Reasonable vehicles rule the road

To get a driving license, humans learn the rules; so should autonomous vehicles

White Paper 2020

Oxford Semantic Technologies

RDFOX

Oxford Semantic Technologies

Table of Contents

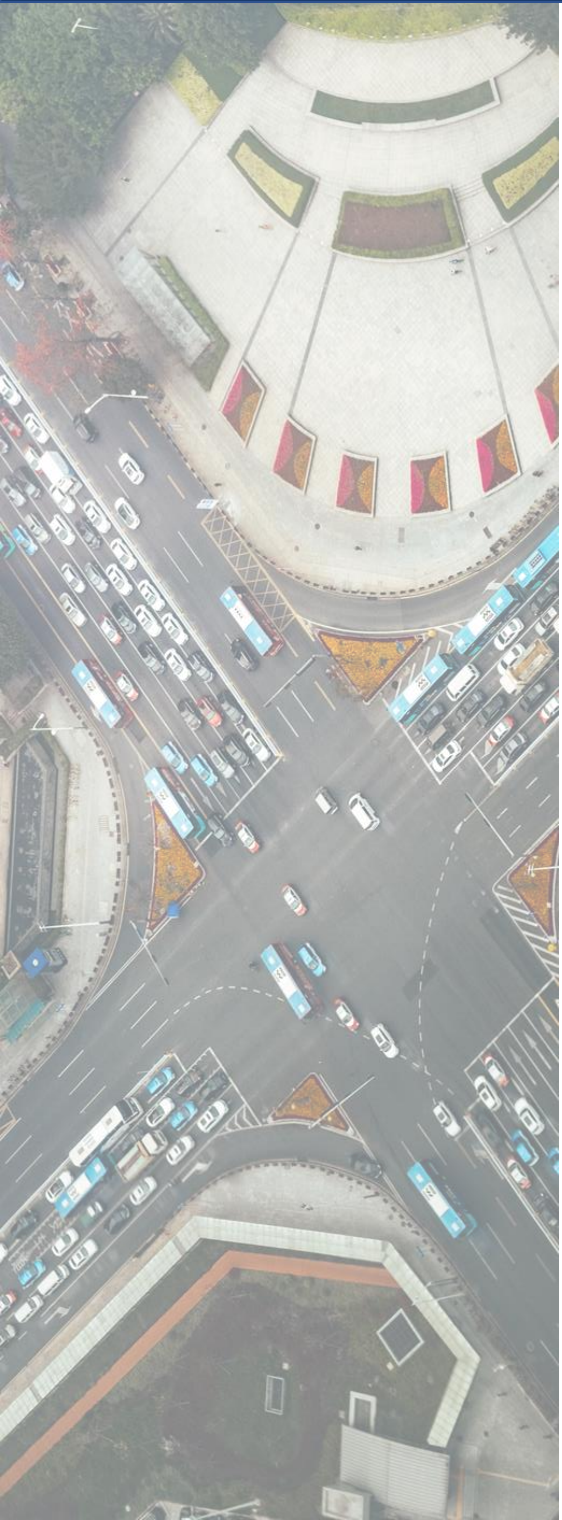

This white paper is inspired by award winning research undertaken by Haonan Qiu,[1] and colleagues at BMW and the University of Ulm. The research is funded by BMW.

The views expressed herein are those of Oxford Semantic Technologies Limited, the developers of RDFox. Oxford Semantic Technologies' mission is to bring cutting-edge research in semantic web technologies to industry.

RDFox is the first semantic reasoning engine which allows flexible incremental addition and retraction of data, and incremental reasoning. It has been optimised for speed and correctness. It is unmatched in power, is an in-memory solution and is mathematically validated at the University of Oxford.

Introduction

It is estimated that the automotive AI Market will reach more than $10.5 billion by 2025[3]

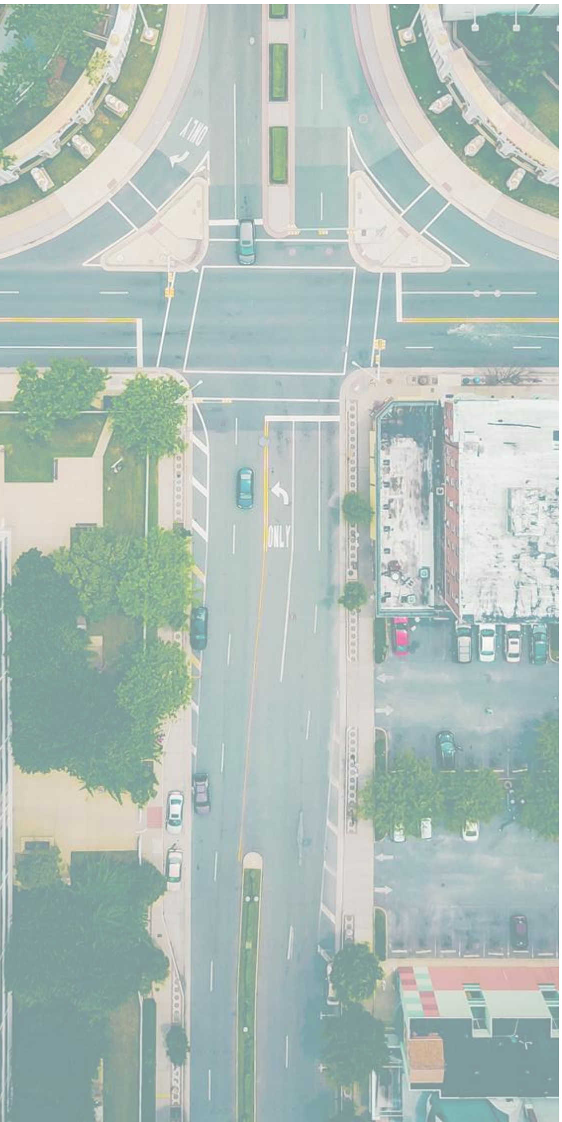

Autonomy

Autonomy is a central pillar of the fourth industrial revolution. However, fully autonomous vehicles remain an area of development. Autonomous vehicles operate in a high-risk environment - if something goes wrong, lives could be at risk. To function effectively, the system needs human-like reasoning and decision-making capabilities.

There are numerous benefits associated with the advent of autonomous vehicles, including increasing the safety on the roads, real-time route optimisation, increased lane capacity, reduced energy consumption and time efficiency for the users.[2] Research and development into autonomous vehicles has boomed in recent years, with most major automotive manufacturers exploring their development. It is estimated that the automotive AI Market will reach more than $10.5 billion by 2025[3] and companies are determined to get their piece of the pie.

Development of autonomous vehicles

The BMW group are amongst leading manufacturers developing AI for autonomous vehicles. They have invested resources into R&D across their operations, including the opening of a 248,000 square feet test site near Munich.[4] They are actively participating in ground-breaking research, and sponsor many PhD projects.[5]

Their research has unveiled a significant technological advancement and introduced a novel application of RDFox, a semantic reasoning engine developed by Oxford Semantic Technologies.[6] This application uses human-like reasoning and sophisticated data processing to combat the challenges faced by the industry. Autonomous vehicle development represents a suitable application of RDFox, as it is a data intensive setting, where speed, performance and reasoning are critical.

The white paper

This white paper will outline the challenges in autonomous vehicle development, research into using ontologies and semantic reasoning, and the benefits of this approach for the industry. RDFox's attributes make it uniquely able to perform in this context.

Challenges for the development of autonomous vehicles

There are many challenges associated with the advent of autonomous vehicles into mainstream road systems, for example: safety, public perception and technological demands.

There are also many different approaches being developed comprising many different, discrete technologies.

Sensors and high definition (HD) maps

Effective autonomy requires sensor information combined with high definition (HD) maps.[6] Manufacturers such as Google, HERE, TomTom, Baidu, BMW and Toyota use HD maps within their research and development of autonomous vehicles.[7]

To navigate and make decisions in real-time, autonomous vehicles must be location-aware, understand their surroundings and be able to anticipate action. The environment the vehicle operates in is highly dynamic and the system must deal with large volumes of streaming data. As such, the speed of processing is directly related to the speed autonomous vehicles will be able to safely travel, thus exceptionally fast data integration and processing is required.

A widespread approaches in development, Simultaneous Localisation and Mapping (SLAM), situates the vehicle and navigates roads by leveraging sensors and HD maps with detailed information; for example, lane data, geometry, traffic signs, road surface, location of objects and speed limits.[8] This provides a multi-leveled, combined data approach. However, until now HD maps have been problematic, resulting in some manufacturers moving towards 'map-less' methods, due to the absence of technology which could cope with the demands.[9]

Creating an HD map

Effective, onboard HD maps require the resources for map generation, significant storage space, power to process and integrate large volumes of data which exists in various formats and the ability to maintain the map in real-time.[10]

Lack of standardisation makes the sharing of maps and map information difficult. The various formats may be incompatible and therefore pose an integration challenge with the larger system architecture. As such, there are concerns over the scalability and interoperability of map processing solutions.[11] Further challenges include the ability to support efficient onboard/offboard development.

Challenges for the development of autonomous vehicles

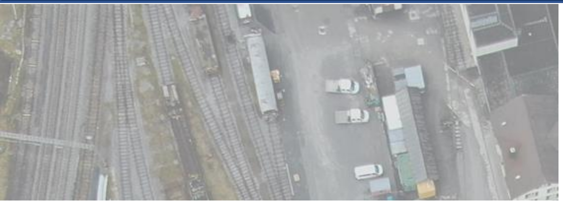

"In an autonomous car, we have to factor in cameras, radar, sonar, GPS and LIDAR –components as essential to this new way of driving as pistons, rings and engine blocks. Cameras will generate 20–60 MB/s, radar upwards of 10 KB/s, sonar 10–100 KB/s, GPS will run at 50 KB/s, and LIDAR will range between 10–70 MB/s. Run those numbers, and each autonomous vehicle will be generating approximately 4,000 GB –or 4 terabytes –of data a day." Brian Krzamich, CEO of Intel, 2016[16]

As a result of these challenges, many researchers and manufacturers have chosen to move away from maps, towards sensor-only methods and machine learning techniques, including MIT's Computer Science Artificial Intelligence Laboratory (MIT CSAIL), Wayve, Apple and Tesla.[12]

If sensors and HD Maps are created effectively, they still face challenges, such as the difficulty of capturing the implicit information which helps human drivers make decisions. For example, a sensor won't know why a car in front has stopped or that it might intend to reverse and park until the car starts to reverse. Understanding the context of the objects identified by sensors is important and this requires meaning to be encoded.

Increasing the number of sensors to provide a complete image of the surroundings and improve clarity may not be a viable solution, as increasing the volume of incoming data exacerbates processing issues. The feasibility of storing and processing information must be considered. On their own, sensors miss the critical knowledge provided by the map.

Future improvements

There have been numerous suggestions for improving the viability of autonomous vehicles, including end-to-end machine learning solutions. However, many machine learning techniques have been unable to meet the challenges of computing time, scale and the ability to react to constantly changing environments.[13]

Proponents of the SLAM approach note the need for a software element which can combine incoming data with localisation and mapping, path planning and motion control.[14] However, recent developments indicate that using HD Maps and sensors effectively can be achieved using semantic reasoning and ontologies.[15]

Moving forward: ontologies and semantic reasoning

Ontologies

To overcome these challenges, research indicates a growing interest in ontology-based road environment modelling within the Intelligent Transportation Systems community.[17] Ontologies offer many benefits and applications for the development of autonomous vehicles; they have been used in representing road intersections,[18] creating driving decision-making systems,[19] and for providing meaning and context, both implicit and explicit, to sensor data, for example the relationships between perceived road objects.[20]

Various levels of ontologies can be integrated using semantic reasoning (rules) to create a “common high-level ontology resulting in a unified knowledge base”.[21] Using an ontological approach allows the road environment to be modelled and road situations to be reasoned over.

The necessary knowledge for decision-making and vehicle motion planning can be deduced from this unified knowledge base.[22] Ontologies provide a flexible architecture, as various lower-level maps can be used to feed information to the high-level ontology.

Semantic Reasoning

Researchers at BMW have suggested a novel approach for the progression of autonomous vehicles, which uses RDFox, a semantic reasoning engine, in its knowledge-spatial architecture.

Semantic reasoning is the “ability to make logical deductions from the information that is explicitly available”.[23] This approach models the road environment, provides a dynamic update mechanism and has semantic reasoning at its core. Reasoning allows information to be aggregated and integrated leading to knowledge of the environment and effective decision-making, e.g. changing lanes at an appropriate time so the car leaves the motorway.[24]

Moving forward: ontologies and semantic reasoning

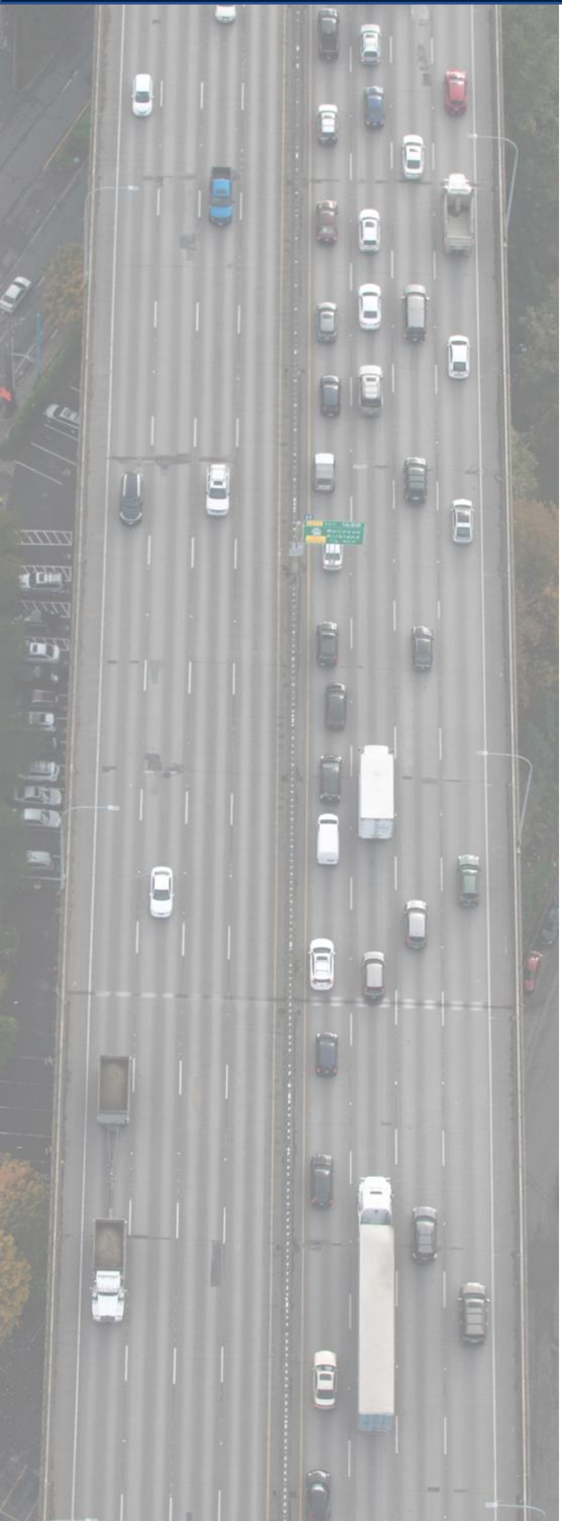

RDFox and semantic reasoning

RDFox provides a comprehensive software component which can be used to process and integrate data, in real-time, with constrained processing power. RDFox can deal with large, constant streams of data due to its incremental processing capabilities. This means that as new data comes in, only the calculations which are affected are updated.

Through rules RDFox can make real-time decisions resulting in motion control. By harnessing the power of semantic reasoning, logic can be integrated into the software. This means the rules and best practices associated with driving can be embedded within the system.

RDFox is suitable for this application because, like a (competent) human driver, it is a sophisticated reasoner, which can scale, and is optimised for fast, incremental reasoning. Speed of thought is critical in an autonomous vehicle, because a difference measured in milliseconds could be an incident on the motorway. In reality, the safe operating parameters of autonomous vehicles will determine whether and how they operate. This will determine, for example, maximum speeds, minimum distances between vehicles, and the minimum conditions that must be satisfied before manoeuvers are permitted, such as lane changes. The faster and more accurate the reasoning, the larger and more flexible the safe operating window, and the greater the benefits of autonomous transportation.

The in-memory solution and economical storage that RDFox offers also bolster its viability, as it provides the ability to hold information on-board the vehicle effectively, whilst also integrating various datatypes. Whilst RDFox has all of the capabilities synonymous with graph databases, its incremental semantic reasoning capability make it uniquely suitable to meet the demands of autonomous driving.

Benefits of semantic reasoning for autonomous vehicles

Semantic reasoning is the key to giving autonomous systems salient information and decision-making capabilities whilst reducing their dependency on big data feeds and huge processing power.

The contribution made by researchers at BMW demonstrates improvements in various areas including the ability to embed domain knowledge on safe driving practices, data volume issues, monitoring of sensors and explainability.

Safe driving practices

Rules and common-sense help humans plan and infer responsible driving behavior for situations on the road. A driver might slow down when driving through a residential neighbourhood on a sunny weekend because they expect more pedestrians and children playing with balls near the road. Semantic reasoning makes it possible to model these implicit concepts so that the appropriate decisions can be applied at scale without having to explicitly state each variation individually.

An autonomous vehicle must also be able to anticipate potential outcomes and respond in real-time. For example, if someone walks out into the road unexpectedly, the car must respond. By modelling domain expertise on safe driving practices into the system we can improve the safety element of autonomous vehicles.

Furthermore, of the mistakes made by humans driving, 98.3% are due to inattention or ineffective driving.[25] When such errors are combined with a deliberate reduction of the safety margin, through excessive speed or driving too close to the vehicle in front, the result is a serious and potentially fatal incident.[26] By removing the human element, which is responsible for lethally combining errors and violations, there will likely be a reduction in road incidents, which result in more than 1.2 million deaths and 50 million injuries a year, globally.[26]

Benefits of semantic reasoning for autonomous vehicles

Data volume

The volume and computation pressure of data processing is a huge concern for those developing methods using HD Maps. Researchers have found that using RDFox, they can store a proportion of the information in memory, and for a limited amount of time.[28] By using spatial reasoning, a 'spatial window' can be created. This defines a range that information will be stored for, using a specific distance around the car. Information pre-loads for areas soon to come into the spatial window. Once beyond the range, the information will be incrementally deleted from the system.

This limits the amount of data stored and processed at any given time, greatly improving processing efficiency. The map is incrementally updated whilst progressing along a route. This is all done onboard, signaling that effective onboard systems are possible.

Data integration

At present the various maps accessible by manufacturers are in incompatible formats due to the lack of standardisation. RDFox is well suited to solving this issue due to its ability to integrate data from various locations and formats, into a unified knowledge graph which can then be reasoned over.

By using RDFox the manufacturer can integrate the existing map knowledge base into their system, and build upon this, rather than starting from scratch. This represents a serious breakthrough in map development.

Sensors

Rules offer the solution to safety concerns over sensors. When information is incomplete because sensors are erroneous or missing, rules can be used to determine decision-making, and can rely on the wealth of other information incorporated into the system. For example, if a sensor fails, then the car will be able to aggregate the information from other sensors to ensure that safe driving continues.

Benefits of semantic reasoning for autonomous vehicles

Additionally, when sensor data is inconsistent, rules can be used to flag this to the user in real-time. For example, rules can prevent false emergency breaks by validating the consistency of the derived knowledge from the fused sensor and map data ontologies.[28]

Explainability

As safety is a huge concern for integrating autonomous vehicles into the mainstream road systems, being able to explain a vehicle's actions is of interest to the developers. At present, machine learning techniques are unable to provide explainability. Should an incident occur, or should the user be interested in why the autonomous vehicle took a specific path, they can't get an explanation from the system. With RDFox, the system can provide explainability. This would allow for accountability, and also provides transparent information to the Manufacturer and investigators.

Just as the black box, together with a learning culture, has helped transform safety in flight, transparency in autonomous vehicle control systems will help achieve the same for road transport. It is, in fact, astonishing to think that aviation is orders of magnitude safer than driving.

Looking forward

There is a growing body of research which uses ontologies and semantic reasoning in autonomous vehicles and other autonomous applications.

The novel approach put forward in the research conducted by Haonan Qiu and colleagues at BMW and the University of Ulm, explored the use of ontological approaches for dynamic map processing, and reasoning for decision-making and motion control.

The inclusion of RDFox in environmental modelling and knowledge-spatial architecture for the facilitation of autonomous vehicle decision-making opens new possibilities for applying human-like reasoning to this field.

The choice of RDFox was integral to this research, due to its unique ability to perform with speed and correctness in the acutely constrained and time-sensitive cockpit of the autonomous vehicle.

The results are encouraging, and suggest that the ontological approach using semantic reasoning could contribute to the strive for full autonomy in vehicles and, of course, other areas of robotics.

Contact Us

Oxford Semantic Technologies

Website: www.oxfordsemantic.tech

Email: info@oxfordsemantic.tech

Medium: https://medium.com/oxford-semantic-technologies

Twitter: @oxfordsemantic

Haonan Qiu
BMW Group
Haonan.Qiu@bmw.de

Adel Ayara
BMW Group
Adel.ayara@bmw.de

References

1. Qiu, Ayara, and Glimm (2020a) A Knowledge Architecture Layer for Map Data in Autonomous Vehicles. Institute of Artificial Intelligence, University of Ulm.; and, Qiu, Ayara and Glimm (2020b). Ontology-Based Processing of Dynamic Maps in Automated Driving. In KEOD.
2. Infiniti Research report: The top-3 trends that will drive the market for autonomous vehicles. Access: https://www.infinitiresearch.com/thoughts/top-3-trends-will-drive-market-autonomous-vehicles
3. Markets and Markets Press Release: Automotive Artificial Intelligence Market worth 10,573.3 Million USD by 2025. Access: https://www.marketsandmarkets.com/PressReleases/automotive-artificial-intelligence.asp
4. MIT Technology Review. Self-driving cars take the wheel. Access: https://www.technologyreview.com/2019/02/15/137381/self-driving-cars-take-the-wheel/
5. Qiu, Ayara, and Glimm (2020a); and, Qiu, Ayara and Glimm (2020b).
6. Oxford Semantic Technologies. Access: www.oxfordsemantic.tech
7. Wevolver: 2020 Autonomous Vehicles Technology Report. Access: https://www.wevolver.com/article/2020.autonomous.vehicle.technology.report
8. Ibid.
9. Qiu, Ayara and Glimm (2020b).
10. Wevolver: 2020 Autonomous Vehicles Technology Report. Qiu, Ayara, and Glimm (2020a); and, Qiu, Ayara and Glimm (2020b).
11. Wevolver: 2020 Autonomous Vehicles Technology Report.
12. Ibid
13. Ibid
14. Qiu, Ayara, and Glimm (2020a); and, Qiu, Ayara and Glimm (2020b)
15. Intel Newsroom. Data is the New Oil in the Future of Automated Driving. Access: https://newsroom.intel.com/editorials/krzanich-the-future-of-automated-driving/#gs.i9vpay
16. Multiple references:
    1. Katsumi and Fox, (2018). Ontologies for transportation research: A survey. *Transportation Research Part C: Emerging Technologies*, vol. 89, pp. 53 – 82.
    2. Chen and Kloul (2018). An ontology-based approach to generate the advanced driver assis- tance use cases of highway traffic. In *KEOD*, pp. 73–81.
    3. Bagschik, Menzel, and Maurer (2018). Ontology based scene creation for the development of automated vehicles. In *IEEE Intelligent Vehi- cles Symposium (IV)*, pp. 1813–1820.
17. Multiple references
    1. Regele (2008). Using ontology-based traffic models for more efficient decision making of autonomous vehicles. In *4th Int. IEEE Conf. on Autonomic and Autonomous Systems (ICAS'08),* pp. 94–99.
    2. Hu¨lsen, Zo¨llner, and Weiss (2011). Traffic intersection situation description ontology for advanced driver assistance. In *IEEE Intelli- gent Vehicles Symposium,* pp. 993–999.
    3. M. Hu¨lsen, J. M. Zo¨llner, N. Haeberlen, and C. Weiss, "Asynchronous real-time framework for knowledge-based intersection assistance," in *14th Int. IEEE Conference on Intelligent Transportation Systems (ITSC)*, 2011, pp. 1680–1685.
18. Qiu, Ayara and Glimm (2020a); and, Qiu, Ayara and Glimm (2020b).
19. Armand, Ibanez-Guzman, and Zinoune (2017). *Digital Maps for Driving Assistance Systems and Autonomous Driving*. Springer, pp. 201–244.
20. Qiu, Ayara, and Glimm (2020a); and, Qiu, Ayara and Glimm (2020b)
21. Belouaer, Bouzid, and Mouaddib. (2010). Ontology based spatial planning for human- robot interaction. In *17th Int. Symposium on Temporal Representation and Reasoning (TIME 2010)*, pages 103–110, Los Alamitos, CA, USA. IEEE Computer Society.
22. Oxford Semantic Technologies. Access: www.oxfordsemantic.tech/glossary
23. Qiu, Ayara and Glimm (2020a). A Knowledge Architecture Layer for Map Data in Autonomous Vehicles . Institute of Artificial Intelligence, University of Ulm
24. Venture Beat. Wave raises 20 million to give autonomous cars better ai brains. Access: https://venturebeat.com/2019/11/17/wayve-raises-20-million-to-give-autonomous-cars-better-ai-brains/
25. Parker and Stradling. Road Safety Research Report No.17. Access: https://trid.trb.org/view/689647
26. FIA Foundation. Road Safety. Access: https://www.fiafoundation.org/our-work/road-safety
27. Qiu, Ayara, and Glimm (2020a); and, Qiu, Ayara and Glimm (2020b)
28. Ibid.

Image References

Page 1 Photo by John Lockwood on Unsplash
Page 2 Photo by Robert Bye on Unsplash
Page 3 Photo by Tim Trad on Unsplash
Page 4 Photo by Matty Powell on Unsplash
Page 5 Photo by Nathan Queloz on Unsplash
Page 6 Photo by Daria Nepriakhina on Unsplash
Page 7 Photo by Patrick Perkins on Unsplash
Page 8 Photo by K. Mitch Hodge on Unsplash
Page 9 Photo by Nathan Queloz on Unsplash
Page 10 Photo by Lachlan Donald on Unsplash
Page 11 Photo by John Salzarulo on Unsplash
Page 12 Photo by Javier Haro on Unsplash