

Article

# Cloud2Edge Elastic AI Framework for Prototyping and Deployment of AI Inference Engines in Autonomous Vehicles

Sorin Grigorescu <sup>1,\*</sup>, Tiberiu Cocias <sup>1</sup>, Bogdan Trasnea <sup>1</sup>, Andrea Margheri <sup>2</sup>,  
Federico Lombardi <sup>2</sup> and Leonardo Aniello <sup>2</sup>

<sup>1</sup> Robotics, Vision and Control Laboratory (ROVIS), Transilvania University of Brasov and Elektrobit Automotive, 500036 Brasov, Romania; tiberiu.cocias@unitbv.ro (T.C.); bogdan.trasnea@unitbv.ro (B.T.)

<sup>2</sup> School of Electronics and Computer Science, University of Southampton, Southampton SO16 7NS, UK; a.margheri@soton.ac.uk (A.M.); f.lombardi@soton.ac.uk (F.L.); l.aniello@soton.ac.uk (L.A.)

\* Correspondence: s.grigorescu@unitbv.ro

Received: 30 August 2020; Accepted: 16 September 2020; Published: 23 September 2020



**Abstract:** Self-driving cars and autonomous vehicles are revolutionizing the automotive sector, shaping the future of mobility altogether. Although the integration of novel technologies such as Artificial Intelligence (AI) and Cloud/Edge computing provides golden opportunities to improve autonomous driving applications, there is the need to modernize accordingly the whole prototyping and deployment cycle of AI components. This paper proposes a novel framework for developing so-called AI Inference Engines for autonomous driving applications based on deep learning modules, where training tasks are deployed elastically over both Cloud and Edge resources, with the purpose of reducing the required network bandwidth, as well as mitigating privacy issues. Based on our proposed data driven V-Model, we introduce a simple yet elegant solution for the AI components development cycle, where prototyping takes place in the cloud according to the Software-in-the-Loop (SiL) paradigm, while deployment and evaluation on the target ECUs (Electronic Control Units) is performed as Hardware-in-the-Loop (HiL) testing. The effectiveness of the proposed framework is demonstrated using two real-world use-cases of AI inference engines for autonomous vehicles, that is environment perception and most probable path prediction.

**Keywords:** autonomous vehicles; self-driving cars; artificial intelligence; deep learning; cloud computing; edge computing

## 1. Introduction

Among the new technological trends arisen in the last decade, Autonomous Driving has gained a lot of attention, with significant effort and resources invested by both academia and enterprises. The breakthroughs in self-driving cars have been made possible by the emergence of novel algorithms and real-time computing systems in the field of Artificial Intelligence (AI), deep learning (DL), Internet-of-Things (IoT) and Cloud computing [1].

An Autonomous Vehicle (AV) is an intelligent agent that observes its environment, makes decisions and performs actions based on these decisions. Autonomous driving applications aim to enable vehicles to automatically control their behavior. Traditionally in classical cars, these aspects and functionalities are the exclusive responsibility of human drivers, including an understanding of the driving scene, path prediction, behavioral planning and control. The development and deployment of accurate AI models, which can automatize these tasks is fundamental to support the progress of autonomous driving. Addressing these challenges lies within the realms of Data Science, with AI and DL proving to be successful tools to generate sufficiently accurate representations

of real-world processes. Deep learning (DL) and Deep Neural Networks (DNN) have become leading technologies in many domains, enabling vehicles to perceive their driving environment and take actions accordingly. Although several cloud-based solutions have been proposed for the automatic reconfiguration and adaptation of distributed platforms in Industry 4.0 [2], automotive grade systems have been scarcely reported.

Nowadays, numerous software tools are available, providing solutions for each particular stage in the development of an AI-based application. A key benefit for a driving function engineer is to have a comprehensive framework for enabling both the prototyping of the AI components, as well as their deployment and evaluation on target Edge devices such as Electronic Control Units (ECUs). General AI tooling systems such as MATLAB [3], Amazon Web Services (AWS) [4] or AI-One [5] provide functionalities to enable users to design and integrate custom AI modules. However, these tooling systems are neither tailored for the automotive sector, nor for the specificity of the autonomous driving applications. One of the major drawbacks is their lack of alignment to the automotive functional safety standard ISO26262 [6], particularly to the ASIL (Automotive Safety Integrity Level) hierarchy. On the other hand, major automotive and chip manufacturing companies are developing their own in-house AI frameworks, such as Tesla's Full Self-Driving and *Dojo* computers Tesla's Dojo—<https://syncedreview.com/2020/08/19/tesla-video-data-processing-supercomputer-dojo-building-a-4d-system-aiming-at-l5-self-driving/>), NVIDIA's *DriveWorks* platform (NVIDIA DriveWorks—<https://developer.nvidia.com/drive/driveworks>), or Uber's Michelangelo training framework (Uber's Michelangelo—<https://eng.uber.com/michelangelo-machine-learning-platform/>). The problem with these tooling systems is that they are either designed to be used solely by the developing company (e.g., Tesla, Uber, etc.), or are built to operate only on specific hardware, as in the case of NVIDIA's AI solutions.

In this paper, we first analyze the development process of AI-based autonomous driving applications and identify its current limitations. Secondly, we propose the *Cloud2Edge Elastic EB-AI Framework* as a revised prototyping and deployment workflow for building AI based solutions for autonomous driving. We introduce the novel concept of an *AI Inference Engine*, which encompasses a simple yet elegant solution to the training, evaluation and deployment of AI-based components for self-driving cars. Throughout the paper, we refer to a real scenario, where autonomous driving applications are developed at Elektrobit (EB) Automotive (Elektrobit Automotive—<https://www.elektrobit.com/>), a world-wide supplier of embedded and connected software products and services for the automotive industry.

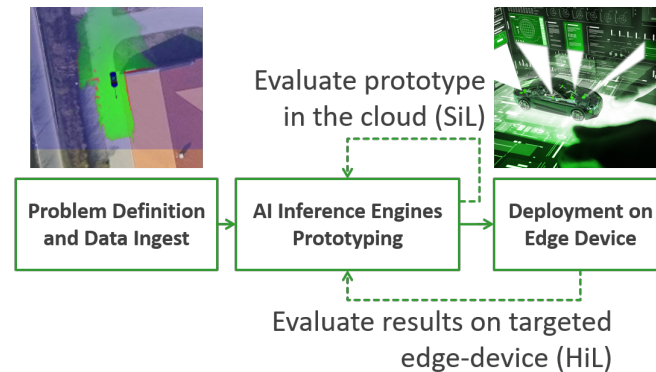
### 1.1. The EB-AI Framework

The key differences that the EB-AI framework provides with respect to tooling systems such as the ones listed above are:

- Ability to digest automotive specific input data such as video streams, LiDAR and/or radar data;
- Prototyping and development of AI Inference Engines in the Cloud using an SiL paradigm, agnostic of low-level AI libraries such as TensorFlow [7], PyTorch [8], or Caffe [9];
- Providing tunable state of the art DNN architectures for a broad spectrum of autonomous driving applications;
- Ability to deploy and evaluate the obtained AI Inference Engines via HiL testing on edge devices (e.g., target ECUs).

The diagram in Figure 1 illustrates the development workflow of an AI Inference Engine based on the EB-AI framework, according to our own data driven V-Model detailed in Section 3. Once the specific problem to solve has been defined and enough data have been collected, that data is ingested at a prototyping level by a DNN, which stands at the core of the inference engine. During this stage, the DNN is trained, evaluated and refined within the Cloud according to the Software-in-the-Loop (SiL) principles. Finally, the inference engine is deployed on a target Edge device inside a vehicle

and evaluated again in real-world scenarios as Hardware-in-the-Loop (HiL). This process allows us to refine the engine even further, according to a continuous feedback loop aimed to keep improving applications over their entire life span. The novelty of the AI Inference Engine concept lies in the effective integration of the DNN's training, evaluation as SiL and testing using the HiL paradigm.



**Figure 1.** AI Inference Engine concept. Deep learning based inference engine evaluated at the prototyping level within the Cloud as Software-in-the-Loop (SiL), as well as on the targeted edge device Electronic Control Unit (ECU), as Hardware-in-the-Loop (HiL).

A key aspect to consider in the EB-AI workflow is the huge amount of data required to train DNNs. This data can be either synthetically generated (first stage) or collected via real sensors mounted on test vehicles (last stage). In both cases, the data has to be made available at training the stage. Although the training of DNNs can be highly parallelized [10], with unprecedented levels of parallelization reached by leveraging the recent wide availability of GPUs, it is still prohibitive to have enough computational power to process the amount of data required by autonomous driving applications. As a consequence, all Cloud computing providers offer commoditized AI solutions (e.g., IBM's Watson AI, Amazon AI, Microsoft's Azure AI and Google's Cloud AI) for scalable computing training facilities. However, the following limitations hold at the AI Inference Engine prototyping stage, when large quantities of data have to be uploaded to the Cloud: (i) data upload impracticality due to bandwidth bottlenecks and latency [11], (ii) difficult enforcement of privacy requirements (e.g., GDPR [12]), since part of the collected raw data can be sensitive and thus cannot be shared with cloud providers.

### 1.2. Contributions

We propose a solution to overcome the presented limitations by exploiting recent developments in Over-The-Air (OTA) systems for AVs [13,14], enabling the use of an AV as an additional computational device. These computing resources localized on the Edge ECUs, also referred to as *Edge computing*, can be integrated with Cloud resources to provide an *Elastic* distributed infrastructure in which the training can be collaboratively carried out [15,16]. Furthermore, Edge devices can limit the data transmitted to the Cloud by pre-processing local raw data or pre-training partial models [17], which in turn can also help mitigate privacy concerns, since raw data sharing can be directly avoided. This kind of hybrid deployment over the Cloud and Edge resources calls for a modular framework for AI-based autonomous driving applications. To this aim, the proposed EB-AI architecture provides a modular toolchain that enables the deployment of autonomous driving applications across Edge and Cloud resources.

Therefore, the main contributions of the paper are:

- A simple yet elegant AI Inference Engine concept, based on the SiL and HiL principles;
- A data-driven V-Model approach guiding the design of AI-based autonomous driving applications;
- A modular Cloud2Edge AI framework for autonomous driving applications coined EB-AI;
- An elastic framework able to overcome network bandwidth and privacy concerns by dynamically deploying deep learning training tasks among Edge and Cloud;

- The development of two real-world AI Inference Engines for *environment perception* and *most probable path prediction*;
- A discussion on the advantages of such a hybrid deployment, in terms of training parallelization, privacy-preservation, fault tolerance and scalability.

*Structure of the Paper.* Section 2 reports the most relevant background concepts and related work. Section 3 presents the V-Model approach and the proposed EB-AI framework. Section 4 discusses the deployment of the architecture across Cloud and Edge devices. Finally, two relevant AI Inference Engines are detailed in Section 5, while the conclusions are stated in Section 6.

## 2. Background and Related Work

In the following, we overview the main concepts of DL, report on general DL tooling systems and highlight the main limitations of available tools and learning environments for automotive applications.

### 2.1. Deep Learning Overview

*Deep learning* is a branch of machine learning that leverages on large DNN architectures to build a layered representation of the input data [18]. DNNs are universal non-linear function approximators, built using multiple hidden layers, and can be classified based on their architecture. *Convolutional Neural Networks* (CNNs) are mainly used for processing spatial information, such as images, and can be viewed as image feature extractors. Recurrent Neural Networks (RNN) are especially good in processing temporal sequence data, such as text, or video streams. Different from conventional neural networks, an RNN contains a time dependent feedback loop in its memory cell. Due to the high number of layers, DNNs are difficult to train using vanilla backpropagation algorithms. As an example, in the case of auto-encoders [19], each layer is trained separately, then backpropagation is used to train the whole network. To avoid overfitting (i.e., trained models that cannot generalize or predict unseen values), common methods such as Dropout [20] are used to regularize the training.

DNNs are effective in domains characterized by a high number of features, such as computer vision and natural language processing. Training of DNNs critically requires large annotated datasets, which have been increasingly released in the last decade, especially in the computer vision field. One of the first large collections of annotated images is the ImageNet database [21], which contains 1.5 mil images representing 1000 object classes in a variety of shapes, poses and illumination conditions. Although images relevant to autonomous driving systems are only a subset of the whole collection, the complete database was used to pre-train the first convolutional layers of object detectors used for driving scene perception [1].

In the last years, mainly due to the increasing research interest in autonomous vehicles, many driving datasets were made public. These vary in size, sensory setup and data format, commonly comprising of synchronized ego-data, video, LiDAR, Radar, ultrasonic, inertial measurements and GPS datastreams. Among others, popular ones are the KITTI Vision Benchmark dataset (KITTI) [22] by the Karlsruhe Institute of Technology (KIT), NuScenes [23] by Aptive, or Cityscapes [24] by Daimler, Max Planck Institute and TU Darmstadt.

Since acquiring large training datasets is a demanding process, usually performed via manual annotation, alternative methods have been explored for synthetic data generation and simulation engines. In our previous work we have proposed and patented a semi-parametric approach to one-shot learning, coined *Generative One-Shot Learning* (GOL) [25]. Given as input single one-shot objects (or generic patterns and templates), together with a small set of regularization samples, GOL uses the sample to drive the generative process by outputting new synthetic data. Specifically, it permits to generalize on unseen data, while increasing the classification accuracy on synthetic data as much as possible.

## 2.2. Deep Learning Libraries and Tools

Some of the well-established low-level AI libraries are Tensorflow [7] by Google, PyTorch [8] by Facebook, the Cognitive Neural ToolKit (CNTK) (CNTK—<https://docs.microsoft.com/en-us/cognitive-toolkit/>) by Microsoft, and Caffe2 [9] from Berkeley University. The APIs are mainly available for C/C++ and Python languages, although other notable implementations exist, such as Weka [26], which is written in Java. For facilitating their usage and large scale adoption, these libraries are integrated into high-level frameworks (e.g., Lasagne (Lasagne—<http://lasagne.readthedocs.org/en/latest/>)).

For fast computation, training commonly takes place on GPUs interconnected as clusters, which are available as a service on the major Cloud platforms, for example, Microsoft Azure and Amazon AWS. After training, the DNN models are downloaded into edge devices for inference. In this case, edge devices can be smartphones, PC commuters, or target ECU devices part of automotive applications.

In order to exploit parallel computation, the training can be distributed based on data partition (e.g., among threads, different machines, GPUs) or model partition [27]. Scalable deployments are enabled by distributed computing platforms such as Storm [28] and Spark [29]. In particular, Trident-ML (Trident-ML—<https://github.com/pmerienne/trident-ml>) is a real-time library for machine learning upon Storm, while CaffeOnSpark [30] and TensorSpark [31] are Spark integrations of Caffe and TensorFlow, respectively.

To overcome bandwidth and privacy issues when uploading data to the Cloud, federated learning [32] offers a technical framework to scale learning horizontally or vertically across devices [33]. Edge computing is therefore exploited to carry out pre-training, the results of which are integrated on the Cloud and then offloaded on the Edge for on-device inference [34].

The above-mentioned frameworks are designed solely for training purposes within the Cloud. In comparison, the main advantage of EB-AI over them is our proposed elastic approach, which considers both training and inference as a joint task distributed onto powerful GPU clusters residing in the Cloud and Edge computing devices available in each autonomous vehicle belonging to a connected fleet of cars. Additionally, as opposed to the more general AI frameworks presented in this subsection, EB-AI has been designed specifically for automotive purposes, following the constraints of Automotive SPICE and the ISO26262 [6] functional safety standard, along with a proposed data driven V-Model for building AI Inference Engines.

## 2.3. AI Frameworks in Autonomous Driving

The application of deep learning, and more general of AI, to the automotive industry has grown significantly in the last few years. AI and deep learning are used to obtain human-like behaviors in automated driving [35], such as environment perception and scene understanding (e.g., detecting traffic signs, traffic participants, road obstacles), or trajectory planning and control. Deploying AI Inference Engines inside autonomous vehicles requires overcoming limitations of platform dependencies and limited computation resources. Therefore, new transportation vehicles are equipped with embedded ECU devices specialized for AI and exploited to ensure adequate performance and energy efficiency [36].

Overall, the increasing interest in AI Inference Engines for the automotive industry has been driven by the field of autonomous vehicles. While a number of AI-enabled functions are already deployed in cars and interacting with drivers [37], it is predicted that full-fledged AVs, that is SAE levels 4 and 5 [38], will require significant additional cost for developing the new required AI solutions [39]. As such, there is the need to optimize the design, deployment and maintenance of automotive AI Inference Engines by reducing their complexity and by improving accuracy via the access to (real-time) data. Lu et al. proposed a work close to ours [40], where they introduced distributed learning based on self-driving cars as Edge computing devices. However, they only train local models on the cars and then aggregate the parameters of the neural network on a Parameter EdgeServer. Conversely,

we propose a solution to split training tasks between Cloud and Edge, in order to deploy training tasks, which process sensitive data only over the Edge while deploying the other over both Cloud and Edge.

Car manufactures and their suppliers are both developing their own in-house AI frameworks. Example of such systems are Tesla's *Full Self-Driving* and *Dojo* computers, NVIDIA's *Drive Works* platform, or Uber's *Michelangelo* training framework. However, these are either closed software platforms, built for in-house operations, or are designed to work only on dedicated hardware, as in the case of NVIDIA's software solutions. Based on Elektrobit's history as a key supplier of automotive grade software and real-time operating systems, we have designed EB-AI for interoperability, incorporating optimizations of AI Inference Engines for different embedded ECU devices, coupled with the necessary functional safety standards required by the automotive industry.

### 3. EB-AI Toolchain

In the following, we describe the proposed EB-AI toolchain used to design, train and evaluate the AI Inference Engines. The engines represent application deployment wrappers for trained DNN models, which are used to build driving functions inside EB Robinos<sup>®</sup>, which is Elektrobit's autonomous driving system (EB Robinos—<https://www.elektrobit.com/products/automated-driving/eb-robinos/>).

#### 3.1. V-Model and Workflow

Automotive software engineering still demands a robust and predictable development cycle. The software development process for the automotive sector is subject to several international standards, namely Automotive SPICE and ISO 26262 [6]. Accepted standards, as far as the software is concerned, rely conceptually on the traditional V-Model development lifecycle. It is important to approach deep learning from a more controlled V-model perspective in order to address a lengthy list of challenges, such as the requirements for the training, validation and test datasets, the criteria for the data definition and pre-processing, as well as the impact of hyperparameters tuning.

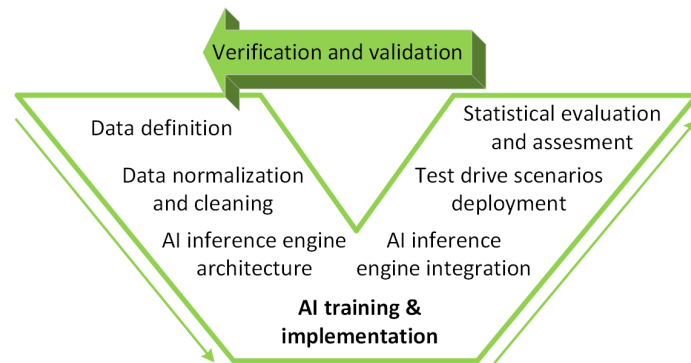
In Figure 2 we propose a data driven V-Model for prototyping and development within EB-AI. The first two steps regard *data definition* and *data normalization and cleaning*. These are used to define the *AI inference engine architecture*, where we properly configure the layers of the DNN and select the input training data according to the pre-processed datastreams. We then have the *AI training and implementation* step, where we train the DNN model to the requirements of the application. The trained model is converted automatically to a format that can be deployed within the autonomous vehicle. The converted model represents the *AI Inference Engine*, which is integrated and tested in different driving scenarios. According to the results obtained at different levels within the data driven V-Model, we can start the *Verification and Validation* process, where we refine the inference engine.

This circular approach allows the EB-AI architecture to employ continuous learning, where DNNs can improve their accuracy overtime. To perform this task, we proposed a generative technique to compute synthetic data that can be used in the training process based on the GOL algorithm [25].

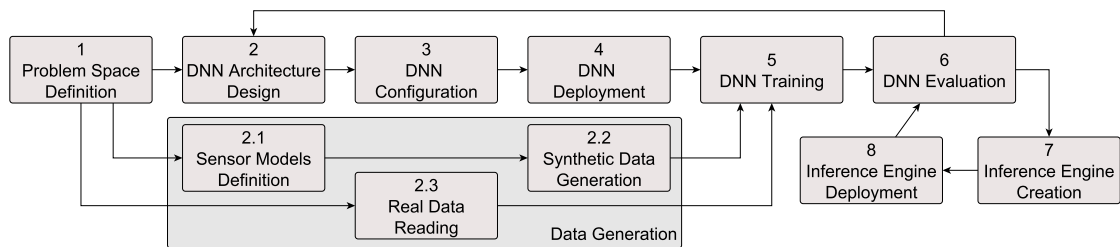
Figure 3 details the workflow for AI Inference Engines design, deployment and testing based on the data driven V-Model. Phase 1 consists of defining the problem space and collecting the necessary data for the training process. In this phase, the data is pre-processed, annotated, normalized and filtered. In phase 2 we design the DNN architecture, activation functions, structure of the hidden layers and output nodes. Phase 3 deals with the tuning of the DNN, that is, all hyperparameters necessary for the training phase, such as the learning rate, loss function, regularization strategy and accuracy metrics.

Once the DNN setting has been completed, we deploy the DNNs among the cloud nodes in phase 4 and start the training in phase 5. In this phase we feed the DNN both with real sensory data, as well as with synthetic data, generated in a parallel workflow branch (phases 2.1–2.3). After training, the DNN is evaluated and redesigned if necessary (phase 6). The output of the training and evaluation

phases is the *AI Inference Engine* (phase 7). This is then deployed among the edge device in the car (phase 8) and evaluated again with respect to in-car performance measures.



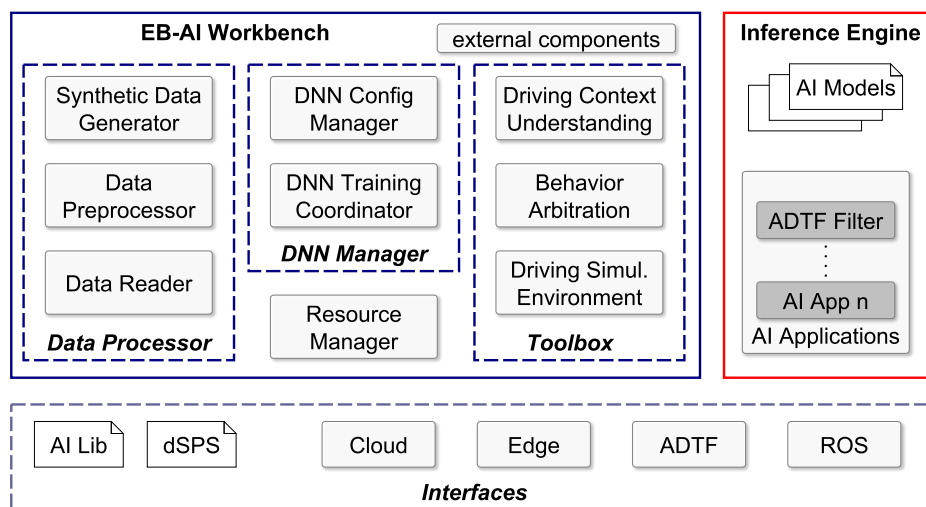
**Figure 2.** Data driven V-Model development of AI Inference Engines. The proposed V-model encompasses a synergy of data driven SiL and HiL design, training and evaluation of AI components.



**Figure 3.** AI Inference Engines design, training and evaluation workflow within Elektrobit (EB)-AI. The workflow builds on top of the proposed data driven V-model from Figure 2.

### 3.2. EB-AI Architecture

The EB-AI cloud architecture is depicted in Figure 4. It is based on two main components: (i) the *EB-AI Workbench* and (ii) the *Inference Engines*. The former generates and trains AI models, while the latter uses such models for prediction and inference on different embedded ECUs.



**Figure 4.** EB-AI cloud architecture. The framework follows the two major operations from the DL workflow: design and training in the EB-AI Workbench and inference. The interfaces allow for flexibility with respect to the core AI libraries, edge device hardware and type of application deployment wrapper.

Specifically, the EB-AI Workbench is responsible for prototyping, training, and development of the AI models and is composed of four main modules: (i) *Data Processor*, (ii) *DNN Manager*, (iii) *Resource Manager* and (iv) *Toolbox*.

**Data Processor.** It reads, pre-processes and analyzes the data provided by the vehicles' sensors. Data pre-processing typically deals with detecting outliers in the input data, visualizing it using feature space reduction techniques such as PCA (Principal Components Analysis) and t-SNE (t-distributed Stochastic Neighbor Embedding), checking its variance and standard deviation across its data distribution domain, splitting it into appropriate train-evaluation-test batches and finally normalizing it. The probable capacity of the DNN model is determined based on the quantity of training samples. Furthermore, this stage is used to generate synthetic data through our patented GOL algorithm [25]. This allows to learn patterns of real data and generate artificial samples to improve the data available when training the DNN.

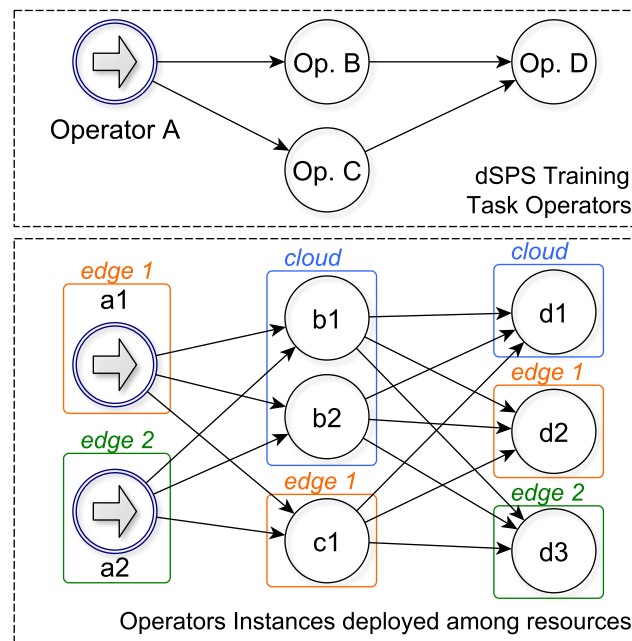
**DNN Manager.** It contains two modules: the *DNN Config Manager* handles the configuration of the DNN, while the *DNN Training Coordinator (TC)* manages the training process. Similar to the work in [2], DNN models are automatically configured based on the results obtained from pre-processing the training data. By analyzing the structure and type of input data, the DNN Config Manager automatically determines the input and output shapes of the DNN architecture, thus enabling the usage of existing network architectures from the workbench. The first step is to automatically define the input and output layers of the model based on the input–output data types. The input layer's shape is given by the structure of the training samples (e.g., images, LiDAR, radar, ultrasonic and IMU measurements, log traces, GPS tracks), while the output layer's shape is computed based on the labels' type. For example, if the training data is composed of image samples and 2D bounding boxes, then different network architecture for spatial object recognition is suggested, as shown in use-case Section 5.1. Similarly, as presented in Section 5.2, different RNN architectures are proposed if the data is structured as temporal sequences. In this first step, training samples and labels are aggregated into single input and output layers, respectively. The automatic configuration of the input–output layers can be manually tuned if the samples and labels are composed of multiple types of sensory measurements (e.g., images and LiDAR) and/or multiple prediction heads (e.g., 2D object detection and 3D object reconstruction).

Once the input–output shapes of the DNN model have been defined, we proceed to computing the architecture of the inner layers. We split the anatomy of a DNN architecture into three main blocks: (i) backbone model, (ii) feature extractor model and (iii) prediction heads. These are chosen from subsets of state-of-the-art DNN architectures, such as VGG16, MobileNetV2 and Darknet for the backbone, or YoloV3–V5, RetinaNet and SSD (Single Shot Detector) for the case of prediction heads in object detection. We compute the DNN structure depending on the model capacity determined at the Data Processor stage. Namely, in order to avoid overfitting if training data is scarce, we use a light-weighted backbone, such as VGG16, coupled to prediction heads having a lower number of convolutional layers. Different DNN models are suggested based on this analysis. The models are subsequently adapted and modified in our SiL and HiL paradigm for obtaining an optimal AI Inference Engine for the task at hand.

The TC employs libraries for DL on top of a Distributed Stream Processing System (dSPS), allowing us to run a distributed *training application* on the dSPS. Such an application is represented as a directed acyclic graph, where vertices represent *operators* and edges represent streams between pairs of operators. Each operator carries out a piece of the overall computation, which is a *training task*, and can be distributed over multiple instances.

**Resource Manager.** It is the elastic manager that dynamically handles the deployment of the DNN training tasks over Cloud and Edge resources. Figure 5 shows an example of training tasks for a dSPS application composed of four operators (A, B, C, D) allocated over the cloud and two edge nodes.





**Figure 5.** Training Tasks Operators of the dSPS. In this example we show an instance where data managed by the operator A is sensitive, thus the RM deploys operator instances a1 and a2 only on the two edge nodes. Operator B can instead be performed in the Cloud, while Operator C, which produces the AI Inference Engine, is available on all nodes.

The Resource Manager (RM) is able to reconfigure over time (i) the number of edge/cloud resources, (ii) the operators' parallelism and (iii) the allocation of operator instances to available resources. Thus, RM can redistribute dynamically the training tasks on a different pool of resources to keep desired performances according to the input workload (e.g., the input rate of values read by the vehicles). The operators of the dSPS application can be distributed among edge and cloud nodes according to how much resources are necessary for tasks, as well as according to the sensitivity of the acquired training data, ensuring that sensitive data is processed where it is produced, without propagating it to other cloud or edge nodes, thus mitigating privacy. Once the learning is completed, the trained model is replicated among all edge and cloud nodes so that vehicles can immediately have the latest version available.

**Toolbox.** It is a set of components specific for automotive AI applications. The training itself feeds the previously defined data to the network, allowing it to learn a new capability by reinforcing correct predictions and correcting the wrong ones. The module *Driving Context Understanding* classifies the driving context from grid-fusion information, the *Behavior Arbitration* understands the driving context and optimizes the strategy from real-world grid representations, while the *Driving Simulation Environment* trains, evaluates and tests AI algorithms in a virtual simulator, such as Microsoft's AirSim.

**Inference Engine.** It contains the trained AI models from the EB-AI Workbench and deployment wrappers (or application) for the models. For example, such a wrapper can be built as a Robotic Operating System (ROS) Node (ROS—<https://www.ros.org/>), or an EB Assist ADTF<sup>®</sup> component. ADTF stands for *Automotive Data and Time-Triggered Framework* (ADTF—<https://www.elektrobit.com/products/automated-driving/eb-assist/adtf/>) and is a tool for the development, validation, visualization and testing of advanced driver assistance systems and automated driving features.

**Interfaces.** Both EB-AI Workbench and Inference Engine use an interface layer for low-level libraries (e.g., CUDA, Tensorflow, PyTorch, CNTK, or Caffe2), integration with the dSPS engine (e.g., Spark) and the APIs for cloud and edge nodes. Furthermore, it provides a specific interface to ADTF and ROS.

## 4. Hybrid Deployment Advantages

This section discusses the advantages brought by the proposed hybrid deployment. Specifically, it discusses how our solution preserves privacy (Section 4.1) and how it improves fault tolerance and scalability (Section 4.2). In Appendix A we also provide an experimental evaluation aimed to show the benefit of increased parallelism in DNN training.

### 4.1. Privacy-Preserving Techniques

Although the Cloud can offer a scalable training infrastructure, some training data should not flow to the Cloud due to data protection requirements, as those on personal data enforced by the GDPR. Both data encryption and obfuscation solutions can be used to address these challenges.

Data obfuscation techniques like Differential Privacy (DP) have been successfully used both in centralized and federated learning [41]. Differently from pre-processing techniques, using DP on Edge devices as part of a federated learning approach guarantees better accuracy [15,32]. In particular, Edge devices are carrying out the training of a number of front layers by using DP on the sensitive data. Then, they transmit the intermediate results to the Cloud for the completion of the training and dissemination to the Edge of the updated layers.

Learning on encrypted data has also been proposed. While homomorphic encryption has been prototyped on centralized Cloud training [42], applications of Multi Party Computation (MPC) offers distributed privacy-preserving training [43].

The EB-AI toolchain supports DP learning by relying on the TensorFlow Privacy library (TensorFlow Privacy—<https://github.com/tensorflow/privacy>). Differently from MPC learning that must rely on the ad-hoc computing framework (e.g., Sharemind [44]), DP learning can be managed by the EB-AI framework to take advantage of Edge devices on self-driving cars. The TensorFlow Privacy library relies on a differentially private stochastic gradient descent (SGD) algorithm [45], which provides privacy protection for deep neural networks while incurring a tolerable overhead in terms of training time.

### 4.2. Fault Tolerance and Elastic Scalability

The EB-AI architecture can support flexible learning deployment across Cloud and Edge, offering scalable execution of learning tasks according to resource and privacy constraints.

Despite the advances in OTA software, the deployment of learning tasks must tolerate faults due to lack of connectivity with AVs. Distributed computing frameworks, such as Apache Spark, can be deployed to enable elastic fault tolerant federated learning via, for example, TensorSpark and CaffeOnSpark.

The EB-AI Resource Manager extends the ELYSIUM [46] autoscaler in order to rebalance the operators of the training application. Specifically, it elastically scales both the dSPS operators' parallelism and resources (cloud and edge nodes), allocating these operators according to workload changes or faults predicted and/or observed.

The flexibility of the task definition provided by existing dSPS platforms, together with their seamless integration with high-level AI frameworks, allows the EB-AI workbench to directly introduce new pre-processing or pre-training tasks at the Edge according to emerging privacy and learning constraints.

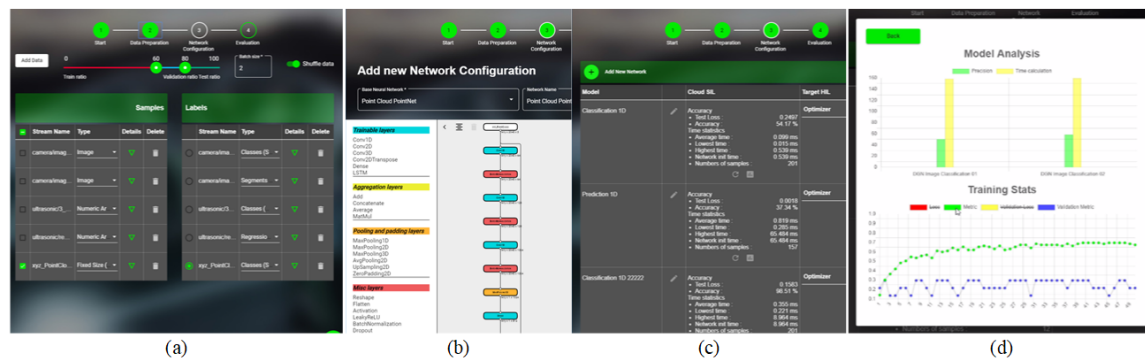
## 5. EB-AI Use Cases

State of the art autonomous driving systems are defined using modular perception-planning-action pipelines, in which the main problem is divided into smaller sub-problems, each module being designed to solve a specific task and deliver the outcome as input to the adjoining component [1].

In the following, we aim to highlight the functionalities of the proposed EB-AI framework in two autonomous driving application use-cases, namely *driving environment perception* and *most probable*

*path prediction*. Since this article focus on the underlying computing framework, as opposed to a specific algorithm, we will not focus on the accuracy of the two use-case methods, but on their designed, deployment and evaluation using the proposed EB-AI framework.

In both use cases presented below, we use the proposed Elastic Cloud2Edge paradigm to perform data ingestion and the configuration of the DNNs using EB-AI's Workbench interface from Figure 6, while the *DNN Manager* and *Resource Manager*, detailed in Section 3, are used to jointly distribute the training on the Cloud and the AVs' Edge devices.



**Figure 6.** Snapshots from the EB-AI Cloud interface. (a) Data upload and pre-processing. (b) DNN architecture design. (c) Deployment and evaluation in the Cloud and Edge ECU as SiL and HiL, respectively. (d) Models comparison.

### 5.1. Use Case 1: Driving Environment Perception

The perception of the driving environment is a key part of any autonomous driving system, characterized as the ability of the inference engine to understand the image scene and accurately represent the dynamic and stationary objects (e.g., cars, pedestrians, traffic signs). The EB-AI framework offers ready to use DNN architectures, such as SegNet for traffic scene image segmentation [47], different versions of the Yolo detector [48] and our own Deep Grid Net (DGN) for driving context understanding [49]. For the use case of environment perception, the specific problem space of object detection in images will be considered.

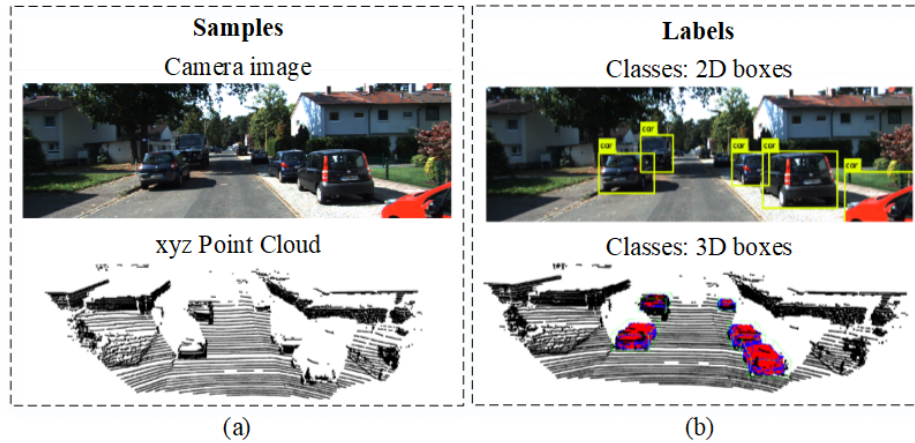
For environment perception, we perform object detection and classification on RGB color images and XYZ point clouds, as illustrated in Figure 7. We used labeled data from the CamVid dataset [50], which we have converted into the ROS bags. The ROS bags are uploaded into the EB-AI Cloud for pre-processing operations such as normalization, encoding or clipping.

Figure 6 shows a couple of snapshots from the EB-AI Cloud interface, depicting the data upload, DNN architecture design and training as SiL, deployment and evaluation via HiL and models comparison. The first steps in building an AI inference engine are the definition of the requirements and collection of the necessary data for training. The data are collected from two sources, namely vehicles' sensors and synthetic generation through our GOL algorithm [25]. This allows the improvement of the patterns learned from real data and the ability of the trained network to generalize. A manual splitting ratio of the data between train, validation and test is available, as well as automatic k-fold variations.

By analyzing the input data the *DNN Config Manager* automatically determines the input and output shape of the DNN architecture, thus enabling the usage of existing network architectures from the workbench. Additional adjustments can be made in order to improve the accuracy of the DNN, based on the task at hand. In this driving environment perception use-case we use a Yolo V5 DNN trained on pairs of camera images and annotated 2D bounding boxes of objects of interest.

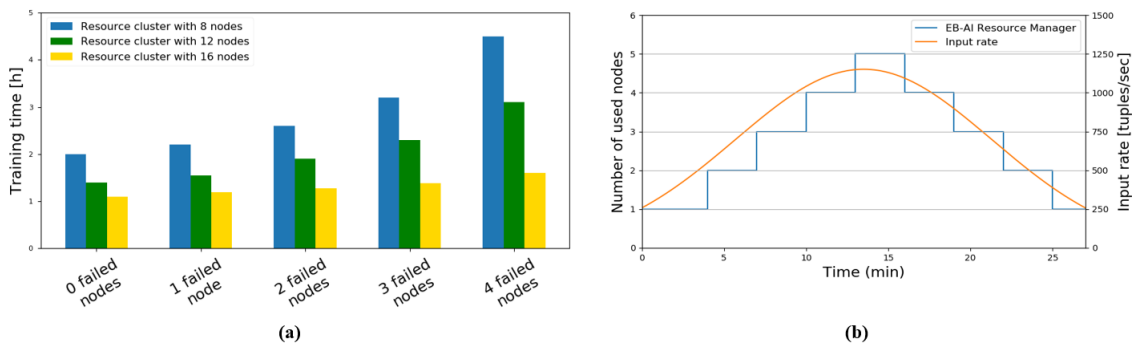
The training is managed by the *DNN Training Coordinator*, which handles the DL libraries on top of a distributed streaming processing system. The Resource Manager handles the available training resources over the cloud and/or edge device. In order to illustrate the fault tolerance and scalability of the training process, we have used different ECU configurations with 8, 12 and 16 nodes respectively.

The ECUs are a baseline desktop computer equipped with an Intel Core i9 9900K CPU, 64 GB RAM, with two high-performance NVIDIA GeForce RTX 2080 Ti graphics cards, an NVIDIA Jetson AGX platform, a Kalray Konic board and a RaspberryPi 3B+ computer.



**Figure 7.** Pairs of samples and labels commonly used in environment perception DNNs. (a) Input RGB image and 3D point cloud. (b) Labels of objects of interest annotated in 2D on images and in 3D onto point cloud data, respectively.

Figure 8 shows how the training resources are used in the context of fault tolerance and elastic scalability. In order to evaluate the fault tolerance, specific resource nodes were deliberately blocked, while the impact on the overall training time was measured. As it can be seen in Figure 8a, with a hardware setup comprised of 16 nodes, the impact of failing one to four nodes is negligible on the training time. However, the impact becomes visible when four nodes (50% of total available nodes) are blocked in a hardware setup of eight total nodes. Figure 8b shows the number of used nodes in the training process while injecting an input stream curve of 250–1200 tuples (subset data) per second. The EB-AI Resource Manager reactively allocates or deallocates nodes in order to compensate for the input demand.

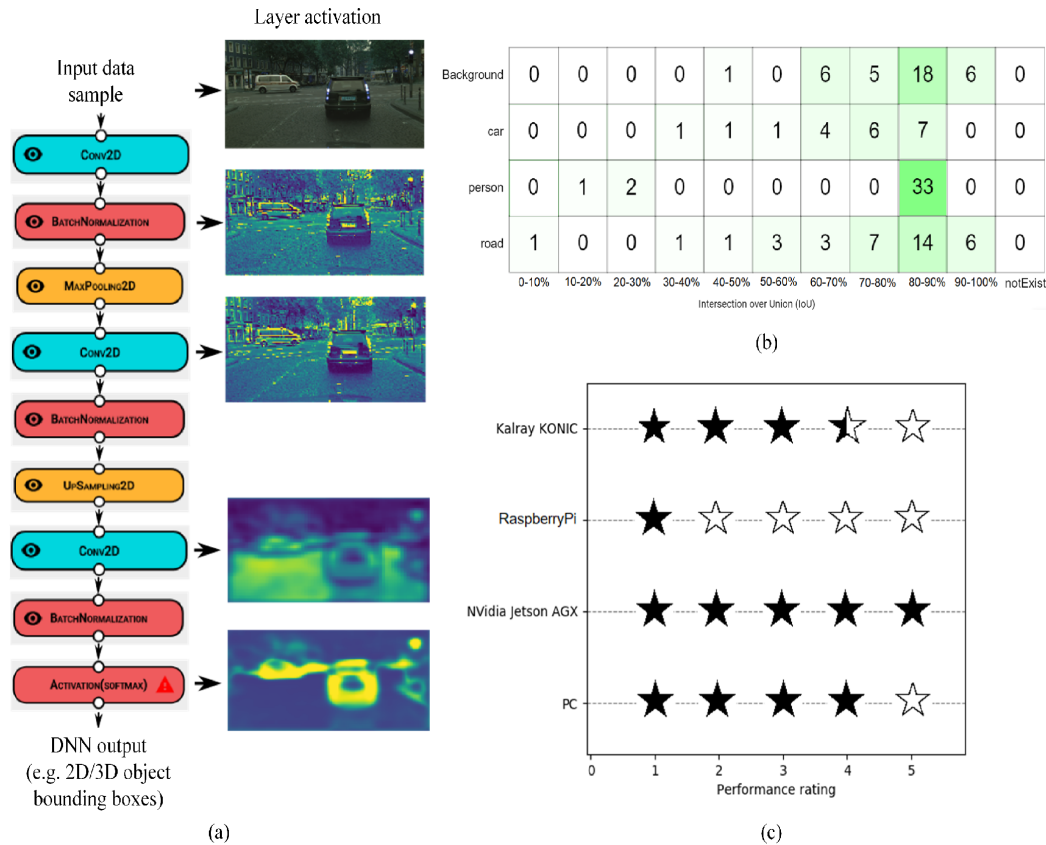


**Figure 8.** Usage of training resources in the context of fault tolerance and elastic scalability. (a) Fault tolerance based on number of failed nodes. (b) Elastic scalability of the EB-AI Resource Manager.

The performance of a DNN can be evaluated based on a series of metrics, depending on its architecture. Figure 9 shows different evaluation approaches for our object detection use-case based on direct visualization of the DNN layers' activations, Intersection over Union (IoU), as well as a performance estimator, which aggregates the overall accuracy of the model into a five-star rating system.

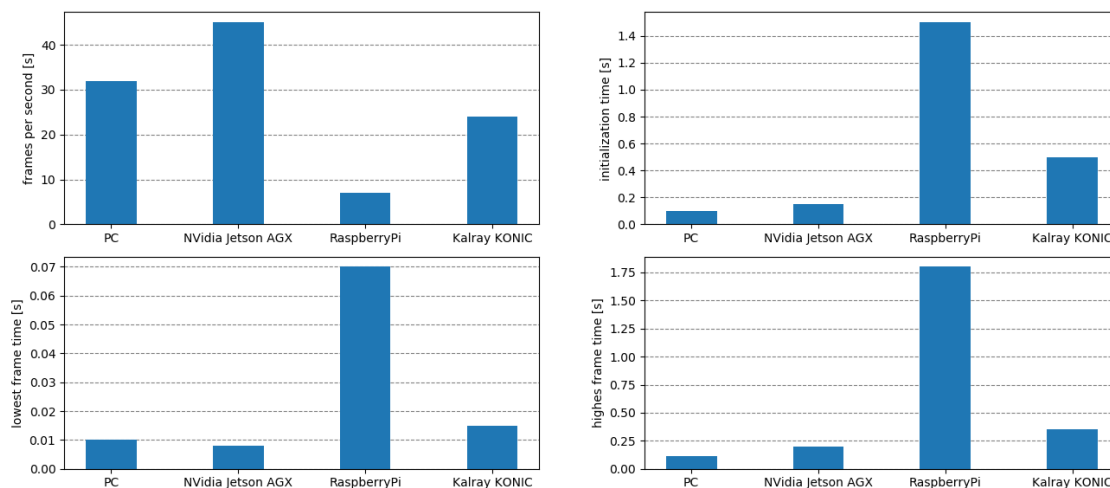
To assess the precision of object detection, we use the IoU quality measure, which is a percentage measure reflecting the amount of overlapping between the predicted bounding box, as outputted from the AI inference engine, and the labeled bounding box. An accurate model has an IoU value close

to 100%, which represents maximum accuracy. Figure 9b illustrates the IoU matrix determined for the trained Yolo V5 model. From the obtained IoU matrix, it can be observed that the majority of the detected objects overlap in a ratio of 80–90% with the ground truth. Featuring easy interpretability over the testing data, the user can trace back the samples where the evaluation failed or was not accurate enough, while iteratively improving the neural network until the desired performance is achieved.



**Figure 9.** Performance evaluation metrics within EB-AI. (a) Direct visualization of the activations of each DNN layer. (b) Intersection over Union (IoU) evaluation matrix describing in percentage the amount of overlapping bounding boxes. (c) Five-star rating aggregates overall accuracy.

The obtained AI Inference Engine contains the trained DNN model encapsulated in a deployment wrapper. In this case, it was deployed as an ADTF wrapper on Elektrobit’s Automotive Data and Time Triggered Framework (ADTF). Figure 10 shows different measure performance indicators for the four considered target ECU devices. The metrics quantify the computation time in frames per second, initialization time, as well as maximum and minimum computation time, respectively. The obtained results are proportional to the computation power of each ECU, yielding low framerates for low-range grade devices, such as the RaspberryPi, as opposed to the high performance of automotive ECUs, such as the NVIDIA AGX platform. The metrics are logged in the ECU device for later importing into the Cloud, where they are optimized by reconfiguring and retraining the DNN using the proposed data driven V-Model.



**Figure 10.** Real-time performance metrics obtained on different deployment ECU devices. The diagrams show the Frames per Second (FPS), AI Inference Engine initialization time, as well as the lowest and highest achieved framerate for four different ECUs.

## 5.2. Use Case 2: Most Probable Path Prediction

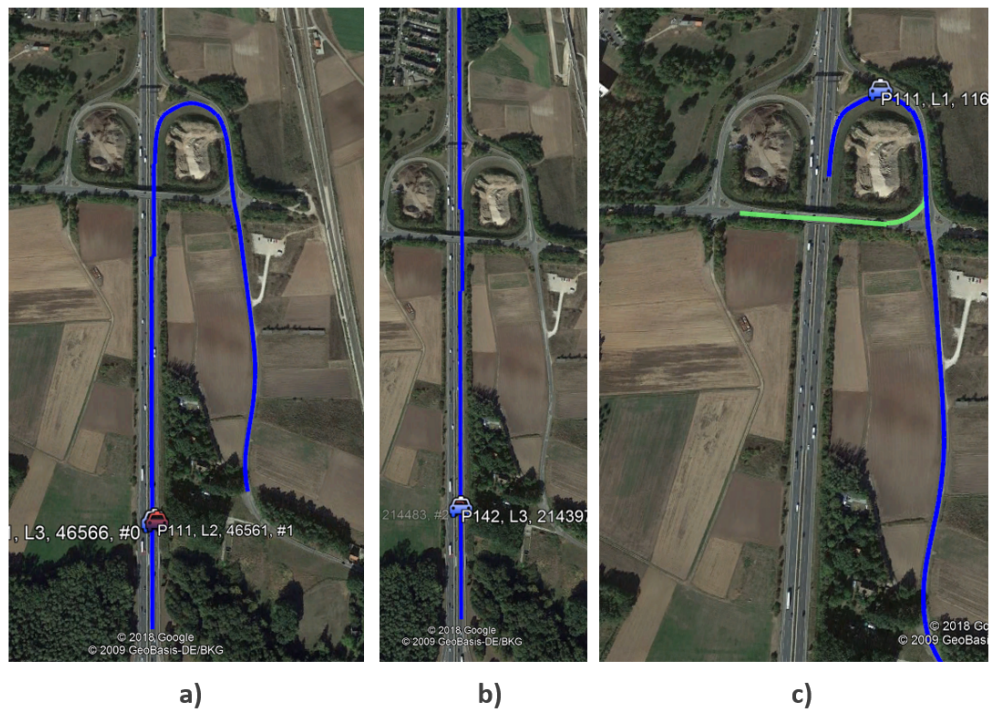
In autonomous driving and driver assistance, in order to prepare the vehicle for upcoming hazards or to warn the driver, it is important to predict in advance the route that the vehicle will follow over a future time horizon. Such algorithms, illustrated in Figure 11, are called Most Probable Path (MPP) estimators. With an increasing level of map details, past trips of vehicles can be used to store information about the road paths that are likely to be reached within a specified distance using GPS observations. Using this information, the MPP algorithm can learn and further predict the probable road network ahead. As the vehicle travels forward, traveled road segments are left behind, while new ones are added to the MPP.

A map is a web of paths implemented using a tree structure, where each node in the tree represents a link between two segments of the road. The root node of the tree represents the link where the vehicle is currently located. Starting from a link, the vehicle can travel to a considerably large number of possible paths. A transition from one link to a reachable successor is represented by adding a child to the proper parent node in the tree. This is important because as the vehicle travels forward, nodes that become unreachable are removed from the tree. Every time the vehicle transitions from one link to the next, at least one node becomes unreachable and will be removed from the graph.

In order to predict where the vehicle is heading in the future, the MPP method needs to calculate the likelihood associated to each possible turn. In the classical way, the weight associated to each turn, or transition, depends on many factors and considers the turning angle and the road class of the following road segment. A smaller turning angle results in a higher probability of driving straight, whereas a bigger turning angle results in a smaller probability of following that specific road segment.

The main drawback of the classical approach in computing the MPP is that in most cases these predictions are performed at a small distance (a few hundred meters) in front of the current position, since the computation time required for computing the MPP probabilities is high. Additionally, the result will always be similar, since the turning angle and road class will be constant on the same map, so there will never be any difference between different drivers or different cars.

In this use-case, we have trained and deployed an AI Inference Engine using GPS data collected with Elektrobot's autonomous test vehicle, shown in Figure 12. We have performed 25 trips of 40 km each. Using the EB-AI framework, we were able to enhance the links using the pre-processing operations. As in the environment perception use-case example, the user can split the ratio of training data between train, validation and test datasets.



**Figure 11.** Comparison between different MPP computation methods. (a) Ground truth. (b) Prediction given by a classical (non-AI) method. (c) Correctly predicted path using an AI Inference Engine.

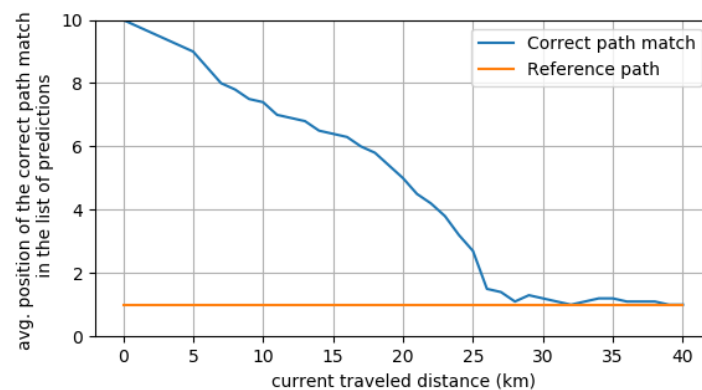


**Figure 12.** Test vehicle used for data acquisition and testing for the Most Probable Path (MPP) AI Inference Engine.

The DNN at the core of the AI Inference Engine is a Recurrent Neural Network (RNN) consisting of 256 Long-Short-Term-Memory (LSTM) units, followed by a Fully Connected layer (SoftMax activation), providing as output a list of predictions of the most probable paths. We have chosen an RNN network architecture due to its robustness in processing time dependent sequences, thus encoding the links relations in its memory cells. EB-AI was used to efficiently design and optimize the neural network. The obtained AI model wrapped in the ADTF format was deployed in a container ready to run on a Nvidia Jetson AGX board mounted in the vehicle from Figure 12.

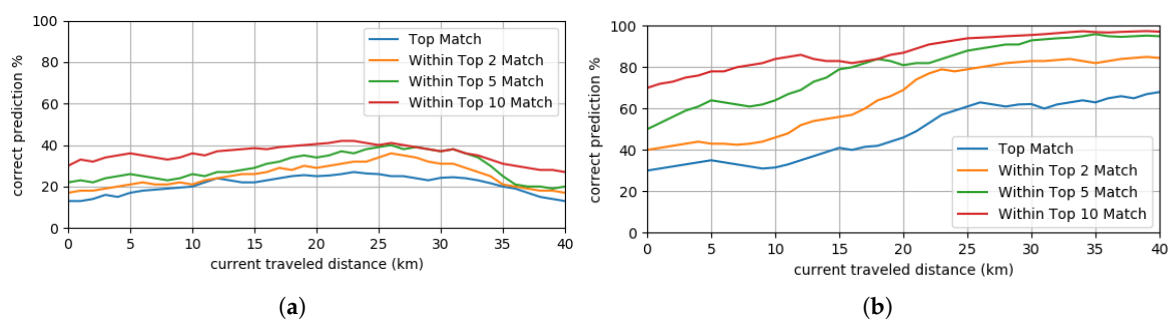
To evaluate the performance of the MPP approach, EB-AI provides the means to measure the similarity between two paths, namely predicted and ground truth. The output of the DNN is a list of paths ordered by the probability of being the correct one. Figure 13 shows the average position of the most probable route as the trip advances. The  $y$  axis in the figure represents the ground-truth reference path and the DNN predicted path, order based on their probabilities. We show the top most probable paths estimated by the neural network. Ideally, the predicted path should be the same as the reference one. In the beginning, when little information is known about the past vehicle's route, the MPP has

a lower accuracy. The accuracy increases gradually as the drive progresses, reaching the top match after kilometer 25.



**Figure 13.** The average correct route match relative to the traveled distance. At halfway through the 40 km trip, the predicted route is within the top 5 matches, while in the last 10 km the predicted route is the top match.

An additional performance metric is the correct prediction rate of the DNN. Figure 14 presents how often the MPP correctly predicted across different trips. The left most graph presents the results obtained when driving all the time on new routes, while the right most graph illustrates the results obtained when driving repetitively on the same routes.



**Figure 14.** Performance of the MPP algorithm across different trips. (a) Driving on new routes. (b) Driving repetitively on the same routes.

## 6. Conclusions

In this paper, we have proposed an elastic AI development framework for autonomous driving applications based on deep learning, which takes into account some of the advances brought about by AI, Cloud and Edge computing. A modular Elastic toolchain providing all the required DL components enables the deployment of training tasks over both Cloud end Edge resources, where the latter can be located on the vehicles themselves. Pre-processing raw data and/or training partial models at the Edge allows to reduce the amount of data to upload to the Cloud and help mitigate privacy issues. We showed the effectiveness of the proposed toolchain in two application cases implemented at Elektrobit Automotive, namely Environment Perception and Most Probable Path Prediction. Furthermore, we showed the convenience of exploiting a distributed setting to parallelize as much as possible the training of the DNNs.

As future work, we aim to prototype and deploy AI Inference Engines at a large scale using the proposed hybrid deployment strategy over a larger fleet of autonomous vehicles.



**Author Contributions:** Conceptualization, S.G., B.T. and F.L.; methodology, S.G., T.C., B.T., A.M., F.L. and L.A.; software, S.G., T.C. and B.T.; validation, S.G., T.C., B.T. and F.L.; formal analysis, A.M., F.L. and L.A.; investigation, S.G. and F.L.; resources, S.G., T.C., B.T. and F.L.; data curation, S.G., T.C., B.T.; writing—original draft preparation, F.L., A.M., L.A. and B.T.; writing—review and editing, F.L.; visualization, F.L.; supervision, L.A.; project administration, S.G.; funding acquisition, S.G. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research was funded by the European Commission under the CyberKit4SMEs project, grant number 883188.

**Conflicts of Interest:** The authors declare no conflict of interest.

## Abbreviations

The following abbreviations are used in this manuscript:

ADAS	Advanced Driver Assistant Systems
ADTF	Automotive Data and Time-Triggered Framework
AI	Artificial Intelligence
AV	Autonomous Vehicle
CNN	Convolutional Neural Network
DL	Deep Learning
dSPS	distributed Stream Processing Systems
DNN	Deep Neural Network
DP	Differential Privacy
EB	EletkroBit
GOL	Generative One-shot Learning
GPU	Graphic Processing Unit
HiL	Hardware-in-the-Loop
IoT	Internet-of-Things
IoU	Intersection over Union
ML	Machine Learning
MLP	Multi-Layer Perceptron
MPC	Multi-Party Computation
OTA	Over-the-Air
ReLU	Rectified Linear Unit
RM	Resource Manager
RNN	Recurrent Neural Network
ROS	Robotic Operative System
SiL	Software-in-the-Loop
TC	Training Coordinator

## Appendix A

A hybrid deployment over Cloud and Edge resources allows to significantly increase the level of parallelism in AI Inference Engine development. To give an idea about the scale of improvement that parallelization can bring to DNN training, we present a toy example with an experimental evaluation where two DNNs, i.e., a Multi-Layer Perceptron (MLP) and a CNN, are trained. The MLP has two hidden layers of 800 units each, followed by a softmax output layer of 10 units with 20% dropout to input data and 50% dropout to both hidden layers. The CNN has two convolutional layers with 32 kernels of size  $5 \times 5$  strided and padded, a fully-connected layer of 256 units with 50% dropout on its inputs and a 10-unit output layer with 50% dropout on its inputs. We compare the time required to train them by using a CPU, i.e., low level of parallelism, and a GPU, higher parallelism.

The experimental evaluation was carried out on a Dell Precision T7610 server, equipped with an Intel Xeon E5-2660 v2 CPU, 8 GB of RAM, and with two GPUs, namely Tesla K20c and Nvidia Quadro K4000. The DNNs were developed in Python with Lasagne, by relying on Theano to exploit the CUDA GPU libraries. The used dataset is based on MNIST, with a training set of 60k samples and a test set of 10k samples. The tests aim at comparing: (i) the training accuracy of the MLP and the CNN, after the

same number of training iterations and (ii) the training time on the available CPU and GPUs. Optimal settings for each network were computed based on a Q-Learning approach [46] that can automatically apply the dropout to minimize overfitting.

Figure A1 reports the accuracy results for the two networks during training. The CNN achieves better accuracy than the MLP, being over 98% in five epochs, while the MLP requires about 50 epochs. The maximum accuracy of the CNN is 99.21% versus 98.56% of the MLP. Figure A2a,b report, respectively for MLP and CNN, the performance of training an epoch on the different hardware devices. As expected, training with both the GPUs is much faster than with the CPU. The epoch training time for MLP (resp., CNN) on the GPU Tesla and Quadro are  $\approx 14$  and  $\approx 6$  (resp.,  $\approx 73$  and  $\approx 38$ ) times faster than on the CPU. Notice how the Tesla is about two times faster than the Quadro.

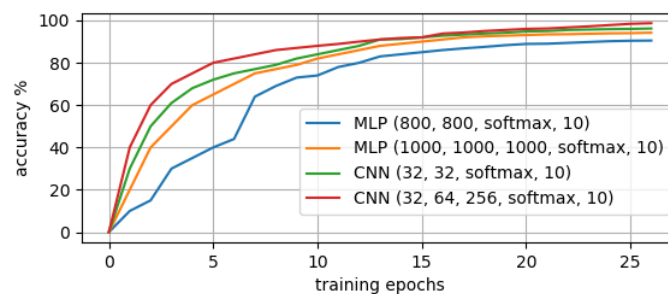


Figure A1. Accuracy comparison between MLP and CNN.

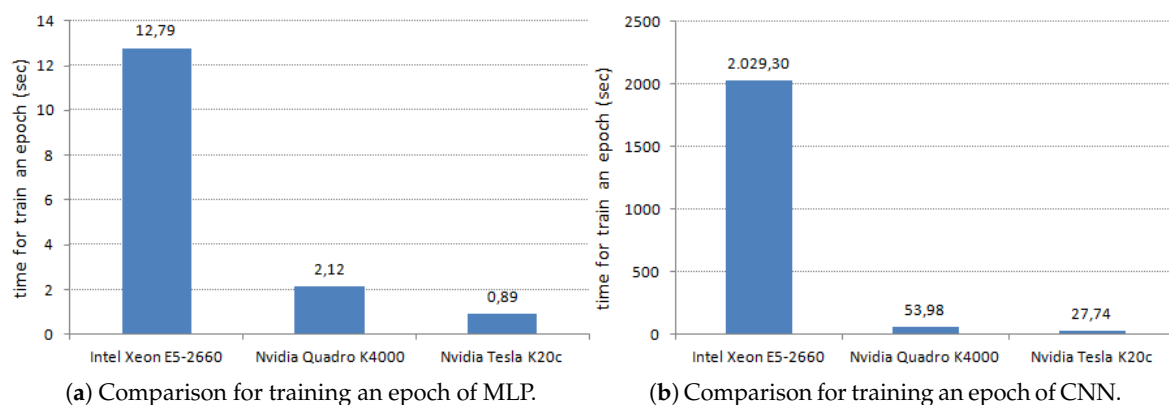


Figure A2. Performance evaluation of the two types of DNNs .

## References

1. Grigorescu, S.; Trasnea, B.; Cocias, T.; Macesanu, G. A Survey of Deep Learning Techniques for Autonomous Driving. *J. Field Robot.* **2019**, *37*, 362–386. [CrossRef]
2. Villalonga, A.; Beruvides, G.; Castaño, F.; Haber, R.E. Cloud-Based Industrial Cyber-Physical System for Data-Driven Reasoning: A Review and Use Case on an Industry 4.0 Pilot Line. *IEEE Trans. Ind. Inform.* **2020**, *16*, 5975–5984. [CrossRef]
3. MATLAB. 9.7.0.1190202 (R2019b); The MathWorks Inc.: Natick, MA, USA, 2018.
4. Deep Learning on AWS. Available online: <https://aws.amazon.com/training/course-descriptions/deep-learning/> (accessed on 12 August 2020).
5. The Analyst Toolbox. Available online: <http://coppelia.io/2014/06/the-analysts-toolbox/> (accessed on 12 August 2020).
6. Salay, R.; Queiroz, R.; Czarnecki, K. An Analysis of ISO 26262: Using Machine Learning Safely in Automotive Software. *arXiv* **2017**, arXiv:1709.02435.

7. Abadi, M.; Agarwal, A.; Barham, P.; Brevdo, E.; Chen, Z.; Citro, C.; Corrado, G.S.; Davis, A.; Dean, J.; Devin, M.; et al. TensorFlow: Large-Scale machine learning on heterogeneous distributed systems. In Proceedings of the 12th Symposium on Operating Systems Design and Implementation, Savannah, GA, USA, 2–4 November 2016; pp. 265–283.
8. Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; Lerer, A. Automatic Differentiation in PyTorch. Available online: <https://openreview.net/forum?id=BjJsrnfCZ> (accessed on 28 October 2017).
9. Jia, Y.; Shelhamer, E.; Donahue, J.; Karayev, S.; Long, J.; Girshick, R.; Guadarrama, S.; Darrell, T. Caffe: Convolutional architecture for fast feature embedding. In Proceedings of the 22Nd ACM International Conference on Multimedia, Orlando, FL, USA, 3–7 November 2014; ACM: New York, NY, USA, 2014; pp. 675–678. [CrossRef]
10. Batiz-Benet, J.; Slack, Q.; Sparks, M.; Yahya, A. Parallelizing machine learning algorithms. In Proceedings of the 24th ACM Symposium on Parallelism in Algorithms and Architectures, Pittsburgh, PA, USA, 25–27 June 2012; pp. 25–27.
11. Nangare, S. Gartner’s Strategic Tech Trends Show the Need for an Empowered Edge and Network for a Smarter World. Available online: <http://cloudcomputing-news.net/news/2018/oct/17/gartners-strategic-tech-trends-show-need-empowered-edge-and-network-smarter-world/> (accessed on 17 October 2018).
12. Wiles, J. Top Risks for Legal and Compliance Leaders in 2018. Available online: <https://www.gartner.com/smarterwithgartner/top-risks-for-legal-and-compliance-leaders-in-2018/> (accessed on 30 March 2018).
13. Greenough, J. The Connected Car Report: Forecasts, Competing Technologies, and Leading Manufacturers. Available online: <https://www.businessinsider.in/THE-CONNECTED-CAR-REPORT-Forecasts-competing-technologies-and-leading-manufacturers/articleshow/46436661.cms> (accessed on 18 August 2016).
14. Khurram, M.; Kumar, H.; Chandak, A.; Sarwade, V.; Arora, N.; Quach, T. Enhancing connected car adoption: Security and over the air update framework. In Proceedings of the 2016 IEEE 3rd World Forum on Internet of Things, Reston, VA, USA, 12–14 December 2016; pp. 194–198.
15. Huang, Y.; Ma, X.; Fan, X.; Liu, J.; Gong, W. When deep learning meets edge computing. In Proceedings of the Computer Society, Toronto, ON, Canada, 10–13 October 2017; pp. 1–2.
16. Li, H.; Ota, K.; Dong, M. Learning IoT in Edge: Deep Learning for the Internet of Things with Edge Computing. *IEEE Netw.* **2018**, *32*, 96–101. [CrossRef]
17. Huang, Y.; Zhu, Y.; Fan, X.; Ma, X.; Wang, F.; Liu, J.; Wang, Z.; Cui, Y. Task scheduling with optimized transmission time in collaborative cloud-edge learning. In Proceedings of the 2018 27th International Conference on Computer Communication and Networks (ICCCN), Hangzhou, China, 30 July–2 August 2018; pp. 1–9. [CrossRef]
18. Bengio, Y.; Courville, A.; Vincent, P. Representation learning: A review and new perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.* **2013**, *35*, 1798–1828. [CrossRef] [PubMed]
19. Hinton, G.E.; Zemel, R.S. Autoencoders, minimum description length and Helmholtz free energy. In Proceedings of the 6th International Conference on Neural Information Processing Systems Morgan, San Mateo, CA, USA, 29 November 1994; Kaufmann Publishers Inc.: San Francisco, CA, USA, 1994; pp. 3–10.
20. Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; Salakhutdinov, R. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **2014**, *15*, 1929–1958.
21. Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. ImageNet Large Scale Visual Recognition Challenge. *Int. J. Comput. Vis. (IJCV)* **2015**, *115*, 211–252. [CrossRef]
22. Geiger, A.; Lenz, P.; Stiller, C.; Urtasun, R. Vision Meets Robotics: The KITTI Dataset. *Int. J. Robot. Res.* **2013**, *32*, 1231–1237. [CrossRef]
23. Caesar, H.; Bankiti, V.; Lang, A.H.; Vora, S.; Liong, V.E.; Xu, Q.; Krishnan, A.; Pan, Y.; Baldan, G.; Beijbom, O. NuScenes: A multimodal Dataset for Autonomous Driving. *arXiv* **2019**, arXiv:1903.11027.
24. Cityscapes. Cityscapes Data Collection. Available online: <https://www.cityscapes-dataset.com/> (accessed on 17 February 2018).
25. Grigorescu, S.M. Generative One-Shot Learning (GOL): A semi-parametric approach to one-shot learning in autonomous vision. In Proceedings of the 2018 IEEE International Conference on Robotics and Automation (ICRA), Brisbane, Australia, 21–25 May 2018; pp. 7127–7134.

26. Hall, M.; Frank, E.; Holmes, G.; Pfahringer, B.; Reutemann, P.; Witten, I.H. The WEKA data mining software: An update. *ACM SIGKDD Explor. Newsl.* **2009**, *11*, 10–18. [[CrossRef](#)]
27. Xing, E.P.; Ho, Q.; Xie, P.; Wei, D. Strategies and principles of distributed machine learning on big data. *Engineering* **2016**, *2*, 179–195. [[CrossRef](#)]
28. Toshniwal, A.; Taneja, S.; Shukla, A.; Ramasamy, K.; Patel, J.M.; Kulkarni, S.; Jackson, J.; Gade, K.; Fu, M.; Donham, J.; et al. Storm@ twitter. In Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data, Snowbird, UT, USA, 22–27 June 2014; ACM: New York, NY, USA, 2014; pp. 147–156.
29. Zaharia, M.; Chowdhury, M.; Franklin, M.J.; Shenker, S.; Stoica, I. Spark: Cluster computing with working sets. *HotCloud* **2010**, *10*, 95.
30. Noel, C.; Shi, J.; Feng, A. Large Scale Distributed Deep Learning on Hadoop Clusters. Available online: <https://on-demand.gputechconf.com/gtc/2016/presentation/s6836-andy-feng-large-scale-distributed-deep-learning-hadoop-clusters.pdf> (accessed on 25 September 2015).
31. Smith, C.; Nguyen, C.; TensorFlow, U.D. Distributed TensorFlow: Scaling Google’s Deep Learning Library on Spark. Available online: <https://arimo.com/machine-learning/deep-learning/2016/arimo-distributed-tensorflow-on-spark/> (accessed on 30 March 2018).
32. Yang, Q.; Liu, Y.; Chen, T.; Tong, Y. Federated Machine Learning: Concept and Applications. *ACM TIST* **2019**, *10*, 12:1–12:19. [[CrossRef](#)]
33. Konečný, J.; McMahan, H.B.; Ramage, D.; Richtárik, P. Federated Optimization: Distributed Machine Learning for On-Device Intelligence. *arXiv* **2016**, arXiv:1610.02527.
34. Chen, J.; Li, K.; Deng, Q.; Li, K.; Yu, P.S. Distributed Deep Learning Model for Intelligent Video Surveillance Systems with Edge Computing. *IEEE Trans. Ind. Inform.* **2019**. [[CrossRef](#)]
35. Luckow, A.; Cook, M.; Ashcraft, N.; Weill, E.; Djerekarov, E.; Vorster, B. Deep learning in the automotive industry: Applications and tools. In Proceedings of the 2016 IEEE International Conference on Big Data (Big Data), Washington, DC, USA, 5–8 December 2016; pp. 3759–3768.
36. Brilli, G.; Burgio, P.; Bertogna, M. Convolutional Neural Networks on embedded automotive platforms: A qualitative comparison. In Proceedings of the 2018 International Conference on High Performance Computing & Simulation (HPCS), Orleans, France, 16–20 July 2018; pp. 496–499.
37. Fridman, L.; Brown, D.E.; Glazer, M.; Angell, W.; Dodd, S.; Jenik, B.; Terwilliger, J.; Kindelsberger, J.; Ding, L.; Seaman, S.; et al. Mit autonomous vehicle technology study: Large-scale deep learning based analysis of driver behavior and interaction with automation. *arXiv* **2017**, arXiv:1711.06976.
38. The Society of Automotive Engineers. Available online: <https://blog.ansi.org/2018/09/sae-levels-driving-automation-j-3016-2018/#gref> (accessed on 12 September 2018).
39. Litman, T. *Autonomous Vehicle Implementation Predictions*; Victoria Transport Policy Institute: Victoria, BC, Canada, 2019.
40. Lu, S.; Yao, Y.; Shi, W. Collaborative learning on the edges: A case study on connected vehicles. In *2nd USENIX Workshop on Hot Topics in Edge Computing (HotEdge 19)*; USENIX Association: Renton, WA, USA, 2019.
41. Jiang, L.; Lou, X.; Tan, R.; Zhao, J. Differentially private collaborative learning for the IoT edge. In Proceedings of the 2019 International Conference on Embedded Wireless Systems and Networks (EWSN’19), Beijing, China, 25–27 February 2019; Junction Publishing: Junction, TX, USA, 2019; pp. 341–346.
42. Yuan, J.; Yu, S. Privacy Preserving Back-Propagation Neural Network Learning Made Practical with Cloud Computing. *IEEE Trans. Parallel Distrib. Syst.* **2014**, *25*, 212–221. [[CrossRef](#)]
43. Mohassel, P.; Zhang, Y. SecureML: A system for scalable privacy-preserving machine learning. In Proceedings of the 2017 IEEE Symposium on Security and Privacy (SP), San Jose, CA, USA, 22–24 May 2017; pp. 19–38. [[CrossRef](#)]
44. Bogdanov, D.; Laur, S.; Willemson, J. Sharemind: A framework for fast privacy-preserving computations. In Proceedings of the 13th European Symposium on Research in Computer Security: Computer Security (ESORICS’08), Egham, UK, 9–13 September 2008; Springer: Berlin/Heidelberg, Germany, 2008; pp. 192–206. [[CrossRef](#)]
45. Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H.B.; Mironov, I.; Talwar, K.; Zhang, L. Deep learning with differential privacy. In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, 24–28 October 2016; pp. 308–318.
46. Lombardi, F.; Aniello, L.; Bonomi, S.; Querzoni, L. Elastic symbiotic scaling of operators and resources in stream processing systems. *IEEE Trans. Parallel Distrib. Syst.* **2017**, *29*, 572–585. [[CrossRef](#)]

47. Badrinarayanan, V.; Kendall, A.; Cipolla, R. SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **2017**, *39*, 2481–2495. [[CrossRef](#)] [[PubMed](#)]
48. Redmon, J.; Divvala, S.; Girshick, R.; Farhadi, A. You only look once: Unified, real-time object detection. In Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 27–30 June 2016; pp. 779–788. [[CrossRef](#)]
49. Marina, L.; Trasnea, B.; Tiberiu, C.; Vasilcoi, A.; Moldoveanu, F.; Grigorescu, S. Deep Grid Net (DGN): A deep learning system for real-time driving context understanding. In Proceedings of the 2019 Third IEEE International Conference on Robotic Computing (IRC), Naples, Italy, 25–27 February 2019.
50. The Cambridge-Driving Labeled Video Database. Available online: <http://mi.eng.cam.ac.uk/research/projects/VideoRec/CamVid/> (accessed on 23 July 2019).



© 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).