. We act as a platform for discussion and pre-negotiation of policy issues across all transport modes. We analyse trends, share knowledge and promote exchange among transport decision makers and civil society. The ITF's Annual Summit is the world's largest gathering of transport ministers and the leading global platform for dialogue on transport policy.

The Members of the Forum are: Albania, Armenia, Argentina, Australia, Austria, Azerbaijan, Belarus, Belgium, Bosnia and Herzegovina, Bulgaria, Canada, Chile, China (People's Republic of), Colombia, Croatia, Czech Republic, Denmark, Estonia, Finland, France, Georgia, Germany, Greece, Hungary, Iceland, India, Ireland, Israel, Italy, Japan, Kazakhstan, Korea, Latvia, Liechtenstein, Lithuania, Luxembourg, Malta, Mexico, Republic of Moldova, Mongolia, Montenegro, Morocco, the Netherlands, New Zealand, North Macedonia, Norway, Poland, Portugal, Romania, Russian Federation, Serbia, Slovak Republic, Slovenia, Spain, Sweden, Switzerland, Tunisia, Turkey, Ukraine, the United Arab Emirates, the United Kingdom, the United States and Uzbekistan.

International Transport Forum
2 rue André Pascal
F-75775 Paris Cedex 16
contact@itf-oecd.org
www.itf-oecd.org

ITF Roundtables

ITF Roundtables bring together international experts to discuss specific topics notably on economic and regulatory aspects of transport policies in ITF member countries. Findings of ITF Roundtables are published in a Summary and Conclusions paper. Any findings, interpretations and conclusions expressed herein are those of the authors and do not necessarily reflect the views of the International Transport Forum or the OECD. Neither the OECD, ITF nor the authors guarantee the accuracy of any data or other information contained in this publication and accept no responsibility whatsoever for any consequence of their use. This document and any map included herein are without prejudice to the status of or sovereignty over any territory, to the delimitation of international frontiers and boundaries and to the name of any territory, city or area.

Cite this work as: ITF (2021), *Artificial Intelligence in Proactive Road Infrastructure Safety Management: Summary and Conclusions*, ITF Roundtable Reports, No. 187, OECD Publishing, Paris.

Acknowledgements

This report reflects an expert discussion held on 10-12 February 2021 at the Roundtable on Artificial Intelligence in Road Traffic Crash Prevention. The report was written by Alexandre Santacreu, who also co-ordinated this piece of research.

The author is grateful to the participants who presented their research at the roundtable:

- Monica Olyslagers (iRAP) for "The Accelerated and intelligent collection of RAP attributes"
- Richard Owen, Craig Smith and George Ursachi (Agilysis) for "Road Danger Prediction: Classic models, AI models and Data Challenges"
- Bas Turpijin (National Road Traffic Data Portal, NDW) for "Probe Vehicle Data"
- Apostolos Ziakopoulos (National Technical University of Athens, NTUA) for "Using AI for Spatial Prediction of Driver Behaviour"
- Luis Fernando Villaça Meyer (Cordial Institute) for "Average Speeds and Road Safety in São Paulo"

The ITF Secretariat would like to thank Professor George Yannis (NTUA) for chairing the discussion and for advising ITF research in the field. Sincere gratitude also goes to the 48 roundtable participants, whose names are listed in Annex A.

The author would like to acknowledge guidance received from a number of experts via email or telephone: Jürgen Beyer (University of Weimar), Jonas Ambeck and Florent Garnier-Follet (Toyota Motor Europe), Caspar Chorus (Delft Technical University), Aliaksei Laureshyn (Lund University), Hlács András (Permanent Delegation of Hungary to the OECD), Jean Pilon-Bignell (Geotab), Henk Stipdonk (KIM), Lee Fawcett (Newcastle University), Marianne Rostoft (TOI) and Dominique Mignot (Université Gustave Eiffel).

This research also benefited from the contributions of members of the ITF Safer City Streets network to an ITF webinar on "Identifying dangerous locations": Rob Viola and Seth Hostetter (New York City Department of Transportation), Remco Smit (City of Rotterdam) and Arjan Knol (DOK Data Netherlands).

Finally, thanks to Stephen Perkins (ITF) for his guidance in the co-ordination of the research, Layla Khalaf (ITF) who assisted in the delivery of the roundtable, Rebecca Brite (independent) who copy-edited the report and Edwina Collins (ITF) who managed the publication process.

Table of contents

Figures

Tables

Executive summary

What we did

This report examines and determines the most relevant cases for artificial intelligence (AI) use in a transport planning context for crash prevention on an entire road network. It explores the possibility of using computer vision to acquire relevant information and the capability of computer models to map high-risk locations. It offers recommendations to stakeholders on the development and appropriate use of life-saving AI solutions.

The report summarises the findings of an ITF roundtable in February 2021 that brought together experts from 33 organisations and 15 countries. Participants represented public authorities, the transport, technology and data industries, research institutes and international organisations.

What we found

AI facilitates proactive traffic safety management in two ways: Through sensors and systems such as computer vision, it helps collect data on road infrastructure condition and traffic events over an entire road network. And, through predictive models, AI learns to identify locations on the network where crash risk is highest. In regions where precise and relevant data exists, AI can identify dangerous locations proactively, before crashes happen.

Decision makers are more familiar with the traditional reactive approach to traffic safety. This involves countermeasures – often revised road geometry – at locations where one or several crashes have occurred. It is straightforward, attracts broad public support and often appears to reduce crash numbers. However, the statistical artefact known as regression to the mean often inflates the claimed benefits and makes the reactive approach appear more efficient than it is.

AI pushes the limits of pattern recognition beyond human capabilities and may thus discover hitherto unknown crash-prone road configurations. However, many consider AI a black box, lacking traditional statistical models' transparency and interpretability of predictions. Fortunately, recent developments in the field of "explainable AI" have begun to fill the gap by disclosing which factors have most influence on predictions overall, which factors explain any given prediction and what would result from changes to design or traffic management on a given road segment.

Computing power being broadly available, what now limits AI is availability of data and of skilled individuals to supervise modelling. Data are often in short supply because they remain in silos instead of being shared, the main barrier being fear of litigation for disclosure of identifiable personal information. The automotive industry holds some of the most precious data for risk prediction, known as "floating car data": They include indicators of traffic volume, speed and incidents of engagement of active safety systems, notably antilock braking systems (ABS), electronic stability programs (ESP) and autonomous emergency braking (AEB).

What we recommend

Develop a competitive market for the sharing and monetising of traffic and mobility data

All stakeholders should seek to stimulate the growth of a diverse, competitive and transparent market for data from vehicles, telematics, smartphone apps and other services. Governments should acknowledge that industrial partners need incentives such as monetisation to lift the barriers to data sharing, develop innovative quality data products and cover the cost of data collection, cleaning, analysis, safe storage and transfer. Collection of road asset data using computer vision is an example of such data products. Road safety could be neglected and lack quality data if data producers do not consider it a core business activity. Relevant privacy safeguards must be in place as an intrinsic part of the legal frame for monetisation of traffic and mobility data.

Do not wait for real-time data before developing risk maps

Road network managers, researchers and consultants should remain aware of hurdles preventing access to real-time data from vehicles, smartphones and wearables. They should start by developing risk prediction models that feed on aggregate instead of real-time data, and map abnormally dangerous locations on the road network. Once aggregated over a year or a week, data have fewer privacy issues and a lower market value but remain highly relevant to risk mapping applications. Rotterdam provides an illustration of how AI can learn and predict crash risk with aggregate input data.

Mandate the sharing of aggregate vehicle data

Governments should consider defining a minimum set of data for all vehicle manufacturers to report, in an anonymous standard aggregate format, to facilitate elaboration of proactive road safety strategies. Such a set could include data on traffic volume, speed distribution and locations where vehicles' active safety systems (ABS/ESP/AEB) engaged. The data's value for producing performance indicators related to the global road safety performance targets set by the United Nations should be a criterion for what should be in the data set. Precise governance aspects covering consent, data collection and data processing should be set by representatives of data protection authorities and privacy organisations, road transport authorities, industry/businesses and academia.

Learn from other fields and best practice for data sharing and privacy protection

Stakeholders should draw parallels with other contexts where government seeks access to crowd-sourced data, such as mobile telecommunication data. Partners should consider more secure alternatives to data exchange, such as exchange of queries and responses instead of raw information. Data providers, integrators and marketplaces should envisage hosting a secure computing workspace to facilitate such an approach.

Support research and innovation towards trusted and explainable AI

Authorities and sponsors should support research on AI use in computer vision and risk prediction. They should build trust in AI by facilitating benchmarking and validation of methods (including machine learning) for proactive road network safety management. New systems' capabilities should be fully assessed. Priorities should include development of "explainable AI" techniques providing road network authorities with clarity on which interventions would bring the most benefits. Funding must also be available to road safety professionals to conduct post-intervention assessment and validate or recalibrate the risk prediction tools.

Align new tools with precise policy objectives

Governments should verify that new tools for proactive safety management do not distract from clear objectives. If the policy goal is to eliminate fatal and serious crashes, governments should commission research to assess the capability of proxy data and risk mapping tools to predict those specific crash types, in both cross-sectional and longitudinal analyses.

Develop new skills and digital infrastructure

Authorities need to develop the skills to become informed consumers in this complex field. They should seek to promote a multidisciplinary approach to road safety combining expertise in data science, technology and safety. Authorities should also create national access points for collecting and reporting transport-related data.

Clarify regulatory frameworks for data protection and digital security

Governments should clarify privacy protection rules where uncertainty on their interpretation deters data sharing. Such clarification would benefit all economic sectors and areas of public action, well beyond the road safety field. Governments should also examine how freedom of information laws interact with data protection laws. Finally, they should review current regulations for protection of proprietary data and intellectual property against misuse of data and infringement of data exchange conditions.

Design user-friendly, risk-mapping tools

To encourage tool uptake and road safety investment, authorities should specify particular features in the risk-mapping tools they develop or procure. Tools should include estimates of the total annual social cost of predicted crashes on each road section and estimates of interventions' benefit/cost ratio. They should also have accessible, user-friendly interfaces. The city of Rotterdam and the International Road Assessment Programme are among entities that have developed examples of such tools.

The proactive approach to road safety

Road traffic injuries kill 1.35 million people each year, and are the top cause of deaths worldwide among those aged 5 to 29 (WHO, 2018). More than half of those killed are vulnerable road users – motorcyclists, bicyclists and pedestrians. In response to this challenge, the UN General Assembly (2020) proclaimed the Decade of Action for Road Safety 2021-2030 and set a target of preventing at least 50% of road traffic deaths by 2030.

Without new tools and methods, the world is unlikely to meet UN targets. An earlier decade of action for road safety stopped the increase in road deaths but failed to reduce their number significantly. Fortunately, considerable change is happening in the road safety field. A growing number of national and local governments adopted Vision Zero and seek to eliminate road traffic deaths. To make it happen, they embrace a holistic approach to traffic safety called the Safe System approach.

The principle of a Safe System

The Safe System approach is considered the best practice among road safety professionals. It acknowledges that humans inevitably make mistakes, and that all parts of the transport system must contribute to avoiding a fatal outcome in the event of a collision. Vehicle design, road geometry and traffic rules should reflect the human body's known limits in withstanding crash forces. Roads and streets should be forgiving. Vehicles should protect both their occupants and vulnerable road users.

In a Safe System, the planning approach is not strictly *reactive* to incidents; such an incremental approach lacks pace, scale and ambition. Instead, as Table 1 shows, it is *proactive*, identifying risk factors in all parts of the system and seeking to address them before serious harm occurs. It puts particular emphasis on building a safe road system, rather than fixing crash accumulation spots.

Table 1. Comparing the traditional road safety approach with a Safe System

Question	Traditional road safety approach	Safe System approach
What is the problem?	All traffic crashes	Crashes resulting in fatal and serious injuries
What is the goal?	Reducing fatal and serious injury numbers	Eliminating fatal and serious injuries
Which planning approaches?	Reacting to incidents Incremental approach to reduce the problem	Proactively targeting and treating risk Systematic approach to build a safe road system
What causes the problem?	Non-compliant road users	People inevitably make mistakes People are fragile
Who is ultimately responsible?	Individual road users	Responsibility is shared by individuals with system designers
How does the system work?	Isolated interventions	Different elements are combined so that if one fails, others provide protection

Source: Adapted from ITF (2016a) and Belin et al. (2012).

The reactive and proactive approaches

Decision makers and citizens are familiar with the traditional reactive approach to traffic safety involving the redesign of a location where one or several crashes have occurred. The rationale is that crashes will continue to happen where they have already happened, especially where a spatial cluster has occurred. The reactive approach is straightforward, attracts broad public support and often appears to reduce crash numbers. However, some crash number reduction would most likely have happened without intervention, through the statistical phenomenon called regression to the mean (Box 1).

Box 1. Regression to the mean

In most before-and-after studies reported in road safety literature, remedial measures have been deployed after a period of collision counts deemed unacceptably high. However, a period of high collision counts at a specific location is often due to random fluctuation in relatively small collision numbers. Due to such fluctuation, it is likely that collision counts will later return to a lower baseline level regardless of the value of the intervention. This natural statistical phenomenon is known as *regression to the mean.*

Figure 1. Separating regression to the mean from the true benefit of an intervention

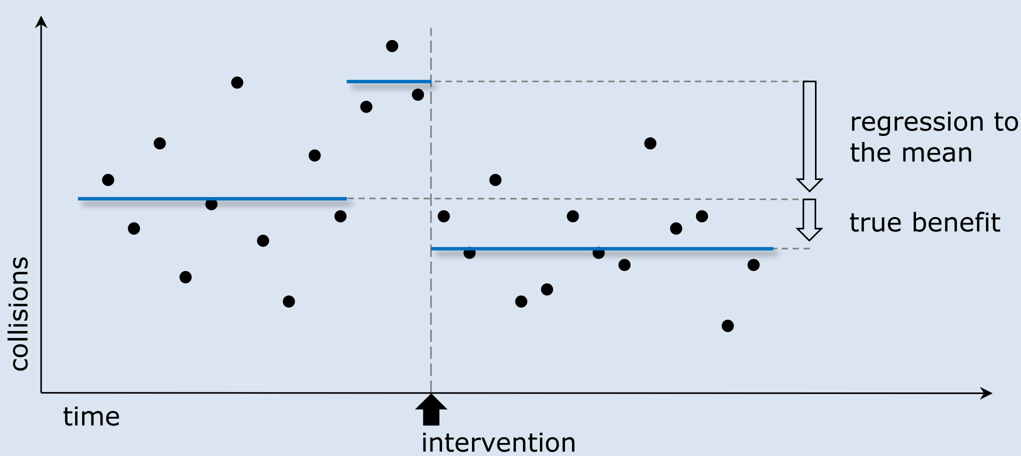

Studies assessing road safety programmes' effectiveness are notoriously bedevilled by the problem of regression to the mean. Figure 1 indicates that most of the change observed immediately after a road safety intervention would have been observed anyway, even without intervention, simply through regression to the mean.

Source: Adapted from Fawcett et al. (2017).

Another weakness of the reactive approach is that crashes must happen at a given location, and be reported to the authorities, for the location to be considered for safety treatment. Not only is this morally questionable, but such an approach can be incomplete or biased due to the vast under-reporting of crashes (ITF, 2011; Aldred, 2018). Under-reporting being greatest in lower-income regions, the approach may be biased against diagnosis of high-risk locations in the regions most needing road safety treatment.

A proactive approach also takes crash data as an input, but complements it with other kind of information and seeks to identify systematic risks and solutions. Its benefits are not limited to the places where crashes have occurred and been reported. The following examples help illustrate the proactive approach to road safety in the context of road design and operations:

- By observing a correlation on many sites between a specific junction design and casualty numbers, one can predict the casualty risk on all parts of the network designed in that way and intervene on all such locations regardless of their crash history.
- Knowing the correlation between speed and crash risk (ITF, 2018), one can map precise speed data to identify areas where speed management solutions could prevent the next serious injury.
- By developing models capable of predicting the number of serious crashes from traffic data (close calls, vehicle speed and mass, emergency braking events, swerving, etc.), one can identify dangerous locations before crashes happen. This is often described as using surrogate safety metrics (Box 2). Such models can be further improved using other types of data, including road user mix, road geometry and weather, to name only a few.

Underpinning the use of surrogate safety metrics is the assumption that correlations exist between crashes and conflicts. This explains why analysis of conflicts helps predict crashes. Hydén (1987) uses a “safety pyramid” to describe this relationship (Figure 2). The pyramid base consists of normal traffic encounters that are quite safe and frequent. The tip of the pyramid represents the most severe events, such as crashes resulting in injuries or fatalities, which are highly infrequent.

Hydén proposed that the number of conflicts could serve as a proxy to predict, and hence prevent, occurrence of rare but more serious crashes. A method for counting the number of conflicts is a surrogate safety metric.

Figure 2. Hydén’s safety pyramid

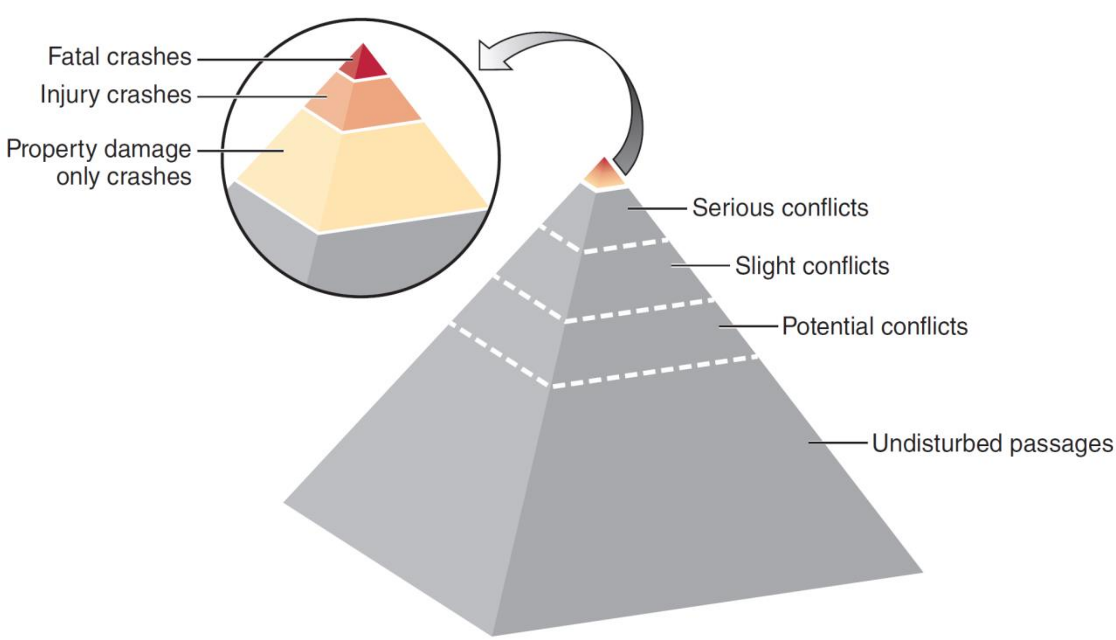

Source: Chang, Saunier and Laureshyn (2017), adapted from Hydén (1987).

Experts have challenged the use of surrogate safety metrics. Some fear such metrics would be elaborated to predict all crashes, not just serious and fatal ones, distracting professionals from focusing on preventing the latter type of crash. The concern is legitimate and is based on the fact that serious and fatal crashes are rare. Training statistical models on rare events is difficult, as it requires longer time frames and a wider road network. Ultimately, misuse of surrogate safety metrics could result in fewer crashes but more fatal and serious injuries.

Accordingly, several developers of surrogate safety metrics have worked to account more precisely for road users' kinetic energy, conflict angle and vulnerability so as to predict serious and fatal crashes. Indeed, one could argue that a risk of head-on collisions at high speeds is enough to predict serious and fatal casualties, regardless of crash history or near misses. The International Road Assessment Programme (iRAP), for instance (**Error! Reference source not found.**), takes kinetic energy into account in its star rating o f road assets.

A major strength of surrogate safety metrics is their ability to estimate a road safety intervention's effects within weeks, whereas analysis of serious and fatal crash data can take three to five years, at best.

Box 2. Surrogate safety metrics

Surrogate safety metrics are increasingly used to diagnose traffic safety problems so that action can be taken before a serious crash happens. They are typically based on identifying the occurrence and severity of traffic conflicts involving evasive actions such as braking or swerving, also known as near misses or close calls. Using artificial intelligence (AI), large amounts of video footage and data points from roadside cameras and sensors can be analysed to identify close calls in which a crash was narrowly avoided.

Street imagery can be used to support assessment of roads' safety characteristics. This is already being done for attribution of star ratings in road assessment programmes. The next generation of star rating programmes will likely benefit from more frequent and broader image data collection as well as automated image analysis using computer vision. Drones and satellites can capture additional data and will play an increasing role in road safety. Such innovations will facilitate road safety investment planning as well as monitoring of results.

Source: ITF (2019).

What is Artificial Intelligence?

Artificial intelligence (AI), a term coined in 1956, is an umbrella term for algorithmic and computer science techniques allowing computer software to learn from experience, perceive, cognise, adapt to situations, reach decisions and act. Machine learning (ML) is an area of AI that involves detecting patterns in data, making predictions or enabling actions to be taken without explicit programming in the form of the usual "if-then" routines and without classic automation and control engineering. Most people are familiar with the use of AI in games and speech recognition. Its use by government is less visible but nonetheless promising. To foster use of trustworthy AI solutions, the G20 has adopted human-centred AI Principles that draw from the OECD AI Principles (Box 3).

The OECD AI Experts Group defines an AI system as "a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments". AI applications are experiencing rapid uptake in sectors where they can detect patterns in large volumes of data and model complex, interdependent systems to improve decision making (OECD, 2019b).

AI has undergone rapid transformation over the last decade. Since 2011, the AI subset called machine learning has dramatically improved machines' ability to make predictions from historical data. The OECD (2019b) cites "the maturity of a ML modelling technique called 'neural networks', along with large datasets and computing power", as underlying the recent expansion in AI development.

Hopes and questions raised

AI facilitates proactive management of a safe road network in two ways. First, through sensors and systems such as computer vision, it helps collect and label data on infrastructure condition and traffic events over an entire road network. Second, through predictive models, AI learns to identify locations where crash risk is highest before crashes happen. For both aspects, this report outlines where and how AI adds value and what policy makers need to know to make the best use of AI.

Trust and public acceptance will determine the adoption of AI-based techniques in the field of road safety, as in other fields. Is AI's performance superior to those of various other techniques? Are AI outputs transparent and explainable, or do AI systems function as an impenetrable black box? Can AI provide guidance on which countermeasure could address a safety problem? Will road authorities embrace the predictive safety management principle? Will authorities have the resources to elaborate and validate crash prediction models and react to these new crash predictions?

AI is notoriously dependant on provision of large amounts of quality data. Will drivers, carmakers, commercial transport operators and telematics companies be willing to share the data they produce? Will they consider their data sensitive for privacy or commercial reasons? What are effective ways to protect privacy? Should authorities ask the transport industry to provide safety-relevant data free of charge? Examples exist of public agencies that have developed risk prediction models with no or limited data from the private sector, but the roundtable discussion left no doubt on the added benefit of using data from the transport industry.

Crash risk prediction could benefit from a vast number of data inputs, such as speed compliance, hard braking events, road star rating and exposure data, to give just a few examples. Which data sets are the most helpful in prediction of road crashes?

Crash data form the pillar of crash prediction models. Are there effective methods for training AI systems where crash data is under-reported to police? Will the use of incomplete crash data lead to AI perpetuating existing biases to the detriment of some user groups, areas or populations?

The following two chapters reflect the main ways AI can support road safety interventions. One involves sensing the road asset and traffic conditions, including near misses. The other identifies risky areas and elaboration of countermeasures.

Box 3. Principles for responsible stewardship of trustworthy AI

The OECD calls on all AI actors to promote and implement, according to their respective roles, the following principles for responsible stewardship of trustworthy AI.

1. Inclusive growth, sustainable development and well-being

Stakeholders should proactively engage in responsible stewardship of trustworthy AI in pursuit of beneficial outcomes for people and the planet, such as augmenting human capabilities and enhancing creativity, advancing inclusion of underrepresented populations, reducing economic, social, gender and other inequalities, and protecting natural environments, thus invigorating inclusive growth, sustainable development and well-being.

2. Human-centred values and fairness

a) AI actors should respect the rule of law, human rights and democratic values, throughout the AI system lifecycle. These include freedom, dignity and autonomy, privacy and data protection, non-discrimination and equality, diversity, fairness, social justice, and internationally recognised labour rights.

b) To this end, AI actors should implement mechanisms and safeguards, such as capacity for human determination, that are appropriate to the context and consistent with the state of art.

3. Transparency and explainability

AI Actors should commit to transparency and responsible disclosure regarding AI systems. To this end, they should provide meaningful information, appropriate to the context, and consistent with the state of art: i) to foster a general understanding of AI systems, ii) to make stakeholders aware of their interactions with AI systems, including in the workplace, iii) to enable those affected by an AI system to understand the outcome, and, iv) to enable those adversely affected by an AI system to challenge its outcome based on plain and easy-to-understand information on the factors, and the logic that served as the basis for the prediction, recommendation or decision.

4. Robustness, security and safety

a) AI systems should be robust, secure and safe throughout their entire lifecycle so that, in conditions of normal use, foreseeable use or misuse, or other adverse conditions, they function appropriately and do not pose unreasonable safety risk.

b) To this end, AI actors should ensure traceability, including in relation to data sets, processes and decisions made during the AI system lifecycle, to enable analysis of the AI system's outcomes and responses to inquiry, appropriate to the context and consistent with the state of art.

c) AI actors should, based on their roles, the context, and their ability to act, apply a systematic risk management approach to each phase of the AI system lifecycle on a continuous basis to address risks related to AI systems, including privacy, digital security, safety and bias.

5. Accountability

AI actors should be accountable for the proper functioning of AI systems and for the respect of the above principles, based on their roles, the context, and consistent with the state of art.

Source: OECD (2019a).

Sensing and sharing safety-relevant data on an entire road network

No proactive approach to road network safety can be developed without provision of quality data on the road asset, traffic and traffic events. Data provision is what limits and will continue limiting the performance of AI models in this area. Hence this section explores the various data sources that are most likely to help predict crash risk. It considers the role of AI in data collection and proposes data governance solutions to facilitate data sharing.

AI is particularly valuable in video processing, e.g. to detect traffic conflicts or infrastructure design features. Fixed equipment, including CCTV cameras and roadside traffic sensors, provides a rich stream of information on traffic conditions. In New York City, for instance, researchers detected occurrences of double parking using an AI computer vision application with a municipal open data feed providing images from 700 CCTV cameras. To collect data over an entire road network however, other data sources, such as probe vehicles, are needed.

Nine out of ten of the world's road deaths occur in low- and middle-income countries (LMICs). Can data-driven crash prediction models work in LMICs, where the market penetration of connected vehicles will remain low for the near future? Yes, in two ways. The first concerns computer vision applications in road surveys commissioned to assess risk level, such as the surveys feeding into iRAP assessments (Box 4). Computer vision helps reduce such surveys' cost and thus facilitates diagnosis of road safety risk across a much wider network. The second way uses the telematics devices mounted on large commercial fleets, regardless of vehicle age.[1] Not only do telematics capture data on vehicle dynamics, but they also offer the opportunity to collect images of the road thanks to the increasingly popular on-board cameras. In LMICs, smartphone apps could also become a key source of data on vehicle dynamics. Partnerships with insurance companies and navigation app developers should be envisaged to this end.

Box 4. Automated coding for the International Road Assessment Programme

The International Road Assessment Programme is a registered charity established to deliver the vision of a world free of high-risk roads. It has partnerships with regional, national and local government, development and financing institutions, research bodies and civil society in more than 100 countries, with over a million kilometres of roads assessed thus far.

iRAP's star rating models predict crash risk for each road section. They are based on from road attributes using well-established crash modification factors (Box 5). Road attributes primarily reflect road design and conditions, but also include supporting data such as traffic speed and flow. Collecting this information typically requires dedicated survey teams and manual coding of attributes, which is a barrier to global deployment and frequent updating of road assessments.

iRAP's *accelerated* and *intelligent* collection and coding of road attribute data, known as AiRAP, has the potential to reduce the time and effort road safety assessments require, reduce costs per unit of road length, improve accuracy and make data available for every road on earth. To unlock this potential, AiRAP's main goal is to source road attribute, traffic flow, and speed data following iRAP's "common global specification" and map safety performance and star ratings.

AiRAP captures advances in artificial intelligence, machine learning, vision systems (street and sky), LIDAR, telematics and other data sources. Providers can propose translation routines to convert machine collected data into iRAP-compliant data (Figure 3). The conversion is certified by iRAP. The data are to be published worldwide, along with any associated conditions of use, as agreed by iRAP and the provider.

Source: iRAP (n.d.).

Figure 3. Source of ai-RAP data and conversion to iRAP attributes

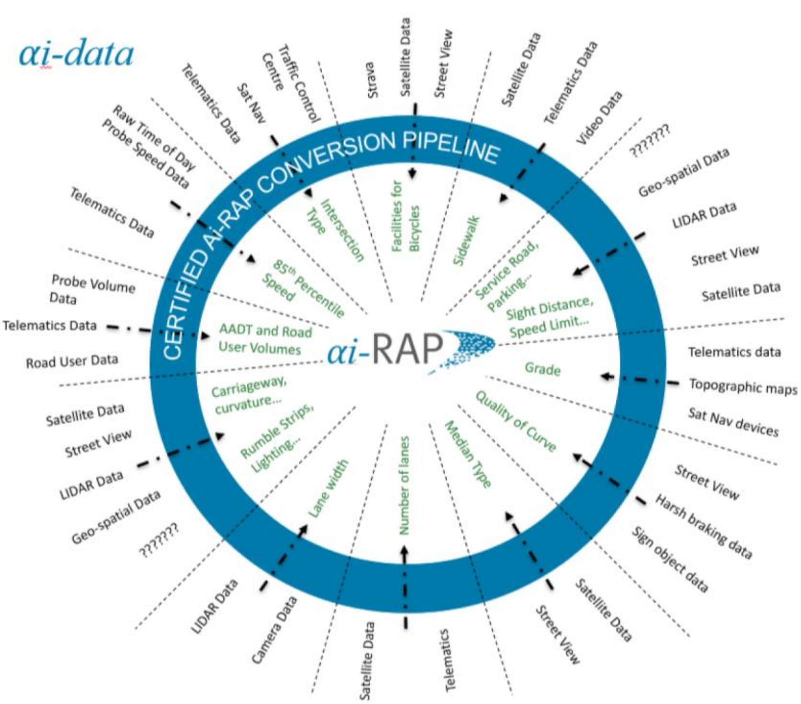

Source: iRAP (n.d.).

Box 5. Crash modification factors

A crash modification factor (CMF) is a multiplicative factor used to compute the expected number of crashes after a given countermeasure at a specific site. For instance, a countermeasure with a CMF of 0.80 is expected to reduce the number of crashes by 20%.

The Highway Safety Manual (AASHTO, 2010) provides information and tools to facilitate roadway design and operational decisions based on a quantitative assessment of their safety consequences. CMFs are one such tool, used either to support an agency's roadway safety management process or as input to safety prediction methods.

CMFs come with guidance that specify in which context the results are transferable. Roundtable participants stressed how important such guidance is to prevent misuse of CMFs.

Source: CMF Clearinghouse (2021).

Roadway attributes

Figure 3 gives an idea of the diversity of road attributes that could contribute to a road safety assessment. Infrastructure star rating, such as that delivered by the International Road Assessment Programme (Box 4) is quite a data hungry process.

One can monitor pavement condition using AI with data from dedicated survey vehicles, from partner fleets or from smartphones. Vehicle vibration, pavement images and LIDAR point-cloud are examples of data sources for pavement maintenance (ITF, 2021). Smartphone motion sensors and cameras can supply vehicle vibration data and pavement images. Specialised companies propose such a service, with smartphones mounted on patrol or survey vehicle dashboards. Alternatively, road users may allow apps to collect such data when they drive.

Street view images are a particularly useful source of information. The state of North Carolina is developing computer vision to automate roadway feature extraction from video logs (Box 8). New vehicles come with cameras and other sensors to read road signs and offer various advanced driver assistance systems (ADAS). Autonomous vehicles will come with even more sensors and greater connectivity. These trends create an opportunity to collect much more information on the road asset layout and condition.

Curb activity and the number of property access points along a road could make a significant contribution to a model’s prediction power. Yet this information is rarely available and needs to be created by skilled individuals from local data sources.

Some attributes, such as slippery road surfaces in winter conditions, are dynamic by nature. Collecting such attributes in real time enables road authorities to manage road risk in a timely manner (ITF, 2019). Data collection could happen in real time where possible, but also in batches to support non-real-time applications, e.g. risk mapping over an entire road network, including segments without mobile data coverage.

Traffic attributes

The traffic attributes of a given road section include traffic volume, speed and events such as near misses. A single pass of a probe vehicle cannot capture traffic attributes, although it can capture most roadway attributes simultaneously. Estimating traffic attributes requires a fleet of instrumented vehicles or road users, raising practical questions about fleet representativeness and data privacy.

There was consensus among roundtable participants on collecting traffic volume, speed and emergency manoeuvre data as a matter of priority to feed into risk prediction models. They suggested collecting this data via vehicles, smartphones, wearables and other devices to achieve whole network coverage, even in LMICs. Analytics companies such as The Floow already supply telematics data to the iRAP. One participant reported using data from a system that not only records vehicle position and dynamics, but also fits two cameras: one facing the driver and the other on the road ahead.[2] There is hardly any limit to the number of traffic attributes that can be collected and might ultimately increase the predictive power of risk prediction models.

Traffic volume

Experts identified traffic volume as being among the most useful input data for risk prediction modelling. Traffic volume, frequently used as exposure data, is the cornerstone of risk analysis and yet is often lacking in risk prediction modelling exercises.

Telematics companies have estimated traffic volumes along entire road networks in several countries. Since the source data come from a sample of the vehicle stock and lacks representativeness, the estimation is particularly challenging. Traffic volume for areas near logistics depots could be overestimated due to the concentration of instrumented commercial vehicles, for instance, while volume in lower-income residential areas could be underestimated due to the lower market penetration of high-end connected vehicles equipped with on-board navigation systems.

The key to estimating traffic volume on an entire road network lies in the availability of ground truth data, such as traffic counts, as open data. The University of Central Florida estimated cycling volumes, using trip data from a fitness app called STRAVA, by modelling the relationship between the volume from the app and that from actual counts. Although incomplete and biased, the app data were adjusted to provide a bicycle traffic estimate for every link of the cycling network. The same technique could be applied to any source of traffic data, such as insurance or navigation telematics. Agilysis, an analytics company, reports having conducted such a protocol to estimate traffic data on every road segment in the UK.

Traffic speed

Speed management is at the core of the Safe System approach, as it determines the amount of harmful kinetic energy in the transport system. Speed limit compliance is one of 12 voluntary targets[3] for the new UN Decade of Action for Road Safety 2021-2030. The target for 2030 is to halve the proportion of vehicles travelling over the posted speed limit. Evidence suggests that a 1% change in speed results in a 4% change in the road fatality number (ITF, 2018).

There is thus a need to draw speeding heat maps as a basis for development of targeted engineering, education and enforcement measures. Frequent updates of such maps would help in assessing such measures' effects. A road safety software provider called VIA reports using vehicle telematics data from the HERE platform to provide most city authorities in the Netherlands with speed maps. VIA's dashboard separates out congested periods to reveal speeds in free flow conditions.

GPS technology could power traffic speed data collection in some circumstances, but is not accurate enough in urban contexts and is unavailable in tunnels. A telematics company called OSeven reports performing a data fusion exercise with GPS and accelerometer signals. Another solution is to read vehicles' odometer data feed.

Biases in the vehicle sample could have a significant effect on speed estimates' accuracy. Insurance telematics nudge drivers to adopt a safer attitude and comply with speed limits, so insurance telematics data may not be representative of average driver behaviour. Navigation telematics, by flagging speed enforcement sites, could facilitate deliberate speeding between them. Yet despite its vulnerabilities, speed data covering an entire road network should be a priority item in the list of safety-related data for road authorities to procure.

Routine provision of speed data by all car manufacturers could alleviate the biases found in insurance and navigation telematics data, although it would introduce a bias towards more recent models of connected vehicles.

Near misses and other potential crash predictors

Naturalistic driving studies often result in mapping of abrupt vehicle movements, but rarely provide a sufficient sample size to cover a complete network or provide robust crash predictors. To identify hotspots of abrupt vehicle movements at scale, it thus seems essential to take advantage of existing technology, such as connected vehicles and smartphones.

Box 6. Road safety-related minimum universal traffic information

The European Union's Directive 2010/40/EU calls for provision of road safety-related minimum universal traffic information (known as safety-related traffic information or SRTI) free of charge to road users. The European Commission specifies that the SRTI service covers eight categories of events or traffic conditions:

(a) temporary slippery road;

(b) animal, people, obstacles, debris on the road;

(c) unprotected accident area;

(d) short-term road works;

(e) reduced visibility;

(f) wrong-way driver;

(g) unmanaged blockage of a road;

(h) exceptional weather conditions.

Source: European Commission (2013).

To develop AI-powered crash risk prediction, roundtable participants considered collection of customer engagement data for anti-lock braking systems (ABS), electronic stability programs (ESP) and autonomous emergency braking (AEB) to be a major opportunity. The automotive industry is not currently required to collect ABS/ESP/AEB engagement data systematically, but some carmakers do collect and monetise such data. Volvo's Connected Safety function, for instance feeds ABS and ESP data to the cloud so that other Volvo vehicles can benefit from them.

Some data processing methods use thresholds to simplify raw signals but the process can destroy information. It may also be difficult to compare or merge data from various providers adopting different thresholds. More generally, data integration always entails information loss. One may lose the most important part of the data without even being aware of it. Ideally, one should start integrating data after having defined the use case for it. Careful selection and processing of sensor signals determine the data usefulness for a specific use case.

Even if signal processing methods are harmonised, differences in the sensor hardware limit the comparability of data collected from different vehicles or devices such as smartphones and telematics equipment. Lime, a micro-mobility company, reports possible variability in data availability and applicable safety metrics between vehicle models. The Floow, a telematics company, observes that accelerometer chipsets can differ between smartphones even for the same make and model. Normalising the signal from each device is a huge task that not all analytics companies do.

Barcelona City Council, looking to increase road safety, has fitted buses and other municipal vehicles with Mobileye's collision avoidance system.[4] The technology constantly scans the environment, including certain blind spots not visible to the driver. Mobileye experts consider near misses and hard braking events to be the most significant crash predictors. They have mapped braking and cornering hotspots using connected vehicles in several countries, including LMICs. With this information, road authorities can undertake proactive and targeted inspections, understand if a serious crash risk exists and take action to address it.

Lime reports using vehicle sensor data to automatically detect potential crashes. The solution involves machine learning to identify a series of safety-critical event signatures. The data thus gathered provides insight that can assist in tailoring safety programmes, local operations and hardware development.

Data sharing and aggregation

Local and national governments have access to a wealth of information through roadside sensors and CCTV cameras. They should not neglect this information source and should consider making it open data. In turn, open data should be well publicised so crash prediction model developers can benefit from it.

However, in most cases the vast majority of the road network is not equipped with sensors. A risk assessment covering an entire road network thus requires access to private sector data. The EU Data Task Force (DTF) created an ecosystem for sharing safety-related traffic data and information between vehicle manufacturers and governments (Box 7). It has been praised for taking the automotive industry in the same direction as the air transport industry, which adopted the principle of sharing *all* safety-related data.

The following sections examine several barriers to the sharing of data:

- the silo effect – the lack of connections between organisations and between teams within organisations
- technical costs (collecting, processing, hosting, etc.), which are not negligible
- privacy protection imperatives and associated fears of litigation, often cited as the #1 barrier
- commercial sensitivity.

Silos and interoperability

Data marketplaces could go some way towards breaking the silos. They also create a financial incentive for disseminating data that would otherwise not be shared spontaneously. HERE, CARUSO and Amazon are among the companies developing traffic data marketplaces. The HERE platform includes free data sets that original equipment manufacturers (OEMs) provide to governments, public agencies and universities, with conditions specified by a creative commons licence or other licence.

With its global presence, iRAP is working to create a marketplace that federates stakeholders and secures data from a wide range of sources and competing suppliers, using specifications tailored for the road and traffic attributes that underpin star rating algorithms.

Caution is needed to avoid creating yet another data silo, one limited to road safety professionals. Most of the data used in road safety analysis could serve other purposes as well, powering research in economic activity, climate change, public health and other fields. Data sharing platforms should facilitate mutual exchange among a range of research fields. Making such connections could also strengthen political support for road safety policies, for instance when speed management reduces greenhouse gas emissions, local air pollution and noise in addition to making the road network safer. Elected officials rarely consider road safety a top priority, which is why co-benefits must be highlighted.

Interoperability greatly facilitates data sharing. A cluster of vehicle manufacturers support a data model called Vehicle Signal Specification (VSS) and have started deploying it in production vehicles.[5] VSS introduces a domain taxonomy for vehicle signals that could become a standard in automotive applications. Similarly, the automotive industry's car-to-cloud data standard, called SENSORIS, ensures that data from separate manufacturers can easily be combined.[6] If widely adopted, such standards will facilitate data sharing within the industry, with government agencies and with other parties. Different fleets could use a single application programming interface, unlike today when data aggregators' approaches vary by manufacturer. Data standards could make it easier for manufacturers to monetise their data, selling it to an aggregator or a telematics company. A common data format from cloud to cloud could also facilitate data merging, for instance when several manufacturers wish to contribute to detection of temporary slippery roads.

To facilitate data set merging, Sustainable Mobility for All (SuM4All), an advocacy platform for international cooperation on transport and mobility issues, recommends developing common data platforms at the regional and national levels (SuM4All, 2021). This requires collaboration between public, private and civil society members to identify priority use cases and corresponding data points.

Creating a digital representation of road network geography is one challenge in designing a platform for sharing and merging road safety data. To what level of detail should the network be described? One solution would be to define a standard for geocoding of the network, that all data providers would adopt. Another would be to retain raw spatial coordinates and let data aggregators perform whatever analysis they need. It is indeed hard to bring data sets together, so this could take more time than doing the actual research with them.

Box 7. An ecosystem for the exchange of safety-related traffic data and information

Established in 2017, the EU Data Task Force was a public-private partnership that enabled collaboration between vehicle manufacturers and countries to enhance traffic safety for all road users. Its members created Data for Road Safety (DFRS) as an 'ecosystem' to bring safety-related traffic information to road users across the EU, in accordance with EU regulations (Box 6). DFRS involves pooling SRTI data from multiple sources, including road infrastructure and various vehicle manufacturers, and making them accessible to all project partners. Within the ecosystem, five roles are identified:

- data sources, which share or provide access to data
- aggregators, which enrich these data, e.g. harmonising and cleansing data from multiple sources
- creators, which uses the available data to create SRTI
- National access points, a regulated role for EU member states, providing access to SRTI
- Service providers, which render and distribute SRTI directly to end users.

The ecosystem relies on a common interpretation of Delegated Regulation 886/2013 and the DFRS Multi-Party Agreement represents practical implementation of EU law. Data are exchanged within the ecosystem for the sole purpose of improving road safety, without any financial compensation between the parties. The agreement prohibits using the data for any other purpose, as the members consider other uses to be commercial use cases.

The Dutch national access point, NDW, received over a million vehicle-generated messages for ABS engagements over the course of a seven-month pilot. This illustrates the potential for sharing of safety-related data in general, and emergency braking data in particular, by vehicle manufacturers. Aggregation of such data to produce ABS engagement heat maps, however, falls outside the scope of pro-bono data sharing; it is considered commercial use, as is road asset management.

Source: EU Data Task Force (2020).

Monetisation

Given vehicle crashes' economic cost – around 3% of GDP in low, middle and high income countries (WHO 2015, Wijnen et al., 2017) – it is not unreasonable to expect governments to fund procurement of quality data in response. The data often come from individuals and their vehicles. The current situation, where drivers act as probes/proxies to collect road network data but have no ownership of them, may seem paradoxical. Yet drivers commonly trade their data for a free service or discounted insurance premium. There is also room for tools that collect personal data and give individuals the choice of selling it or donating it to third parties such as governments.

One rationale for monetising data from individuals is that it creates the financial incentive without which much data collection, cleaning, processing and hosting would simply not happen, and the road safety community would lack a precious source of information.

Governments are considering the option of asking the transport industry to provide data free of charge. The burden on the industry would remain less than proportionate to the cost of road crashes, an externality largely borne by society. Pro bono data sharing is already part of e-scooter operators' permit conditions. The experts did not consider this approach to be the most conducive to deployment of accurate risk prediction tools, however. Transmitting, translating, processing, aggregating, anonymising and disseminating data have a considerable financial impact. Mandatory data sharing could be a disincentive to collect data in the first place. Participants felt that data quality and availability would go down if policy makers required the automotive industry to supply data free of charge to governments. Instances of such requirements should be limited to the simplest cases of data reuse.

The availability of data at a fair cost will depend on the level of competition among data producers and data platforms. As in other industries, monopolistic positions could lead to higher prices and hinder adoption of data-driven road safety tools. In AiRAP (**Error! Reference source not found.**), data product s tandardisation seeks to facilitate competition across providers. It is based on the Common Global Specifications for coding of road features. The AiRAP protocol also includes independent confirmation that whoever provides data adheres to quality levels and repeatability criteria. Such accreditation offers transparency on where the provider matches the criteria and where it does not (e.g. certain countries or road types). On the other hand, standardisation could delay the data sharing process. For instance, precisely defining each SRTI (e.g. what is a slippery road) is time consuming, especially at the international level, whereas the automotive industry is willing to share raw data immediately.

Privacy

Privacy regulations such as GDPR in Europe are a major obstacle to sharing of micro-level data, e.g. from telematics. Yet GDPR makes people more confident that their data will not be misused: it creates the trust that facilitates data collection. People might be more willing to give away their personal data if they knew exactly who used it, how they used it and for what purpose. Experts believe many opportunities for automated data collection are missed due to high levels of mistrust in government and in private companies. Where trust in government is low, the roll-out of high-resolution video cameras, for instance, becomes impossible, preventing the wider roll-out of surrogate safety analysis such as that demonstrated in the US city of Bellevue, Washington.

Data protection regulators should work with transport and interior ministers to review and refine the trade-offs between protecting privacy and eliminating road traffic deaths. Setting clear rules on this matter would lift a major barrier currently deterring many OEMs and telematics companies from sharing vehicle data. Privacy concerns pose a risk of leading authorities to use inadequate, second-best data sources whose likely biases lead to wrong conclusions.

It is important to bear in mind that there are multiple ways to aggregate vehicle data, primarily (a) by road segment/by area across many vehicles and (b) by driver/by vehicle across many road segments. Aggregating driver behaviour over a segment of roadway across many vehicles can resolve privacy issues as it is one way to eliminate personal identifiable information and unlock data sharing perspectives. It would thus be possible to produce data on traffic volume, speed, speed percentiles, ABS/ESP/AEB engagements, etc. over a whole week or month. Roundtable participants called for the adoption of such “privacy by design”.

Aggregation, however, destroys some information. For instance, it makes it difficult to investigate any possible bias in the data. Having a specific use case in mind is important before performing data aggregation.

To address businesses’ fear that data sharing will result in lawsuits on privacy grounds, legislators should review the current liability regime. Some participants called for a Good Samaritan approach that encourages data sharing, especially in cases where the data is not monetised.

Governments should follow simple data governance principles, such as purpose specificity and data minimisation, and comply with privacy protection regulation. Location and trajectory data are the most vulnerable, so they should be covered by the most robust protection methods (ITF, 2021).

Finally, it must be acknowledged that digital security is fundamental to privacy protection. To promote greater uptake and potential for data sharing, governments should continue elaborating digital security policies.

Commercial sensitivity

Drivers and operators of commercial fleets may be reluctant to share telematics data for fear that, in a highly competitive market, commercially sensitive information could be mined from the data. Lifting this barrier could require awareness-raising messages about AI's use in traffic safety and transparency about who precisely will have access to the data. Telematics companies themselves fear for their intellectual property, worrying that their algorithms could be reverse engineered, for instance. Ownership of the intellectual property represented by the data exchanged in the SRTI ecosystem remains an open issue, several experts said.

Partners in any data exchange may fear that their data are ultimately misused, that the other party is not following the licence terms limiting use of the data. Among innovations to enforce data licences, the International Data Spaces (IDS) initiative aims at data sovereignty and interoperability. By specifying data use constraints, it defines the terms and conditions for the data economy. By describing an open software architecture and publishing open source software codes, it ensures maximum adoption. The IDS Association has more than 110 members from more than 20 countries (IDSA, 2020).

Open government

Many countries adopted legislation allowing public access to government-held data. They are known as freedom of information acts, access to information acts or open records legislation and apply to any data set that a government owns. Such open-ended access to data containing private citizen information can represent a major hurdle to data use for policy making.

Through a freedom of information act request, the New York City Taxi & Limousine Commission's entire trip records data set became public, potentially supporting any number of research projects related to transport planning. But, although anonymised, this data set could be matched with auxiliary data sets to reveal the likely religion of particular cab drivers, which rides celebrities took and the likely identity of individuals frequenting strip clubs. This shows that anonymisation is not always meaningfully protective or may not be technically viable in ordinary circumstances (Accenture, 2016).

Freedom of information regulations create two kinds of disincentives or barriers to data use by government agencies:

- They are a disincentive to the use of data sets containing personal identifiable information, since open access to such information would make government liable for privacy breach.
- They would be a barrier to government purchases of data, since data suppliers may refuse to let their data sets become public.

Both issues could be somewhat addressed by procurement of dashboards or query-based "safe answer" solutions whereby government procures only insights and neither owns nor hosts any data. In the Data for Road Safety initiative (Box 7), the DFRS ecosystem for SRTI exchange includes a multiparty agreement barring governments from turning vehicle-generated raw data into open data (Data Task Force, 2020).

Accurate risk prediction and guidance for proactive road network safety management

AI models could outperform traditional risk prediction models, coping with more input variables and discovering unexpected interactions between variables that enhance their predictive power. Will the right conditions be in place for this prospect to be realised? Will AI models be sufficiently transparent to be trusted? Will they deliver actionable insights? The answers to these questions will determine AI models' acceptance and adoption by road authorities.

Real-time applications

Real-time risk prediction models provide frequent updates (every minute to every week, depending on the system specifications) on high-risk locations. These alerts result in traffic management responses such as lowering of dynamic speed limits, a heightened or more visible police presence, treatment of a slippery road surface or reinforced surveillance by a control centre. On the other hand, non-real-time models can deliver a map of the permanent high-risk locations a road authority must investigate as part of a road safety audit.

Slippery road alerts are one example of real-time risk prediction. Other examples are the pilot projects described in Box 8, part of the US Department of Transport's Safety Data Initiative (SDI).

Volvo shares a continuous stream of slippery roads alerts and hazard warnings from its connected vehicles in Europe and North America through its Connected Safety function. Customers include the Swedish Road Authority and an AI-based traffic management company called Waycare. Other OEMs also share and monetise data, e.g. on hard braking events; no OEM is known for not doing so.

The rest of this report explores the questions of accuracy and biases, explainability and acceptance applicable to all AI-based risk prediction models, whether they run in real time or not.

Accuracy and biases

Data-driven approaches are highly dependent on domain knowledge, i.e. an understanding of the particularities of an industrial sector or other area of activity. For traffic safety, domain knowledge can describe the combination of system failures that leads to a crash. The understanding of such failures helps in selecting relevant data sets and the most appropriate algorithms (ITF, 2021).

AI benefits from domain knowledge, such as an understanding of crash contributing factors, in the selection of input variables. The deepest possible understanding of the root causes of crashes is what in-depth crash investigation programs seek to deliver.

Box 8. US Safety Data Initiative

Led by the US Department of Transport (USDOT) since 2017, the Safety Data Initiative takes advantage of the advent of big data and powerful analytics to provide actionable road safety insights. It involves a number of federal agencies, US states, local governments, private sector bodies and non-governmental organisations. Two pilot projects under the SDI illustrate the value of AI in crash risk prediction and mapping:

- The USDOT's Volpe National Transportation Systems Center employed machine learning techniques to estimate crash risk by combining crowdsourced crash data from Waze with crash history. This helped the **state of Tennessee** predict crash risk in one-hour time blocks on a one-square-mile grid. The Tennessee Highway Patrol used the model's output in prioritizing patrol locations.
- The **University of Central Florida** delivered real-time crash risk visualisations, with predictive analytics and real-time traffic safety diagnostics, through a tool using AI to suggest real-time interventions and long-term countermeasures to decision makers and operators and to inform the public of ZIP code-level safety conditions.

In 2020, the USDOT awarded over USD 3 million in funding to SDI projects, including the following:

- The **North Carolina Department of Transportation** will develop an AI tool for automated analysis of existing video log data that would extract roadside hazards – e.g. trees, embankments, steep slopes – on all rural roads in the state to help identify roadway segments in need of infrastructure safety improvements.
- The **Massachusetts Department of Transportation** will expand an existing crash data portal to help transport practitioners identify higher-risk roadways and risk factors so as to target roadway safety improvements and develop publicly available analytic tools and data visualizations.

The next step for the SDI is to expand the range of data sources supplied by the private sector. Data types will include satellite and street view imagery, on-board vehicle sensor data, road user movements and land use attributes.

Source: USDOT Volpe (2019), USDOT (2020).

Biases in Artificial Intelligence

Road authorities' adoption of data-driven approaches arguably increases objectivity, equity and fairness. Hence AI could make tasks such as infrastructure maintenance and network safety management more effective. Used inappropriately, however, AI technology could produce misleading results. AI algorithms are not generally biased, but they can unintentionally perpetuate biases if the source data are biased (ITF, 2021).

Excessive reliance on motor vehicle fleet data may result in biases, such as analysts neglecting road safety problems affecting people walking and cycling. The solution is to collect data from a wide range of devices, including shared bikes, personal bikes, e-scooters, phones and wearables. For instance, connected equipment such as motion sensing bike lights would allow collection of data from personal bicycles (ITF, 2019). Wearables would permit collection of pedestrian exposure, crash and fall data. It should also be recalled, however, that only the more tech-savvy members of the population use such devices, leading to further biases that need to be corrected.

Computer vision is another technology with biases that must be identified and removed. There may be racial and gender biases in AI, either accidental or due to AI reproducing biases already present in society.

Lack of data on crashes and road safety interventions

Whether AI-based or using more traditional statistical methods, crash risk prediction modelling takes more than just road and traffic attributes as input data. It also requires accurate crash data and a record of road safety countermeasures and when they were deployed. Crash data are often under-reported and inaccurate. The use of incomplete crash data could lead to AI perpetuating existing biases to the detriment of neglected areas or populations. Thus collection of good quality crash data is of the utmost importance.

Solutions to improve crash data completeness and accuracy were discussed in ITF (2019). Improving the quality of police and health records is important, as is occasional or permanent linkage of data from the transport and health spheres. Self-reported traffic injury surveys can also complement other data sets. Governments should use such solutions in ensuring the completeness and precision of their crash data – especially on crashes resulting in serious and fatal injuries – and publish this information as micro-level open data.

One difficulty in setting up risk prediction tools is that high-risk locations often receive safety treatment investments soon after diagnostic data reveal the problem. An accurate log of road safety countermeasures is needed, including indications of policing, targeted education campaigns, engineering solutions, etc. This is particularly challenging.

Calibration/training

Model calibration is not a one-off effort but an iterative process. Even in the case of non-real-time prediction tools, model developers should continually collect new data and assess whether the prediction power holds or could be improved through new model training. Behaviours differ in different parts of the world. Will explanatory factors found in one region apply to another region? Developers of risk prediction tools should constantly seek to verify their predictions' accuracy against recent and local data, and retrain their models when needed. As the UNECE (2020) ML project report emphasised, ML must be maintained to improve as it continues to learn (Figure 4).

Models may not be transferable between regions due to factors including behaviour differences and inconsistent input data. Training of data scientists around the world needs to take these factors into account.

Ideally, models should be update and retrained frequently and locally. Under the Safe System approach, the road safety community focuses on predicting and preventing fatal and serious crashes, which represent a minority of crashes and may not follow the same pattern as other crashes. Training a model requires a large number of events, yet fatal and serious crashes are relatively infrequent. Hence the spatial and temporal granularity of risk prediction modelling is limited. This simple statistical barrier is one of the most challenging technical aspects of risk prediction modelling.

Like weather models, risk prediction models may disagree, in which case taking the average output across several models can help. Each individual data set or algorithm has shortcomings, so it is usually better to use several. To prevent overfitting, Amazon Web Services proposes tools that develop several rough AI models and then take an average of the result of, say, 100 such models.

Performance

Will AI outperform classic models? With a given set of input variables, ML rarely outperforms traditional modelling by an extent that justifies the added complexity and lower transparency. However, ML can cope with more input variables and could significantly push the boundaries of crash risk prediction.

There is a risk of misusing AI models if trust is placed blindly in their abilities. Governments should commission robust evaluations of the benefits delivered by new tools and adjust trust levels accordingly to safeguard against assumed accomplishments. At some locations, models will make predictions that are far off recent crash numbers. An example is a location where several fatal crashes occurred within three years but a crash prediction model finds it to be perfectly safe. Without the most robust benchmarking of model performance, road authorities would lose trust in predictive safety.

Conversely, not using AI models due to lack of trust represents a missed opportunity. The road safety community should communicate clearly on new techniques' performance. To this end, governments, universities, industry and other stakeholders should organise or support competitions to benchmark crash prediction models' performance. In such events, teams of researchers compete to develop the best prediction model using the same input data and prediction scope. The HERE data marketplace organises AI competitions powered by the road traffic data available on its platform.[7] Competitions or benchmark efforts should always verify whether a model performs better than a simple extrapolation from observed crash numbers.

Explainable models for diagnostic and guidance

The strength of AI in crash prediction is its discovery of patterns the human eye would have missed, especially when combining many seemingly unrelated variables. But surprising and counterintuitive patterns call for sanity checks. A consensus emerged from the roundtable in favour of human supervision by data scientists, statisticians and road safety experts.

Experts noted that both AI and classic models sometimes associate crash risk "wrongly" with variables that may not pose any crash risk themselves but are influenced by missing data that would best represent the root causes of crash risk.

The conceptual framework, methodology and protocol for AI must be carefully defined to deliver meaningful results. AI can only identify the root causes of crash risk if it recognises the right variables of problematic patterns. These data are crucial to predicting possible causes of crash risk (e.g. behaviour, road design). AI has many other areas of application besides traffic safety, so it should be possible to transfer knowledge from other sectors.

An inadequate choice of input variables, together with a rapid interpretation with insufficient domain knowledge, could have dangerous consequences. For instance, due to confounding factors, a model could identify urban congestion as the root cause of crashes and recommend road widening as the cure, a policy known to in fact cause more motor vehicle traffic and greater crash risk over the whole city. This is an example of how AI could mislead authorities if used blindly. The problem often lies in hidden correlations between variables and the opportunistic choice of variables that do not capture root causes. In short, having some data is not enough: the devil is in the details.

Government officials are unlikely to be trained to detect weaknesses in predictive risk analysis and may not question results, especially if the conclusion goes in a politically practical way. Domain experts need to deliver sanity checks based on decades of established knowledge.

The most appreciated risk prediction models are those that explain why a given location is flagged as risky and suggest how to resolve the problem. Traditional modelling – with techniques such as regressions and generalised linear modelling – is considered stronger than AI in this area. But that is about to change as explainable AI models become available.

Explainability/interpretability

Although AI will probably do a lot to help identify dangerous locations, the lack of explainability that often characterises AI techniques could impede their acceptance and use. Modellers often hesitate to use an artificial neural network (ANN), a popular ML method, because it is hard to explain the results and how they were arrived at. It is indeed difficult to determine whether a trained ANN has learned intuitively reasonable relationships or ones that are spurious, inexplicable or otherwise undesirable. Choice modellers thus often find it hard to trust even an ANN whose predictive performance is strong (Alwosheel, van Cranenburgh and Chorus, 2019).

Many argue that ML should have no role in life and death decisions. Interpretable models, transparently relating source information with model outcome, seem to be the preferred choice when it comes to such decisions. Black box machine learning models are nevertheless used for high-stakes decision making throughout society, in healthcare, criminal justice and other domains. Some methods for explaining black box models to make them more transparent have been developed, yet for Rudin (2019), trying to explain black boxes, rather than creating models that are interpretable in the first place, is likely to perpetuate bad practice and could cause great harm to society. Rudin recommends instead designing models that are inherently interpretable.

There are ways to help derive meaning and explanations from AI models, or at least to diagnose them and check whether they have learned intuitive relations. Alwosheel, van Cranenburgh and Chorus (2019) propose a method involving synthesising prototypical examples that expose the fundamental relationships the model has learned. An analyst can then evaluate these to see whether they make sense and are desirable or not. In the context of crash risk prediction, a prototypical example might be the introduction of a lower speed limit.

Table 2. Strengths and weaknesses of classic and AI models

Strengths and weaknesses	Classic models	AI models
Expose coefficients	Yes	No
Expose the decision process for the inclusion or exclusion of variables	Yes	No
Allow for interrogation	Yes	Rarely
Find the best fit	Rarely	Yes
Allow for more input variables and interaction between input variables	Rarely	Yes
Could assign effects “wrongly” to variables that are collinear or co/inter-dependent with missing variables	Yes	Yes

Source: Based on Agilysis (2020).

Experts noted that explainable AI, or XAI, is becoming more popular and accessible. It bridges the gap between the most accurate (but obscure) AI models and classic models that are more easily interpreted.

Unless AI models become explainable, they are unlikely to be trusted to assist in decision making. Considering the huge benefits these new techniques could offer, however, it is important for society to persist in producing such models, scrutinising them in detail and learning their limitations, rather than simply sticking with familiar techniques because of the challenges AI poses.

Acceptance and adoption by road authorities

Despite the growing emphasis on tackling risky behaviours and adopting a proactive approach to road safety, most elected officials continue giving priority to reactive treatment of crash hotspots. They have little faith in black box solutions for risk mapping. To foster a new vision, road safety professionals should build trust in risk mapping and risk prediction tools. A good way to build trust in a model would be to identify crash risks that experts did not foresee but later acknowledged as real. Models can also identify crash risks that residents have identified but that no crash data had yet revealed. Attractive, legible and colourful maps can work well for policy makers and constituents in such efforts.

AI needs to show a substantial improvement over statistical models, and a significant level of interpretability, before governments will widely use it for road safety applications. Some experts consider established techniques, such as crash modification factors or functions (Box 5), to be both robust and geographically transposable enough to negate the need for AI. Others argue that a given action or measure can have different safety effects in different countries (ITF, 2012).

Figure 4. Acceptance and facilitation of machine learning for official statistics

Source: UNECE (2020).

Overcoming obstacles to acceptance

Roundtable participants observed that governments do not always follow the conclusions of road safety audits and star-rating programmes. AI is unlikely to lead to increased commitment to road safety. Road authorities will be reluctant to invest in AI if unsure of the return on this investment. The road safety community needs to assess the number of casualties a proactive approach will avoid and reassure decision makers of the value of their investment. For example, iRAP computes the predicted economic benefit of a road safety intervention. Integration of such business case information into operational road management systems is the next step forward. It would supplement one-off political decisions with a more reliable maintenance and investment programme.

Will authorities have the resources to act on these new crash predictions? Can they redouble their efforts, fixing not only crash hotspots, as they used, to but also predicted hotspots? Awareness of proactive safety management tools is a necessary but not sufficient condition for successful implementation of infrastructure-oriented safety measures. Procedures that are not obligatory are easily skipped to save time and money. Some road authorities reject risk prediction modelling for fear of the legal liability that could result from knowing the location of high risk road segments and not acting in time. ITF (2015) therefore recommends making road infrastructure safety management compulsory. The European Commission, for instance, requires roads of the Trans-European network to undergo regular road safety inspections.

Some authorities refuse to share even aggregated crash data on data protection grounds. One case mentioned at the roundtable involved an authority maintaining that a collision's exact time and place could not be disclosed, to preserve privacy. Participants concluded that policy makers should clarify regulatory frameworks on how to anonymise and aggregate data.

Software availability is not seen as the main challenge in developing AI solutions for crash risk prediction. Tools already exist in statistical software and cloud services for the calibration and use of AI models. These may not make AI models as accessible as spreadsheet models, but go some way towards closing the gap.

Skills

Predicting road injury risk requires both quality data and professional skills. Data scientists should join forces with subject matter experts for appropriate interpretation of often complex and incomplete data sets. Data without context rarely provide much information.

Smaller government agencies have limited resources, and their staffs are spread thin across many public issues. That is why local governments often rely on university students to analyse data. Government needs a simple and intuitive way to access essential data sets. VIA in the Netherlands has developed standardised software to help municipalities get access to speeding data.

Because governments' analytical capabilities are limited, civil society's capacity to mine relevant data and deliver insights should not be neglected. Opening up government data, including crash data and a log of countermeasures, can deliver tremendous benefits by facilitating civil society research efforts.

Notes

1 In a previous ITF workshop, industry experts said telematics in professional vehicle fleets were easier to use in countries with lower levels of privacy protection regulations. For instance, the capture of video images inside truck cabins is developing in Africa. See http://www.itf-oecd.org/sites/default/files/docs/new-directions-data-driven-transport-safety_0.pdf.

2 See https://www.lytx.com/en-us/

3 See https://www.who.int/violence_injury_prevention/road_traffic/12GlobalRoadSafetyTargets.pdf

4 See https://www.mobileye.com/uk/fleets/blog/barcelona-introduces-mobileye-city-streets/

5 See https://covesa.github.io/vehicle_signal_specification/introduction/

6 See https://sensoris.org/objectives/

7 See https://www.globenewswire.com/news-release/2020/06/04/2043438/0/en/Traffic4cast-competition-calls-on-AI-community-to-better-predict-urban-traffic-flows.html

References

AASHTO (2010), Highway safety manual. American Association of State Highway and Transportation Officials: Washington, DC, USA.

Accenture (2016), The Ethics of Data Sharing: A guide to best practices and governance, https://www.accenture.com/_acnmedia/pdf-35/accenture-the-ethics-of-data-sharing.pdf

Agilysis (2020), Road Danger Prediction - Classic models, AI models and Data Challenges, Presentation delivered on 11 February 2011, https://www.itf-oecd.org/road-danger-prediction-classic-models-ai-models-data-challenges-agilysis

Aldred, R. (2018), "Inequalities in self-report road injury risk in Britain: A new analysis of National Travel Survey data, focusing on pedestrian injuries", Journal of Transport and Health, Vol. 9, pp. 96-104, https://doi.org/10.1016/j.jth.2018.03.006.

Alwosheel, A., S. van Cranenburgh and C. G. Chorus, (2019), 'Computer says no' is not enough: Using prototypical examples to diagnose artificial neural networks for discrete choice analysis, Journal of Choice Modelling, https://doi.org/10.1016/j.jocm.2019.100186.

Belin M.-Å., P. Tillgren and E. Vedung (2012), Vision Zero – a road safety policy innovation, International Journal of Injury Control and Safety Promotion, 19:2, 171-179, DOI: 10.1080/17457300.2011.635213

Chang, A., N. Saunier and A. Laureshyn (2017), "Proactive Methods for Road Safety Analysis", SAE Technical Paper, WP-0005, SAE International.

CMF Clearinghouse (2021), Crash Modification Factors Clearinghouse website, www.cmfclearinghouse.org (accessed 18 November 2021).

Data Task Force (2020), Final report and recommendations, October, https://www.dataforroadsafety.eu/images/Documenten/DTF-REPORT-OCTOBER-2020-021020.pdf.

European Commission (2013), Commission Delegated Regulation (EU) No 886/2013, European Commission, Brussels, https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32013R0886&from=EN.

Fawcett, L. et al. (2017), "A novel Bayesian hierarchical model for road safety hotspot prediction", Accident Analysis and Prevention, Vol. 99/A, pp. 262-271, https://doi.org/10.1016/j.aap.2016.11.021.

Hydén, C. (1987), The Development of a Method for Traffic Safety Evaluation: the Swedish Traffic Conflict Technique, Lund University, retrieved from UC Berkeley Transportation Library (Accession No. 00422429).

IDSA (2020), Implementing the European Strategy on Data Role of the International Data Spaces (IDS), Position Paper, v1.0, April, https://internationaldataspaces.org/wp-content/uploads/IDSA-Position-Paper-Implementing-European-Data-Strategy-Role-of-IDS1.pdf.

iRAP (n.d.), Ai-RAP and Automated Coding for iRAP, https://resources.irap.org/General/Ai-RAP_Intelligent_Collection_of_RAP_Data_Worldwide.pdf.

ITF (2021), "Data-driven Transport Infrastructure Maintenance", International Transport Forum Policy Papers, No. 95, OECD Publishing, Paris, https://doi.org/10.1787/a1cc71cc-en.

ITF (2019), "New Directions for Data-Driven Transport Safety", *International Transport Forum Policy Papers*, No. 83, OECD Publishing, Paris, https://doi.org/10.1787/2b2f71bc-en.

ITF (2018), Speed and Crash Risk, OECD Publishing, Paris, www.itf-oecd.org/speed-crash-risk.

ITF (2016a), *Zero Road Deaths and Serious Injuries: Leading a Paradigm Shift to a Safe System*, OECD Publishing, Paris, https://doi.org/10.1787/9789282108055-en

ITF (2015), Road Infrastructure Safety Management, OECD Publishing, Paris.

ITF (2012), Sharing Road Safety: Developing an International Framework for Crash Modification Functions, ITF Research Reports, OECD Publishing, Paris, https://doi.org/10.1787/9789282103760-en.

ITF (2011), Reporting on Serious Road Traffic Casualties, IRTAD research report, International Transport Forum, OECD Publishing, Paris, https://www.itf-oecd.org/reporting-serious-road-trafficcasualties .

OECD (2019a), Recommendation of the Council on Artificial Intelligence, OECD/LEGAL/0449, https://legalinstruments.oecd.org/api/print?ids=648&lang=en.

OECD (2019b), Artificial Intelligence in Society, OECD Publishing, Paris, https://doi.org/10.1787/eedfee77-en.

Rudin, C. (2019), Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead, . Nat Mach Intell 1, 206–215. https://doi.org/10.1038/s42256-019-0048-x.

SuM4All (2021), Sustainable Mobility: Policy Making for Data Sharing, Sustainable Mobility for All, Washington DC, License: Creative Commons Attribution CC BY 3.0, https://www.wbcsd.org/contentwbc/download/11663/176259/1 .

UN General Assembly (2020), Resolution 74/299 "Improving global road safety"

USDOT (2020), U.S. Department of Transportation Announces Over $3 Million for Roadway Safety Tools, Newsroom, June 10, https://www.transportation.gov/briefing-room/us-department-transportation-announces-over-3-million-roadway-safety-tools

UNECE (2020), UNECE – HLG-MOS Machine Learning Project, Project Report, https://statswiki.unece.org/display/ML/Machine+Learning+Project+Report.

Volpe (2019), Safety Applications of Crowdsourced Data, Tennessee Highway Patrol Case Study Report, https://www.transportation.gov/sites/dot.gov/files/docs/subdoc/836/2019-tn-highway-patrol-case-study-brief.pdf

WHO (2015), Global Status Report on Road Safety 2015, World Health Organization, Geneva, https://www.who.int/violence_injury_prevention/road_safety_status/2015/Executive_summary_GSRRS2015.pdf

WHO (2018), Global Status Report on Road Safety 2018, World Health Organization, Geneva, www.who.int/violence_injury_prevention/road_safety_status/2018/en/ (accessed on 20 March 2019).

Wijnen, W. et al. (2017), Crash cost estimates for European countries, Deliverable 3.2 of the H2020 project SafetyCube, https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e5b1e92ba3&appId=PPGMS

Annex A. List of Roundtable participants

George YANNIS (Chair), Professor, National Technical University of Athens (NTUA), Greece

Mohamed ABTEL-ATY, Trustee Chair, Dept. of Civil, Environmental & Construction Engineering, University of Central Florida (UCF), USA

Johan AMORUSO WENNERBY, Volvo cars, Sweden

Glenn ATKINSON, Vice President, Product Safety, Geotab, Canada

Mohamadreza “Mohamad” BANIHASHEMI, Transportation Specialist, Federal Highway Administration (FHWA), USA

Matts-Ake BELIN, Director Vision Zero Academy, Trafikverket (Swedish Transport Administration), Sweden

Ana Isabel BLANCO BERGARECHE, Deputy Director for Mobility, DGT (Transport Ministry), Spain

James BRADFORD, Global Technical Director, iRAP

Dave BRAUNSTEIN, President, Together for Safer Roads (TSR), USA

Alina BURLACU, Senior Transport Specialist, World Bank

Annie CHANG, Director of Safety Programs, Lime, USA

Sam CHAPMAN, CIO, Co-Founder, Director, The Floow, UK

Erik DONKERS, Senior Road Safety Expert, VIA Traffic Solutions Software, Netherlands

Véronique FEYPELL, Manager, Road Safety, ITF / OECD

Carlo GIRO, Manager, Research and Innovation, International Road Transport Union (IRU)

Andrew HOFFMAN, Director, Strategy & Innovation, Together for Safer Roads (TSR), USA

Pedro HOMEM DE GOUVEIA, Senior Policy & Project Manager, Coordinator for Road Safety & Transport Security, POLIS Network

Martin IVARSON, Strategist, Volvo cars, Sweden

Jonas JANSSON, (Dr) Research Director, National Road and Transport Research Institute (VTI), Sweden

Lina KONSTANTINOPOULOU, Secretary General, EuroRAP

Jens LANDVOGT, Senior Principal Business Solution Architect, Here Technologies, Germany

Vincent LEDOUX, Projects Director, Road Safety and New Technology, Cerema, France

Rob MCINERNEY, CEO, iRAP

Ashley MILLER, Senior Solutions Architect, Amazon Web Services, USA

John MILTON, Chair of TC 3.1 Road Safety Committee, PIARC, USA

Luis MIRANDA-MORENO, Associate Professor, McGill University, Canada

Sudeshna MITRA, Transport Specialist, World Bank

Ahmed NASR, Data Monetization and Marketplace, Here Technologies, Belgium

Monica OLYSLAGERS, Safe Cities and Innovation Specialist, iRAP

Richard OWEN, CEO, Agilysis, UK

Kaan OZBAY, Director, C2SMART Center, New York University (NYU), USA

Eleornora PAPADIMITRIOU, Assistant Professor, Delft University of Technology, Netherlands

Dimitrios PAPAIOANNOU, Modeller-Analyst, ITF / OECD

Stephen PERKINS, Head of Research and Policy Analysis, ITF / OECD

Rachele POGGI, Statistician, ITF / OECD

Veronica Ines RAFFO, Senior Infrastructure Specialist, World Bank

Himanshu RAJ, Officer, Sustainable Mobility, Local Governments for Sustainability (ICLEI), Germany

Alexandre SANTACREU, Policy Analyst, Road Safety, ITF / OECD

Barnaby SIMKIN, Manager, Regulatory Affairs, NVIDIA, Germany

Craig SMITH, Data Scientist, Agilysis, UK

Per-Olof SVENSK, Project Manager, Development Manager and Advisor, Trafikverket (Swedish Transport Administration), Sweden

Bas TURPIJN, Traffic Engineer, NDW (National Dataportal Road Traffic), Netherlands

George URSACHI, Senior Research Analyst, Agilysis, UK

Luis Fernando VILLAÇA MEYER, Director of Operations, Cordial Institute, Brazil

Adam WEISS, Strategy & Regulation Manager, MobilEye, Israel

Ivo WENGRAF, Research & Data Manager, RAC Foundation, UK

Björn ZACHRISSON, Business Development Manager, NIRA Dynamics AB, Sweden

Apostolos ZIAKOPOULOS, Research Associate, National Technical University of Athens (NTUA), Greece

Artificial Intelligence in Proactive Road Infrastructure Safety Management

This report examines and determines the most relevant cases for artificial intelligence (AI) use in a transport planning context for crash prevention on an entire road network. It explores the possibility of using computer vision to acquire relevant information and the capability of computer models to map high-risk locations. It offers recommendations to stakeholders on the development and appropriate use of life-saving AI solutions.

International Transport Forum
2 rue André Pascal
F-75775 Paris Cedex 16
+33 (0)1 73 31 25 00
contact@itf-oecd.org
www.itf-oecd.org

Published: 12/2021 | Photo credit: watophotography / Shutterstock