Explainable AI Reasoning for Urban Traffic Understanding

Domenico Redavid[a,*,1], **Eleonora Bernasconi**[b,**,2] **and Stefano Ferilli**[b,***,3]

[a]Economics and Finance Department, University of Bari, Italy
[b]Computer Science Department, University of Bari, Italy
ORCID (Domenico Redavid): https://orcid.org/0000-0003-2196-7598, ORCID (Eleonora Bernasconi): https://orcid.org/0000-0003-3142-3084, ORCID (Stefano Ferilli): https://orcid.org/0000-0003-1118-0601

Abstract. Efficient traffic analysis and management are essential for supporting various stakeholders, enabling the automation of real-time monitoring, anomaly detection, and strategic planning for urban and suburban road networks. Manual approaches are impractical due to the complexity and scale of traffic systems. Automated solutions, however, must prioritize transparency and interpretability to avoid reliance on opaque black box models. In complex urban scenarios, automated reasoning techniques were effective in interpreting traffic patterns, allowing real-time identification of anomalies and hazards, contributing to more informed decision-making. This paper discusses possible uses of multi-strategy automatic reasoning for traffic interpretation to build useful support systems for more mindful road traffic management.

1 Introduction

Effective traffic management is crucial for everyday life in modern societies. Road crashes represent the 8th leading cause of death worldwide [1]. Dealing with massive traffic crowds is urgent and essential [29]. Infrastructures have not kept up with the increased traffic load, but effectively managing traffic on existing roads is more profitable than building new ones [6]. It has also connections with social issues, pollution and the green economy [40].

Both academia and industry have endeavoured to traffic prediction, analysis, and understanding. The demand for real-time, automated traffic solutions is increasing, yet the sheer complexity, scale, and heterogeneity of the data involved make traditional methods insufficient. As such, there is a critical need for advanced AI-based frameworks capable of dealing with these complexities while providing actionable, interpretable insights.

Current traffic analysis tools must contend with various data challenges, often falling within the realm of Big Data [16] and Data Mining [31]. The primary challenges include:

- **Volume**: The vast number of vehicles results in an enormous amount of traffic data.
- **Variety**: Traffic data comes in multiple forms, from video feeds to GPS logs, with varying structures and formats.
- **Velocity**: Traffic data is generated continuously and at a rapid pace, requiring near-instantaneous processing.

In addition to these general Big Data characteristics, traffic data has unique traits such as:

- **Vinculation**: Traffic data streams are highly interconnected; events like traffic light changes, accidents, or pedestrian movements affect the entire system.
- **Validity**: The accuracy of traffic data depends heavily on the technologies used for collection, such as cameras and sensors, which require regular validation to ensure quality.

Big Data fuels black-box models-the more data available, the more complex the models can become, improving accuracy but reducing interpretability. Simpler models (such as linear regression) are interpretable, but may not perform as well on Big Data as deep learning models. In deep learning, black-box reasoning can be applied. This refers to AI/ML models in which decision making is not easily interpretable by humans. These models receive input and produce output, but the internal logic is complex and opaque. To address black box problems, researchers are developing Explainable AI (XAI) techniques like:

- LIME (Local Interpretable Model-agnostic Explanations)[4]: Approximates black box models with simpler, interpretable ones [27].
- SHAP (SHapley Additive exPlanations)[5]: Quantifies feature importance [22].
- Decision Trees & Rule-Based Models[6]: More transparent alternatives [23].

However, all explanations obtained by these methods will always have some degree of uncertainty. In some areas, including road traffic management, regulatory problems may arise as regulators require concrete explanations (e.g., insurance companies in the case of a serious accident).

In a more general view, there is a need to develop a comprehensive framework for road traffic analysis that meets the needs of the

* Corresponding Author. Email: domenico.redavid1@uniba.it.
** Corresponding Author. Email: eleonora.bernasconi@uniba.it.
*** Corresponding Author. Email: stefano.ferilli@uniba.it.
[1] Equal contribution.
[2] Equal contribution.
[3] Equal contribution.

[4] https://github.com/marcotcr/lime
[5] https://github.com/shap/
[6] https://github.com/adaa-polsl/RuleXAI

various stakeholders involved in traffic planning, management, safety and security. Early detection or prediction of traffic problems to facilitate proactive measures to prevent or mitigate their impact assumes a key role in this framework. To ensure reliability and enable users to make informed decisions, in this paper we will analyze one of the main characteristics that a traffic management framework must embody: Explainability, i.e., the ability to explicitly illustrate the reasoning behind each result, offering a full explanation that outlines the logical steps and relationships among the data that determine the results. For this task we address the use of explainable automated multi-strategy reasoning for traffic understanding as a module of a traffic management tool. The rest of this paper is organised as follows. After discussing related work in the next section, we describe the proposed module in Section 3, while in Section 4 we describe and discuss the application of the module to a sample scenario. Finally, we provide the conclusions and future work directions in Section 5.

2 Related Works

In the urban traffic setting, anomalies may vary according to specific scenarios: an anomaly may refer to a specific vehicle or a specific dangerous situation (e.g. accidents, fires, congestion).

A common need for traffic management is the prevention, or at least forecasting, of accidents. Many factors come into play when dealing with traffic accidents, such as age [34] or environmental factors [15]; however, most works do not take these factors into account. The possibility of taking action against road accidents can have a sort of big impact on security and health, but also from an economic point of view. In 2014 Australia was estimated to spend $27 billion on traffic analysis [13]. This investment encompasses various costs related to road crashes, including the lifetime care cost of young drivers involved in crashes, indicating a broad approach to traffic management that spans from immediate accident response to long-term care and prevention strategies. For example, the work by Low and Odgers discusses the rethinking of the cost of traffic congestion and includes lessons from Melbourne's City Link Toll Roads, indirectly touching on the broader economic implications of traffic management [21]. Similarly, Gordon's analysis on applying benefit-cost analysis to Intelligent Transportation Systems (ITS) in the Australian context suggests a comprehensive view on managing traffic through advanced technological systems [12]. Furthermore, the analysis by Buckis, Lenné, and Fitzharris underscores the economic dimensions of traffic accidents, highlighting the extensive financial impact these incidents have on society [3].

Traffic understanding resides in the Big Data domain; hence, the format of data available varies. As an example, in many scenarios, the detection of congested areas or accidents in a city requires the use of GPS. D'Andrea et al. [5] proposed a segment traffic classification to distinguish dangerous (or suspect) areas. More recently, satellite video-based solutions have been gaining momentum. Zhang et al. [39] proposed strategies to recognise a vehicle through Computer Vision (CV) techniques and an improvement of the traditional Adaboost [26].

Following this trend, many applications in the field of traffic understanding have a visual approach, which is eager to be combined with embedded Machine Learning[24], using a camera and recording and/or shooting traffic roads. In some applications, the visual approach is fundamental due to the calibration of instruments based on vehicle and lane dimensions as in [33]. Understanding of traffic may be related to the forecasting of traffic in specific areas in some periods. For this reason, they are mainly modelled with time series approaches [14].

As Xu et al. [37] showed, mobile traffic and urban traffic are linked and some strategies can be interchanged between the two domains. Trinh et al. [35] designated a LSTM-based solution to deal with the typical problem of vanishing gradients, while Feng et al. [7] went even further by using end-to-end neural networks. Many techniques are also borrowed from the field of Swarm Intelligence [42] such as Federation-Imitating Learning [38].

Introducing explanations is essential when decisions have a strong social impact. Given the popularity of neural networks, it is a relevant issue today to introduce explanations in neural-based algorithms, leading to the so-called hyrid models [41].

The extent to which this can work is still under examination and far from being solved.

The decision about the kind of approach enormously affects the instruments to be employed, performance, kinds of explanations, and costs. Current data analytics approaches revolve around the idea of using data to characterize and predict traffic risk in order to prescribe better (safer) routes, driver assignments, rest breaks, etc. With advances in information technology, it is possible to collect ever more relevant data, such as comprehensive incident databases, real-time driving data feeds or relevant factor characteristics [32].

Among the non-explainable solutions, traffic understanding in videos plays a key role. The technological support is provided by cameras. Computer vision algorithms allow us to detect relevant road elements and visualise objects in the scene. Visual multiple object tracking (VMOT) aims to locate multiple targets of interest, infer their trajectories, and maintain their identities in a video sequence [17], or in real-time traffic scenarios [18]. Today, many videos or images recorded are available thanks to visual surveillance and displays in autonomous vehicles. Santhosh et al. [28] provided a survey on anomaly detection techniques by taking photos from video surveillance systems. One of the main points of the detection is represented by the tracking system. Further applications concern motion planning in autonomous vehicles. Apart from recordings, one may find himself surprised by the variety of simulated data that can be generated with videos.

This opportunity is given by computer games, urban visualisation, and urban planning simulation in autonomous driving settings. The design process has been reconsidered. Umbrello et al. [36] tried to mitigate the issue thanks to the Value Sensitive Design (VSD) [11] in combination with the Belief-Desire-Intention (BDI) [10] model. The combination leads to making the vehicle behave by following some design principles which, in turn, follow our ethical schemes and values on the road. Other approaches rely on cutting-edge technologies, such as transformer-based architecture [20]. Also, it should be evaluated when explanations are desired or needed. Shen et al. [30] investigated this need to better understand user expectations and increase their reliability in non-explainable complex systems.

Symbolic solutions, which the proposed approach falls under, are part of what is called an umbrella term, the field of Knowledge Representation and Reasoning (KRR). This approach may require a technological change in the data acquisition process because such data must be structured and, as is well known, are not prone to be extracted from images or videos. But this disadvantage could be overcome by the increased explainability of their models considering also the fact that another key data source is sensors that can produce, manipulate and transfer information in a structured way. The first approaches, such as Cuena et al. [4], adopted classical logic and the Prolog programming language. The aim was to improve traffic or reduce the severity of existing problems. They were able to recommend increasing the duration of a traffic light phase (e.g.: green), or suggest displaying certain messages on some Variable Message Panels to divert traffic. Hülsen et al. [19] extended the state-of-the-art by introducing ontologies into the traditional logic programming setting. In the context of visual surveillance, the description of object behaviour requires real-time analysis [25].

The natural upsides of the new paradigm emerged from the user experts' evaluation, which made them associate the semantics of rules with scenes.

3 Automated MultiStrategy Reasoning for Traffic Interpretation

In addition to learning patterns of standard or abnormal traffic behavior that can be used to supervise future traffic situations [9], a useful functionality is traffic interpretation. This function is designed in the form of an expert decision support module that uses automatic reasoning to reproduce the inferences that an expert continuously watching traffic videos would make. In this way, dangerous or otherwise relevant situations can be detected in real time and notified to relevant stakeholders, possibly even suggesting associated actions that they might accept, modify, or reject. For this purpose, relevant domain knowledge must be expressed in a logical formalism and stored in a KB. It can be provided by domain experts and formalized by knowledge engineers, or learned automatically using First Order Logic (FOL) approaches. Using the KB, a traffic-related video can be continuously "watched," and the information it provides can be formalized and provided to the automated system that interprets it, issuing warnings if relevant situations arise. Furthermore, the system can respond to specific questions from stakeholders.

In particular, our module uses GEAR (an acronym for 'General Engine for Automated Reasoning') [8], a logical inference engine capable of MultiStrategy Reasoning, i.e. of integrating and bringing to cooperation several inference strategies in order to cope with the several complexities posed by real-world tasks and problems. The current GEAR prototype includes the following strategies:

- **Deduction** aims at making explicit knowledge that is implicit in the available knowledge but is a necessary consequence thereof.
- **Abstraction** reduces the amount of information conveyed by a set of facts, abstracting from unnecessary details.
- **Abduction** is devoted to coping with missing information, by guessing unknown facts that are not stated in the available knowledge, but are needed to solve a given problem, provided that they satisfy some integrity constraints. Of course, there may be many plausible explanations for a given observation.
- **Uncertainty** is the possibility that uncertainty can dramatically improve the flexibility and robustness of reasoning.
- **Argumentation** deals with inconsistent knowledge, to distinguish which of several contrasting, but internally consistent, positions are justified.
- **Induction** is the inference of general knowledge starting from specific observations.
- **Ontological**. An ontology defines and describes the kinds of entity that are of interest in a domain, their properties, and relationships. Typical ontology-based reasoning tasks are inheritance and consistency checks.
- **Similarity-based computation** between FOL descriptions is complex due to non-unique mapping between the descriptions.
- **Analogy** matches the characterizing features of two subjects, objects, situations, etc. even if they use different descriptors.

Reasoning operators act on the content of a so-called *Knowledge Base* (KB). Knowledge bases handled by GEAR may include various kinds of knowledge items, including Facts (simple statements), Rules (implication-like formulas), Integrity Constraints, Abstraction Operators, and Argumentative relationships. Rules may have a priority (a number used to determine which rule should be executed first in case of conflicts). The premises of rules can use any, possibly nested, composition of conjunction, disjunction, and negation. Its conclusion can be a single atom or a conjunction of atoms or negated atoms. Additional components are available to express abducibles and integrity constraints for abduction, abstraction operators, or argumentative relationships. Knowledge can be organised in modular way. Other predicates can be used to specify system settings (e.g.: to set flags that direct the system's behaviour), information related to user interaction (e.g.: to specify the information that can be asked to the user if missing in the KB), calls to pre-defined procedures (e.g.: to call Prolog to carry out some computations), etc.

4 Sample Scenario

We evaluated the proposed method on a use case drawn from a hypothetical scenario involving urban traffic areas. We assume that we use information from cameras placed at critical points in the chosen road area and interpret the behavior using an automatic reasoning system. A number of state-of-the-art techniques have been employed to achieve vehicle identification in real time; various types of vehicles can be identified, including cars, trucks, buses, and motorcycles [2]. For this task, we used continuous streaming videos provided by the SkylineWebcams streaming platform[7]. In particular, we chose an urban setting in Rome, namely *Piazza Venezia*.

Figure 1 shows a screen shot illustrating an instrument we made for this purpose in action. The right side of the screen presents real-time logs, with information such as the number of vehicles detected, the accuracy of detection, and the duration of identification of each vehicle. The left side of the screen displays frames illustrating the bounding boxes surrounding the identified vehicles and their accuracy rates. While the video flows, the video analysis tool generates vehicle-related information and asserts corresponding facts in the KB. Note that throughout the case study, we used real-world data for all tasks, so our experimental outcomes can be considered as meaningful for practical application of our framework.

The experiments were run on a state-of-the-art high-performance computing (HPC) system with a 64 bit architecture, endowed with an Intel64 Family 6 Model 85 Stepping 4 GenuineIntel CPU with base clock speed of 3.312GHz, 32GB RAM, and an NVIDIA GeForce RTX 2080 Ti GPU featuring 4352 CUDA cores.

4.1 Automated Reasoning for Traffic Interpretation

We provide a demonstration use case for this feature in relation to the location of Piazza Venezia in Rome. While the video flows, the video analysis module generates vehicle-related information and asserts corresponding facts in the KB based on the following predicates:

`object(O,X0,X1,Y0,Y1,T)` : object with identifier O is recognised in the scene at time T, enclosed in a bounding box with coordinates $X0, X1, Y0, Y1$.

`next(T',T")` : time T'' follows time T'.

[7] Live Cams in Italy - SkylineWebcams, available at: https://www.skylinewebcams.com/en/webcam/italia.html

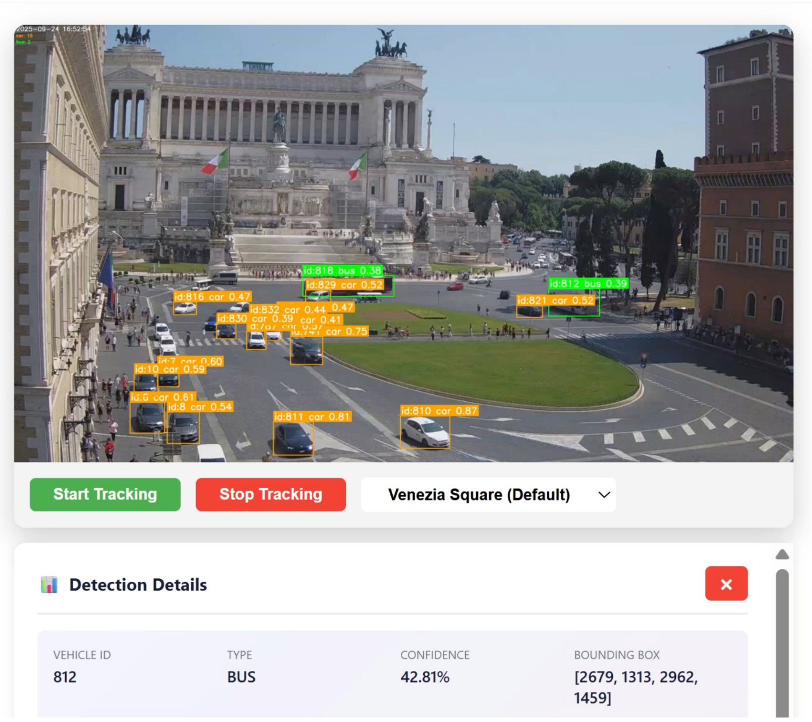

Figure 1. Vehicle detection application

Whenever appropriate or useful (e.g.: periodically every k frames processed), the reasoning engine GEAR is started to interpret what happened and return relevant notifications. It starts by deriving simple knowledge, as expressed by the predicates that we have defined for this sample application reported in Table 1.

Rules for defining these predicates and for determining when a time period or displacement is ‘relevant are stored in the KB using GEAR’s formalism. These concepts are interrelated, meaning that some are defined upon others (e.g.: `move/6` is defined in terms of `distance/9`).

Then, the reasoning may proceed with further concepts or situations of interest that the system is expected to identify (e.g.: road traffic offences, etc.), at higher and higher levels of abstraction (i.e., not associated with a simple position occupied by a vehicle but determined according to the overall vehicle behaviour and to its relationships to the road features and the behaviour of other vehicles). In our sample use case, we focus on the following situations:

- traffic jam;
- vehicle going faster than the maximum allowed speed;
- vehicle passing from forbidden zones of the road;
- vehicle stopping in places where the stop is forbidden;
- vehicle taking a wrong turn;
- vehicle going around the square in a loop;

and provided detailed descriptions thereof and of how they can be detected, possibly based on simpler situations that may not be relevant by themselves. Multistrategy reasoning comes in handy for many such tasks. Specifically:

- abstraction can simplify the scenario by removing noise or irrelevant background information;
- deduction can infer the behavior of drivers or categories of drivers (e.g., dangerous, safe, distracted etc);
- induction can generalize the behavior of drivers for some category or during some specific conditions (e.g. day vs night driving);
- abduction can infer unseen events (e.g., when one vehicle in the scene hides another);
- uncertainty can make reasoning in complex scenarios more flexible;
- argumentation can fix possible inconsistencies due to errors in sensor data (e.g., two cars in the same place or two distinct cars identified as the same).

4.2 *An example of deduction rule*

As a simple example of a deduction rule, we can define the concept of *fast* driver by the rule reported in Algorithm 1.

Algorithm 1 Fast Movement Detection Rule

Input: X, T_1, T_2
Output: fast(X, T_1, T_2, C)

if for some X,
  C is the number of facts move$(X, S_1, S_2, S_3, S_4, T)$ such that:
  $T \geq T_1 \wedge T \leq T_2 \wedge C \geq 2$
then
  fast(X, T_1, T_2, C)
  WITH PRIORITY 1.0
  AND CERTAINTY 1.0
end if

In other words, if X moves more than twice in the time interval $[T1, T2]$, then X is *fast* and C is an indicator of how fast it is.

Table 1. Predicates defined for the Piazza Venezia use case.

Predicate	Description
`move(O,X0,X1,Y0,Y1,T)`	Object O moved (by a considerable distance) at time T, where $X0, X1, Y0, Y1$ are the displacements of each coordinate of its bounding box.
`enter(O,P,T)`	Object O entered place P (a RoI) at time T.
`leave(O,P,T)`	Object O left place P (a RoI) at time T.
`still(O,T)`	Object O stopped at time T.
`halt(O,L,T)`	Object O stopped for a certain period of time L starting from time T.
`stay(O,P)`	Object O stayed in place P (a RoI).
`placetime(O,P,T)`	Object O was in place P (a RoI) for a considerable time T.
`status(O,T,S)`	S is the status of the object O at time T, where S can be 'moving', 'still', or someplace (RoI) identifier.
`meet(L,T,P)`	Objects in the list L were in the same place P at time T.
`wait(X,Y,T)`	Object X was still at time T, but is now in the same place as object Y.
`distance(X00,X01,Y00,Y01,X10,X11,Y10,Y11,D)`	D is the Euclidean distance between the coordinates $X00, X01, Y00, Y01$ and $X10, X11, Y10, Y11$ of two bounding boxes.
`closetimes(X,Y,T,L)`	L is the last timestamp, starting from T, in which X and Y were close to each other.
`close(X,Y,T,L)`	L is the amount of time for which X and Y were close to each other, starting from timestamp T.
`accomplices(X,Y,T,D)`	Objects X and Y are close to each other for a certain time D in the 'halt' state, and so still for a considerable time, at time T.
`fast(O,T1,T2,D)`	Object O moved many times between timestamps $T1$ and $T2$ with distances greater than or equal to D.

In GEAR's formalism, this is expressed as:

```
rule(14,
  fast(X,T1,T2,Count),
  call_p((aggregate_all(count,
    (move(X,S1,S2,S3,S4,T),T≥T1,T≤T2),
    Count),Count≥2)),
  1,1).
```

These descriptions were formalized by a knowledge engineer and used to create a KB, that consisted of several dozen rules. Applied to short videos taken from the Piazza Venezia location, this knowledge could allow the system to successfully identify several occurrences of those situations while the video was running and raise associated alarms, as well as to identify those situations upon specific requests of the stakeholder, such as:

- "Did any vehicle go around the square in a loop from time X to time Y in the video?".
- "Did any vehicles stand stop and occlude the passage of the square from X hour to Y hour in the video?".

5 Conclusion and Future Work Directions

Efficient traffic management remains a critical challenge for urban transportation systems globally. Traffic anomalies, such as accidents, congestion, and roadblocks, result in substantial social and economic costs. Timely detection of these anomalies is essential for improving traffic safety, reducing delays, and mitigating environmental impacts. While numerous methods have been proposed, ranging from traditional statistical approaches to modern machine learning and computer vision techniques, an important aspect is the explanability.

In this paper, we introduced a novel AI MultiStrategy Reasoning approach able to apply automated reasoning to detect both sudden and gradual traffic anomalies, such as accidents, congestion, and lane violations. The approach was evaluated using real-world traffic data from a major urban area, demonstrating its capability to explain by means of logical rules various types of anomalies.

While the proposed approach is promising, it also has limitations. The efficacy and accuracy of the detection methods highly depend on the quality and completeness of the input data.

Future work may proceed with the definition of a complete framework for Traffic Management. Starting having a solution for the interpretation of the urban-suburban Traffic, we can proceed to develop or integrate the most useful solution for:

- Information Extraction from Traffic Videos. In particular, two fundamental functionality are "Vehicle Detection and Tracking" and "Trajectory analysis and event detection".
- Identification of Noteworthy Areas, i.e., given a map or video frame, how identify areas that are relevant for traffic interpretation and understanding?
- Behavioral Modeling, i.e, model over which apply Supervision, Prediction and Classification.

Obtained a framework with these characteristics it will be possible to address several specific problems:

- Cross-Domain data integration: Expanding the framework to incorporate data from multiple sources, including IoT devices, social media, and weather reports, could provide a more comprehensive view of traffic conditions, taking into account external factors that influence road behavior.
- Human behavior analysis: Integrating the analysis of human behavior—both drivers and pedestrians—can offer deeper insights into traffic dynamics. Understanding driver behavior and pedestrian movements could help develop more adaptive traffic management systems that cater to the needs of all road users.
- Integration with smart city infrastructures: Connecting the traffic management framework with broader smart city systems, such as

smart lighting, parking, and emergency services, could enhance urban responsiveness and create more cohesive city-wide traffic management solutions.

- Sustainability and environmental impact: By optimizing traffic flows and reducing congestion, the framework can contribute to reducing greenhouse gas emissions and fuel consumption, promoting sustainability in urban transportation systems.

Acknowledgements

This research was partially supported by project FAIR – Future AI Research (PE00000013), spoke 6 – Symbiotic AI, under the NRRP MUR program funded by the NextGenerationEU.

References

[1] The top 10 causes of death. https://www.who.int/en/news-room/fact-sheets/detail/the-top-10-causes-of-death. Accessed: 2023-02-27.

[2] Y. Amit, P. Felzenszwalb, and R. Girshick. Object detection. *Computer Vision: A Reference Guide*, pages 1–9, 2020.

[3] S. Buckis, M. G. Lenné, and M. Fitzharris. An analysis of young driver crash types and the associated lifetime care cost in victoria, australia. *Traffic injury prevention*, 16(sup2):S66–S76, 2015.

[4] J. Cuena, J. Hernández, and M. Molina. Knowledge-based models for adaptive traffic management systems. *Transportation Research Part C: Emerging Technologies*, 3(5):311–337, 1995.

[5] E. D'Andrea and F. Marcelloni. Detection of traffic congestion and incidents from gps trace analysis. *Expert Systems with Applications*, 73:43–56, 2017.

[6] D. Elkin and V. Vyatkin. Iot in traffic management: Review of existing methods of road traffic regulation. In R. Silhavy, editor, *Applied Informatics and Cybernetics in Intelligent Systems, Volume 3*, volume 1226 of *Advances in Intelligent Systems and Computing*, pages 536–551. Springer, 2020. doi: 10.1007/978-3-030-51974-2_50.

[7] J. Feng, X. Chen, R. Gao, M. Zeng, and Y. Li. Deeptp: An end-to-end neural network for mobile cellular traffic prediction. *IEEE Network*, 32 (6):108–115, 2018.

[8] S. Ferilli. Gear: A general inference engine for automated multistrategy reasoning. *Electronics*, 12(2), 2023.

[9] S. Ferilli and S. Angelastro. Activity prediction in process mining using the woman framework. *Journal of Intelligent Information Systems*, 53: 93–112, 2019.

[10] L. Fichera, D. Marletta, V. Nicosia, and C. Santoro. Flexible robot strategy design using belief-desire-intention model. In *Research and Education in Robotics-EUROBOT 2010: International Conference, Rapperswil-Jona, Switzerland, May 27-30, 2010, Revised Selected Papers*, pages 57–71. Springer, 2011.

[11] B. Friedman. Value-sensitive design. *interactions*, 3(6):16–23, 1996.

[12] C. Gordon. Applying benefit-cost analysis to intelligent transportation systems (its) and the australian context. *Available at SSRN 2389844*, 2013.

[13] A. Government. *Impact of road trauma and measures to improve outcomes / Bureau of Infrastructure, Transport and Regional Economics*. Department of Infrastructure and Regional Development, Bureau of Infrastructure, Transport and Regional Economics Canberra, ACT, 2014. ISBN 9781925216103. URL http://www.bitre.gov.au/publications/2014/files/report_140.pdf.

[14] J. D. Hamilton. *Time series analysis*. Princeton university press, 2020.

[15] H. M. Hammad, M. Ashraf, F. Abbas, H. F. Bakhat, S. A. Qaisrani, M. Mubeen, S. Fahad, and M. Awais. Environmental factors affecting the frequency of road traffic accidents: a case study of sub-urban area of pakistan. *Environmental Science and Pollution Research*, 26:11674–11685, 2019.

[16] Y. Himeur, M. Elnour, F. Fadli, N. Meskin, I. Petri, Y. Rezgui, F. Bensaali, and A. Amira. Ai-big data analytics for building automation and management systems: a survey, actual challenges and future perspectives. *Artificial Intelligence Review*, 56(6):4929–5021, 2023.

[17] C. Huang, Y. Li, and R. Nevatia. Multiple target tracking by learning-based hierarchical association of detection responses. *IEEE transactions on pattern analysis and machine intelligence*, 35(4):898–910, 2012.

[18] Y. Huangfu, M. Ahrabi, R. Tahal, J. Huang, A. Mohammad-Alikhani, S. Reymann, B. Nahid-Mobarakeh, S. Shirani, and S. Habibi. Efficient edge computing device for traffic monitoring using deep learning detectors. *IEEE Canadian Journal of Electrical and Computer Engineering*, 46(4):371–379, 2023.

[19] M. Hülsen, J. M. Zöllner, and C. Weiss. Traffic intersection situation description ontology for advanced driver assistance. In *2011 IEEE Intelligent Vehicles Symposium (IV)*, pages 993–999, 2011. doi: 10.1109/IVS.2011.5940415.

[20] N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu. Neural speech synthesis with transformer network. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 6706–6713. AAAI Press, 2019. doi: 10.1609/aaai.v33i01.33016706.

[21] N. Low and J. Odgers. Rethinking the cost of traffic congestion, lessons from melbourne's city link toll roads. *Urban policy and research*, 30 (2):189–205, 2012.

[22] S. M. Lundberg and S.-I. Lee. A unified approach to interpreting model predictions. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4765–4774. Curran Associates, Inc., 2017. URL http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf.

[23] S. M. Lundberg, G. Erion, H. Chen, A. DeGrave, J. M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S.-I. Lee. From local explanations to global understanding with explainable ai for trees. *Nature Machine Intelligence*, 2(1):2522–5839, 2020.

[24] F. Maccantelli, G. Peruzzi, and A. Pozzebon. Traffic level monitoring in urban scenarios with virtual sensing techniques enabled by embedded machine learning. In *2023 IEEE Sensors Applications Symposium (SAS)*, pages 1–6. IEEE, 2023.

[25] A. Parashar, A. Parashar, A. F. Abate, R. S. Shekhawat, and I. Rida. Real-time gait biometrics for surveillance applications: A review. *Image and Vision Computing*, page 104784, 2023.

[26] M. T. Ramakrishna, V. K. Venkatesan, I. Izonin, M. Havryliuk, and C. R. Bhat. Homogeneous adaboost ensemble machine learning algorithms with reduced entropy on balanced data. *Entropy*, 25(2):245, 2023.

[27] M. T. Ribeiro, S. Singh, and C. Guestrin. "why should i trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, page 1135–1144, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450342322. doi: 10.1145/2939672.2939778.

[28] K. K. Santhosh, D. P. Dogra, and P. P. Roy. Anomaly detection in road traffic using visual surveillance: A survey. *ACM Computing Surveys (CSUR)*, 53(6):1–26, 2020.

[29] R. Savithramma, R. Sumathi, and H. Sudhira. Reinforcement learning based traffic signal controller with state reduction. *Journal of Engineering Research*, 11(1):100017, 2023.

[30] Y. Shen, S. Jiang, Y. Chen, E. Yang, X. Jin, Y. Fan, and K. D. Campbell. To explain or not to explain: A study on the necessity of explanations for autonomous vehicles. *arXiv preprint arXiv:2006.11684*, 2020.

[31] X. Shu and Y. Ye. Knowledge discovery: Methods from data mining and machine learning. *Social Science Research*, 110:102817, 2023.

[32] A. Sohail, M. A. Cheema, M. E. Ali, A. N. Toosi, and H. A. Rakha. Data-driven approaches for road safety: A comprehensive systematic literature review. *Safety science*, 158:105949, 2023.

[33] X. Tang, W. Wang, H. Song, and Y. Li. Novel optimization approach for camera calibration in traffic scenes. *Transportation research record*, 2677(3):1048–1066, 2023.

[34] F. Tarlochan, M. I. M. Ibrahim, and B. Gaben. Understanding traffic accidents among young drivers in qatar. *International journal of environmental research and public health*, 19(1):514, 2022.

[35] H. D. Trinh, L. Giupponi, and P. Dini. Mobile traffic prediction from raw data using lstm networks. In *2018 IEEE 29th annual international symposium on personal, indoor and mobile radio communications (PIMRC)*, pages 1827–1832. IEEE, 2018.

[36] S. Umbrello and R. V. Yampolskiy. Designing ai for explainability and verifiability: a value sensitive design approach to avoid artificial stupidity in autonomous vehicles. *International Journal of Social Robotics*, 14(2):313–322, 2022.

[37] F. Xu, Y. Lin, J. Huang, D. Wu, H. Shi, J. Song, and Y. Li. Big data driven mobile traffic understanding and forecasting: A time series approach. *IEEE Transactions on Services Computing*, 9(5):796–805, 2016. doi: 10.1109/TSC.2016.2599878.

[38] B. Yang, H. Shi, and X. Xia. Federated imitation learning for uav swarm coordination in urban traffic monitoring. *IEEE Transactions on Industrial Informatics*, 19(4):6037–6046, 2022.

[39] L. Zhang, J. Wang, and Z. An. Vehicle recognition algorithm based on haar-like features and improved adaboost classifier. *Journal of Ambient Intelligence and Humanized Computing*, 14(2):807–815, 2023.

[40] L. Zhang, M. Xu, H. Chen, Y. Li, and S. Chen. Globalization, green economy and environmental challenges: State of the art review for practical implications. front. *Financial and Trade Globalization, Greener Technologies and Energy Transition*, 16648714:133, 2023.

[41] G. Zheng, W. K. Chai, J.-L. Duanmu, and V. Katos. Hybrid deep learning models for traffic prediction in large-scale road networks. *Information Fusion*, 92:93–114, 2023.

[42] Y. Zhou, B. Rao, and W. Wang. Uav swarm intelligence: Recent advances and future trends. *Ieee Access*, 8:183856–183878, 2020.