KI WISSEN

Automotive AI Powered by Knowledge

Funded by the European Union
NextGenerationEU

Supported by:
Federal Ministry for Economic Affairs and Climate Action
on the basis of a decision by the German Bundestag

KI WISSEN
Automotive AI Powered by Knowledge

KI FAMILIE

KI Wissen is a project of the KI Familie. It was initiated and developed by the VDA Leitinitiative autonomous and connected driving and is funded by the Federal Ministry for Economic Affairs and Climate Action.

www.kiwissen.de @KI_Familie KI Familie

KI WISSEN
Automotive AI Powered by Knowledge

Welcome

Dear Reader,
We are pleased to present the final KI Wissen result booklet.

As a vital part of the „KI-Familie“ within the VDA Leitinitiative autonomous and connected driving, KI Wissen adheres to clear safety standards for automated systems, proving crucial for the German automotive industry in an era of global competition and rapid digitization in mobility.

The new structuring as project families efficiently utilizes research and development resources, mobilizing synergies within the project families to address overarching requirements and share results. It significantly contributes to the overall goal of accelerating the joint innovation agenda.

At the forefront of current AI development is a data-centric methodology leading to black-box models with limitations. KI Wissen boldly tackles these challenges by exploring the seamless integration of established knowledge into AI systems, redefining the foundation for training, and validating AI functions.

The transformation from data-based to information-based AI in KI Wissen confronts key hurdles: generalizing AI to phenomena with limited data, enhancing stability against data disturbances, ensuring data efficiency, validating and securing AI-assisted functions, and elevating functional quality.

The overarching aim of the KI Wissen Project was to address critical challenges in AI development through four key sub-projects: „Knowledge Integration,“ „Knowledge Extraction,“ „Knowledge Conformity,“ and „Enabler, Integration, and Demonstration.“

„Knowledge Integration“ aimed to optimize training processes by reducing data volume and computational demands while simultaneously enhancing generalization capability and functional quality. This approach laid the foundation to improve machine learning models by integrating relevant domain knowledge. The success in this area paved the way for enhanced data efficiency, generalization, and safety, providing a viable path forward to overcome limitations associated with purely data-driven models that rely on extensive training data and may exhibit unpredictable, safety-critical behavior.

The focus of „Knowledge Extraction“ was multifaceted, centering on concepts, measures for deriving insights, and models delivering structured outputs. This effort not only facilitated

the recognition and utilization of newly acquired information but also contributed to the explainability and analyzability of AI systems. Successful development of methods for the interface between the output of ML models and human interpretation enhanced the safety and traceability of the system.

„Knowledge Conformity“ emphasized the development of methods to assess the conformity of AI system outputs with existing knowledge. This aspect was crucial for ensuring the validation and safety of autonomous vehicle functions, with findings influencing further training or real-time enhancements. The methods developed in this discipline checked the behavior and decisions of knowledge-infused AI driving functions for conformity with existing knowledge. Various approaches were demonstrated to formalize and utilize knowledge, verifying the knowledge conformity of an AI decision. These results found applications in improving the efficiency of AI model training, the reliability of AI abductions, and the safety of AI driving functions during operational phases.

Fulfilling the holistic approach of KI Wissen, „Enabler, Integration, and Demonstration“ evaluated the developed functions, components, and methods through three selected use cases within a comprehensive system. This evaluation went beyond assessing individual components, allowing for an examination of their suitability within the broader system

architecture. The integration of newly developed AI modules from the aforementioned sub-disciplines into the KI Wissen demonstrator created a collaborative synergy between the expertise of various sub-projects and work packages. This integration not only improved the performance of the AD stack but also laid a robust foundation for future advancements, ensuring a cohesive and interoperable framework for autonomous driving functions.

To sum it up: The solutions and methods developed in KI Wissen for the integration of domain knowledge in AI systems are a significant step forward in the development of hybrid AI technologies tailored to the automotive environment. The ability to extract new knowledge, check AI decisions for conformity, and reduce training efforts through domain knowledge integration is a substantial contribution to a new foundation for safer, better, and more efficient AI systems. The possibilities unlocked by KI Wissen represent a major advancement in the field of AI and the entire automotive sector.

The success of KI Wissen wouldn‘t have been possible without the 15 project partners‘ unwavering commitment and contributions. The results presented in this booklet are a testament to their collective expertise and dedication.

We extend our sincere thanks to the Federal Ministry for Economic Affairs and Climate Action for supporting AI research, shaping the direction of our project. Furthermore, we extend our heartfelt thanks to TÜV Rheinland Consulting for their dedicated monitoring and unwavering support throughout the project. Particular recognition goes to our project management partner EICT for invaluable guidance and organizational support, significantly enhancing our efforts.

Thank you to ALL for being the driving force behind the success of KI Wissen!

Dr. Jörg Dietrich
Continental AG
Project Coordinator KI Wissen 2021-2023

Simon Heinz
Continental AG
Project Coordinator KI Wissen 2023-2024

Greeting from the Federal Ministry for Economic Affairs and Climate Action

Artificial Intelligence plays an increasingly and impactful role within research on automated driving. At an early stage the VDA Leitinitiative was established to address, among other topics, various AI-related challenges. It is supposed to advance AI competencies in helping to develop industry-wide and concrete R&D-projects specifically in the complex realm of autonomous driving. The Federal Ministry for Economic Affairs and Climate Action supports these efforts, e. g., through substantial funding of the KI Familie. Commencing with confidence, KI Wissen, as part of the KI Familie, faced an uncertain starting point amid the contentious debate that prevailed at the project‘s initiation. In the discussions and questions regarding hybrid AI, the merge of data- and rule-based knowledge was questioned and discussed. KI Wissen faced the challenging goal of bringing together two previously separated worlds. Modern AI functions are powerful in the domain for which they were developed, but at the same time they lack an understanding of what they have learned. They are not able to abstract and draw conclusions; they can only reproduce what they have learned. KI Wissen has made the connection, and the project has achieved its goal of integrating different types of knowledge into machine learning. Modern data-driven AI methods have been combined with knowledge-based approaches. KI Wissen has successfully created a comprehensive ecosystem for integrating knowledge into training and safety validation of AI functions. The BMWK focuses on funding collaborative research approaches to bring together the expertise of diverse participants. This also helps to foster the transfer of knowledge between partners and beyond. The KI Wissen consortium consists of 15 partners from industry, research and science and is led by Continental AG. We, the Federal Ministry for Economic Affairs and Climate Action, would like to thank all participants for successfully completing the KI Familie projects. The impactful results of all four projects are clearly essential for the industry to take a leading role in safe autonomous driving. Please enjoy the KI Wissen booklet and find out about the research results. The project outcome which will contribute to come another important step closer to the realization of numerous opportunities of the diverse and dynamic field of autonomous driving to the benefit of our society and the competitiveness of our industry.

Ernst Stöckl-Pukall
Head of Division for Digitalisation and Industry 4.0, Federal Ministry of Economic Affairs and Climate Action

Federal Ministry for Economic Affairs and Climate Action

Collaboration in Artificial Intelligence

Classic automotive questions are re-emerging with regard to AI. AI technology know-how and its safe use in modern vehicles will determine the leading role in the mobility markets of the future. The German automotive industry addresses this challenge with the projects of the **KI Familie**. The KI Familie was initiated and developed by the **VDA Leitinitiative Connected and Autonomous Driving**. 80 leading partners from science and industry are involved receiving funding from the Federal Ministry for Economic Affairs and Climate Action (BMWK).

In this unique setting, all KI Familie projects are working together. The partners are sharing knowledge while fostering pre-competitive collaboration which is essential in an ever more competitive and complex environment with fast pace innovations. Exchanging findings across project boundaries accelerates the knowledge buildup in cutting edge technologies for the good of industries, research institutions and society. The joint commitment to share pre-competitive knowledge helps each partner to stay technologically ahead and multiplies resources and investments of each partner.

The KI Familie has four sibling projects which all are focusing on special AI topics.

KI ABSICHERUNG
Methods and measures to safeguard AI-based perception functions for automated driving.
www.ki-absicherung-projekt.de

KI WISSEN
Development of methods for the integration of knowledge into machine learning.
www.kiwissen.de

KI DELTA LEARNING
Development of methods and tools for the efficient expansion and transformation of existing AI modules in autonomous vehicles to meet the challenges of new domains.
www.ki-deltalearning.de

KI DATA TOOLING
Methods and tools for the generation and refinement of training, validation and safeguarding data for AI functions in autonomous vehicles.
www.ki-datatooling.de

Facts & Figures

Simon Heinz

Continental AG

Project Coordinator

39 Months

Project Duration: (01/01/2021 - 31/03/2024)

15 Project Partners

8 Industry Partners

1 Academic Partner

6 Research Institutes

€ 25.9 M

Project Budget

€ 17.4 M

Funding Budget

Funding Body

Federal Ministry for Economic Affairs and Climate Action (BMWK)

KI Wissen – Automotive AI Powered by Knowledge

Challenge and Goal

AI-based processes are paving the way to fully automated driving. Up until now, the development of AI solutions has been purely driven by data, which is needed for training and validation. Furthermore, data-based AI functions are often black box models, their decision making cannot always be retraced. In the research project KI Wissen, we addressed these challenges and developed methods for integrating existing knowledge in order to make the training process more efficient and to significantly increase the comprehensibility of AI decision making. The goal of the project was to create a comprehensive ecosystem for the integration of knowledge into the training and safeguarding of AI functions. By combining conventional data-based AI methods with the knowledge- or rule-based methods developed in the project, the basis for training and validating of AI functions were completely redefined: This basis includes not only data, but information, i.e., data and knowledge. The development from data- to information-based AI carried out in the project addressed the central challenges towards autonomous driving: the generalization of AI to phenomena with small data bases, the increase of the stability of the trained AI to disturbances in the data, the data efficiency, the plausibility check and the validation of AI-supported functions as well as the increase of the functional quality.

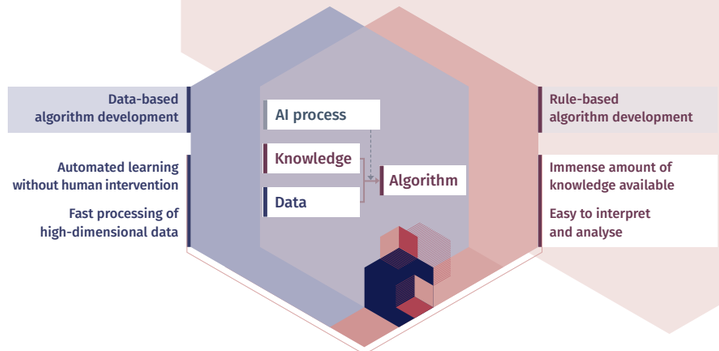

Combine the advantages and eliminate the disadvantages of both worlds in one approach

Approach, Development and Outcome

KI Wissen has fostered a collaborative environment, leveraging the extensive competence of a diverse consortium comprising leading suppliers and a robust group of technology providers, research institutes, universities, federal research institutions, and supporting partners. Within the KI Wissen project, a collaborative synergy was successfully established between the expertise of the various subprojects and work packages. The ‘Knowledge Integration‘ team aimed to enhance machine learning models by integrating relevant domain knowledge. This approach seeks to improve data efficiency, generalization, and safety, addressing the limitations of purely data-driven models that require extensive training data and can exhibit unpredictable, safety-critical behavior. In ‘Knowledge Extraction‘, KI Wissen developed methods for the interface between the output of the ML models and human interpretation, thereby improving the safety and traceability of the system. KI Wissen‘s ‘Knowledge Conformity‘ team has successfully developed methods to check the behavior and decisions of knowledge-infused AI driving functions to ensure knowledge conformity. It was demonstrated in various ways how knowledge can be formalized and utilized to check the knowledge conformity of an AI‘s decision. These findings have applications in enhancing the efficiency of training AI models, the reliability of AI inference, and the safety of AI driving functions during operation phase. The Enabler, Integration, and Demonstration team of KI Wissen was building on the integration of newly developed AI modules into the KI Wissen demonstrator. This integration has not only enhanced the performance of the AD Stack but also established a robust foundation for future advancements, ensuring a cohesive and interoperable framework for autonomous driving functionality. In parallel with the integration of AI modules, KI Wissen has made significant progress in simulation solutions. The establishment of an advanced simulation platform in the project has allowed the accurate representation of various use case scenarios, enabling thorough testing and validation of the integrated AI modules. This serves as a critical tool for identifying potential challenges, fine-tuning algorithms, and strengthening the overall reliability and adaptability of the project‘s deliverables. Furthermore, a meticulous approach was taken to collect data by including both synthetic and real-world data. The rigorous selection process, along with thorough post-processing, has also ensured the high quality and suitability of the real-world data for labeling purposes. This strategy of dual-source data provisioning not only enriches the datasets for the AI modules, but also enhances the system‘s robustness by exposing it to a collection of scenarios, both simulated and real, contributing to a comprehensive and effective validation process.

The collaborative efforts of the project have yielded significant advancements in AI-driven technologies for autonomous driving.

Detection

Knowledge-Augmented Object Detection with Transformer 22
Model-Agnostic Body Part Relevance Assessment for Pedestrian Detection 24
Investigating the Effects of Augmentation Techniques on Interpretability and Perfor-mance of Pedestrian Detection Models 26
Static Scene Knowledge for Pedestrian Detection 28
Method to Fuse Map Layer and Pedestrian State Space to Improve 3D Detector Output 30
Knowledge Guided Occluded Pedestrian Detection 32
Knowledge Aware Pedestrian Detection 34
Geo-Informed Conformity Check of Pedestrian Detection Models 36
Efficient Pedestrian Detection for Autonomous Driving 38
Prior-Knowledge Based Deep-Learning Approach for Pedestrian Detection 40
Performance Impact of Input Image Characteristics and Prior Knowledge Conformity in Prior Mask R-CNN 42
Plausibility Verification For 3D Object Detectors Using Energy-Based Optimization 44
Object Detection Plausibility with Concept-Bottleneck Models 46
Digital Twin Creation and Evaluation of Color Dependent Object Detection 48
Multi-Modal Sensor Fusion for Robust Obstacle Detection 50
Leveraging Knowledge for Traffic Sign Detection 52

Detection

The task of detection in the context of autonomous driving involves identifying and categorizing objects in the vehicle‘s surroundings using sensor data, such as cameras, LiDAR, and RADAR. The primary goal is to recognize other traffic participants and especially vulnerable road users like pedestrians in realtime. Typically, deep neural networks are utilized to analyse sensor data generating detailed understanding of the environment and thus enabling the autonomous vehicle to make informed decisions for navigating safely. Thereby, accurate detections are crucial for tasks like collision avoidance and path planning, overall ensuring the safety and efficiency of autonomous driving systems. Current approaches lack robustness, because of generalization insufficiencies in real world driving environments.

To overcome that challenge, the partners in KI Wissen developed methods to explain the decision-making process as well as including prior knowledge into ML-Models for object classification and detection. For explaining model decisions and extracting knowledge, we have evaluated augmentation techniques, state-of-the-art architectures like visual transformers and concept relevance, e.g. for body-parts.

Furthermore, also the plausibilization of object detections using knowledge (e.g. semantic concept, spatial perspective, or map information). To integrate knowledge into models, the training process was adapted to include explicit information, e.g. static scene knowledge, pedestrian body-parts or the co-occurrence of objects. Different strategies were evaluated on how to integrate the knowledge by utilizing for example meta learning approaches and attention mechanisms. In addition, we evaluated methods for knowledge integration by utilizing simulations and sensor data fusion.

The results show that the training process clearly benefits from the use of knowledge: the amount of training data can be reduced significantly while the robustness and interpretability of the system increases. Finally, it could be shown that knowledge can efficiently be used for testing and to validate object detection models.

Knowledge-Augmented Object Detection with Transformer

Tianming Qiu, Julian Wörmann, fortiss

We present a model that integrates synthetic data as prior knowledge into a Transformer-based object detection pipeline. With features extracted from synthetic support image patches, our model improves the efficiency of the Transformer-based framework. It enhances the quality of region proposal initialization to obtain a better detection performance and achieves faster detection model training under limited available data. Our approach outperforms the baseline model on the Cityscapes and KITTI datasets, suggesting a potential benefit for use cases such as domain adaptation when dealing with limited data.

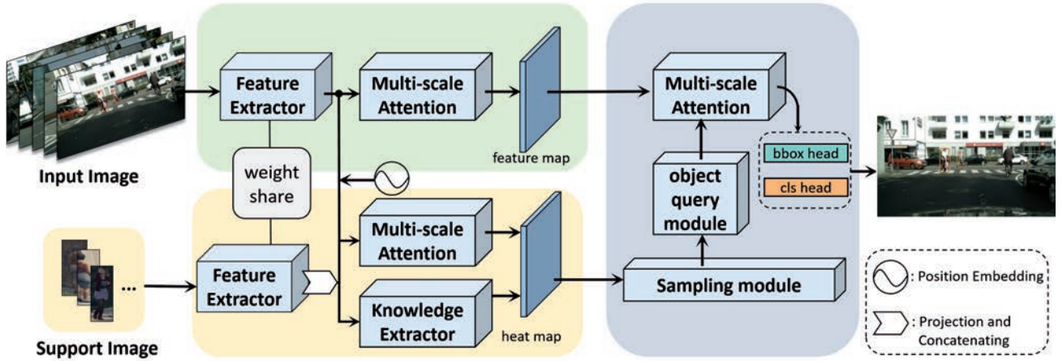

Our method's framework comprises three components: the Encoder for input images, the Encoder for support images, and the Decoder for object detection. The weight-shared feature extraction network processes the augmented support image and input training image. Leveraging support images from the Encoder allows the disentanglement of semantic information, enabling the attention to input images to guide the object queries in the Decoder. This, in turn, facilitates more target-oriented initial queries. (© fortiss GmbH)

Model-Agnostic Body Part Relevance Assessment for Pedestrian Detection

Maurice Günder, FhG IAIS

Model-agnostic explanation methods for deep learning models are flexible regarding usability and availability but struggle with complex model architectures. Sampling-based methods like KernelSHAP are inefficient for large inputs, such as images in object detection use cases. We propose a framework for using sampling-based explanation models in computer vision, specifically for pedestrian detection and demonstrate this by assessing body part relevances for detection. Our novel method, similar to KernelSHAP, is more efficient with lower sampling sizes for explainability analysis on large-scale datasets.

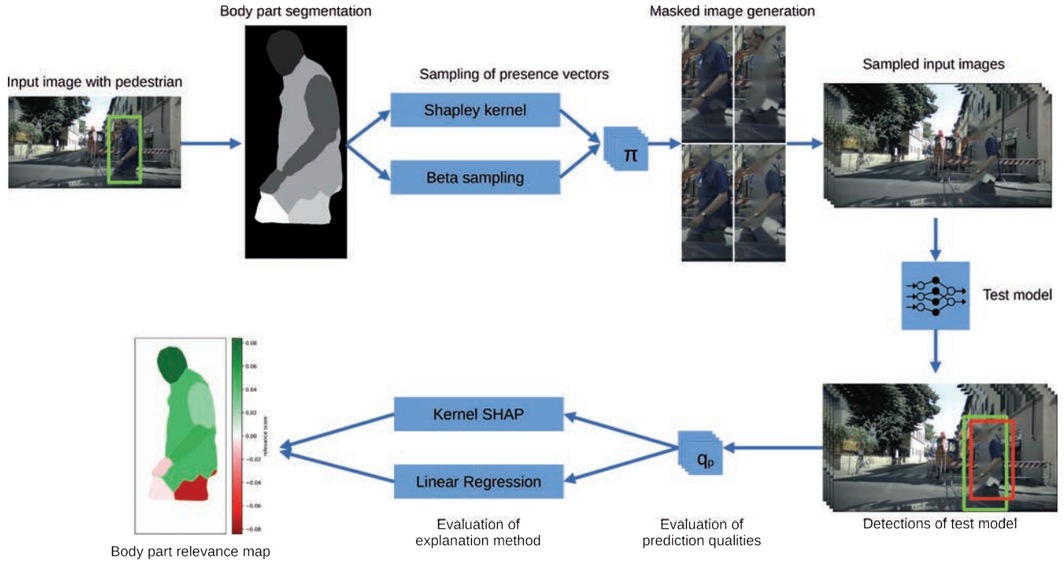

Overview sketch of our body part relevance assessment method. The pedestrian bounding boxes of a real-world street scene image is segmented into the visible body parts. Our model takes the image regions (superpixels) and blends them with a masked version of the individual body parts. Now, different instances of the pedestrian with different presences of body parts can be drawn and placed into the original image. Our test model under investigation then predicts the pedestrian again and we observe the quality of the output regarding classification confidence and bounding box overlap. KernelSHAP or our adapted explanation method then assesses the relevance of each body part visualized as a so-called relevance map.

Investigating the Effects of Augmentation Techniques on Interpretability and Performance of Pedestrian Detection Models

Kostadin Cvejoski, FhG IAIS

Pedestrian detection is crucial in computer vision, impacting vehicles, robotics, and security. Challenges like occlusion persist. We explored data augmentation's impact on models like Faster R-CNN, SSD, and RetinaNet. Techniques included Horizontal Flip, AugMix, CutMix, Cutout, and Gaussian Noise on datasets like CityPersons, EuroCity Persons, and Caltech. Findings show no consistent link between model performance and interpretability. Some augmentations enhance performance but not interpretability, and vice versa. This highlights the need to balance performance and interpretability, directing future research in pedestrian detection.

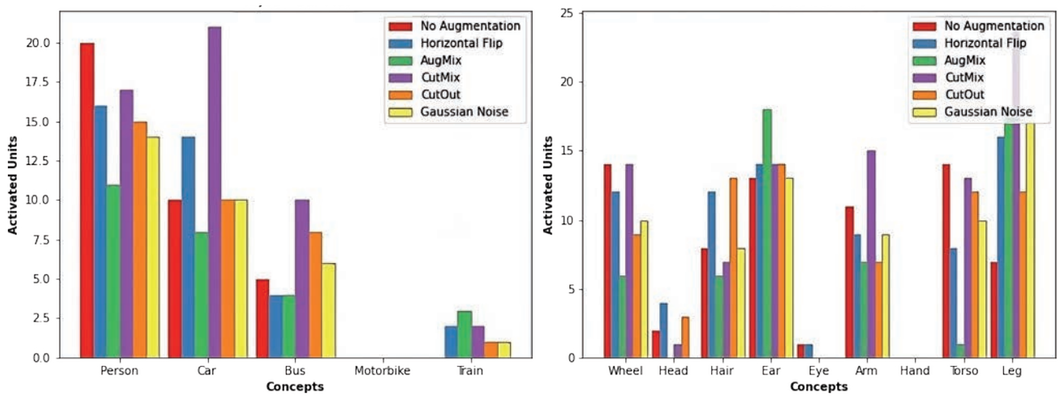

Number of units per concept for different augmentation techniques. (© Fraunhofer IAIS)

Static Scene Knowledge for Pedestrian Detection

Vera Stehr, Valeo

For some applications (e. g., Automated Valet Parking in a parking facility) cameras can be mounted statically on walls and ceilings. We took advantage of the static field of view of the camera and used an empty camera image without any objects to be detected as knowledge. A network architecture was designed which incorporates this knowledge to enhance pedestrian detection. As an alternative application, the empty image was used in a post processing step for a conformity check to reduce the false positive rate. We tested normalization techniques to diminish photometric differences between the empty and the live camera image.

Detection in the camera image (left), the same corresponding empty image (right) and the difference of the cutout calculated with the structural similarity (SSIM). (© Valeo)

Method to Fuse Map Layer and Pedestrian State Space to Improve 3D Detector Output

Ravi Kothari, Aditya Kumar Agarwal, Federica Paolicelli, AVL

We propose a method that combines the location of objects, specifically pedestrians, with their detected state space attributes to establish a decision boundary. The enhanced detections are deemed more reliable, leading to improved performance in real-world applications. Plausibility scores were calculated using attributes such as 3D detector output, map layers, and state space features. Results indicate improvements in mean average precision (mAP) values for both the datasets (nuScenes and KI Wissen UC1).

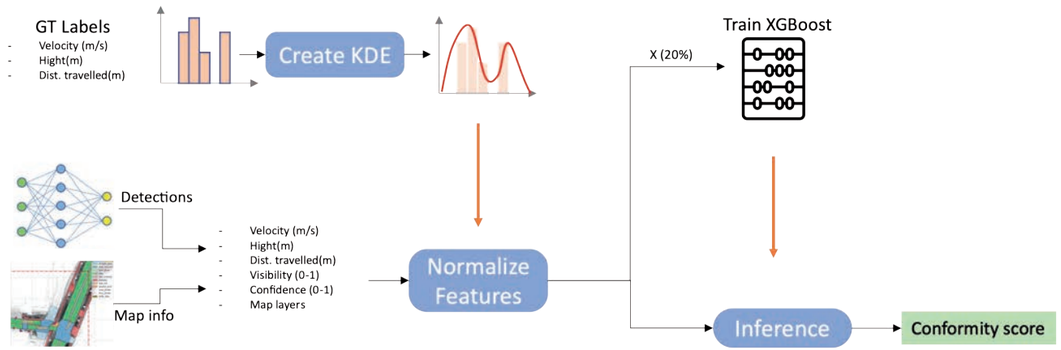

Overview of the Conformity Pipeline. (© AVL Software and Functions GmbH)

Knowledge Guided Occluded Pedestrian Detection

Julian Wörmann, fortiss

Reliable detection of safety critical occluded pedestrians is still challenging. In this work, we utilized concepts of attention to emphasize selected features. Attention is realized as a hypernetwork, that predicts reweighting coefficients for the feature maps. To associate semantics to the feature maps, an additional loss is incorporated in the training. By means of this penalty, we injected knowledge about relations between the features of different sub-classes into the model. Eventually, the conformance to this knowledge is exploited to scale the set of reweighting coefficients. During inference, the detector model can be used without any additional input required.

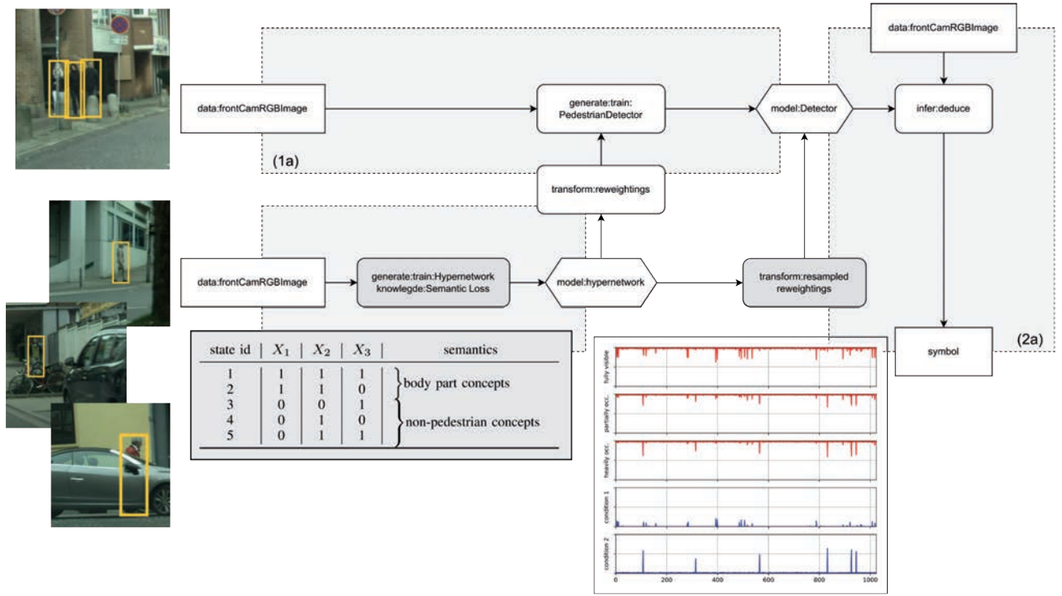

Knowledge Guided Detector. Semantic knowledge is used to penalize reweighting coefficients. The reweightings guide the feature selection process. (© fortiss GmbH)

Knowledge Aware Pedestrian Detection

Muhammad Ali Chattha, DFKI

Pedestrian detection is crucial for the safety and efficiency of autonomous and smart vehicles. While Deep Neural Networks (DNNs) have shown remarkable success, they cannot incorporate explicit information. We propose a knowledge-aware pedestrian detection framework that enables DNNs to integrate explicitly defined information in their learning process. We estimated Pedestrian Probability Distribution (PPD), identifying areas in an image with a high probability of pedestrians. These PPDs are integrated into the baseline DNN network at the intermediate and inference layers. This dual-stage fusion enhances overall accuracy and also expedites model convergence.

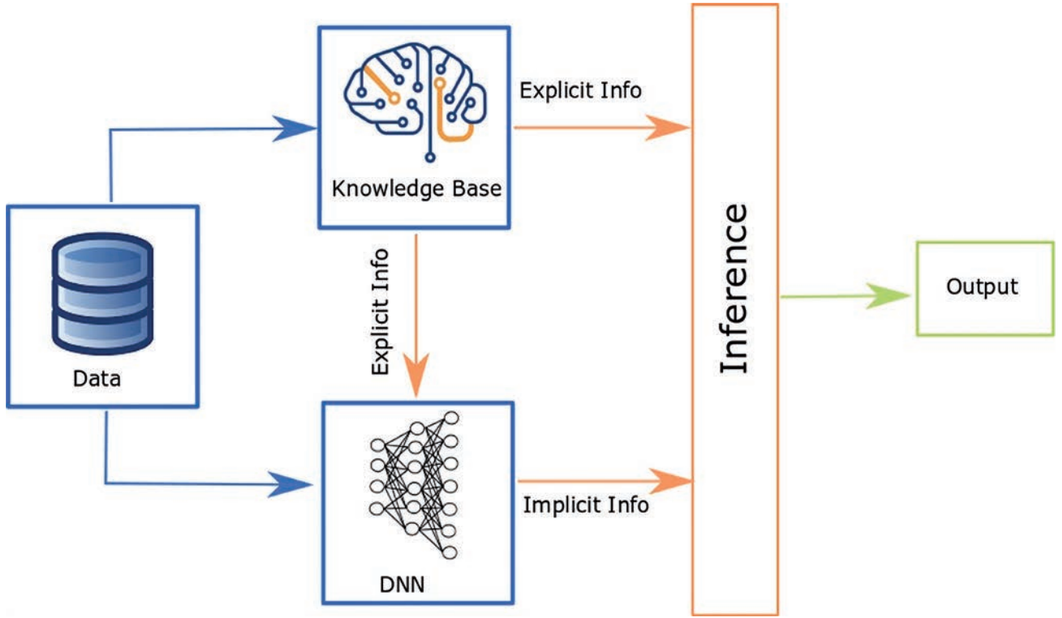

High-level overview of Knowledge Aware Pedestrian Detection Framework. The framework is agnostic to underlying knowledge and neural architecture and can work with a wide set of architectures.

Geo-Informed Conformity Check of Pedestrian Detection Models

Laura von Rueden, FhG IAIS

Although AI-based pedestrian detection models are already quite good, there is still a challenging trade-off between a low miss rate and few false positives. To improve this, we propose to check if the model output is consistent with geospatial knowledge. We have developed three modules for conformity checks: 1. Spatial perspective and pedestrian sizes, 2. Street map and pedestrian positions, 3. Combination.We find that all modules can help to identify false positive predictions, which in turn allows to fine-tune models for a lower miss rate.

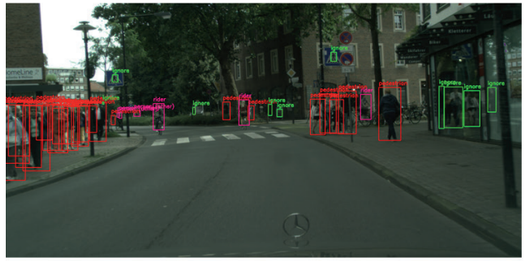

Figure 1

Figure 2

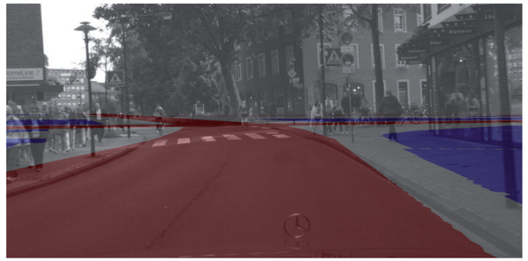

Figure 3

Figure 1: Example image from CityPersons data with GT BBoxes for pedestrian detection. Figure 2: Knowledge Module 1 - Spatial perspective to check pedestrian sizes. Figure 3: Knowledge Module 2 - Street map projected into camera perspective to check pedestrian positions.

Efficient Pedestrian Detection for Autonomous Driving

Abdul Hannan Khan, DFKI

Pedestrian Detection is vital for autonomous driving. It needs to be performant and efficient to provide effective and in-time perception. We revisited the architecture of the feature pyramid network and detection head, to remove the components that are computationally expensive and don‘t contribute directly to the learning capability of the network. We designed a novel detector that uses MLP-Mixers and works in batches of patches to boost performance. Furthermore, we benchmarked our proposed method on four well-established pedestrian detection datasets, where it beats the state-of-the-art as well as the human baseline on the Caltech dataset.

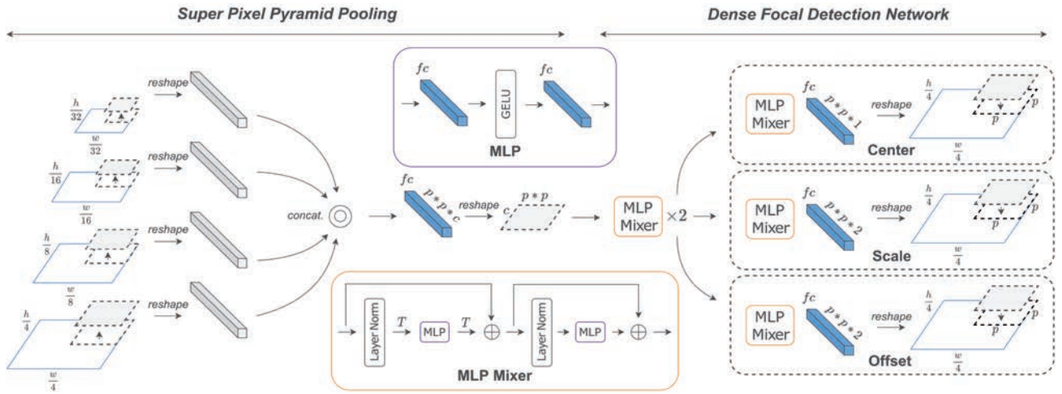

Model architecture of Localized Semantic Feature Mixers. (© DFKI GmbH)

Prior-Knowledge Based Deep-Learning Approach for Pedestrian Detection

Madhu Ameneni, Claas Brüß, Aiman Hsino, Ravikumar Ramaiah Chikkade, Leonard Schroven, Megha Vijayendra Rao, Nils Worzyk, Capgemini

We propose a deep learning model based on Mask R-CNN integrating prior knowledge inputs into the neural network architecture to increase training data efficiency, classification performance and reliability in challenging scenarios. The prior knowledge is represented by a semantic segmentation image, which is obtained as a 2D perspective on a previously optained 3D semantic segmentation map containing only static objects. A scenario dependent increase in object detection performance was observed. Evaluating system level metrics in full loop application, we observed interventions and readjustments earlier in comparison to a classic Mask R-CNN.

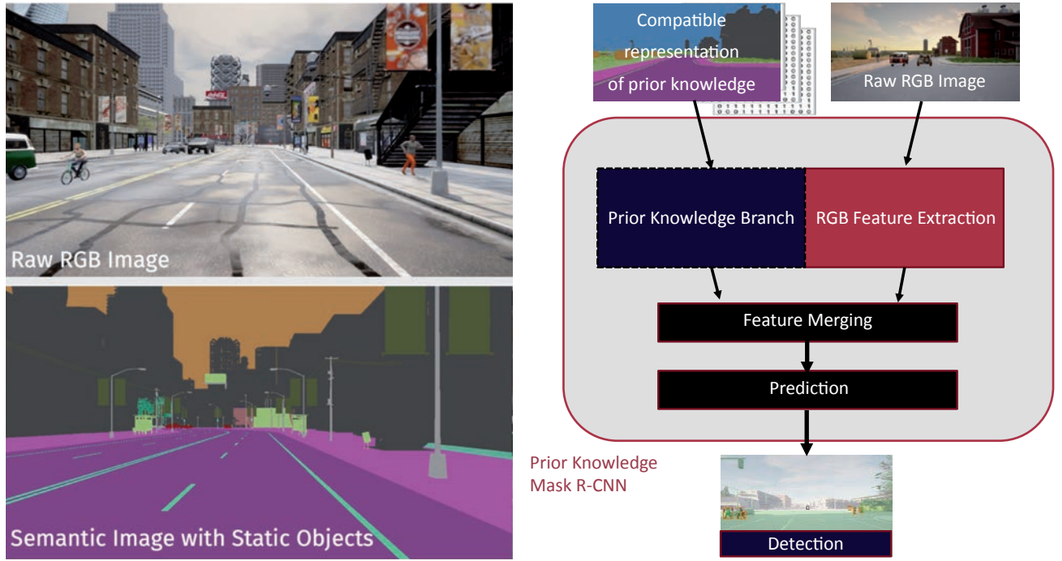

Semantic Prior Image to RGB Image Matching and Prior Mask R-CNN Architecture.

Performance Impact of Input Image Characteristics and Prior Knowledge Conformity in Prior Mask R-CNN

Madhu Ameneni, Claas Brüß, Aiman Hsino, Ravikumar Ramaiah Chikkade, Leonard Schroven, Megha Vijayendra Rao, Luis Wiedmann, Nils Worzyk, Capgemini

To understand the influence of various aspects of a scenario, a range of performance metrics for object detection were cross-correlated to image metrics. The impact of the applying of prior knowledge on performance was investigated to determine whether specific scenarios could be addressed with more reliability. High performance sensitivity to geometric misalignment of prior images and RGB sensor images was observed in case of altered RGB sensor perspectives. The performance gains through the application of prior knowledge in object detection can thus only be ensured in application scenarios with high prior knowledge conformity being maintained.

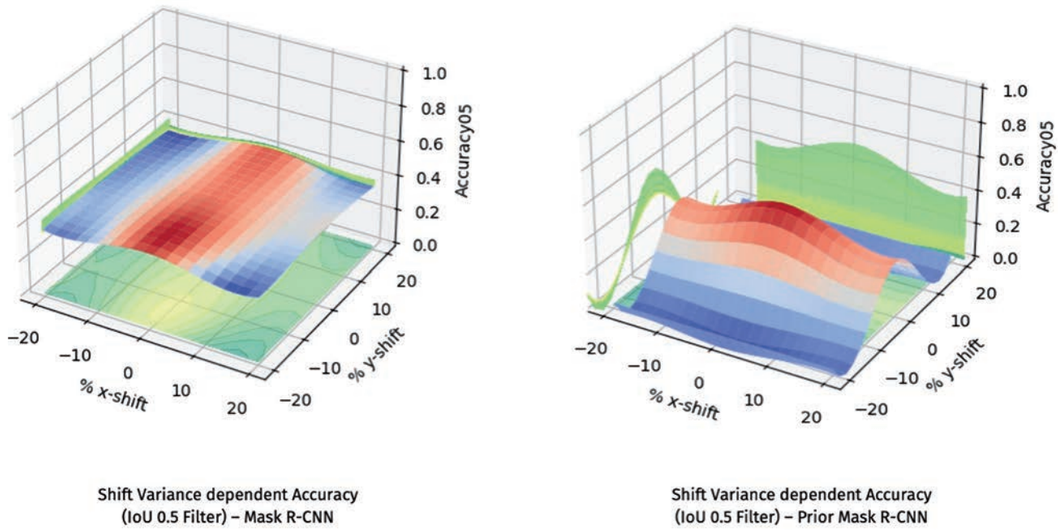

Shift Variance dependent Accuracy (IoU 0.5 Filter) – Mask R-CNN

Shift Variance dependent Accuracy (IoU 0.5 Filter) – Prior Mask R-CNN

Impact of degraded Prior Knowledge Conformity on Accuracy in classic Mask R-CNN and Prior Mask R-CNN.

Plausibility Verification For 3D Object Detectors Using Energy-Based Optimization

Abhishek Vivekanandan, FZI

Environmental perception obtained via object detectors have no predictable safety layer encoded into their model schema, which creates the question of trustworthiness about the system's prediction. This work aimed to verify 3D object proposals from a Neural Network by proposing a plausibility framework that leverages cross sensor streams to reduce false positives. The verification metric being proposed uses prior knowledge in the form of four different energy functions, each utilizing a certain prior to output an energy value leading to a plausibility justification for the hypothesis under consideration.

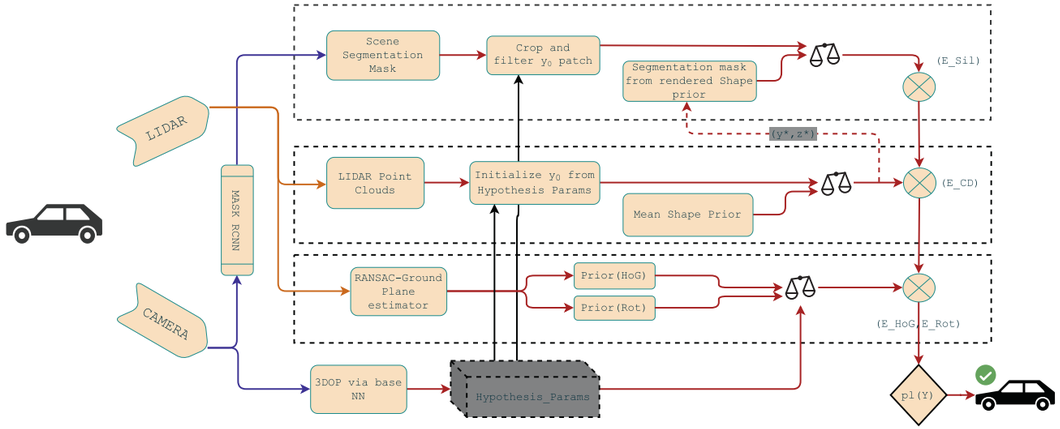

General architecture describing the flow of energy values which are combined to argue about an object's presence. y_0 represents the initial hypothesis. LiDAR and Camera inputs are primarily used to provide cross sensor data streams. Top: E_Sil, compares segmentation mask from MaskRCNN and rendered segmentation mask from the optimized outputs of CD energy function. Middle: E_CD, compares and optimizes a mean 3D shape prior with raw segmented point clouds within the hypothesis space. Bottom: E_HoG, E_Rot, compares ground estimates from a RANSAC regressor with the height and rotation of an initial hypothesis. (© FZI)

Object Detection Plausibility with Concept-Bottleneck Models

Mert Keser, Gesina Schwalbe, Azarm Nowzad, Continental
Alois Knoll, Technische Universität München

For automated driving (AD) safety, KI Wissen introduced a method for validating object detector (OD) outputs using a small, model-agnostic, robust, interpretable image classification model. The concept bottleneck model (CBM) used here is a deep neural network (DNN) with interpretable intermediate outputs, ensuring interpretability and robustness. We enhanced CBM's performance in domain-shift situations, by integrating trainable color-invariance filters. This upgraded CBM effectively detects hallucinated objects and false positives in AD OD, demonstrating zero-shot and few-shot domain adaptation capabilities, marking it as a viable solution for error monitoring in safety-critical applications.

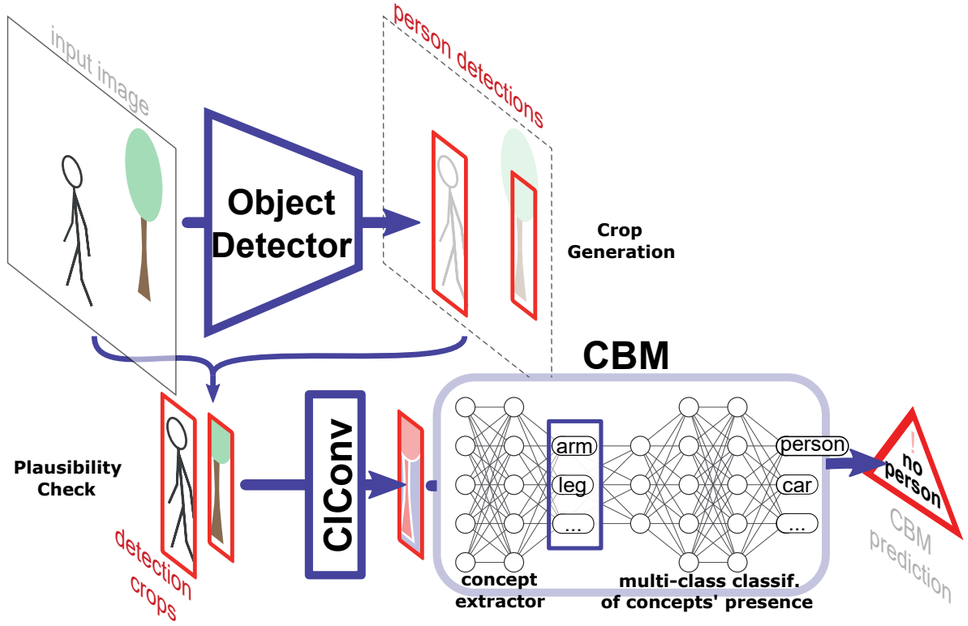

Interpretable, model-agnostic monitoring approach for identification of false positive person detections. It uses an interpretable concept bottleneck model (CBM) as independent classifier on the color-invariant representations (CIConv) of the object detections (here: for class person). (© Continental AG)

Digital Twin Creation and Evaluation of Color Dependent Object Detection

Toni Baric, Valeo

Through simulation, we can explore scenarios that result in an AI system failure and pinpoint the exact root cause of the failure. To conduct a valid simulation, a proper digital twin has been created from real-world assets: environment, road, car, and fisheye camera sensor. One scenario where an AI system produces inaccurate or false results is linked to different car colors. The same car model at the same location is detected with varying detection confidences. The findings highlight parametric weaknesses (such as color) of the Deep neural network (DNN) and enable optimization of the training process.

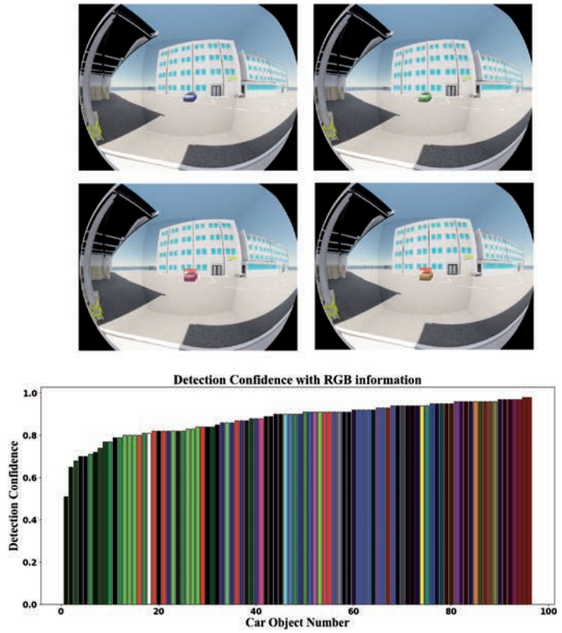

Car with blue and green colors detected using Yolact on a simulated fisheye camera view. Analyzing the object confidence distribution of the detected car across different color variations (RGB values) provides valuable insights. (© Valeo)

Multi-Modal Sensor Fusion for Robust Obstacle Detection

Syed Tahseen Raza Rizvi, Abdul Hannan Khan, DFKI

Integrating visual and LiDAR sensors enhances navigation, obstacle detection, and real-time decision-making, this work proposes a solution using multi-modal sensor data to identify obstacles. The approach leverages 3D LiDAR points to establish the road plane and mapping all the LiDAR points to identify objects. The detections from cross modalities are validated to avoid missing or false detections. The evaluation on the nuScenes dataset demonstrates improved performance, highlighting the method's robustness, ability to identify unknown objects, and a novel object detection metric ($mAP_{critical}$) addressing the criticality of objects in autonomous driving scenarios.

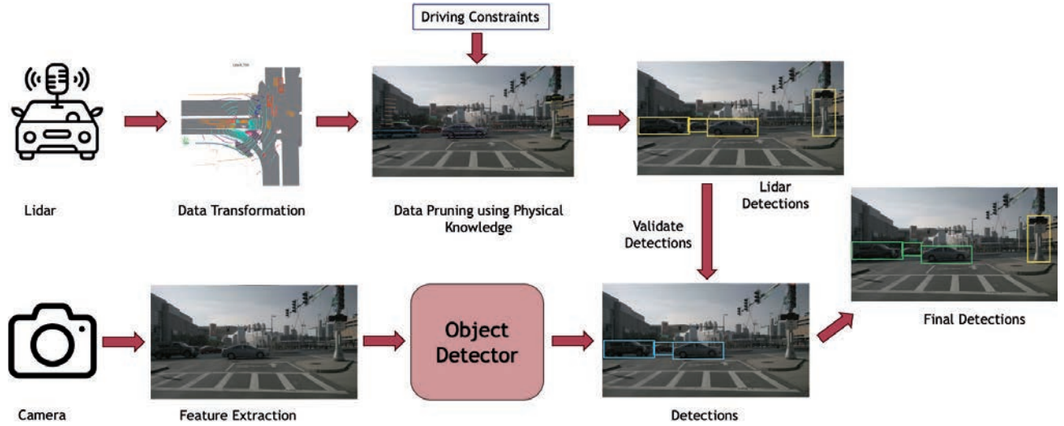

Multi-modal Sensor Fusion Pipeline for Obstacle Detection. (© DFKI GmbH)

Leveraging Knowledge for Traffic Sign Detection

Gurucharan Srinivas, DLR

Our proposed hybrid object detector utilizes semantic correlations, like co-occurrence measures between traffic sign pairs, extracted from training annotations as an undirected knowledge graph. The results reveal its detection metrics is on par with the baseline. To mitigate the incompleteness in our knowledge graph and enhance its efficacy we delved into more sophisticated knowledge representation. We mined association rules with the a priori algorithm, translating them into First-Order Logic (FOL) statements. Our ongoing efforts aim to further integrate FOL with co-occurrence insights via Logic Tensor Networks, enhancing detection capabilities.

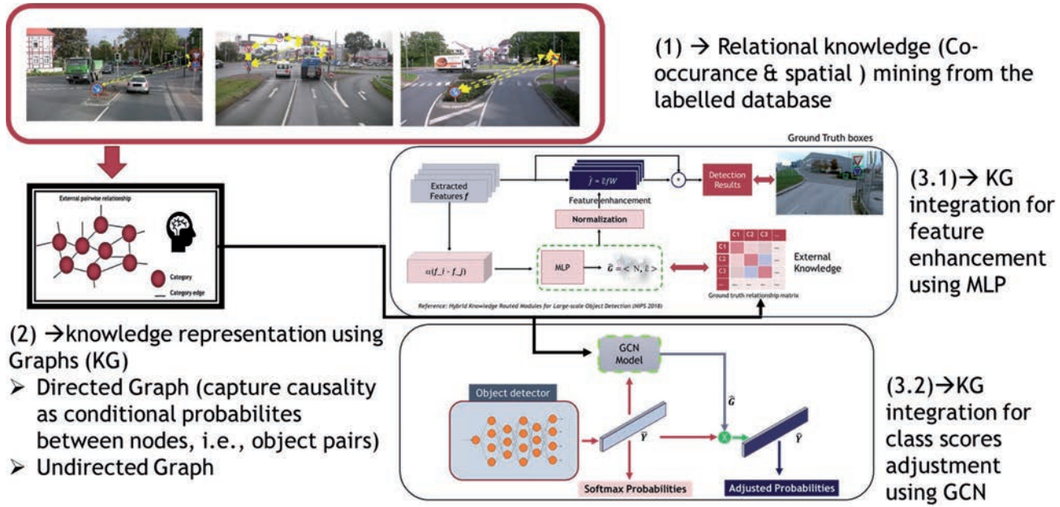

The image illustrates the process of enhancing object detection features, specifically for traffic sign recognition, through the utilization of relational knowledge. Initially, step one involves analyzing the dataset to identify semantic relationships, focusing on patterns of how objects appear together. Following this, step two presents the formulation of this knowledge into two types of graphs: directed and undirected, to visually represent these relationships. The final step, step three, integrates this knowledge to refine and improve the detection capabilities, both by enhancing the features of the detected objects and adjusting the scores of detection accuracy.

Trajectory Prediction and Motion Planning

Informed Motion Planning 58
Knowledge Integrated Plausible Motion Forecasting 60
An Empirical Bayes Analysis of Object Trajectory Representation Models 62
An Efficient and Robust Multi-Modal Trajectory Predictor Baseline for Autonomous Driving 64
Expert Informed Trajectory Prediction 66
Integration of A Priori Knowledge Using a Causal Model of Vehicle Trajectories 68
Causality-Driven Checks of the Physical Conformity of Vehicle Trajectories 70
Tracking and Trajectory Prediction Using Scene Semantics 72
Learning-Aided Warmstart of Model Predictive Control in Uncertain Fast-Changing Traffic 74
Model Predictive Control Under Temporal Logic Specifications 76
An Outlier-Robust and Efficient Bayesian Filter and Smoother 78
Behavior-Conditioned Driving Policies 80

Trajectory Prediction and Motion Planning

In a modular software stack for autonomous driving, the perception module periodically provides a sensor abstracted representation of the current environment. This information is fused into an environment model that forms the basis of situation interpretation, prediction of the future development of the traffic scene and subsequent planning of the vehicle's future path and motion. Several unique challenges arise in this task. First, the inputs coming from the perception module are inherently uncertain due to technical limitations of the sensors and software. Second, even given perfect sensors and software, traffic participants are (human) agents with unobservable intentions who may behave irrationally and can therefore be described only in the language of probability. Third, the interactions between all traffic participants are governed by complex feedback loops as agents communicate their intentions and negotiate conflicts via their behaviors.

Traffic is a well-regulated and reasonably orderly affair, most of the time at least. First, all traffic participants must obey the laws of physics which constrain and limit their motion options. Second, traffic is governed by traffic laws that most traffic participants obey most of the time. Third, traffic takes place in a highly structured environment that is designed for efficiency and reduced conflict.

Within KI Wissen, we developed methods that leverage this kind of knowledge and integrate it into machine learning models. A wide range of different approaches could be identified and assessed. The project focused on methods based on statistical prior information that help make tracking more robust or improve the accuracy of prediction tasks. Furthermore it employed formal methods that leverage scene semantics, causal relationships, or temporal logic, while others tackled the motion planning problem more directly through behavior conditioning and physics informed reinforcement learning or a direct combination of learning and optimization-based methods.

Common to all approaches is their success in showing that combining model based and data driven meth-ods can make the task of situation interpre-tation, prediction, and motion planning more accurate, more efficient, and ultimately safer.

Informed Motion Planning

Daniel Bogdoll, Tim Joseph, Christian Hubschneider, FZI

We integrated knowledge about traffic rules and physical constraints for improved trajectory planning based on Reinforcement Learning (RL). We analyzed real-world driving data to determine a plausible action space for our RL Agent and integrated knowledge into the action space and reward design. In situations where controlled rule exceptions were necessary our agent outperformed our original baseline.

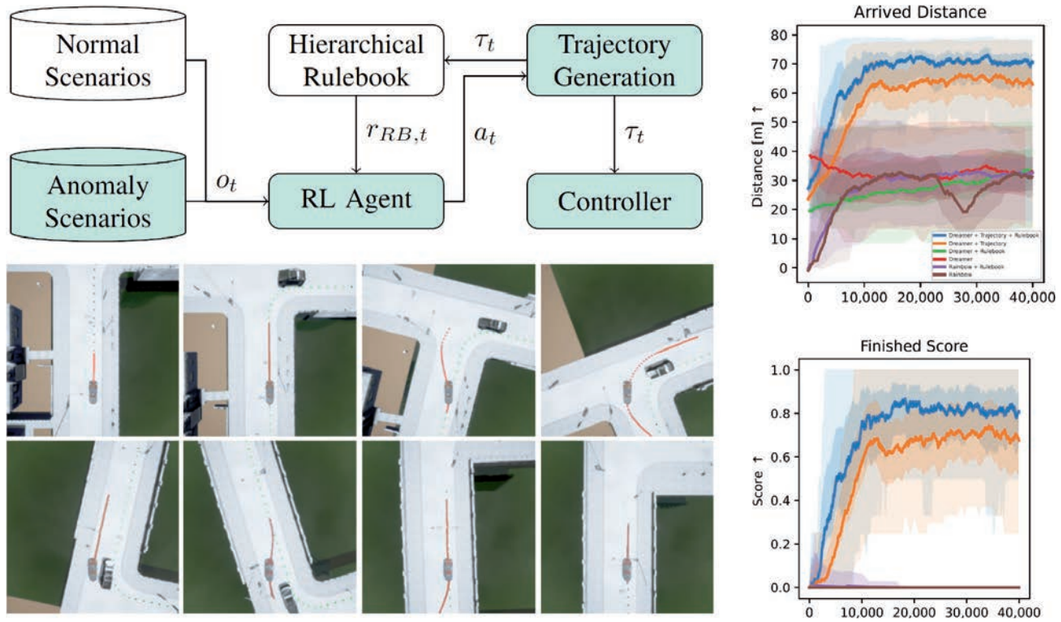

Informed Reinforcement Learning for Situation-Aware Traffic Rule Exceptions.

Knowledge Integrated Plausible Motion Forecasting

Abhishek Vivekanandan, FZI

Current trajectory forecasting approaches primarily focus on optimizing a loss function with a specific metric, which can result in predictions that do not adhere to physical laws or violate external constraints. Our work incorporates explicit knowledge priors that allow a network to forecast future trajectories in compliance with both the kinematic constraints of a vehicle and the geometry of the driving environment. Our proposed method is designed to ensure reachability guarantees for traffic actors in both complex and dynamic situations by incorporating knowledge priors into the training process.

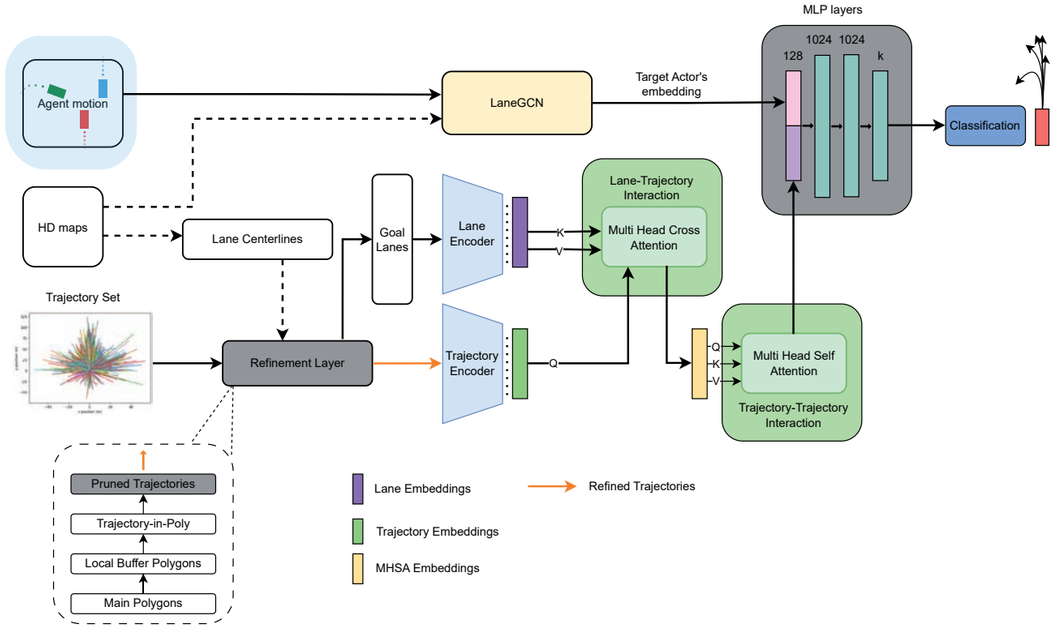

High-level overview of the architecture. The Refinement layer takes two inputs: 1. Lane centerline points in the map coordinate system for a given query window (in meters) and 2. Trajectory sets with ϵ = 2m coverage. The refinement layer produces feasible/pruned trajectories by constructing a lane boundary given the lane-centerline points. It also produces a list of possible goal positions or reachable lanes where the target actor could reach for within a given prediction horizon. LaneGCN is used as a backbone architecture and outputs an embedding representation which is then concatenated with the lower parts of the network. (© FZI)

An Empirical Bayes Analysis of Object Trajectory Representation Models

Yue Yao, Jörg Reichardt, Continental
Daniel Göhring, Freie Universität Berlin

We present an in-depth empirical analysis of the trade-off between model complexity and fit error in modelling object trajectories via simple linear combinations of polynomial basis functions. Our methodology estimates observation noise and prior distributions over model parameters from several large datasets. Incorporating these priors can then regularize prediction models. Analyzing vehicles, cyclists, and pedestrians, results show that linear models do represent real-world trajectories with high fidelity at very moderate model complexity. This suggests the feasibility of using linear trajectory models in motion prediction systems with inherent mathematical advantages.

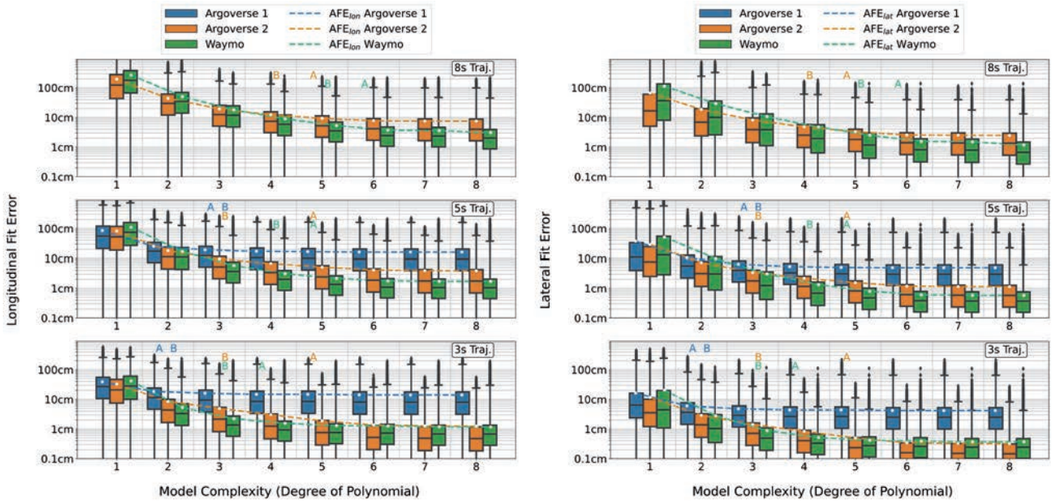

The longitudinal (left) and lateral (right) fit error of models for vehicle trajectories with T ∈ [3s, 5s, 8s]. "A, B" denote the model complexity n=n̂ that maximizes AIC and BIC, respectively. The upper whisker denotes the 99.9% percentile. (Ref: Yao et al., An Empirical Bayes Analysis of Object Trajectory Representation models, in proc. IEEE ITSC, 2023, arXiv:2211.01696, © Continental AG)

An Efficient and Robust Multi-Modal Trajectory Predictor Baseline for Autonomous Driving

Yue Yao, Joerg Reichardt, Continental | **Shengchao Yan,** Universität Freiburg
Daniel Goehring, Freie Universität Berlin

Safe and comfortable driving requires to predict future actions of other traffic participants. Prediction algorithms should combine accuracy and speed with robustness against changes in sensor-setup and map representation. To this end, we present a trajectory prediction algorithm based on polynomial representations for road geometry and object trajectories on both the input and the output side of the model. With much smaller model size, training effort, and execution time, we reached near state-of-the-art performance (Rank 24 in Argoverse2 Leaderboard) for in-dataset validation and significantly improved robustness in out-of-distribution, i.e. cross-dataset, testing.

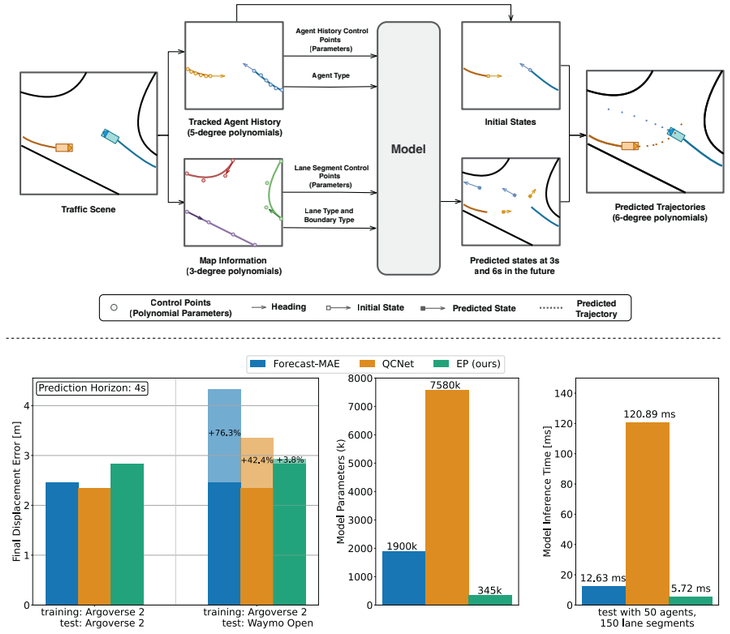

Model architecture and quantitative performance measures of our approach.

For further details, see Yao et al., Everything Polynomial: An Efficient and Robust Multi-Modal Trajectory Predictor Baseline for Autonomous Driving, forthcoming (© Continental AG)

Expert Informed Trajectory Prediction

Christian Schlauch, Christian Wirth, Continental

Expert knowledge, e.g. about traffic regulations, can be highly informative for trajectory predictions. However, its qualitative nature allows for many high-risk exceptions. Our proposed probabilistic informed learning approach integrates expert knowledge into probabilistic deep learning models, without disregarding exceptions, by leveraging regularization-based continual learning techniques. By integrating prior drivability knowledge into two state-of-the-art trajectory prediction models, we demonstrated substantial increases in data-efficiency, outperforming non-informed and informed baselines in low data regimes on two public datasets.

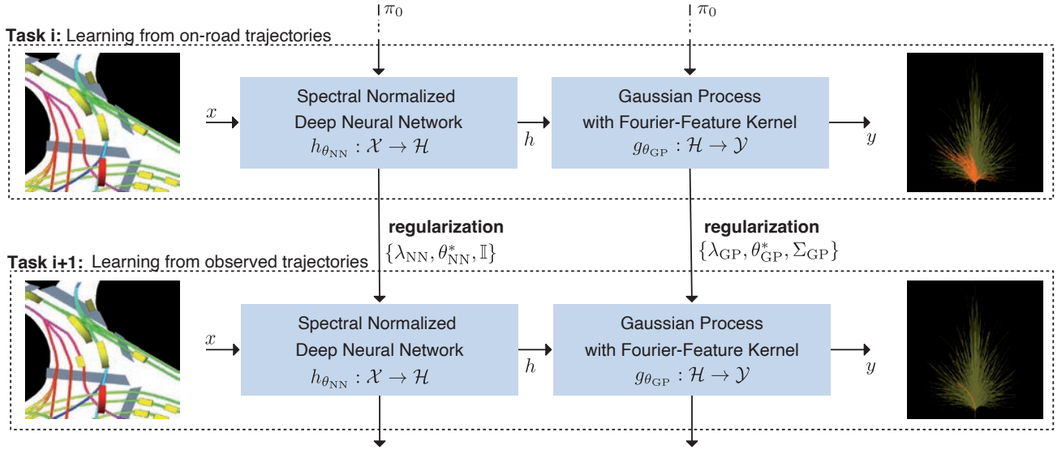

Our approach applied to the CoverNet-SNGP model. In the knowledge task, a multi-label classification is trained on on-road trajectories, with labels derived from drivable area map data. In the conventional task, a multi-class classification is trained on observed trajectories, regularized by the previously learned prior distribution. (©Continental AG). Further details in Schlauch et. al. Informed Spectral Normalized Gaussian Processes for Trajectory Prediction 2024.

Integration of A Priori Knowledge Using a Causal Model of Vehicle Trajectories

Christian Brunner, Himanshu Agarwal, e:fs

Automated driving functions based on machine learning face the challenge of reacting robustly to scenarios that are scarce in the training data. Causal models that encode physical knowledge offer potential to address this issue. We developed a structural causal model (SCM) to formalize the generative process of vehicle trajectories. Causal queries are used to augment the training data for supervised or reinforcement learning (RL) algorithms. We demonstrated our concept on a model-based RL-agent by embedding our SCM into its training. Our evaluations show that knowledge integration significantly improves the performance on scenario variations compared to the baseline agent.

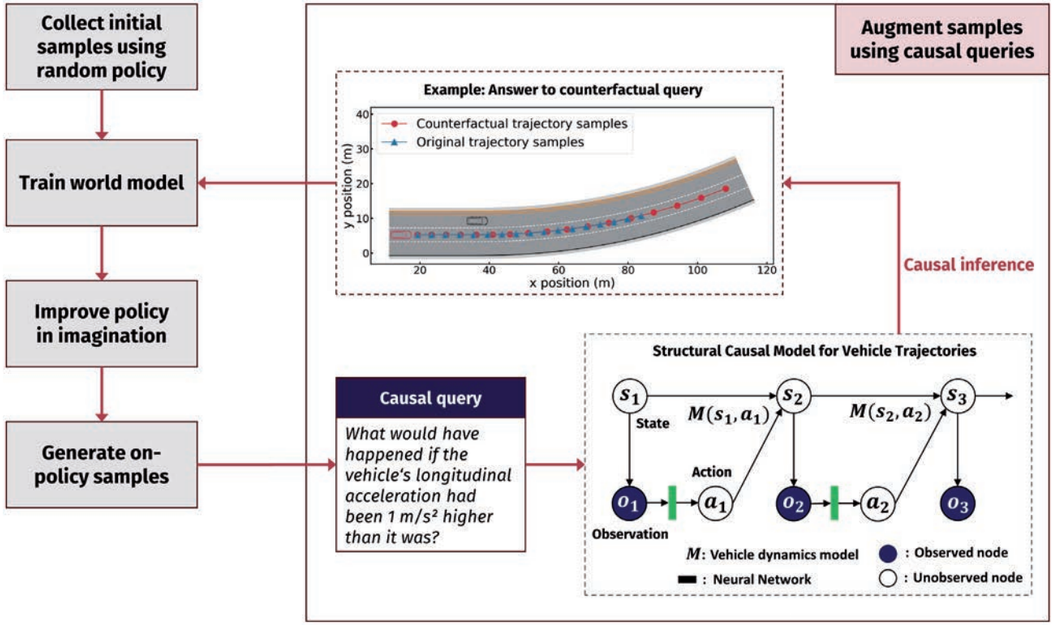

Knowledge integration into the training workflow of a model-based reinforcement learning agent using causal queries for data augmentation. (© e:fs TechHub GmbH)

Causality-Driven Checks of the Physical Conformity of Vehicle Trajectories

Himanshu Agarwal, Christian Brunner, e:fs

Robustness and explainability of machine learning modules are of paramount importance in automated driving. Erroneous perception of vehicle trajectories can trigger dangerous driving decisions. We developed a structural causal model (SCM) to check the physical conformity of perceived trajectories. The knowledge of vehicle dynamics is explicitly integrated into the SCM via equations governing the laws of physics. We evaluated our conformity check methods on real-world trajectories with artificial perturbations added to the observations. Our methods offer reliable performance in detecting unreasonable trajectories even when they are just partially perceived.

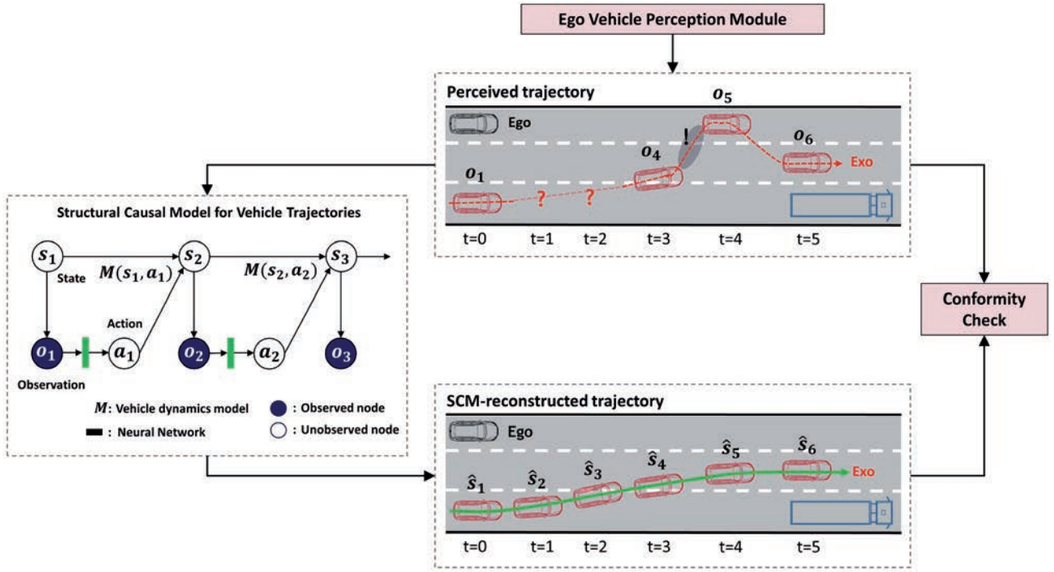

Illustration of our proposed conformity check concept to detect physically implausible vehicle trajectories by means of a structural causal model. (© e:fs TechHub GmbH)

Tracking and Trajectory Prediction Using Scene Semantics

Nils Kornfeld, DLR

To be able to predict the future motion of road users it is important to create a reliable model of the previous motion of all dynamic objects present in a traffic scene. Systems used for tracking the previous motion and predict the future evolution of trajectories show typical shortcomings in challenging situations like full or partial occlusion. Incorporating knowledge like scene semantics about the environment at hand into the multiple object tracking algorithm can help to overcome some of the system's inherent shortcomings.

Exemplary Tracking and trajectory prediction results of one partially and one fully occluded object. (© Deutsches Zentrum für Luft- und Raumfahrt, e. V.)

Learning-Aided Warmstart of Model Predictive Control in Uncertain Fast-Changing Traffic

Mohamed-Khalil Bouzidi, Yue Yao, Jörg Reichardt, Continental | **Daniel Göhring,** Freie Universität Berlin

Purely learning-based motion planning and control lack reliability for safety-critical applications. Model Predictive Control (MPC) offers safety and feasibility guarantees but faces issues: converging to undesired local optimal trajectories and non-convergence in unknown fast-changing environments. To address this, we propose a Learning-aided Warmstart Framework. It employs a neural-network-based multimodal predictor which is anyways needed to forecast traffic participants' trajectories and reutilizes this predictor to generates multiple proposals for the ego vehicle in parallel. The different ego trajectory modes are used to identify multiple distinct local minima.

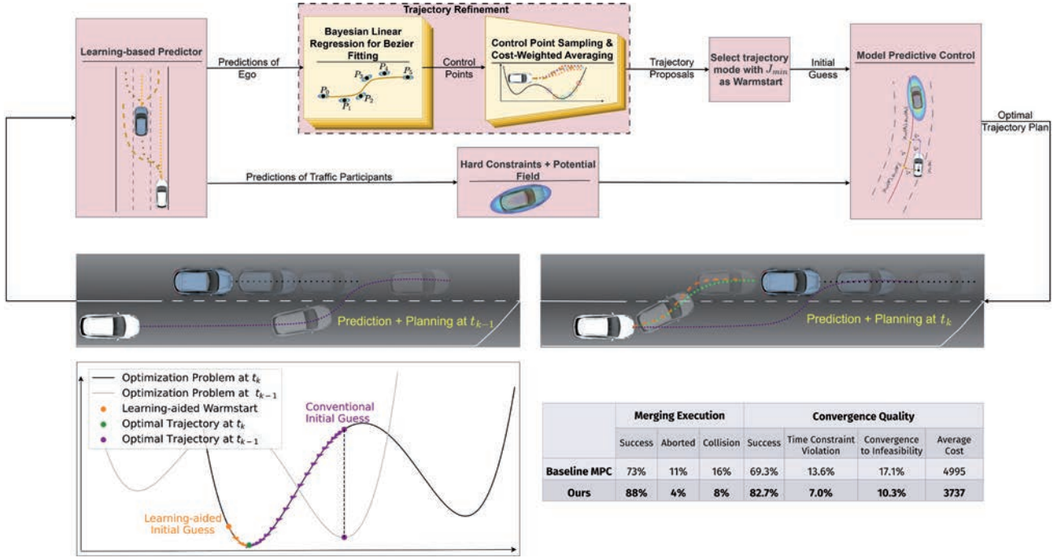

	Merging Execution			Convergence Quality			
	Success	Aborted	Collision	Success	Time Constraint Violation	Convergence to Infeasibility	Average Cost
Baseline MPC	73%	11%	16%	69.3%	13.6%	17.1%	4995
Ours	**88%**	**4%**	**8%**	**82.7%**	**7.0%**	**10.3%**	**3737**

Architecture of our Learning-aided Warmstart Framework (top), Intuition behind our approach (bottom left) and results of a quantitative evaluation over test scenarios (bottom right). For further details, see Bouzidi et al., Learning-Aided Warmstart of Model Predictive Control in Uncertain Fast-Changing Traffic, ICRA 2024, arXiv:2310.02918, © Continental AG)

Model Predictive Control Under Temporal Logic Specifications

Etienne Bührle, Ömer Sahin Tas, Hendrik Königshof, FZI

Urban intersections exhibit continuous agent dynamics as well as discrete logical constraints, making the control task challenging. Prior work has put forward the use of temporal logics for formalization and monitoring of traffic rules. We propose a model predictive control scheme that maintains formal correctness guarantees. We constructed a hybrid automaton by taking the product of the continuous multi-agent system and the Büchi automaton of the temporal logic rule set. We formulated the resulting control problem as a mixed-integer optimal control problem and solved it for various intersection scenarios with different combinatorial behavior modes.

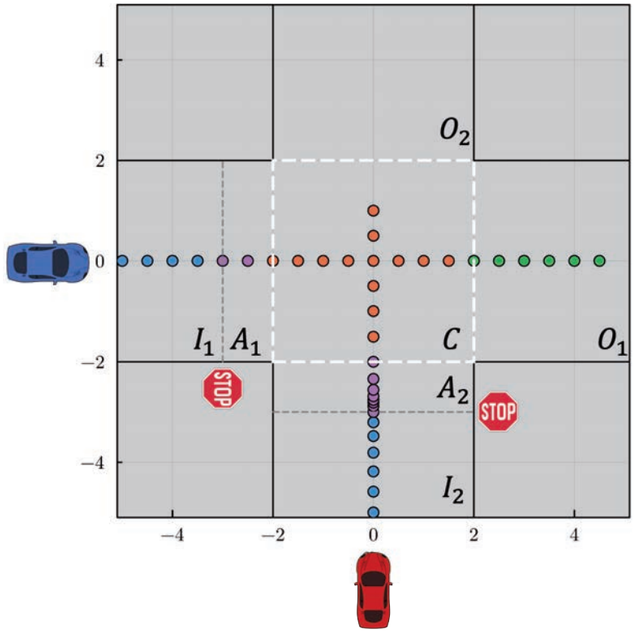

Agent interaction on an uncontrolled four-way stop. (© FZI Forschungszentrum Informatik)

An Outlier-Robust and Efficient Bayesian Filter and Smoother

Jonas Neuhöfer, Jörg Reichardt, Continental

Bayesian filters are essential for many applications in the fields of robotics and sensor-data fusion. Closed-form analytic filtering equations, however, only exist under the idealized assumption of linear dynamic and observation models plus Gaussian noise. This fails to account for outliers often present in real-world applications. We mathematically derived a robust filter and smoother based on heavy-tailed Student's t-distributions that exhibit low time complexity and require only a single hyperparameter adapting to the expected frequency of outliers. We demonstrated an initial implementation of our methods that outperforms comparable state-of-the-art methods.

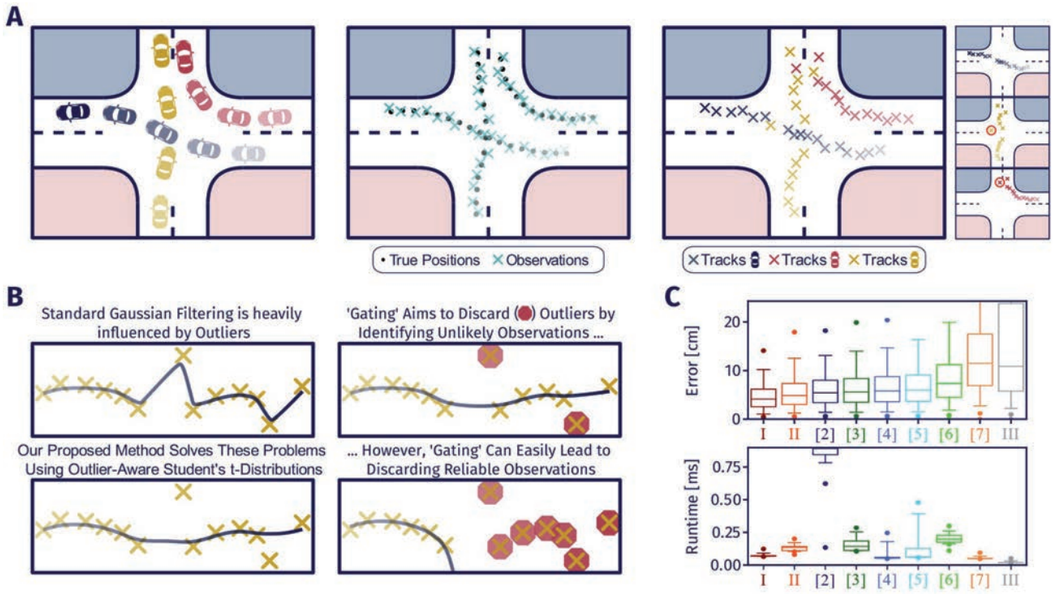

A) An Example of how outliers can easily be produced when observing a complex environment. B) Problems with existing methods in use C) Comparison of our method to other comparable state-of-the-art methods. I is our method when two usually unknown values are known, II is our general method, and III is the general Kalman Filter. [2-7] are state of the art methods from 2009-2020. (© Continental AG)

Behavior-Conditioned Driving Policies

Tim Joseph, FZI

Learning driving policies to control autonomous vehicles via reinforcement learning (RL) offers a solution to learn optimal driving behavior directly from sensor data. However, designing a reward function that leads to a driving policy that works in any situation has not yet been achieved. Instead, we suggest a different direction: we propose a model-based RL agent that learns a conditional driving policy by simulating behavior for many different reward functions in imagination using a world model. We do so by randomly sampling parameters that shape the reward function and optimizing an actor-critic policy that is conditioned on these parameters. The result is a policy that can be controlled at run-time to execute different behaviors dependent on external input.

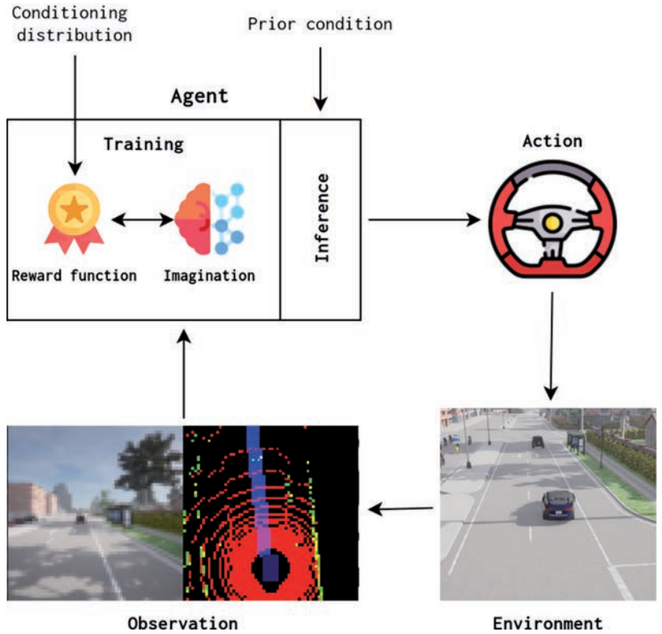

Our agent receives RGB and LIDAR sensor data from the environment and generates driving-decision in an end-to-end fashion. The agent „imagines“ outcomes and scores them according to different reward functions. Thus, it can drive according to any behavior-conditioning that it has imagined in training.

Formal Methods and Knowledge Representation

Harnessing Symbolic Knowledge Extraction and Utilization for Informed Decision-Making 86
Model Agnostic Local Analysis with Latent Attacks (MALALA) 88
A Concept to Support AI Models by Using Ontologies 90
Knowledge-based Comparison of CNNs 92
AI Net Observer 94
Legal Norm Formalization – An Overview 96
From Legal Documents to Formalized Rules with Large Language Models 98
Data-Efficient Automated Rule Formalization Framework. 100
Traffic Sequence Charts for Knowledge Formalization and AI Application 102
Risk Assessment and Adaptive Decision-Making for Autonomous Vehicle Safety 104
Monitoring Traffic Rule Conformance and Integrated Knowledge Impact 106
Runtime Monitoring using TSC-Based Knowledge Building Blocks 108
Scene Understanding for Autonomous Driving Using Visual Question Answering 110
Consistent Joint Action and Explanation Prediction for Autonomous Driving 112
Generating Training Data from Formalized Traffic Dynamics 114
TSC-Based Generation of Concrete Scenarios for Synthetic Data 116
Local Concept-Based Explanations for Object Detectors 118

Formal Methods and Knowledge Representation

The goal of KI Wissen is to integrate prior domain knowledge into AI models to overcome limitations of purely data-driven approaches. The chapter highlights methods addressing knowledge extraction, representation, formalization, and conformance checking. In terms of knowledge extraction, KI Wissen investigates methods that extract concepts to improve the understanding and identify potential weaknesses of AI models. This includes extending datasets for consistent joint action and explanation prediction, utilizing AI net observers to provide feedback, and using concept activation vectors to compare neural network models based on semantic concepts. Additionally, a model agnostic local analysis with latent attacks is used to improve pedestrian detection in rare situations. Knowledge formalization aim to represent knowledge in a machine-readable format to support AI training or conformity checks. Handcrafted approaches extend legal concepts and rules based on explicit and implicit knowledge, while AI-based approaches use large language models to translate legal documents and traffic rules into a machine-readable format. Rule compliance is achieved through the integration of semantic role and chain-of-thought reasoning. To improve AI training, multiple methods formalize traffic scenarios. These methods enable the generation of synthetic training data by formalizing physical and mathematical knowledge within traffic scenarios. A method based on Traffic Sequence Charts demonstrates that the infusion of such knowledge reduces training efforts and improves AI's performance. Addressing conformity, one method uses a reasoner equipped with ASAM OpenXOntology and formalized traffic rules to obtain a robust adaptive decision-making. Another method propagates a knowledge conformance framework to automate the runtime verification of knowledge conformity, leading to accelerated training and improved performance in trajectory/behavior planning. A neuro-symbolic architecture combines LLM agents, symbolic reasoners, and neural networks to improve decision-making in au- tonomous driving systems, ensuring compliance with traffic rules and regulations.

Overall, the pre-sented methods increase the generalization and comprehensibility of AI functions, reduce the amount of required training data, and improve efficiency of training in terms of duration and performance.

Harnessing Symbolic Knowledge Extraction and Utilization for Informed Decision-Making

Ya Wang, Adrian Paschke, FhG FOKUS

Decision-making is a safety-critical component of autonomous driving systems, requiring a comprehensive understanding of traffic scenarios and skilled reasoning based on world and normative knowledge. To address neural networks' issues with transparency and reliability, we introduce a neuro-symbolic architecture comprising two sections: Rule-Compliant Decision Making, which uses Large Language Model agents for data queries and symbolic reasoning for actions, and Symbolic Rule Extraction, analyzing the activations within trained neural networks to effectively extract rules for explainability.

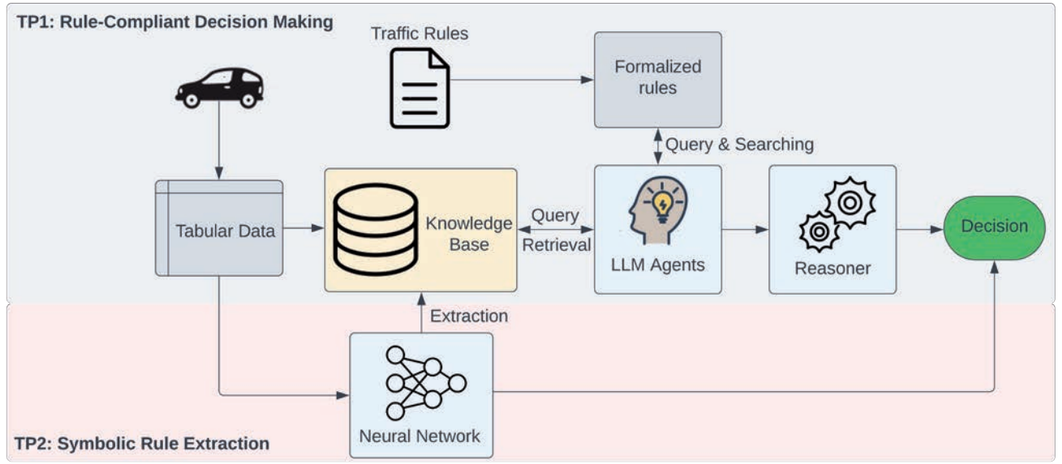

The neuro-symbolic architecture for rule-compliant decision making and rule extraction in autonomous driving. (© Fraunhofer FOKUS)

Model Agnostic Local Analysis with Latent Attacks (MALALA)

Daniel Kaulbach, Luca Bruder, Dina Krayzler, Leonard Rosen, Jan-Hendrik Clausen, Alexander Thamm

Pedestrian detection can fail in unusual situations that were not seen before in the training data, leading to critical errors. Our MALALA approach can be used for detecting these situations. We are harnessing the generative characteristics of variational autoencoders, to generate exactly the borderline cases of pedestrian images similar to existing ones that are confusing the model. This shows for every pedestrian image how stable its detection is, or how quickly it will not be detected any more if small changes are applied. Saliency maps accompany this process, spotlighting areas crucial for model failures. Together, they aid in evaluating problematic pedestrian detection cases.

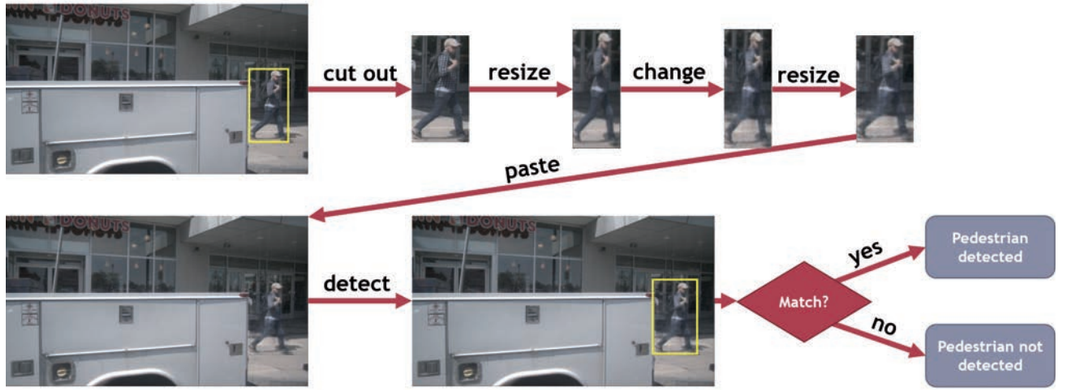

MALALA is changing pedestrians in images to check, if this affects their detection. (© Alexander Thamm GmbH)

A Concept to Support AI Models by Using Ontologies

Maximilian Grabowski, BASt | **Ya Wang,** FhG FOKUS

We propose a concept to support the decision making of the AI of automated vehicles by adding expert knowledge to the AI system. If it is strengthened with formalized knowledge using symbolic reasoning, the tracking of decision making is possible for many situations. The German Technical Specification on Lane Markings from the Road and Transportation Research Association is a suitable candidate to have its knowledge formalized into an ontology. Integrating this ontology into the ASAM OpenXOntology and modifying it with scenario specific concepts and simple rules allows to execute pre-defined queries for an improved situational awareness rule compliance and transparency.

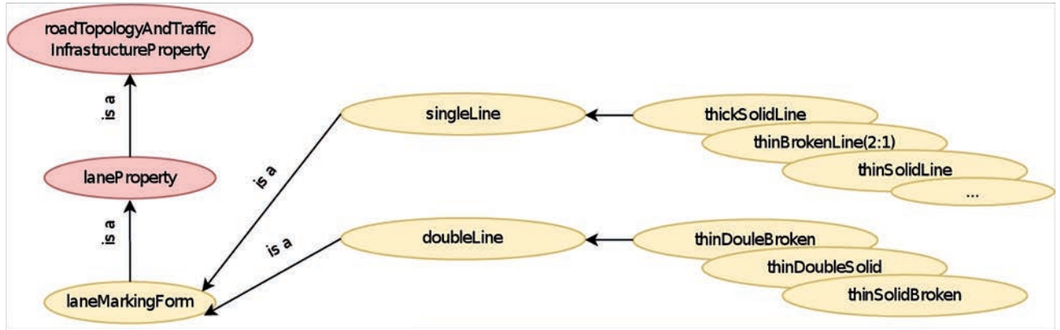

Modified structure of the lane marking form. (© BASt | Fraunhofer FOKUS)

Knowledge-based Comparison of CNNs

Georgii Mikriukov, Christian Hellert, Continental

We explored the potential of comparing neural network models using semantic concepts, diverging from conventional methods that rely on performance or error metrics. Instead, we employ Concept Activation Vectors (CAVs), which correspond to internal representations of models, for the comparative analysis. Utilizing CAVs, we quantified similarities in the feature space surrounding these concepts. The approach involves measuring concept attributions across layer pairs and entire networks, enabling the assessment of knowledge similarity and taking a promising step towards informed, knowledge-based model selection.

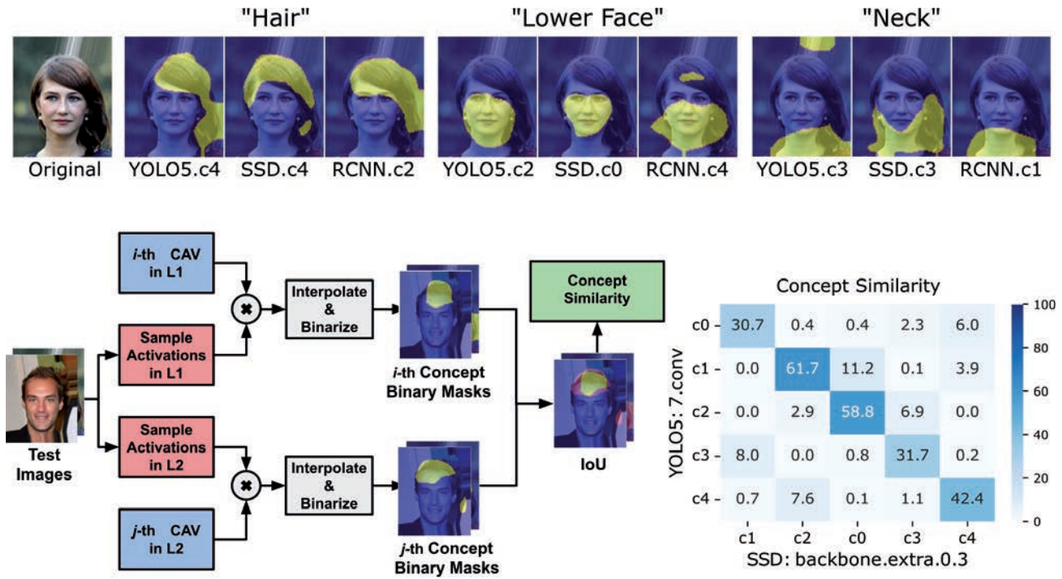

1. Examples of semantic concepts discovered in CelebA dataset (top), 2. The comparison pipeline for semantic concepts in different layers of CNN models (bottom left). 3. The concept similarity matrix for layers 7.conv and backbone.extra.0.3 of YOLOv5 and SSD networks (bottom right). (© Continental AG)

AI Net Observer

Johann Kelsch, DLR

In order to optimize AI Training in a Human-AI-Training System, it is considered necessary to support the human developer‘s decision-making in respect to the choice of data and machine learning methods by using an AI Net Observer. Within the Human-AI-Training System, the AI Net Observer is connected to the human developer via its frontend and can provide feedback on the efficiency of the training and after that of the system usage and therefore optimize the AI Training.

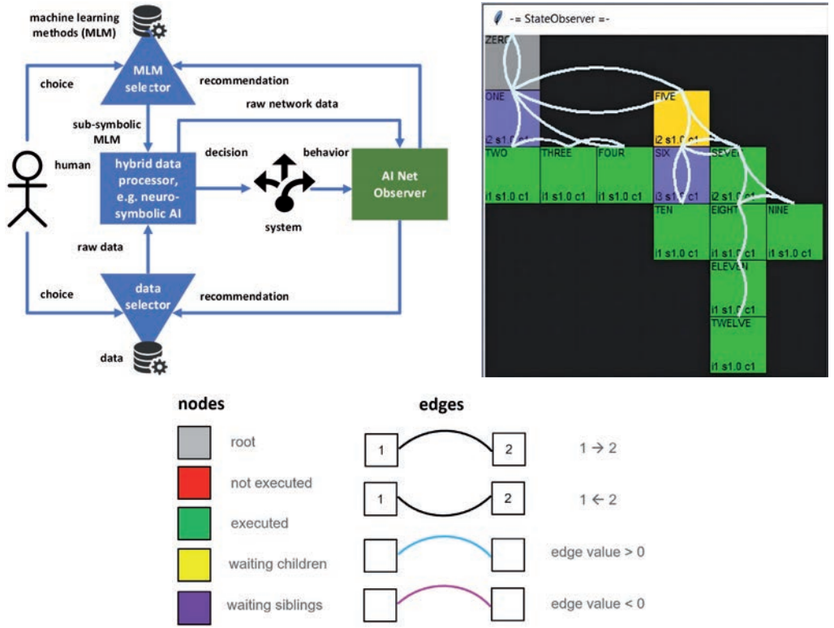

1. AI Net Observer within Human-AI Training System (left), 2. Frontend design (mid), 3. Semantic design (right). (© DLR e.V.)

Legal Norm Formalization – An Overview

Stefan Griesche, Bosch | **Maximilian Grabowski,** BASt | **Kumar Manas,** Continental
Carmen Martin, Universität des Saarlandes | **Mohsin Munir,** DFKI | **Ya Wang,** FhG FOKUS

Legal norm formalization as one building block of a hybrid AI model was realized by two approaches: handcrafted and AI based. The former focused on extracting additional rules and concepts from court decisions, legal commentaries, and technical specifications by experts. The latter employed large language models (LLM) to extract relevant sections from legal documents and translate them into a machine-readable format. Collaboration between natural language processing (NLP) and legal experts enhanced the formalization process. The formalized rules were applied on conformity checks and behavior adaptation. Both resulted in the automated vehicle handling situations more human-like.

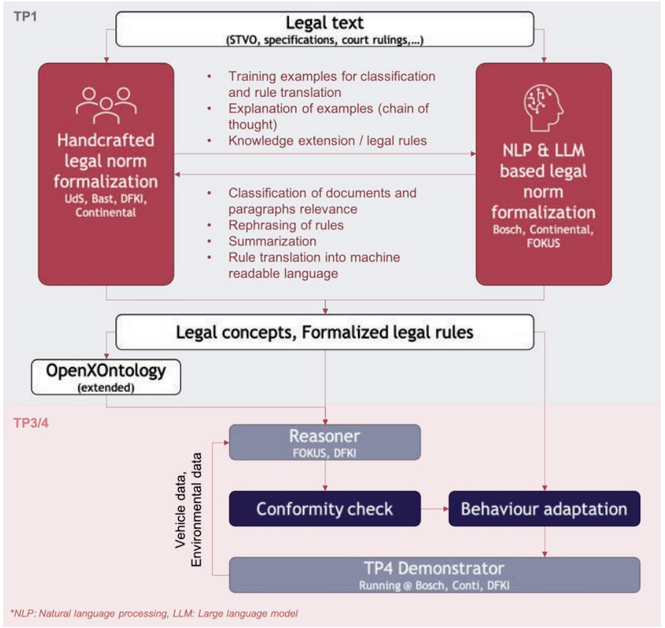

Legal norm formalization approaches and their interaction in KI Wissen. (© Robert Bosch GmbH)

From Legal Documents to Formalized Rules with Large Language Models

Stefan Griesche, Moritz Nekolla, Bosch

Bosch develops products for driving functions worldwide, therefore every market needs to analyzed, compared and checked for legal requirements. To gain efficiency in this process a natural language pipeline (NLP) including generative AI was introduced. The NLP pipeline classified documents and paragraphs based on relevance to product areas. Relevant passages with AI enriched metadata were stored in a vector database. A semantic search retrieved context specific entries (such as sections related to solid lane crossing). A large language model (LLM) was used to extract and formalize rules with zero-shot learning. The rules could be integrated into CARLA and checked in real time on simulation data.

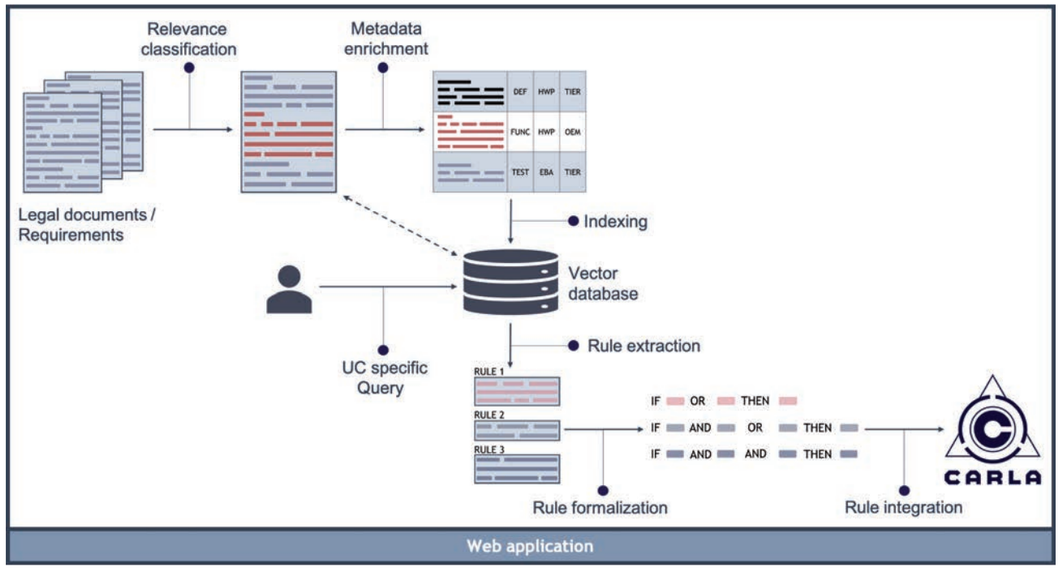

End2End approach for legal rule formalization with GenAI. (© Robert Bosch GmbH)

Data-Efficient Automated Rule Formalization Framework

Kumar Manas, Alexandar Kopte, Stefan Zwicklbauer, Continental

Traffic rules and specifications represented as temporal logic aid for the safety and rule compliance of the driving system. However, such formalization in a handcrafted manner increases turnaround time for the formalization and its integration into prediction and planning. We propose a human-in-loop automated rule formalization framework based on semantic role and chain-of-thought reasoning to assist the formalization of a diverse set of rules from legal and world knowledge sources. We achieved promising performance in traffic rules formalization and increased performance over key trajectory prediction module metrics (ADE, class_acc,..) after rule integration.

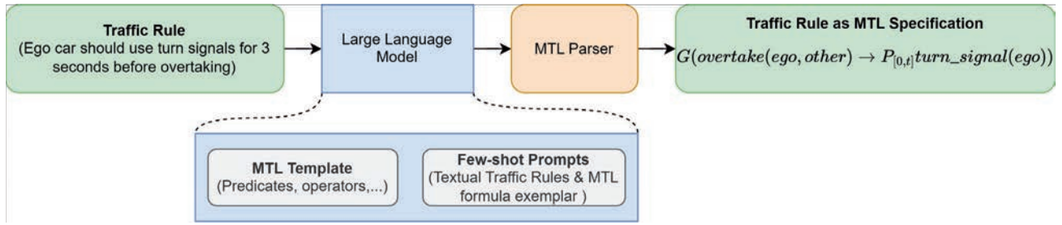

Traffic Rule Translation from legal text to temporal logic.

Traffic Sequence Charts for Knowledge Formalization and AI Application

Philipp Borchers, Dominik Grundt, Tino Werner, Eike Möhlmann, DLR

Infusing knowledge into AI driving functions is regarded as means to reduce training efforts and increase AI's performance. Hence, we merged different knowledge sources into a unified formal representation. Using a single representation reduces the sources of errors and misunderstandings. We categorized relevant knowledge describing what (Ontology) traffic objects should or must (Script, Deontic) do under which (Physics) physical possibilities. Further, we investigated the capabilities of Traffic Sequence Charts (TSCs) to formalize such multimodal knowledge. Finally, we trained a Reinforcement Learning (RL) agent in order to perform maneuvers satisfying TSC-specified knowledge.

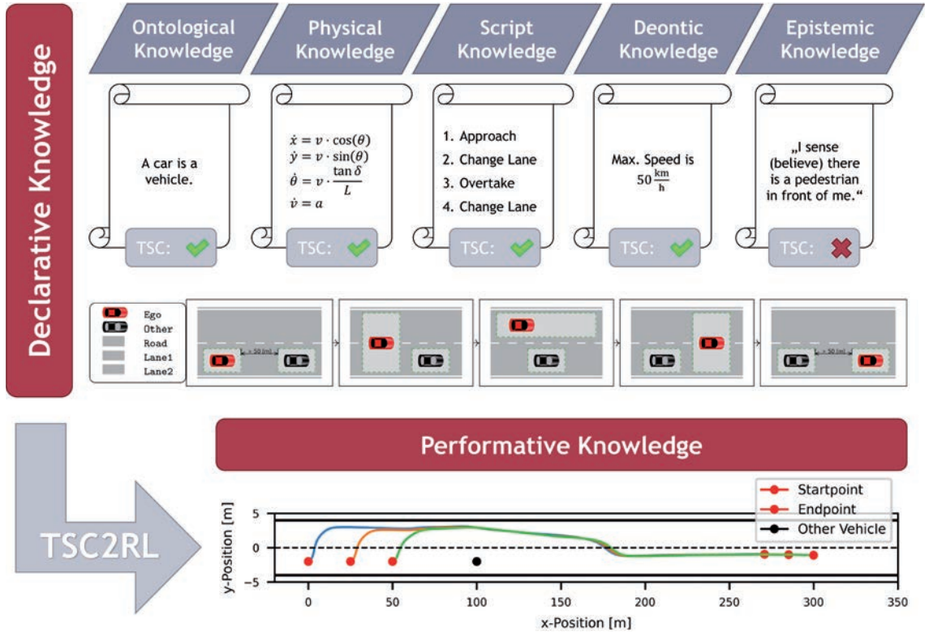

Declarative multi-modal knowledge classification formalized in TSCs to be transferred to performative knowledge in reinforcement learning agents. (© DLR e.V.)

Risk Assessment and Adaptive Decision-Making for Autonomous Vehicle Safety

Syed Tahseen Raza Rizvi, DFKI

This work focuses on optimizing autonomous vehicle safety through a systematic approach, starting with sensor data collection on diverse traffic participants in a simulated environment generated by CARLA and estimating various attributes of all traffic participants. The reasoner, equipped with an ASAM OpenXOntology and STVO-based separation rules, assesses adherence to safety constraints and assigns risk levels (0 to 3) to potential interactions, facilitating proactive decision-making. The proposed framework integrates STVO-based reasoning, compliance checks, risk identification, and quantification, ensuring a robust safety protocol in autonomous driving systems.

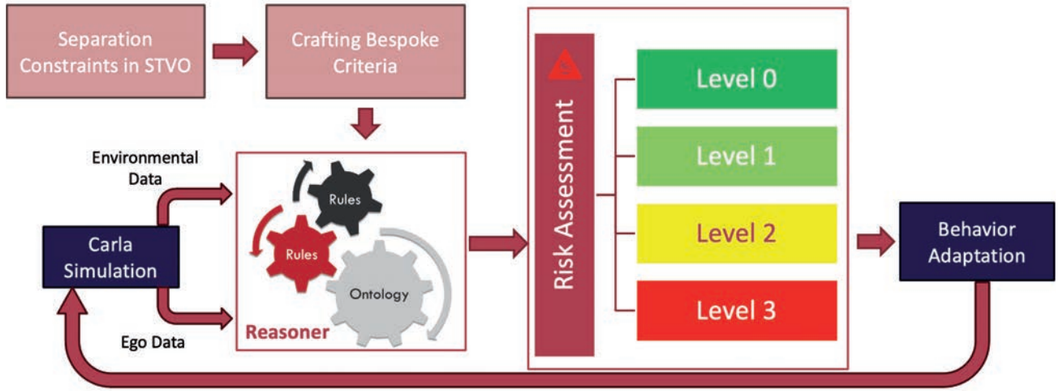

Detailed overview of the Pipeline for the Traffic Scene Safety Analysis. (© DFKI GmbH)

Monitoring Traffic Rule Conformance and Integrated Knowledge Impact

Esra Acar-Celik, Stefan Matthes, Xiangzhong Liu, fortiss

We have developed a „Knowledge Conformance Framework" to aid in developing and testing trajectory/behavior planning for autonomous driving. The framework, based on our in-house developed simulator BARK, simplifies the integration of diverse knowledge sources (e.g., traffic rules and map information) and automates the runtime verification of knowledge conformity. Experiments show promising results of knowledge integration: Involving traffic rules in reward design significantly accelerates the training of deep reinforcement learning and stabilizes model performance, while map knowledge significantly improves the driving score on the CARLA benchmark for motion planning.

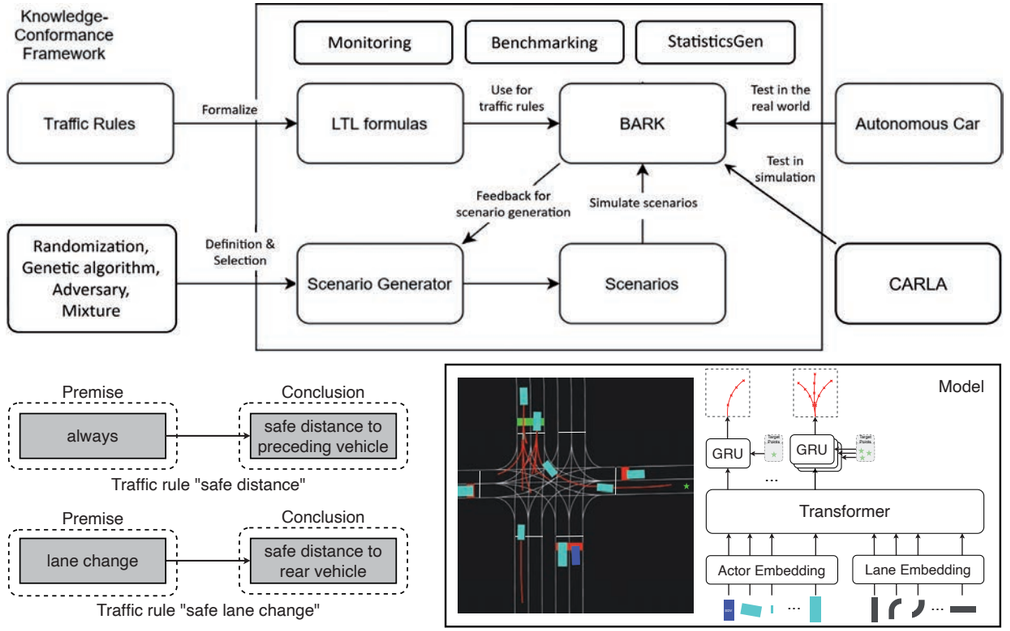

Our „Knowledge Conformance Framework" (top), safe distance and safe lane change rules in premise-conclusion format (bottom left), and our transformer model for integrating map knowledge (bottom right). (© fortiss GmbH)

Runtime Monitoring Using TSC-Based Knowledge Building Blocks

Dominik Grundt, Ishan Saxena, Anna Köhne, Ralf Stemmer, Bernd Westphal, Eike Möhlmann, DLR

Runtime conformity checks of AI driving functions against integrated knowledge can both reduce the training duration and be a safety mechanism in the operation phase. In the project KI Wissen, we have developed a runtime monitoring for conformity checks of knowledge-infused AI driving functions based on formalized multimodal knowledge in the specification language Traffic Sequence Charts (TSCs). Our monitoring exploits the separation of spatial and temporal properties by the TSC language in order to enhance an efficient conformity check during runtime.

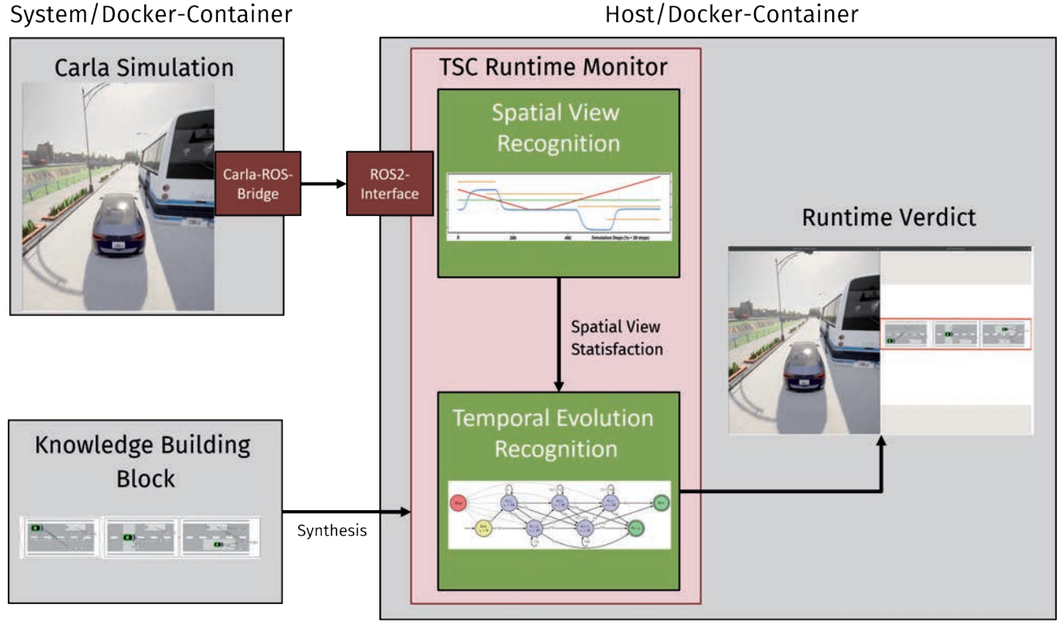

„TSC-Runtime Monitoring" (right), connected via ROS communication to a CARLA simulation, analyzing if the AI driving function behaves conformant to the TSC-specified knowledge on passing-by a stopping bus. (© DLR e.V.)

Scene Understanding for Autonomous Driving Using Visual Question Answering

Tianming Qiu, Stefan Matthes, fortiss

We build a new dataset by extending BDD100K with annotations from 4000 factual answers and improved response performance by modifying a Transformer-based cross-modal model. Furthermore, we extracted cross-attention heatmaps and showed a remarkable correlation between the input question and the visual elements in the image. Both qualitative and quantitative experimental results confirm that our task design offers a practical, interactive approach to image scene interpretation.

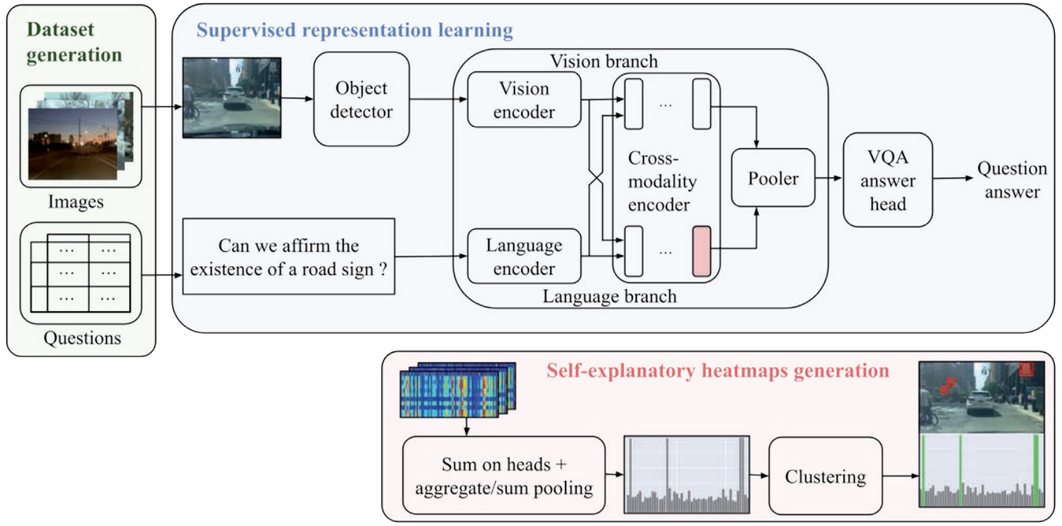

Our pipeline consists of three parts, dataset generation, supervised representation learning, and self-explanatory heatmaps generation. (© fortiss GmbH)

Consistent Joint Action and Explanation Prediction for Autonomous Driving

Leon Luithlen, Marsida Toska, Michelle Reckstadt, Daniel Kurta, Nikita Volkov, Erwin Rudi, Alexander Thamm

We developed a new variant of the BDD-OIA dataset, called BDD-OIA-v2, which addresses several shortcomings of the original, namely potential data leakage between train, validation and test sets, frame selection, and incomplete or partially inconsistent labels. We used this dataset to evaluate several adaptations to the original BDD-OIA classification architecture that reward or enforce consistent action-reason predictions. Preliminary results show that adding consistency constraints can improve classification performance for explanations as measured by F1 score.

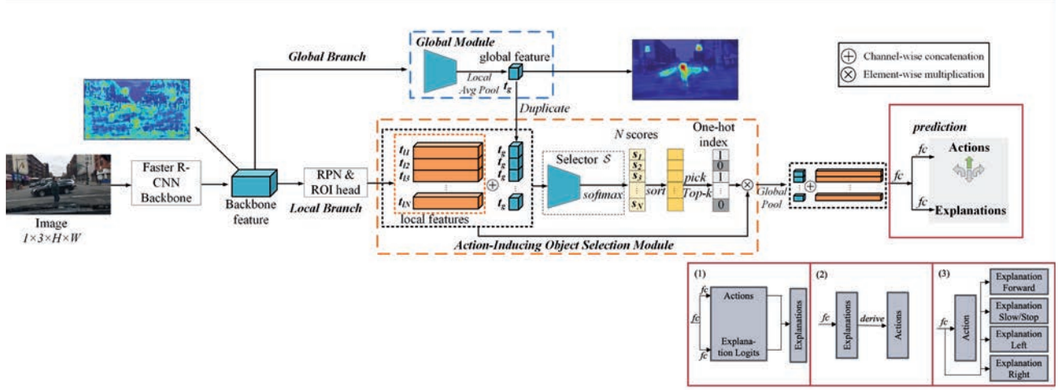

The architecture from BDD-OIA, with the final classification module highlighted. The novel classification approaches are illustrated below. They consist in (1) deriving consistent explanations from the predicted actions and the explanation logits (2) predicting explanations only and deriving the actions (3) training a classifier for each set of explanations consistent with one action.

Generating Training Data from Formalized Traffic Dynamics

Artem Oppermann, Tino Teige, BTC

In autonomous driving, diverse training data is important to ensure the accuracy and reliability of the vehicle‘s decision-making algorithms. Our approach uses BTC‘s scenario specification language Formal Traffic Scenarios (FTS) and a Scenario Concretization Solver to generate artificial training data from formalized physical and mathematical knowledge of traffic scenarios. FTS allows us to specify a wide range of traffic situations, encompassing both common driving scenarios and infrequent corner cases. The solver uses constraint-solving techniques to generate training data in the form of trajectories for the traffic participants in the scenario formalized with FTS.

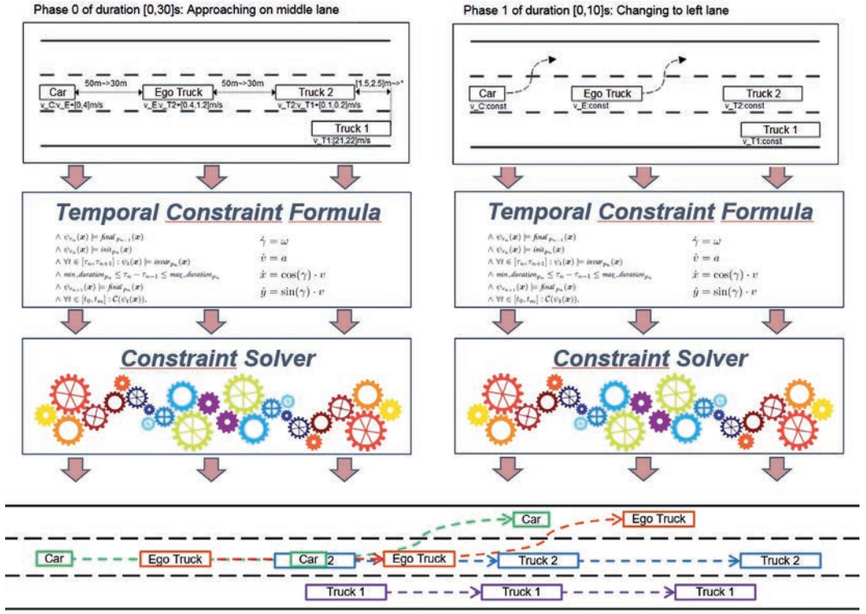

Trajectory generation from FTS. (© BTC Embedded Systems AG)

TSC-Based Generation of Concrete Scenarios for Synthetic Data

Philipp Borchers, Thies de Graaff, DLR

Synthetic data not only greatly reduces the amount of real-world data required for training and testing AI driving functions, but also allows to obtain data for yet unseen rare scenarios. The original intention behind Traffic Sequence Charts (TSC) is to specify scenarios abstractly. A satisfiability modulo theories (SMT)-based play-out is used to derive concrete scenarios. In KI Wissen, we developed techniques solving the SMT formula multiple times resulting in various concrete scenarios. To this end, we generated 300 concrete scenarios (OpenSCENARIO and OpenDRIVE) and simulated them using CARLA. We recorded a dataset of 300 GB with images from six cameras, which can serve as knowledge base for AI training.

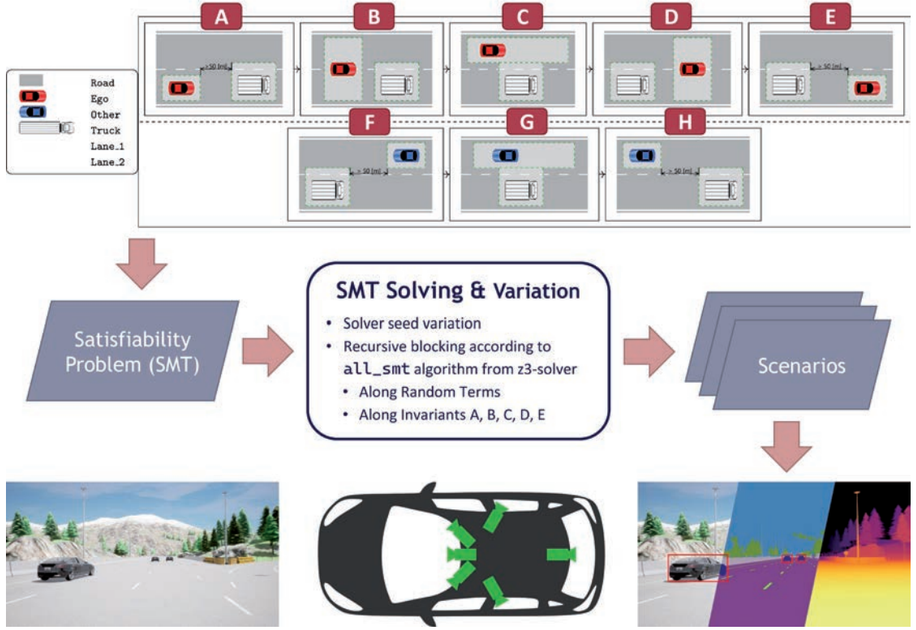

Synthetic Data Generation Processes based on TSCs - The abstract scenario is transferred into a satisfiability problem to be solved by SMT solving and variation methods such that data can be recorded via simulation. (© DLR e.V.)

Local Concept-Based Explanations for Object Detectors

Franz Motzkus, Christian Hellert, Continental

Combining insights into the model processing from local attribution methods (“how was the sample processed”) and global xAI methods (“what does the model generally encode”) provides substantive knowledge about the model functioning. This work extends current methods in local-to-global xAI approaches and applies them to real-world object detection settings. The methods comprise concept-based attributions to test for global concepts, and the dataset-wide extraction of object detection strategies.

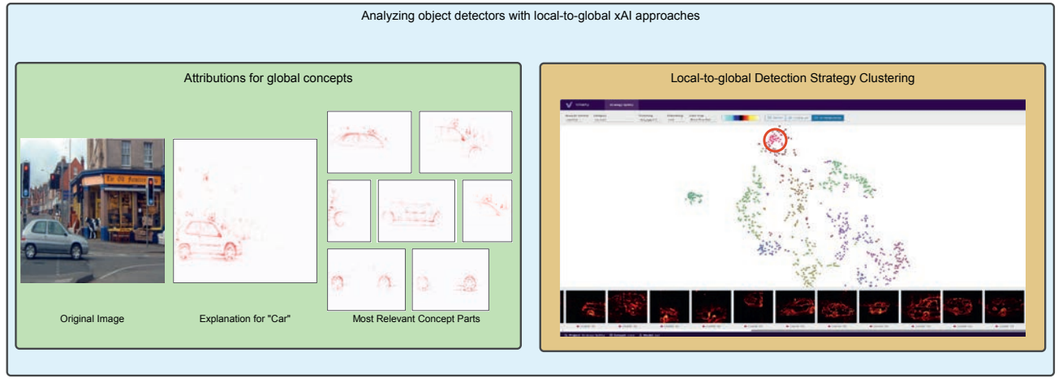

Local-to-global approaches for explaining object detection models: While classical attribution methods show what has been used by the model for a certain detection, concept-based decompositions (left) combine model-internal concept encodings with local attributions. Alternatively, the detection strategies for a single class can be revealed by clustering local attributions (right). Nearby points are likely to show the same features.

Partner Contributions

Continental Automotive GmbH 122
Alexander Thamm GmbH. 124
Capgemini Engineering 126
AVL Software and Functions GmbH 128
BTC Embedded Systems AG 130
Deutsches Forschungszentrum für Künstliche Intelligenz GmbH 132
Deutsches Zentrum für Luft- und Raumfahrt 134
e:fs TechHub GmbH. 136
fortiss GmbH. 138
Fraunhofer-Gesellschaft zur Förderung der angewandten Forschung e.V. FOKUS 140
Fraunhofer-Gesellschaft zur Förderung der angewandten Forschung e.V. IAIS 142
FZI Forschungszentrum Informatik. 144
Robert Bosch GmbH 146
Universität des Saarlandes 148
Valeo Schalter und Sensoren GmbH 150
Bundesanstalt für Straßenwesen (BASt) 152

Continental Automotive GmbH

Company Overview

The Continental AG, as the parent company of Continental Automotive Technologies GmbH, develops technologies and services for the sustainable and connected mobility of people and goods. Continental offers secure, efficient, intelligent, and affordable solutions for vehicles, machinery, traffic, and transportation. Continental Automotive Technologies GmbH focuses on innovative products for a modern automotive future where individual mobility and driving pleasure coexist with driving safety, environmental responsibility, and economic efficiency.

Key Contributions

- Assessment of existing and newly devised methods for knowledge extraction to enhance existing AI models with respect to their traceability and safety
- Development of plausibility component to detect errors in an AI system at runtime
- In the area of knowledge integration: Developing an error robust tracker and an efficient trajectory predictor and based on that, a complete system for trajectory prediction and improved planning algorithms

Exploitation Story

The methods and solutions developed by Continental and its partners significantly exceed the state-of-the-art. They are poised for future integration within Continental, ensuring both relevance and applicability in a product-oriented context. In the automotive context, trustworthy AI systems are paramount, and we have evaluated and applied several techniques in explainable AI, promoting transparency and traceability. Furthermore, our focus on integrating knowledge into situation interpretation and traffic prediction has yielded a comprehensible trajectory prediction system and enhanced planning algorithms, laying a robust foundation for ongoing advancements.

Impact of KI Wissen

The solutions and methods developed in KI Wissen for the integration of domain knowledge in AI systems are a valuable step forward in the development of hybrid AI technologies. The ability to extract new knowledge, check AI decisions for conformity, and reduce training efforts through domain knowledge integration is a significant contribution to the creation of a new foundation for safer, better, and more efficient AI systems of the future. The possibilities unlocked by KI Wissen represent a major advancement in the field of AI and the entire automotive sector, which is increasingly driven by software and AI. Automotive applications can greatly benefit from these achievements.

Alexander Thamm GmbH

Company Overview

Alexander Thamm GmbH is a German consulting company specializing in data and AI. Its mission is an economically strong Europe with digital products and services that are based on European values. The company has leveraged data to generate real value in more than 1000 projects across various industries. Experts from Alexander Thamm can accompany every step on the data journey of customers, ranging from the initial data strategy to the permanent operation of complex AI systems.

Key Contributions

- Leading role in work package about Explainable AI methods for autonomous driving
- Development of an Explainable AI method for exploring critical situations for pedestrian detection models
- Extension of an autonomous driving dataset and improvements on a driving model with a transparent decision process
- Methods for integrating knowledge like the location of objects or similarity to other projects into computer vision models

Exploitation Story

We are planning to incorporate our gained experience from our research into future projects. Our development of new Explainable AI methods and the exchange with research partners help us extend our expertise in building transparent AI systems with decision processes that are understandable for humans. Our work with state-of-the-art computer vision models is applicable to other use cases. Especially our work on object detection with the GLIP model can be valuable in the automotive industry, specifically for autonomous driving.

Impact of KI Wissen

Safety and trustworthiness are crucial requirements for reaching the goal of self-driving vehicles on our streets. Ensuring failure-free driving decisions is a difficult task because many unpredictable and unprecedented situations can occur where AI algorithms still need to make good decisions without having seen them in their training data. The explainability methods that were developed in KI Wissen can help understand if the driving models make their decisions based on the right reasons and in which situations their logic might fail. With its methods for integrating the gained knowledge into models, KI Wissen is providing important research for making autonomous driving safer.

Capgemini Engineering

Company Overview

Capgemini is one of the world's leading providers of technology, consulting, and engineering services. As a pioneer of innovation, we support our customers in the complex challenges surrounding digital transformation. Based on 50 years of experience and extensive industry-specific know-how, we help you achieve your business goals. Capgemini Engineering is a global leader in engineering and R&D, supporting its customers from concept to industrialization in the development of tomorrow's products and services.

Key Contributions

- Early integration into full loop stack as well as repeated feedback and support for improvement of demonstrator software assets available to all consortium partners
- Leading function in AP1.3 and AP4.1 as well as hosting multiple on site workshops for TP1 and TP3 project partners
- Development of neural network based objection detection method directly incoorperating sensor data as well as prior knowledge originating from 3D semantic maps

Exploitation Story

The exploration and continuous work on object detection methods incorporating prior knowledge in addition to sensor data will contribute to the continuous optimization of 2D and 3D data labeling solutions and broaden Capgemini Engineering offerings to customers with needs in this space. The developed evaluation pipeline implemented during KI Wissen highlights the model and training state specific strengths and weaknesses based on scenario types as well as environmental factors. This will yield further improvements in the automatic labeling capabilities in current solutions, resulting in higher labeling speeds while reducing quality control overhead.

Impact of KI Wissen

KI Wissen developed a new neural network architecture directly incorporating prior knowledge in the neural information processing during training and inference. The development of a new evaluation pipeline focusing on multi-variate statistic evaluation methods of distributions in performance metrics enabled a detailed and differentiated comparison of classic neural network architectures and their prior knowledge-enhanced model architecture versions. The research on the impact of degraded prior knowledge conformity on object detection performance emphasized the importance of ensuring the alignment of prior knowledge to current sensor data input.

AVL Software and Functions GmbH

Company Overview

AVL, the world's largest independent company for development, simulation, and testing in the automotive industry, is a pioneer of future mobility trends. Founded in 2008 in Regensburg, Germany, AVL Software and Functions has experienced strong growth with more than 600 employees from over 40 countries. As a specialist for future-oriented software and system solutions, the company focuses on intelligent and environmentally friendly mobility, system integration, and electronics development.

Key Contributions

- Establishing a demonstrator to accurately simulate predefined driving scenarios for use cases
- Providing ongoing support to project partners for integrating into the AD Stack,utilizing new AD Stack features, and using the demonstrator
- Providing accurately post-processed real-world driving data
- Granting free licenses for AVL Scenario Designer to project partners to create simulation scenarios
- Leading roles in subproject and work packages

Exploitation Story

Our experience with KI Wissen has given us valuable insight into various aspects relevant to developing and testing advanced AD functionality. Creating a comprehensive AD stack that enables external user access and integration of their own software components has provided us with the necessary expertise for future collaboration and growth. Investigations into the AVL Scenario Designer and CARLA simulator for interoperability have further enhanced our expertise in this area. Moreover, the knowledge gained from extracting data from a geodatabase like OpenStreetMap and converting it to a simulation-friendly format, such as ASAM OpenDRIVE, lays the groundwork for future development and applications.

Impact of KI Wissen

The project's approach in incorporating diverse sources of knowledge into machine learning is pivotal to advancing AD capabilities. While traffic data is essential for training AI-based functions, real-world data alone is insufficient to expand the capabilities of AI modules to handle unusual situations that may pose a risk to AD vehicles or other road users. In certain situations, limited and carefully considered non-compliance with traffic rules may be necessary. Synthetic data and simulation scenarios that reproduce various and uncommon driving situations, along with the integration of diverse knowledge sources into AI-based functions, establish a comprehensive framework for enhancing future AD development.

BTC Embedded Systems AG

Company Overview

BTC Embedded Systems AG was founded in 1999 with an ambitious goal: to take mathematical verification and test methods out of their academic hiding place and make them available for real-life software projects. We still feel dedicated to this goal and proudly look back on 20 years of technological market leadership in areas like formal verification and automatic test generation. Today, our ISO 26262-certified products are used by car makers and suppliers worldwide for the automated test and verification of safety-critical embedded software.

Key Contributions

- Formalization of mathematical and physical knowledge of traffic dynamics including corner cases and rare events
- Generating synthetic training data for AI systems from this formalized knowledge using constraint solving
- Generating monitors from this formalized knowledge to observe the behavior of AI systems

Exploitation Story

For some years, BTC has been engaged in scenario-based testing of autonomous vehicles. In this context, traffic situations and physical knowledge are described by means of a formal scenario specification language. The experience in formalizing traffic dynamics made within the KI-Wissen project will be considered for integration into BTC's products for verification and testing of autonomous driving functions.

Impact of KI Wissen

Generating artificial training data from formalized traffic scenarios not only reduces the effort and costs involved in obtaining and annotating the data but also makes it possible to include in the training dataset corner cases that are difficult to capture in real-world conditions. Therefore, this approach aims at significantly increasing the size of the dataset and thus further improving the quality of the predictive models.

Deutsches Forschungszentrum für Künstliche Intelligenz GmbH

dfki Deutsches Forschungszentrum für Künstliche Intelligenz
German Research Center for Artificial Intelligence

Company Overview

As the largest independent Artificial Intelligence research center worldwide, the DFKI initiates, realizes, and supports many activities to develop reliable and trustworthy AI and secures a leading role for Germany and Europe in international competition. Based on application-oriented basic research, the DFKI develops product functions, prototypes, and patentable solutions in the field of information and communication technology. DFKI combines scientific excellence with business-oriented research in socially relevant fields of application.

Key Contributions

- Robust pedestrian detection architecture, considering factors such as visibility, height variations, and occlusion
- Efficient anchor-free pedestrian detection architecture, reducing inference time by 55%
- Multi-modal sensor fusion (image, LiDAR) for enhanced robustness and reliability in obstacle detection
- Method that uncovers high-level concepts (e.g. body parts) learned by Deep Neural Networks for pedestrian detection

Exploitation Story

DFKI's mission is to bridge the gap between application-oriented basic research and practical industrial use. Therefore, we strive to refine the knowledge gained in KI Wissen and transfer it to industrial use. The focus here is on three specific exploitation paths: communicating the results in university and industrial training, translating outcomes into practical components through collaboration with industrial users, and further development in follow-up research projects. Various forms are conceivable for this, including independent consulting, R&D projects, and integration into existing collaborations such as joint labs with partners from the automotive industry in particular.

Impact of KI Wissen

Integrating explicit, symbolic knowledge and data-driven machine learning methods into hybrid AI systems represents a promising evolution in developing AI applications. This approach enhances training efficiency and yields more reliable, explainable, and understandable models. These attributes are crucial for establishing trust and acceptance of AI solutions, not only in automotive applications but also across various socially relevant domains. The project outcomes underscore the effectiveness of this hybrid approach, highlighting its potential to establish advanced object detection systems in autonomous driving and serving as a blueprint for innovative solutions in image-centric domains in general.

Deutsches Zentrum für Luft- und Raumfahrt

Company Overview

DLR is the national aeronautics and space research center of the Federal Republic of Germany. Its extensive research and development work in aeronautics, space, energy, transport, and security is integrated into national and international cooperative ventures. In KI Wissen the DLR Institute of Transportation Systems and the Institute of Systems Engineering for Future Mobility have combined their efforts for AI research. Our scientists contribute to increasing the safety and efficiency of traffic on roads and railways.

Key Contributions

- Leading Focus group on knowledge sources to cluster the approaches and design the knowledge building block catalog
- Formalization of multimodal knowledge and integration into AI training, runtime monitoring of knowledge conformity
- Leveraging a priori knowledge for object detection in hybrid AI architectures
- Tracking and prediction of vehicle trajectories by using AI results and semantic information in a multiple object tracking pipeline
- Development of an AI observer for comprehensible AI

Exploitation Story

Based on our proposed knowledge formalization and integration using Traffic Sequence Charts, we want to specify more complex traffic scenarios with multimodal knowledge and scale our methods so that a framework arises that deals with the tasks of formalizing knowledge, integrating it into the training process, and automatically generating monitors for a conformity check. Through hybrid AI, we aim to enhance the capabilities of object detection and build more robust AI modules for our autonomous driving functions. The trajectory prediction will be an integral part of the driving algorithms to reliably forecast other traffic participants movements.

Impact of KI Wissen

Reducing the amount of needed real-world data for AI training and increasing confidence in AI decisions is of great importance for enabling the application of AI in driving tasks. KI Wissen allows for the use of a priori knowledge for training and conformity checks, paving the way for a trustworthy application of AI driving functions. The AI observer that is under development will be used as a basis for future research to make AI more comprehensible and therefore increase the acceptance of AI applications in the population.

e:fs TechHub GmbH

Company Overview

Having started as an automotive engineering supplier, e:fs TechHub GmbH has evolved into an independent partner for developing cutting-edge technology solutions in the field of mobility and beyond. The expertise spans from software development to big data management and AI. e:fs TechHub is dedicated to combining the highest quality with innovative strength to translate future technologies into real-life use cases with series-production quality. The focus extends to the automotive industry, as well as projects in aerospace technology, greentech, and smart cities of the future.

Key Contributions

- Concept for integration of causality in AI models
- Formalization of physical-mathematical knowledge about vehicle trajectories as a causal model and contribution to the toolbox of knowledge building blocks
- Development of methods for integrating the causal model into training processes of machine learning models
- Evaluation and demonstration of methods applying the causal model to detect and correct physically implausible vehicle trajectories

Exploitation Story

We are planning to reduce the time-to-market of comprehensible AI-based driving functions using our concepts in the form of modular, reusable, and standardized AI building blocks for knowledge formalization and integration. This will help our customers to improve training and data efficiency. For example, our modules can make the development and validation of automated driving functions more time- and cost-effective by generating high-value training data at a very low cost. Beyond the automotive use case, the insights from KI Wissen will contribute to our recently launched products and new innovations, facilitating our broader strategy towards the application of safe and explainable AI.

Impact of KI Wissen

Enhancing data efficiency and safeguarding AI-based functions are of great importance in automated driving (AD). The current black-box models of data-driven AI solutions require enormous amounts of data for training and are prone to unexpected and dangerous failures in novel scenarios. The work within KI Wissen provides substantial and convincing arguments to highlight that the combination of a priori knowledge and formalized rules with data-based AI solutions enables the development of more data-efficient and secure AD functions. The combined efforts in developing a comprehensive ecosystem for knowledge integration provide a strong foundation for future collaboration opportunities.

fortiss GmbH

fortiss

Company Overview

fortiss is the research institute of the Free State of Bavaria for software-intensive systems and services, based in Munich. The institute currently employs around 130 people who are involved in research, development, and transfer projects with universities, scientific institutions, and technology companies in Bavaria, Germany, and Europe. fortiss researches and develops methods, techniques, and tools for reliable, secure, and reproducible software solutions and AI and IoT applications.

Key Contributions

- Development of knowledge-augmented perception models that facilitate robust pedestrian detection
- Human-centered explainability concepts based on visual question answering for scene and behavior understanding
- Development of neural planning and forecasting models that integrate map knowledge
- Extension of the BARK framework with a component that can be used to check knowledge-based planning algorithms for conformity with predefined rules

Exploitation Story

In its role as a research and transfer institute, fortiss will use its research results to support local and national partners in future collaborations. Digitalization and the advancement of AI-based methods will be promoted through training courses, publicly funded projects, and contracted research. fortiss will consolidate its gathered expertise into a solution portfolio that will benefit both, global corporations and small to medium-sized enterprises (SMEs) in the automotive sector. We to establish knowledge-augmented machine learning as a new paradigm for creating reliable and sustainable AI models.

Impact of KI Wissen

In many application domains, limited data hinders the use of conventional learning strategies or complex architectures that often require large amounts of training instances. Integrating and leveraging existing domain and expert knowledge is key to developing reliable models, which eventually contribute to acceptance and trust in artificial intelligence processes. We see great potential in models that leverage knowledge about the context in which they are used. The concepts developed in KI Wissen offer the unique opportunity to be applied to real-world problems, as opposed to solutions developed solely on well-curated benchmark datasets.

Fraunhofer-Gesellschaft zur Förderung der angewandten Forschung e.V. FOKUS

Company Overview

The Fraunhofer Institute for Open Communication Systems (FOKUS) is a research institution focused on digital transformation and its impact on society, economy, and technology. It explores and develops solutions for the future digital world, including topics like security, quantum computing, sustainability, artificial intelligence, digital life, and digital networking. FOKUS actively collaborates with partners from various sectors and promotes innovation through its diverse research areas and projects.

Key Contributions

- Innovating in AI with a neuro-symbolic architecture that enhances safety, reliability, and interpretability
- Developing techniques to extract and integrate symbolic knowledge from and into machine learning models, advancing the field of AI explainability
- Collaborating in the development and maintenance of an ontology, ensuring its continued relevance and utility within the project

Exploitation Story

In the realm of autonomous driving, we fuse AI innovation with safety and transparency. Our approach in KI Wissen serves as both an AI decision explainer and an enhancer. We integrate a symbolic solver for reliable, interpretable decision-making and (legal) rule compliance. We extract symbolic knowledge that can be used to improve model performance, reducing data requirements. This adaptable architecture can easily transfer to safety-critical systems demanding transparent, traceable decision-making processes.

Impact of KI Wissen

The neuro-symbolic architecture in autonomous driving can significantly enhance safety, reliability, and regulatory compliance. It merges neural network adaptability with symbolic AI‘s rule-based reasoning, ensuring decisions are data-driven yet law-abiding. Our approach also unravels the complexity of AI‘s ‚black-box‘ through symbolic knowledge extraction, leading to decisions that are both traceable and transparent. This advancement aids in aligning AI with ethical standards and human values such as legal norms, marking a pivotal step in responsible AI development.

Fraunhofer-Gesellschaft zur Förderung der angewandten Forschung e.V. IAIS

Company Overview

As part of the largest organization for application-oriented research in Europe, the Fraunhofer Institute for Intelligent Analysis and Information Systems IAIS is a leading scientific institute in the fields of Artificial Intelligence, Machine Learning and Big Data. With a team of ca. 350 employees, the Institute supports companies in the optimization of products, services, processes, and structures and in the development of new digital business models. IAIS is thus helping to shape the digital transformation of our working and living environment.

Key Contributions

- Development of methods for the extraction of knowledge concepts from black-box ML models
- Support and maintenance of CI/CD pipelines for source code deliverables
- Geo-informed knowledge conformity checks with street maps
- Development and maintenance of a “Visual Analytics Tool” as an interface for knowledge transfer between ML model and human
- Collaboration with other partners for integration of methods into the Visual Analytics Tool

Exploitation Story

Our focus is on the development of trustworthy AI systems, particularly within safety-critical application fields. The KI Wissen project enables us to gain expertise in knowledge transfer related to the interface between machines and humans. Our developed Visual Analytics Tool can be a pivotal element in further projects regarding knowledge transfer and conformity checks. Furthermore, we will apply our „Informed ML“ taxonomy to related real-world use cases and advance our knowledge extraction techniques, facilitating a transparent understanding of decision-making processes and contributing to the ongoing pursuit of trustworthy and reliable AI applications.

Impact of KI Wissen

Trustworthy and transparent decisions by intelligent systems are becoming increasingly important, especially in times when „General AI“ seems to be on the horizon. In a safety-critical application such as autonomous driving, KI Wissen therefore creates a crucial framework for research into the harmonization of black-box AI systems with diverse forms of human, legal, and world knowledge. The insights and experiences gained will help to make future autonomous systems more reliable, more comprehensible, and, finally, safer.

FZI Forschungszentrum Informatik

Company Overview

The FZI Forschungszentrum Informatik is an independent and non-profit foundation for applied cutting-edge research and technology transfer. For over 35 years, the FZI has been researching and developing innovations for the benefit of society and bringing the latest scientific findings in information technology as practical solutions to companies and public institutions. In doing so, the FZI qualifies people for an academic career, a professional start in business, or even the leap into self-employment.

Key Contributions

- Formalizing and integrating legal and social norms into AI maneuver decisions for compliance assurance
- Multisensory AI-based perception and cooperative maneuver planning
- Emphasis on linking explicit knowledge representations with AI architectures in maneuver and trajectory planning
- Development of AI architectures for autonomous driving focusing on safety and multi-input/output architectures
- Approach for direct vehicle control through model-based Reinforcement Learning that enables conditioning the driving behavior on external knowledge

Exploitation Story

KI Wissen played a pivotal role in laying the foundation for subsequent initiatives exploring the application of AI in the context of autonomous driving. The insights garnered have significantly influenced the inception and development of fresh industrial endeavors. Throughout the project, numerous theses were supervised, contributing to the advancement of the upcoming generation of researchers. The scientific discoveries have been disseminated through multiple publications in esteemed international conferences and are integral to various Ph.D. projects undertaken by the research team.

Impact of KI Wissen

KI Wissen formed the basis for substantial research regarding the integration and extraction of knowledge into and from AI systems. Numerous novel models for AI-based perception and maneuver planning were explored. Furthermore, approaches for checking the plausibility and safety of AI architectures in the domains of trajectory planning and vehicle control were researched and engineered.
KI Wissen facilitated the attainment of numerous publications at significant automotive and AI conferences, leading to an enhanced reputation for the FZI and substantial progress in the dissertations of the scientific staff.

Robert Bosch GmbH

Company Overview

The Bosch Group is a leading international technology and service company. In the mobility area the Cross-Domain Computing Solutions division bundles software and electronics expertise for vehicle computers, sensors and control units for numerous applications of driver assistance, automated driving and parking as well as infotainment. Methods of AI and data-driven development are strategic focus areas. Bosch has contributed its extensive know-how to all AI projects within the VDA Leitinitiative.

Key Contributions

- Set up a retrieval augmented generation for legal documents
- Implemented a toolbox of NLP methods to extract, compare, and enrich relevant knowledge from various sources
- Evaluated the use of different large language models, such as GPT4, for rule formulation and formalization
- Bundled the methods into a user frontend to create an end2end pipeline capable of taking legal documents to generate machine-readable rules that can be tested in CARLA

Exploitation Story

We will incorporate our NLP toolbox into our legal requirement management process. Our solutions will increase efficiency by supporting legal experts in formulating unified legal requirements. The LLM based rule formulation and formalization will serve as a framework to extract and formulate target market specific legal requirements. With the formalization of those the approach will foster the traceability between legal requirements and implemented rules. This results in an improved time to market. Furthermore, we will explore how our methods can support domain experts throughout various stages of the system engineering process, such as an AI-generated test case generator at system level.

Impact of KI Wissen

With the proof-of-concepts (POC) conducted in KI Wissen, we have introduced and demonstrated an AI and large language model (LLM) based solution for the system engineering process of automated driving functions.
The quality of the generated results and the accuracy of the AI models have illustrated the immense potential of AI and LLM in this field. We have observed time savings and a reduction in manual effort. The collaboration between NLP experts and legal/domain experts at KI Wissen has paved the way for internal cooperation. As a result, the outcomes of KI Wissen are considered a blueprint for integrating AI and LLM into processes across other product areas of Bosch.

Universität des Saarlandes

Company Overview

Established in 1948, UdS is today a leading university in research in computer technology. The Institute of Legal Informatics (IfR) adds a legal perspective to this research by investigating questions in IT law and legal informatics. Our research and teaching cover topical questions arising from new legislation and its reception in jurisprudence, new technologies, or the clash of the two. Our current research focuses on the legal aspects of autonomous systems, IT security, and data protection, as well as electronic legal proceedings (eJustice) and administration (eGovernment).

Key Contributions

- Implicit norm compilation
- Formalization of normative knowledge
- Alignment of observed behavior with desired (normative) behavior

Exploitation Story

The fascinating aspect of law is that legal norms are not exhaustive representations of principles, but need ongoing interpretation and commentary, which is a major obstacle to their application by AI. Our work covers the formalization of legal norms into machine-readable forms through the development of a normal form in first-order logic, emphasizing violation conditions, and offering flexibility in addressing conflicts between norms. Focusing on German traffic law (StVO), we develop sets of implicit norms based on jurisprudence and scholarly comment. In the next step, we aim to create a generalized approach to replace detailed instructions with a variable-like reference to a body of norms.

Impact of KI Wissen

KI Wissen gives us the opportunity to contribute to the development of Automated Driving Systems (ADS) and make an impact on their safety and reliability. ADS need to blend into non-automated traffic, including all kinds of human and vehicle interaction, which means they must be designed to adhere to current laws and rules. KI Wissen's comprehensive ecosystem for the integration of knowledge into the training and safeguarding of AI functions aids in seamlessly connecting human- and machine-readable concepts and rules. The methodology on how to derive machine-readable traffic rules from court rulings will allow for the establishment of a comprehensive body of traffic rules for autonomous driving.

Valeo Schalter und Sensoren GmbH

Company Overview

Valeo is an automotive supplier and partner to all automakers worldwide. As a technology company, Valeo offers one of the largest ranges of smart sensors and features that improve vehicle safety and comfort. The business group Valeo Comfort & Driving Assistance designs ultrasonic sensors, LiDARS, radars, and cameras that are used in conjunction with apps to detect obstacles around vehicles, making driving and parking more enjoyable and, above all, safer.

Key Contributions

- Analysis of the impact of color variations on detection confidences with a digital twin from real-world assets
- Enhanced 3D object detector for LiDAR by leveraging semantic information in ground truth sampling augmentation
- Recording of use case scenarios and specification of annotation format for the KI Wissen dataset
- Enhanced pedestrian detection and reduced false positive rate with static cameras by incorporating empty scene knowledge

Exploitation Story

The results obtained in KI Wissen and their integration into the development of future vehicles ensure that Valeo remains competitive in the field of highly automated vehicle concepts with high customer value in the long term. The insights on camera perception with a static background and the augmentation of LiDAR datasets can be directly incorporated into the pre-development and product development of AI functions for automated driving functions. The expertise and technologies developed as part of KI Wissen can be incorporated into future follow-up projects, for example, in further use cases or in the development of new sensor technologies.

Impact of KI Wissen

The results of KI Wissen lead to a better understanding of AI in general and allow new developments in this field to progress faster or be used immediately. They find applications in intelligent sensors for driver assistance systems and thereby constitute another building block toward highly automated driving functions. The methods developed for integrating different types of knowledge into machine learning and their combination with modern data-based AI processes lead to an increase in functional quality and improved data efficiency. In addition, the conformity checks that assess AI decisions against formalized knowledge enable a novel type of safeguarding AI-supported functions.

Bundesanstalt für Straßenwesen (BASt)

Company Overview

The Federal Highway Research Institute (BASt) is the practice-oriented, technical-scientific research institute of the German Federal Ministry for Digital and Transport. It covers the whole field of the road transport system and provides scientifically based advice and support for future-proof standards and regulations. It is dedicated to the diverse tasks resulting from the interactions between humans, transport, infrastructure, and environment. Its mission is to improve traffic safety, environmental compatibility, efficiency, and performance of roads.

Key Contributions

- Identification of relevant regulatory documents
- Extraction of expert knowledge
- Formalization and modification into an ontology
- Integration into Reasoner to showcase targeted AI learning

Exploitation Story

The extraction and formalization of the knowledge from the Technical Specification on Lane Markings will help identify other relevant and similar documents to utilize. Further Technical Specifications can be used to extend ontologies, which will support AI systems in decision-making and situational awareness. Additionally, the results from KI Wissen will help us guide regulatory work with respect to the use of AI systems for automated and autonomous driving. The recorded traffic data will be used to further shape type approval regarding virtual simulations applied for testing.

Impact of KI Wissen

Ensuring adherence to laws, rules, and technical specifications is a crucial aspect when designing AI-based Automated Driving Systems (ADS). KI Wissen highlights the significance of incorporating ontologies as a viable approach to seamlessly connecting human- and machine-readable concepts and rules. Additionally, leveraging KI Wissen aids in creating AI systems that ensure transparency and traceability, enhance their acceptance among the public, and facilitate ongoing research in this field.

Table of contents

Welcome **2**

Greeting **8**

Collaboration in Artificial Intelligence **10**

KI Wissen – Automotive AI Powered by Knowledge **14**

Detection **18**

Knowledge-Augmented Object Detection with Transformer **22**

Model-Agnostic Body Part Relevance Assessment for Pedestrian Detection **24**

Investigating the Effects of Augmentation Techniques on Interpretability and Performance of Pedestrian Detection Models **26**

Static Scene Knowledge for Pedestrian Detection **28**

Method to Fuse Map Layer and Pedestrian State Space to Improve 3D Detector Output **30**

Knowledge Guided Occluded Pedestrian Detection **32**

Knowledge Aware Pedestrian Detection **34**

Geo-Informed Conformity Check of Pedestrian Detection Models **36**

Efficient Pedestrian Detection for Autonomous Driving **38**

Prior-Knowledge Based Deep-Learning Approach for Pedestrian Detection **40**

Performance Impact of Input Image Characteristics and Prior Knowledge Conformity in Prior Mask R-CNN **42**

Plausibility Verification For 3D Object Detectors Using Energy-Based Optimization **44**

Object Detection Plausibility with Concept-Bottleneck Models **46**

Digital Twin Creation and Evaluation of Color Dependent Object Detection **48**

Multi-Modal Sensor Fusion for Robust Obstacle Detection **50**

Leveraging Knowledge for Traffic Sign Detection **52**

Trajectory Prediction and Motion Planning **54**

Informed Motion Planning **58**

Knowledge Integrated Plausible Motion Forecasting **60**

An Empirical Bayes Analysis of Object Trajectory Representation Models **62**

An Efficient and Robust Multi-Modal Trajectory Predictor Baseline for Autonomous Driving **64**

Expert Informed Trajectory Prediction **66**

Integration of A Priori Knowledge Using a Causal Model of Vehicle Trajectories **68**

Causality-Driven Checks of the Physical Conformity of Vehicle Trajectories **70**

Tracking and Trajectory Prediction Using Scene Semantics **72**

Learning-Aided Warmstart of Model Predictive Control in Uncertain Fast-Changing Traffic **74**

Model Predictive Control Under Temporal Logic Specifications **76**

An Outlier-Robust and Efficient Bayesian Filter and Smoother **78**

Behavior-Conditioned Driving Policies **80**

Formal Methods and Knowledge Representation **82**

Harnessing Symbolic Knowledge Extraction and Utilization for Informed Decision-Making **86**

Model Agnostic Local Analysis with Latent Attacks (MALALA) **88**

A Concept to Support AI Models by Using Ontologies **90**

Knowledge-based Comparison of CNNs **92**

AI Net Observer **94**

Legal Norm Formalization – An Overview **96**

From Legal Documents to Formalized Rules with Large Language Models **98**

Data-Efficient Automated Rule Formalization Framework **100**

Traffic Sequence Charts for Knowledge Formalization and AI Application **102**

Risk Assessment and Adaptive Decision-Making for Autonomous Vehicle Safety **104**

Monitoring Traffic Rule Conformance and Integrated Knowledge Impact **106**

Runtime Monitoring Using TSC-Based Knowledge Building Blocks **108**

Scene Understanding for Autonomous Driving Using Visual Question Answering **110**

Consistent Joint Action and Explanation Prediction for Autonomous Driving **112**

Generating Training Data from Formalized Traffic Dynamics **114**

TSC-Based Generation of Concrete Scenarios for Synthetic Data **116**

Local Concept-Based Explanations for Object Detectors **118**

Partner Contributions **120**

Continental Automotive GmbH **122**

Alexander Thamm GmbH **124**

Capgemini Engineering **126**

AVL Software and Functions GmbH **128**

BTC Embedded Systems AG **130**

Deutsches Forschungszentrum für Künstliche Intelligenz GmbH **132**

Deutsches Zentrum für Luft- und Raumfahrt **134**

e:fs TechHub GmbH **136**

fortiss GmbH **138**

Fraunhofer-Gesellschaft zur Förderung der angewandten Forschung e.V. FOKUS **140**

Fraunhofer-Gesellschaft zur Förderung der angewandten Forschung e.V. IAIS **142**

FZI Forschungszentrum Informatik **144**

Robert Bosch GmbH **146**

Universität des Saarlandes **148**

Valeo Schalter und Sensoren GmbH **150**

Bundesanstalt für Straßenwesen (BASt) **152**

Table of contents **154**

Contact & Further Information **156**

Contact & Further Information

Project Coordinator

Simon Heinz
Continental AG
AI Enterprise Applications
Guerickestraße 7
60488 Frankfurt am Main, Germany
simon.2.heinz@continental.com

Project Management

European Center for Information and
Communication Technologies – EICT GmbH
EUREF Campus Haus 13
Torgauer Straße 12-15
10829 Berlin, Germany
info@eict.de

Poster Download

www.kiwissen.de

Note on the legal form of the cooperation
The cooperation between the partners within the project has no independent legal personality. In fact a scientific exchange is conducted between the research centers, organizations and universities listed as cooperation partners. A legal or similar relationship under company law, an association or similar is not established by the scientific cooperation. No cooperation partner is entitled to represent individual other cooperation partners or all cooperation partners together towards third parties.

Project Consortium

Continental · BOSCH · Valeo · [at] alexanderthamm
AVL · BTC embedded systems · Capgemini engineering · e:fs TechHub GmbH
dfki Deutsches Forschungszentrum für Künstliche Intelligenz German Research Center for Artificial Intelligence · DLR · fortiss · Fraunhofer FOKUS
Fraunhofer IAIS · FZI · Universität des Saarlandes · bast Bundesanstalt für Straßenwesen

External Partner

eict