



UNIVERSITY OF LEEDS

This is a repository copy of *Orchestrating the Development Lifecycle of Machine Learning-based IoT Applications: A Taxonomy and Survey*.

White Rose Research Online URL for this paper:
<https://eprints.whiterose.ac.uk/166897/>

Version: Accepted Version

Article:

Qian, B, Su, J, Wen, Z et al. (11 more authors) (2020) Orchestrating the Development Lifecycle of Machine Learning-based IoT Applications: A Taxonomy and Survey. *ACM Computing Surveys*, 53 (4). 82. ISSN 0360-0300

<https://doi.org/10.1145/3398020>

© 2020 ACM. This is an author produced version of an article published in *ACM Computing Surveys*. Uploaded in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Orchestrating the Development Lifecycle of Machine Learning-Based IoT Applications: A Taxonomy and Survey

BIN QIAN, JIE SU, ZHENYU WEN*, DEVKI NANDAN JHA, YINHAO LI, YU GUAN, DEEPAK PUTHAL, and PHILIP JAMES, Newcastle University, UK

RENYU YANG, University of Leeds, UK

ALBERT Y. ZOMAYA, The University of Sydney, Australia

OMER RANA, Cardiff University, UK

LIZHE WANG, China University of Geoscience (Wuhan), China

MACIEJ KOUTNY and RAJIV RANJAN, Newcastle University, UK

Machine Learning (ML) and Internet of Things (IoT) are complementary advances: ML techniques unlock the potential of IoT with intelligence, and IoT applications increasingly feed data collected by sensors into ML models, thereby employing results to improve their business processes and services. Hence, orchestrating ML pipelines that encompass model training and implication involved in the holistic development lifecycle of an IoT application often leads to complex system integration. This paper provides a comprehensive and systematic survey of the development lifecycle of ML-based IoT applications. We outline the core roadmap and taxonomy, and subsequently assess and compare existing standard techniques used at individual stages.

Additional Key Words and Phrases: IoT, Machine learning, Deep learning, Orchestration

1 INTRODUCTION

Rapid development of hardware, software and communication technologies boosts the speed of connection of the physical world to the Internet via Internet of Things (IoT). A report ¹ shows that about 75.44 billion IoT devices will be connected to the Internet by 2025. These devices generate a massive amount of data with various modalities. Processing and analyzing such big data is essential for developing smart IoT applications. Machine Learning (ML) plays a vital role in data intelligence which aims to understand and explore the real world. ML + IoT type applications thus are experiencing explosive growth. However, there are unfilled gaps between current solutions and the demands of orchestrating the development lifecycle of ML-based IoT applications. Existing orchestration frameworks for example *Ubuntu Juju*, *Puppet* and *Chef* are flexible in providing solutions for deploying and running applications over public or private clouds. These frameworks, however, neglect the heterogeneity of IoT environments that encompasses various hardwares, communication protocols and operating systems. More importantly, none of them are able to completely orchestrate a holistic development lifecycle of ML-based IoT applications. The development lifecycle must cover the following factors: 1) *how* the target application is specified and developed, 2) *where* the target application is deployed, (3) *what* kind of information the target application is being audited. Application specification defines the requirements including the ML tasks, performance, accuracy and execution workflow. Based on the specification and the available computing resources, the ML models are developed to meet the specified requirements while optimizing the training processes in terms of the cost of time and computing resources. Next, the model deployment considers the difficulty of

*Zhenyu is the corresponding author

¹<https://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide/>

the heterogeneity of the IoT environment for running a set of composed ML models. Finally, ML-based IoT applications closely connect with people’s lives and some applications such as autopilot require high reliability. Therefore, essential monitoring information has to be collected to improve the performance of the application in the next iteration of the lifecycle.

In this survey, we present comprehensive research on orchestrating the development lifecycle of ML-based IoT applications. We first present the core roadmap and taxonomy, and subsequently summarize, compare, and assess the variety of techniques used in each step of the lifecycle. Previous efforts provided broad knowledge that can drive us to build the taxonomy. For instance, [317] discussed encountered challenges of developing the next generation of AI systems. [239, 381] gave comprehensive reviews of available deep learning architectures and algorithms in IoT domain. To the best of our knowledge, this is the first work that presents a comprehensive survey to illustrate the whole development lifecycle of ML-based IoT application, which paves the way for developing an agile, robust and reliable smart IoT application. Before introducing the roadmap and taxonomy, we provide a smart city example in the next subsection that illustrates ML-based IoT applications in real-world.

1.1 Smart City Applications

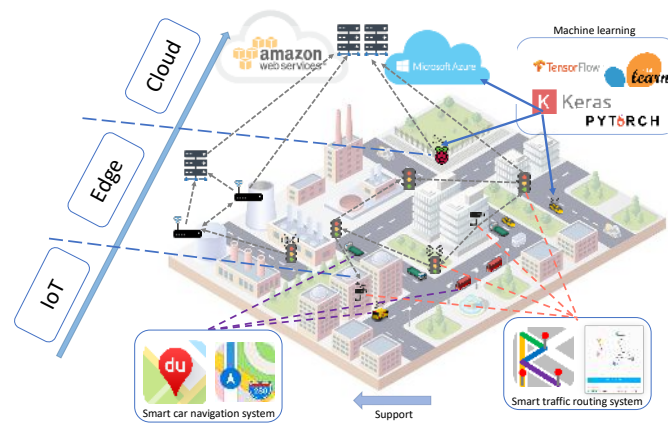


Fig. 1. Smart City

Smart city uses modern communication and information techniques to monitor, integrate and analyze the data collected from core systems running across cities. Meanwhile, smart city makes intelligent responses to various use cases, such as traffic control, weather forecasting, industrial and commercial activities. Fig. 1 represents a smart city which consists of various IoT applications with many of them using Machine Learning (ML) techniques. For example, a *smart traffic routing* system consists of a large number of cameras monitoring the road traffic and a smart algorithm running on the cloud recommending the optimal routes for users [395]. On the other hand, a *smart car navigation system* [161] allows the passengers to set and change destinations via built-in car audio devices. The two systems work together to provide real-time interactive routing services. More specifically, the user’s voice commands are translated in the car edge side and sent to the cloud where the *smart traffic routing* system works. The best route is translated back to voice guiding the users to their destinations. The above-mentioned applications involve various computing resources (e.g., cloud, edge, and IoT devices) and ML techniques, making the development of these ML-based IoT applications very challenging both for the ML models and the IoT system. To fill this gap, we orchestrate the development lifecycle

of an ML-based IoT application. In the next subsection, we present a roadmap for the development lifecycle along with a comprehensive taxonomy that surveys the techniques relevant for developing the application.

1.2 Roadmap and Taxonomy




Fig. 2. The development lifecycle of an ML-based IoT application

Roadmap. Fig 2 shows the roadmap of developing an ML-based IoT application. The roadmap starts with the requirements specification where the required computing resources (hardware and software) and ML models are specified. Based on the specification, we carefully design the infrastructure protocol, data acquisition approach and machine learning model development pipeline. Next, we implement and train the model with various ML algorithms. We also evaluate and optimize the models to achieve high efficiency without sacrificing too much accuracy. After the model development, an optimized deployment plan is generated based on the specified ML models and infrastructures. The deployed application must be audited while it is running on real IoT environments; the audit aims to explore the performance issues in terms of security, reliability and other QoS metrics. Finally, the audited issues will guide the corrections of orchestration details in the next iteration of the application development.

Taxonomy. Fig. 3 depicts our taxonomy which systematically analyzes the core components in the orchestration of the development lifecycle of a ML-based IoT application. Note that the survey in [353] has reviewed cloud resource orchestration techniques. It outlines the key infrastructure orchestration challenges for cloud-based application as well as being extendable for IoT applications. Thus, in this survey, we focus more on the challenges of implementing ML models and orchestrating their IoT application development lifecycle. To this end, we extract the core building blocks of the development lifecycle relevant to ML and identify *four* main categories based on their specific functionality during the development process. The outline of the paper follows the structure of the taxonomy as well.

- (1) **Model Development.** We propose a general pipeline for developing a ready-to-deploy ML model. We investigate the ML techniques to build each block of the pipeline (refer to §2).
- (2) **Model Deployment.** In our work, we review the software deployment techniques and analyze the challenges of applying such techniques to deploy the ML models in IoT environments (refer to §3).
- (3) **Model Audit.** Audit is one of the important dimensions in building a robust application. We survey the main security, reliability and performance issues in ML-based IoT applications (refer to §4).
- (4) **Data Acquisition.** Data quality is important in building ML models. We identify three dimensions that are important throughout the data acquisition pipeline: data collection, data fusion and data preprocessing (refer to §5).

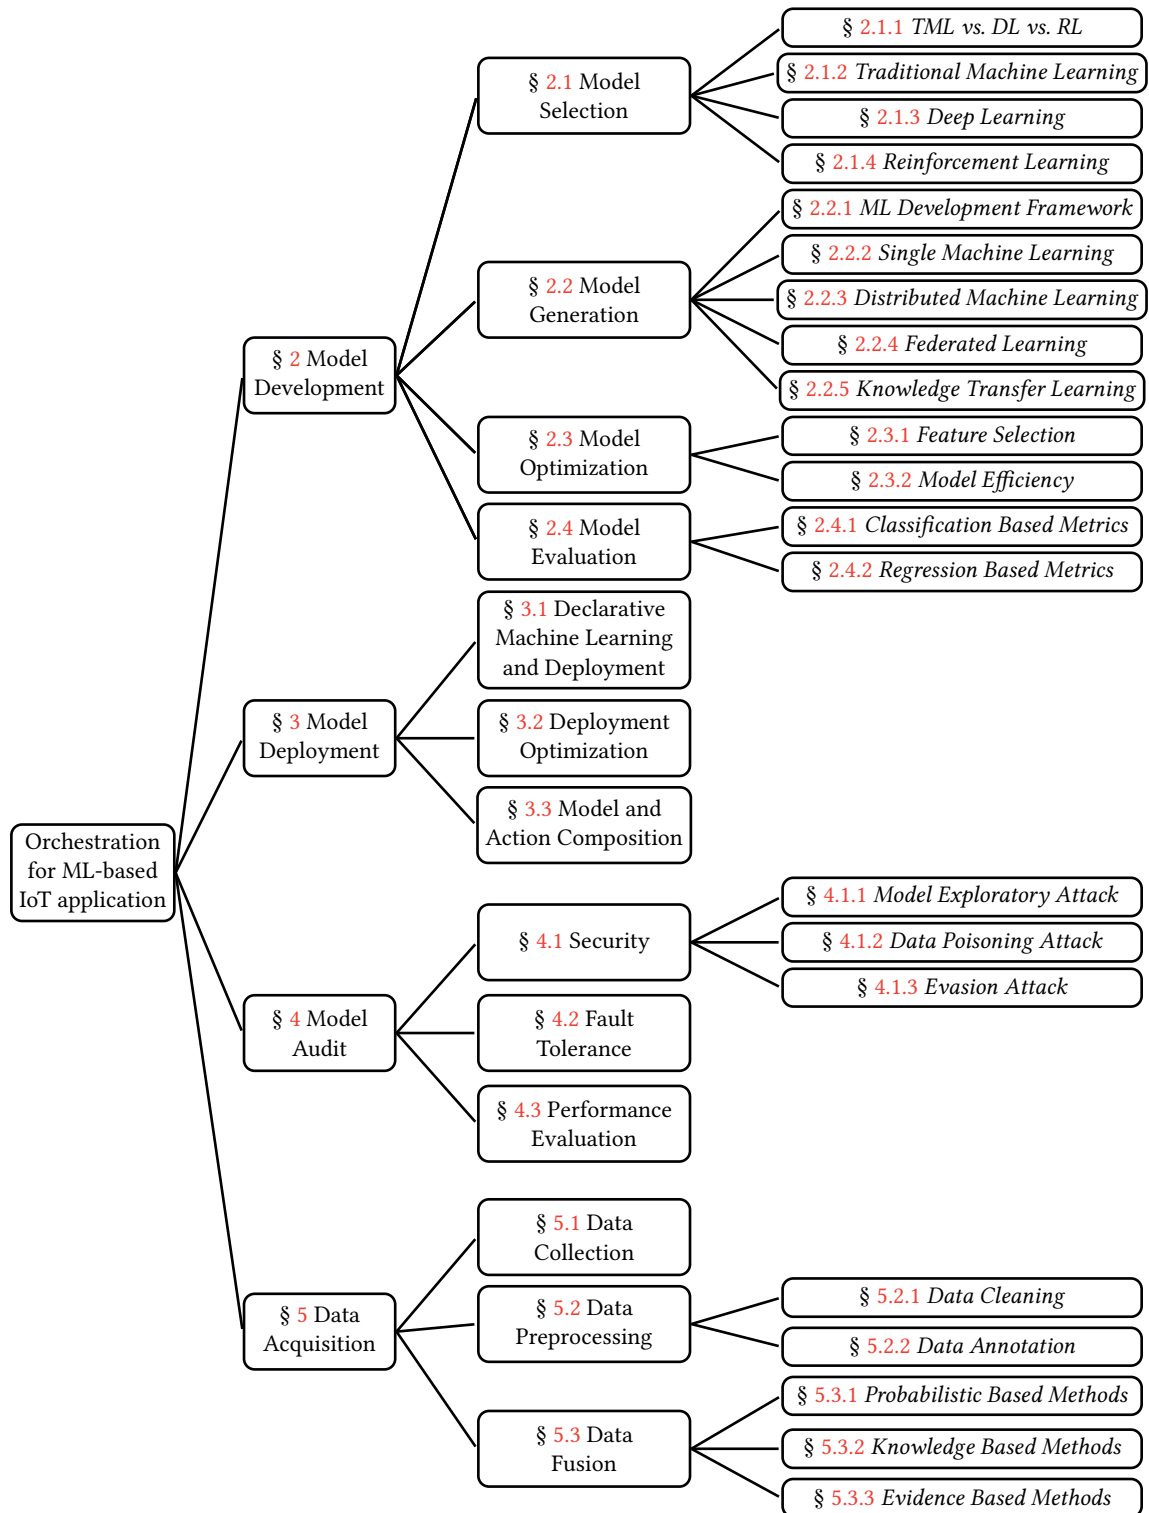


Fig. 3. A taxonomy for orchestrating ML-based IoT application development lifecycle




Fig. 4. A general pipeline of model development

2 MODEL DEVELOPMENT

One of the core components in this paper is machine learning (ML) models, which may be roughly divided into three categories: Traditional Machine Learning (TML), Deep Learning (DL) and Reinforcement Learning (RL). To develop ML models in the IoT environment, we propose a generic pipeline (see Fig. 4) including *model selection*, *model generation*, *model optimization* and *model evaluation*. We initially introduce the generic pipeline by presenting adaptive video streaming [222] as an example.

Adaptive video streaming. Video transmission between server and mobile devices employs http-based adaptive streaming techniques. In a typical video server (e.g., DASH²), videos are encoded and stored as multiple chunks at different bitrates. One video usually consists of several chunks with each containing seconds of content. To maximize video quality, the video player in a client (e.g., mobile device) usually employs adaptive bitrate (ABR) algorithms aiming to pull high-bitrate chunks from the server without compromising the latency. As shown in Fig. 5, ABR algorithms use simple heuristics to make bitrate decisions based on various observations such as the estimated network throughput and playback buffer occupancy. ABR algorithms require fine-grained tuning and can be hardly generalized to handle various network conditions that fluctuate across time and different environments. Thus we are seeking to solve the problem using modern ML technologies.




Fig. 5. Adaptive video streaming

To this end, we first need to perform **model selection** (§2.1) to find a subset of suitable models. In this scenario, the server must give a bitrate decision so that the client can return feedback that conveys whether the decision is satisfactory. Such interaction problems necessitate further use of RL and we will present proper choice of RL algorithms based on different selection criteria (§2.1.4). Next, we will choose a suitable development framework to implement the model and utilize different acceleration techniques to reduce the latency of **model generation** (§2.2). In this example, Tensorflow and A3C algorithm [237] are used as the development framework and distributed training protocol respectively for faster convergence. Once generated, the model has to be adapted into the real environment. Considering heterogeneity of IoT infrastructure, models need to be optimized according to the computing resources. This procedure is called **model optimization** (§2.3). In **model evaluation** (§2.4), model performance is measured to validate whether the model meets expected results. Particularly in this case, performance is evaluated by the total reward obtained from the simulated environment. The following subsections will discuss the the pipeline in detail.

²<https://github.com/Dash-Industry-Forum/dash.js>

2.1 Model Selection

Model selection aims to find the *optimal* ML model to perform a user’s specified tasks, whilst adapting to the complexity of IoT environments. In this section, we first discuss the model selection from three main categories i.e., TML, DL and RL, followed by a survey of well-known models (or algorithms) in each category and their corresponding criteria for model selection.

2.1.1 TML vs. DL vs. RL. In this work we roughly divide the ML approaches/concepts into TML, DL and RL. Compared with the most popular DL, TML is relatively lightweight. It is a set of algorithms that directly transform the input data (to output), according to certain criteria. For supervised cases when a class label is available for training, TML aims to map the input data to the labels by optimising a model, which can be used to infer unseen data at the test stage. However, since the relationship between raw data and label might be highly non-linear, feature engineering— a heuristic trial-and-error process — is normally required to construct the appropriate input feature. The TML model is relatively simple, the interpretability (e.g., the relationship between the engineered features and the labels) tends to be high.

DL has become popular in recent years. Consisting of multiple layers, DL is powerful for modeling complex non-linear relationships (between the input and output) and thus does not require the aforementioned heuristic (and expensive) feature engineering process, making it a popular modelling approach in many fields such as computer vision and natural language processing. Compared with TML, DL models tend to have more parameters (to be estimated) and generally they require more data for reliable representation learning. However, it is crucial to guarantee the data quality and a recent empirical study[247] suggested the increasing number of noisy/less-representative training samples may harm DL’s performance, making it less generalizable to unseen test data. Moreover, DL’s multilayer structures make it difficult to interpret the complex relationship between input (i.e., raw features) and output. However, more and more visualisation techniques (e.g., attention map [385]) were used, which play an important role in understanding DL’s decision-making process.

RL has become increasingly popular due to its success in addressing challenging sequential decision-making problems [324]. Some of these achievements are based on the combination of DL and RL, i.e., Deep Reinforcement Learning. It has shown its considerable performance in natural language processing [197, 365], computer vision [11, 55, 278, 323, 376], robotics [268] and IoT systems [221, 222, 392] and related applications like video games [11], visual tracking [278, 323, 376], action prediction [55], robotic grasping [268], question answering [365], dialogue generation [197], etc. In RL, there is usually one or more agent(s) interacting with the outside environment, where optimal control policies are learnt through experience. Fig. 6 illustrates the iterative interaction circle, where the agent starts without knowing anything about environment or task. Each time the agent takes action based on the environment states, and it receives a reward from the environment. RL optimises this process such that it learns to make decisions with higher rewards received.




Fig. 6. Reinforcement Learning Paradigm

Discussion. In IoT environments, a variety of problems can be modelled by using the aforementioned three approaches. The applications range from system and networking [222] [221], smart city [392] [198], to smart grid [358] [285], etc. To

begin with modeling, it is essential for users to choose a suitable learning concept at the first stage. The main selection criteria can be divided into two categories: *Function-based selection* and *Power Consumption-based selection*.

Function-based selection aims to choose an appropriate concept based on their functional difference. For example, RL benefits from its iterative environment \leftrightarrow agent interaction property, and can be applied to various applications which need interaction with environment or system such as smart temperature control systems, or recommendation systems (with cold start problem). On the other hand, TML algorithms are more suitable for modelling structured data (with high-level semantic attributes), especially when interpretability is required. DL models are typically used to model complex unstructured data, e.g., images, audios, time-series data, etc. and are an ideal choice especially with high amount of training data and low requirement on interpretability.

Power Consumption-based selection aims to choose an appropriate model given constraints in computational power or latency. In contrast to TML, the powerful RL/DL models are normally computationally expensive with high overhead. Recently, model compression techniques were developed, which may provide a relatively efficient solution for using RL/DL models for some IoT applications. However, on some mobile platforms with very limited hardware resources (e.g., power, memory, storage), it is still challenging to employ compressed RL/DL models, especially when there are some performance requirements (e.g., accuracy, or real-time inference) [59]. On the other hand, lightweight TML may be more efficient, yet reasonable accuracy can only be achieved with appropriate features (e.g., high level attributes derived from the time-consuming feature engineering).




Fig. 7. Examples of Supervised Learning (Linear Regression) and Unsupervised Learning (Clustering)

2.1.2 Traditional Machine Learning. Herein we demonstrate several popular TML algorithms (algorithm details are available in TML method **Appendix B**), and discuss the criteria for choosing the TML algorithms. Given different tasks, TML can be further divided into *Supervised Learning* and *Unsupervised Learning*.

Supervised Learning. Supervised learning algorithm (i.e., Fig. 7) can be used when both the input data X and the corresponding labels Y are provided (for training), and it aims to learn a mapping function such that $Y \leftarrow f(X)$. Supervised learning algorithms have been widely used in IoT applications, we introduce the most representative classifiers below.

Perceptron and *Logistic Regression (LR)* are probably the simplest linear classifiers. For both models, the model (i.e., weights and bias) is basically a simple linear transformation. Perceptron can perform binary classification simply based on the sign of the (linearly) transformed input data, while LR will further scale the transformed value into probability (via *sigmoid* function), before a thresholding function is applied for the binary classification decision. LR can also be extended to process multi-class classification scenarios by using *softmax* as the scaling function, with class-wise probabilities as output.

Artificial Neural Networks (ANN) is a general extension of the aforementioned linear classifiers. Compared with *Perceptron* or *LR* which linearly project input data to the output, *ANN* has an additional “hidden layer” (with a non-linear activation function), which enables *ANN* to model non-linearity. However, in contrast to linear classifiers, this additional hidden layer makes it more difficult to see the relationship between the input and output data (i.e., low interpretability). Although in theory, with one hidden layer *ANN* can model any complex non-linear functions, in practice it has limited generalization capabilities when facing unseen data. *ANN* with more layers, also referred to as deep neural networks, tend to have better modelling capability, which will be introduced in the next subsection.

Decision Tree (DT) [269] and *Random Forest (RF)* [41] are two tree-structure based non-linear classifiers. Based on certain attribute-splitting criteria (e.g., *Information Gain* or *Gini Impurity*), *DT* can analyse the most informative attributes sequentially (i.e., splitting) until the final decision can be made. The tree structure makes it interpretable and it has reasonable accuracy with low-dimensional semantic attributes. However, it faces “the curse of dimensionality” problem and does not generalize well when the input feature quality is low. *RF*, on the other hand, can effectively address this overfitting issue. *RF* is an ensemble approach on aggregating different small-scale *DTs*, which are derived based on random sampling of the features/datasets. The random sampling mechanism can effectively reduce the dimensionality (for each individual *DT*) while the aggregation function can smooth the uncertainty of individual *DTs*, making *RF* a powerful model with great generalisation capabilities. However, the interpretability of *RF* tends to be less obvious than that of *DT*, owing to the random sampling and aggregation mechanisms.

Support Vector Machine (SVM) [73] is another popular supervised learning method. It is also called large margin classifier as it aims at finding a hyperplane that is capable of separating the data points (belonging to different classes) with the largest margin. For non-linearly separable datasets, various kernels (e.g., *RBF (Radial Basis Function)*) can be applied into the *SVM* framework with good generalization ability. Yet the time complexity for training this algorithm can be very high (i.e., $O(N^3)$ [2], where N represents the dataset size), making it less suitable for big datasets. On the other hand, *K-Nearest Neighbour (KNN)*[76], which does not require a training process (also referred to as lazy learning), is another powerful non-linear classifier. The classification is performed by distance calculation (between query and all the training examples), distance ranking, and majority voting among the (K) nearest neighbours. So selecting suitable distance functions/metrics (for different tasks) is one of the key issues in *KNN*. Since for any query sample, the distance calculation has to be performed for every sample in the whole training set, it can be time-consuming and thus less scalable to large datasets. Different from the aforementioned methods, *Naive Bayesian (NB)* algorithm [110] takes the prior knowledge of the class distribution into account. Based on the assumption that the features are conditionally independent, the likelihood of each feature can be calculated independently, before being combined with the prior probability according to the Bayes’ rule. If the feature-independence assumption is not significantly violated (e.g., low-dimensional structured data), it can be a very effective and efficient tool.

Unsupervised Learning. The unsupervised learning algorithm (see Fig. 7 right) aims at learning the inherent relationship between the data when only input data X exists (without class label Y). For example, the clustering algorithm can be used to find the potential patterns of some unlabelled data and the obtained results can be used for future analysis. *K-Means*[139] and *Principal Component Analysis (PCA)* [302] are the two most popular unsupervised learning algorithms. *K-means* aims to find K group patterns from data by iteratively assigning each sample to different clusters based on the distance between the sample and the centroid of each cluster. *PCA* is normally used for dimensionality reduction, which can de-correlate the raw features before selecting the most informative ones.

Discussion. For IoT applications, a common principle is to select the algorithm with the highest performance in terms of effectiveness and efficiency. One can run all related algorithms (e.g., supervised, or unsupervised), before selecting the most appropriate one. For effectiveness, one has to define the most suitable evaluation metrics, which can be task-dependent, e.g., accuracy or mean-f1 score for classification tasks, or mean squared errors for regression, etc. Before model selection, a number of factors should be taken into account: data structure (structured data, or unstructured data which may need additional preprocessing), data size (small or large), prior knowledge (e.g., class distribution), data separability (linearly, or non-linearly separable which may require additional feature engineering), dimensionality (low, or high which may require dimensionality reduction), etc. There may also exist additional requirements from the users/stakeholders, e.g., interpretability for health diagnosis. Additionally, it is necessary to understand the efficiency requirement specific to an IoT application and one has to consider how the training/testing time grows with respect to data size. Time complexity shown in Table 2 in **Appendix B** provides more insights. Take *KNN* as an example: although no training time is taken, *KNN*'s inference time can be very high (especially with a large training set), and thus presumably unsuitable for certain time-critical IoT applications. Also, the deployment environment is another non-negligible factor when developing IoT applications since many applications run (or partially run) on low power computing resources.

2.1.3 Deep Learning. In this section, we primarily introduce three classical deep models (i.e., *Deep Neural Networks* (DNN)/ *Multilayer Perceptron* (MLP), *Convolutional Neural Networks* (CNN) and *Recurrent Neural Networks* (RNN)) for supervised learning tasks on unstructured data such as image, video, text, time-series data, etc. We also brief two popular unsupervised models: *Autoencoder* (AE), and *Generative Adversarial Networks* (GAN).

Supervised DL. Next, we will introduce three basic supervised DL models: DNN, CNN and RNN, which require both the data and label for training.

Deep Neural Networks (DNN). As previously mentioned, a deep neural network (*DNN*) is an *ANN* with more than one hidden layer, and hence it is also called multilayer perceptron (*MLP*). Compared with *ANN* with a single hidden layer, *DNN* has more powerful modelling capabilities and its deep structure makes it easier for it to learn higher-level semantic features, which is crucial for classification tasks on complex data. However, for high-dimensional unstructured input data (such as images), there may be many model parameters to be estimated, and in this case, overfitting may occur if there is not enough labelled data. Nevertheless, generally *DNN* has decent performance when input dimensionality is not extremely high, and it has been successfully applied to various applications, for example human action recognition [336], traffic congestion prediction[88] and healthcare[248].

Convolutional Neural Network (CNN). When it comes to high-dimensional unstructured data such as images, in visual recognition tasks it is hard to directly map the raw image pixels into target labels due to the complex non-linear relationship. The traditional way is to perform feature engineering, which is normally a trial-and-error process, and may require domain knowledge in certain circumstances, before TML is applied. This heuristic approach is normally time-consuming, and there exist substantial recognition errors even in simple tasks since it is very challenging to hand-engineer the high-level semantic features. *CNN*, a deep neural network with convolutional layers and pooling layers, can address this issue effectively. The convolution operation can extract the higher level features while the pooling operation can keep the most informative responses and reduce the dimensionality. Compared with *DNN*, the weight sharing concept (of the convolution operation) enables *CNN* to capture the local pattern without suffering from the “curse of high-dimensionality” from the input. These operations and the hierarchical nature make *CNN* a powerful tool for extracting high-level semantic representations from raw image pixels directly, and successfully applied to

various recognition tasks such as object recognition, image segmentation [106] and object detection [140]. Because of the decent performance on various visual analysis tasks, *CNN* is usually considered as the first choice for some camera-based IoT applications, for example traffic sign detection [305].

Recurrent Neural Networks (RNN). Nowadays, with the increasing amount of generated stream and sequential data from various sensors, time series analysis has become popular among the machine learning (ML) community. *RNN* is a sequential modelling technique that can effectively combine the temporal information and current signal into the hidden units for time-series classification/prediction. An improved *RNN* named Long Short Term Memory (*LSTM*) [147], including complex gates and memory cells within the hidden units for “better memories”, became popular in various applications such as speech recognition [126], video analysis [335], language translation [215], activity recognition [130] etc. Since data streaming is most common in the IoT environment, *RNN (LSTM)* is deemed as one of the most powerful modelling techniques, and there are various IoT applications such as smart assistant [109, 336], *smart car navigator system* [161], malware threat hunting [134], network traffic forecasting [272], equipment condition forecasting [387], energy demand prediction system [243], load forecasting [181], etc.

Unsupervised DL. We will also introduce two unsupervised DL models: *Autoencoder (AE)* [17] and *Generative Adversarial Network (GAN)* [122].

Without requiring any label information, *AE* can extract compact features and reconstruct the original (high-dimensional) data with the extracted features. It is normally used for dimensionality reduction, latent distribution analysis or outlier detection. *GAN*, on the other hand, applies an adversarial process to learn the “real” distribution from the input data. More precisely, *GAN* consists of two parts, namely generator and discriminator. The generator aims at generating indistinguishable samples compared to the real data. While the discriminator works adversarially to distinguish the generated fake samples from the real data. It is an iterative competition process that will eventually lead to a state where the generated samples are indistinguishable from the real data. With the learnt “real” distribution, one can generate various samples for different purposes. *AE* and *GAN* are both powerful tools in the computer vision field, and their properties make them promising approaches for IoT applications. *AE* can be used for diagnosis/fault detection tasks [65, 250] or simply as a preprocessing tool (i.e., feature extraction/dimensionality reduction). *GAN* has been used for studies on generating rare category samples, and this upsampling approach may further improve the model performance [390, 391].

Discussion. The aforementioned DL models can be effective tools for processing different unstructured data types. The way of applying them is generally very flexible, and they can be used jointly to process the complex data from various sources in the IoT environments. For example, although *CNN/RNN* could be used in an end-to-end manner (e.g., as image/time-series classifiers), they could also be used as feature extractors, based on which one can easily aggregate features extracted from different sources (e.g., audio, images, sensor data). With high-dimensional video data, one can either model by training *CNN+ LSTM* jointly [342], or use *CNN/AE* as feature extractors, before the sequential modelling (e.g., using *LSTM*). However, when modelling the data with limited labels (e.g., rare event), one needs to consider the potential overfitting effect when using DL directly. One may go back to the TML approaches or use some upsampling techniques (e.g., *GAN*) to alleviate this effect.

2.1.4 Reinforcement Learning (RL). In this section, we first introduce the strategies used to formulate the aforementioned video streaming example (see §2) with Reinforcement Learning (RL). As mentioned earlier, in RL an *agent* interacts with the *environment*, learning an optimal control policy through experience. It requires three key elements, *observation*, *action*, and *reward*. Based on these, we can formulate the adaptive bitrate streaming problem. Specifically, *observation*

can be the buffer occupancy, network throughput, etc. At each step, the *agent* decides the bitrate of the next chunk. A *reward* (for example the quality of service feedback from the user) is received after the *agent* takes *action* (chunk bitrate). The algorithm proposed in [222] collects and generalizes the results of performing the past decisions and optimizes its policy from different network conditions. This RL-based algorithm can also make the system robust to various environmental noises such as unseen network conditions, video properties, etc.

As shown in Fig. 8, there is a plethora of algorithms in the whole reinforcement learning family. More details of these RL algorithms can be found in the **RL methods in Appendix B**, and here we focus on selecting appropriate RL algorithms based on different selection criteria




Fig. 8. Reinforcement Learning Categorization

Environment Modelling Cost In RL modelling, sample efficiency is one of the major challenges. Normally the RL agent can interact either with the real world or a simulated environment during training. However, it can be difficult to simulate the heterogeneous IoT environments and complex IoT devices. RL models can also be trained directly in real world IoT environments, yet one major limitation is the heavy training cost, which may range from seconds to minutes for each step. The model-based RL method, a method that can reduce the sample complexity, can decrease the training time significantly. It first learns a predictive model of the real world, based on which the decisions can be made. When compared with model-free approaches, model-based methods are still in their infancy, and because of the efficiency property, they may attract more attention in the near future.

Action Space: The action space of RL algorithms can be either continuous or discrete. For those RL algorithms with discrete action space, they choose from a finite number of actions at runtime. Take the video streaming task for example, the action space is different bitrates for each chunk. Another task formulated in discrete action space can be found in [221], where the action space is the “schedule of the job at i -th slot”. Available algorithms for discrete action space tasks most reside in the policy gradient group, for example DQN, DDQN. The continuous action space, on the other hand, is infinite for all possible actions. Relationships exist between the actions that are usually sampled from certain distributions such as Gaussian distribution. For example, in an energy-harvesting management system, PPO algorithm [292] is used to control IoT nodes for power allocation. The action space, as stated in [245], is sampled from a Gaussian distribution to denote the load of each node ranging from 0% to 100%. Similarly, in another work [9] that studied energy harvesting WSNs, the Actor-Critic [179] algorithm is implemented to control the packet rate during transmission. One advantage of continuous action space lies in its ability to accurately control the system, thus a higher QoE is expected.

2.2 Model Generation

Based on the user requirement and task specification, we have selected a variety of models. Next, the models need to be developed and implemented. In this section, we will introduce the available tools for the development. We will also present the approaches that can be utilized to accelerate the training process.

2.2.1 Machine Learning Development Framework. The training and execution of ML models can be tricky and it may require numerous engineering efforts. Efforts have been devoted to developing frameworks to support the model

DL frameworks	Core language	Interface	Pros	Cons
Tensorflow (2)	C++	Python, Javascript, C++, Java, Go	<ul style="list-style-type: none"> - Effective data visualization - Distributed learning - Efficient model serving - On-device inference with low latency for mobile devices - Eager Execution with TF2, easy to debug 	<ul style="list-style-type: none"> - Steep learning curve (migration from TF 1 to TF 2) - Poor results for speed
Pytorch	C/C++	Python, C++	<ul style="list-style-type: none"> - Simple and transparent modeling - Eager execution 	<ul style="list-style-type: none"> - Hard to serve even with ONNX support
Caffe (2)	C++	Python, C++	<ul style="list-style-type: none"> - Fast, scalable, and lightweight - Server optimized inference 	<ul style="list-style-type: none"> - Limited community support - Limited in implementing complex networks
Mxnet	C++	Python, C++, Java, Julia, R, Perl, Clojure	<ul style="list-style-type: none"> - Fast, flexible, and efficient in terms of running DL algorithms - Run on any device - Easy model serving - Highly scalable 	<ul style="list-style-type: none"> - Smaller community compared with Tensorflow or Pytorch
DL4J	Java	Java, Clojure, Kotlin	<ul style="list-style-type: none"> - Robust, flexible and effective - Works with Apache Hadoop and Spark 	<ul style="list-style-type: none"> - Robust, flexible and effective - Works with Apache Hadoop and Spark

Table 1. Comparison of Deep Learning Frameworks

development. These frameworks have their own strengths and weaknesses in terms of the supported models, usability, scalability, etc. In this section, we will review several development frameworks.

For TML, the most famous development framework is Sci-kit learn. It is a free ML library with Python interface. Sci-kit learn supports almost all main-stream machine learning models and is a popular tool for fast prototyping. For DL, we list some of the most popular DL frameworks and discuss their pros and cons in Table 1. Users can choose the most suitable frameworks based on their needs.

When IoT comes into context, more challenges arise with edge computing as it is trying to move the computation close where the data is generated [300]. The device heterogeneity of edge computing has made the development of the DL models more complicated. There are many portable Edge computing devices, each optimized with different inference engines. For example Nvidia Jetson series GPU computing unit compiles models with TensorRT³ inference engine while TensorFlow Lite⁴ is specially optimized for Google coral TPU. These inference engines optimize the model graph and quantize the model parameters to lower precision, thus delivering low latency and high-throughput for on-device inference. Some attempts [330] have been made to integrate both inference engines but the compatibility issue still exists. TVM [57] breaks the boundaries among diverse hardware, aiming at cross-framework and cross-device end-to-end optimization of DL models.

2.2.2 Single Machine Learning (Centralized). Model training via single machine is a common strategy for ML model generation. By placing the learning-related computation in the same place, the model learns from the data and updates its parameters. In this subsection, we highlight two approaches that leverage hardware for the training process acceleration: *Computation Optimization, Algorithm Optimization.*

Computation Optimization The basic computation unit of neural networks consists of vector-vector, vector-matrix and matrix-matrix operations. Efficient implementation of computations can accelerate the training and inference process. The Basic Linear Algebra Subprogram (BLAS)⁵ standardizes the building blocks for basic vector, matrix operations. A higher level linear algebra library such as cuBLAS⁶ implements BLAS on top of NVIDIA CUDA and is

³<https://developer.nvidia.com/tensorrt>

⁴<https://www.tensorflow.org/lite>

⁵<http://www.netlib.org/blas>

⁶<https://docs.nvidia.com/cuda/cublas/>

efficient in utilizing the GPU computation resource. Intel Math Kernel Library (MKL)⁷, on the other hand, maximizes performance on Intel processors and is compatible with BLAS without the change of code.

Different DL architectures (e.g., DNNs, CNNs and RNNs) may require different optimizations in terms of basic computations. The DNN computation is usually basic matrix-matrix multiplication and the aforementioned BLAS libraries can efficiently accelerate the computations with GPU resources. The CNNs and RNNs are different in their convolution and recurrent computations. Convolutions can not fully utilize the multi-core processors and the acceleration can be achieved by unrolling the convolution [52] to matrix-matrix computation or computing convolutions as point-wise product [224]. For RNN (LSTM), the complex gate structures and consecutive recurrent layers differ from the DNNs and CNNs in that these computation units can not be split and deployed directly at different devices. This has made parallel computation difficult to apply. Optimization is possible though, with implementations on top of NVIDIA cuDNN [60]. Computations among the same gates can be grouped into larger matrix operations [10] and save intermediate steps. We can also accelerate by caching RNN units' weights with the GPU's inverted memory hierarchy [89]. The weights are reusable between time steps, making a maximum 30× speed up on a TitanX GPU.

Algorithm Optimization Apart from the resource utilization optimization, the algorithmic level optimization is another important research direction for efficient model training and faster convergence. Optimization algorithms aim at minimizing/maximizing a loss function that varies for different machine learning tasks. They can be divided into two categories: *First Order Optimization* and *Second Order Optimization*.

First Order Optimization methods minimizing/maximizing the loss function with the gradient values with respect to the model parameters. Gradient Descent is one of the most important algorithms for neural networks. After back-propagation from the loss function, the model parameters are updated towards the opposite direction of the gradient. Gradient descent approaches fall into local optima when the absolute value is either too big or too small. Also it updates the gradient of the whole data set at one time, memory limitation is always a big problem. Variants have been proposed to address the aforementioned problems, including Stochastic gradient descent [36], mini-batch gradient descent [86]. Also, much famous research enables faster model convergence: Momentum [265], AdaGrad [97], RMSProp [144], ADAM [175]. *Second Order Optimization* methods take second order derivative for minimizing/maximizing loss function. Compared to the *First Order Optimization*, it consumes more computation power and is less popular for machine learning model training. However, *Second Order Optimization* considers the surface curvature performance and is less likely to get stuck on saddle points. Thus it sometimes outperforms the *First Order Optimization*. Famous *Second Order Optimization* methods include [14, 46, 141, 240, 251]. For more systematic survey on the optimization methods for machine learning training, one can refer to [37].

2.2.3 Distributed Machine Learning. Modern ML models such as neural networks require a substantial amount of data for the training process. These data are usually aggregated and stored in the cloud server where training happens. However, when the training process of large volume data outpaces the computing power of a single machine, we need to leverage multiple machines available in the server cluster. This requires the development of novel distributed ML systems and parallel training mechanisms which distribute and accelerate the machine learning workload.

Fig 9 shows the schematic diagram of a distributed ML pipeline. It has multiple components which are engaged in *Training Concurrency*, *Single Machine Optimization*, and *Distributed System*. In *Training Concurrency*, either the models or the data are split into small chunks and placed on different devices. With *Single Machine Optimization* (which shares

⁷<https://software.intel.com/en-us/mkl>




Fig. 9. Distributed Machine Learning Pipeline

similar techniques as conventional ML, see § 2.2.2) that accelerates the training process, we get all local gradient updates. Finally, *Distributed System* discusses strategies that efficiently aggregate the gradient updates.

Training Concurrency in Distributed ML. In the distributed machine learning, the selection of parallel strategy depends on two factors: data size and model size. When either the datasets or the model parameters are too big for single-machine processing, it is straightforward to consider partitioning them into smaller chunks for processing at different places. Here we first introduce two basic methods *data parallel*, *model parallel*. We also introduce *pipeline parallel* and other hybrid approaches that take advantage of both approaches.

Data Parallel. In a multi-core system where a single core can not store all the data, data parallel is considered by either splitting the data samples or the features. Data parallel has been successfully applied to numerous machine learning algorithms [66] with each core working independently on a subset of data. It can be used for training ML algorithm for example decision trees and other linear models where the features are relatively independent. Parallel with the split of data features, though, it can not be used directly with neural networks because different dimensions of the features are highly correlated.

In deep learning, data parallel works by distributing the training dataset across different GPU units. The dominant data parallel approach is *batch parallelism* where mini-batch SGD is employed to compute local gradient updates on a subset of the data. A central server is responsible for aggregating all local updates to global parameter and pushing new models back to the working units. One of the earliest works trained with GPUs can be found in [271] where the authors implemented distributed mini-batch SGD unsupervised learning concurrently with thousands of threads in a single GPU. By varying the batch size [124, 312, 373], this method is effective in reducing the communication cost without too much accuracy loss. In the next paragraph, we will discuss more about the parallel SGD algorithms [226, 316, 375, 383, 399] for improving the communication efficiency, which can be seen as one way of improving the performance of data parallelism. Another type of data parallel that addresses the memory limit on single GPU is *spatial parallelism* [167]. *Spatial parallelism* considers partitioning spatial tensors into smaller subdivisions and allocating them to separate processing units. It thus differs from *batch parallelism* in that the latter puts the groups of data in the same process. *Spatial parallelism* approach has proven to show near linear speedup on modern multi-GPU systems.

Model Parallel. Data parallel suffers from the infeasibility of dealing with very large models especially when it exceeds the capacity of a single node. Model parallel addresses this problem by splitting the model with only a subset of the whole model running on each node [39, 84, 174, 192]. The computation graphs can be divided within the layers (horizontal) or across the layers (vertical). Mesh-tensorflow [299] allows linear within-layer scaling of model parameters across multiple devices after compiling a computation graph into a SPMD program. However, this approach requires high communication cost as it needs to split and combine model updates across a large number of units. [159] introduced

decoupled parallel backpropagation to break the sequential limitation of the back-propagation between the nodes, greatly increasing the training speed without much accuracy loss. For CNN, as each layer can be specified as five dimensions including: samples, height, width, channels, and filters, existing literature [93, 94] studies the split of models among dimensions. Another research direction that optimizes the communication overhead is by searching the optimal partition and device placement of computation graphs via reinforcement learning [234]. The literature [166, 349] followed this idea and shows interest in automatic search of optimal parallel strategies.

Pipeline Parallel. Although model parallel has proven successful in training extremely large models, the implementation is complicated due to the complexity of the neural network structure. This is especially true for CNNs since the convolution operators are highly correlated. Also, GPU utilization is low for model parallel. Due to the gradient interdependence between different partitions, usually only one GPU is in use each time. To solve the aforementioned problems, pipelining has been studied [260, 362] for speeding up the model training. With pipeline parallel, models are partitioned and displayed across different GPUs. Then mini-batches of training data are injected to the pipeline for concurrent processing of different inputs at the same time. Fewer worker GPUs are idle in the pipeline parallel setting as each node is allocated jobs, without waiting for other nodes to finish their work. According to the synchronization strategy we discussed earlier, gradients are aggregated by either synchronous pipeline model (GPipe [157]) or asynchronous pipeline model (PipeDream [138], SpecTrain [53], XPipe [129]). Theoretical analysis of pipeline parallel optimization has also been studied and with Pipeline Parallel Random Smoothing (PPRS) [70], convergence rates can be further accelerated.

Hybrid. Data and model parallel are not mutually exclusive. Hybrid approaches that combine the benefits of both methods are effective in further accelerating the training process. Pipeline parallel [138, 157] can be seen as an approach built on top of data parallel and model parallel. Apart from that, [183] proposed combining data parallel and model parallel for different types of operators. With data parallel for CNN layers and model parallel for DNN layers, it achieved a 6.25 \times speed up with only 1% of accuracy loss on eight GPUs. Another implementation MAPS-Multi [20] borrows the idea of [183] and automates the partitioning of workload among multiple GPUs, achieving 3.12 \times speed up on four GTX 780 GPUs. Other forms of data parallel and model parallel hybrids exist in the literature [61, 84, 116, 117] that reduce the overall communication and computation overhead.

Distributed ML System. When we have acquired a local model update with partial data slice, multi-node and multi-thread collaboration are important for effectively updating the model. Network communication plays an important role in sharing the information across the nodes. In this section, we present the three most important features in network communication: 1) network topology, 2) synchronization strategy and 3) communication efficiency.

Network Topology. The network topology defines the node connection approach in the distributed machine learning system. When the data and models are relatively simple, it is common to utilize existing Message Passing Interface (MPI) or MapReduce infrastructure for the training. Later when the systems are becoming more and more complex, new topologies should be designed to facilitate the parameter update.

The Iterative MapReduce (IMR) or AllReduce approaches are commonly used for synchronous data parallel training. Typical IMR engines (for example the Spark MLlib [230]) generalizes MapReduce and enables the iterative training required by most ML algorithms. Synchronous training can also be implemented by AllReduce topology. MPI⁸ (Message Passing Interface) supports AllReduce and is efficient for CPU-CPU communication. Many researchers implement their

⁸<https://computing.llnl.gov/tutorials/mpi/>

own version of AllReduce for example Caffe2 Gloo⁹, Baidu Ring AllReduce¹⁰. In the ring-Allreduce topology, all nodes connect to each other without a central server, just like a ring. The training gradients are aggregated through their neighbors on the ring. To provide more efficient communication for DL workload in the GPU cluster, libraries such as Nvidia NCCL [72] are developed and support the AllReduce topology. In NCCL2 [164], the multi-node distribution feature is also introduced. Horovod [294] replaces the Baidu ring-Allreduce backend with NCCL2 for efficient distribution.

A Parameter Server (PS) infrastructure [200] is usually composed of a set of worker nodes and a server node which gathers and distributes computation from worker nodes. As asynchronous training of PS neglects stragglers, it provides better fault tolerance capability when some of the nodes break down. Parameter server also features high scalability and flexibility. Users can add nodes to the cluster without restarting the cluster. Famous projects such as DMTK Microsoft Multiverso [99], Petuum [364] and DistBelief [84] enable training of even larger networks.

Synchronization Strategy. In distributed ML, model parameter synchronization between worker nodes is cost-extensive. The trade-off between the communication and the fresher updates has great impact on the parallelism efficiency.

Bulk Synchronous Parallel (BSP) [225] is the simplest strategy for ensuring model consistency of all worker nodes. For each training iteration, all nodes wait for the last (slowest) node to finish the computation and the next iteration does not start before the all the model updates are aggregated. Total Asynchronous Parallel (TAP) [84] approaches are proposed to address the problem of the stragglers within the network. With TAP, all worker nodes access the global model via a shared memory. They can pull and update global model parameters any time when the training is finished. As there is no update barrier for this approach, the system fault tolerance is greatly improved. However, stale model updates can not guarantee convergence to global optimum. Many famous frameworks use the TAP strategy, including Hogwild! [277] and Cyclades [252].

Stale Synchronous Parallel (SSP) [146] compromises between fully-synchronous and asynchronous schemes. It allows a maximum staleness by allowing faster working nodes to read global parameters without waiting for slower nodes. As a result, the workers spend more time doing valuable computation, thereby improving the training speed greatly. But when there is too much staleness within the system, the convergence speed can be significantly reduced. Many state-of-the-art distributed training systems implement BSP and SSP for efficient parallelism, for example tensorflow [1], Geeps [79], Petuum [364].

In contrast to the SSP which limits the staleness of the model update, the Approximate Synchronous Parallel [152] (ASP) limits the correctness. In Gaia [152], for each local model updates, the global parameter is aggregated only if the parameter change exceeds a predefined threshold. This “significance” only strategy eliminates unnecessary model update and is efficient in utilizing the limited bandwidth. However, the empirical determination of threshold only considers the network traffic and is insufficient for dealing with dynamics in the IoT environment. [350] has addressed this problem by also considering resource constraints for efficient parallelism.

Communication Efficiency. Communication overhead is the key and often the bottleneck in distributed machine learning [201]. The sequential optimization algorithms implemented in the worker nodes require frequent read and write from the global shared parameters which poses great challenge on balancing network bandwidth and communication frequency. To increase the communication efficiency, we can either reduce the size of the model gradient (communication content) or the communication frequency.

⁹<https://github.com/facebookincubator/gloo>

¹⁰<https://github.com/baidu-research/baidu-allreduce>

Communication content. The gradient size between working nodes is correlated to both the model size itself and the gradient compression rate. We have reviewed four types of *model compression* techniques in §2.3.2 which are effective in reducing the overall gradient size. Hereby we focus on the techniques that compress the gradient before transmission, discusses the gradient *quantization* and *sparsification*.

Gradient quantization differs from the weight quantization (§2.3.2) as the former compresses the gradient transmission between worker nodes while the latter focuses on faster inference via smaller model size. Works that reduce the gradient precision [83] have been proposed and 1-bit quantization [293, 318] is effective in greatly reducing the computation overhead. Based on the idea, QSGD [7] and Terngrad [356] consider stochastic quantization where gradients are randomly rounded to lower precision. Additionally, weight quantization and gradient quantization can also be combined [150, 158, 361, 382, 397] for efficient on device acceleration.

The weights of the DNNs are usually sparse and due to the large number of unchanged weights in each iteration, the gradient updates are even more sparse. This sparsification nature of the gradient transmission has been utilized for more efficient communication. Gradient sparsification works by sending only important gradients when exceeding a fixed threshold [318] or adaptive threshold [95]. Gradient Dropping [4] uses layer normalization to keep the convergence speed. DGC [207] uses local gradient clipping for sending important gradients first while the less important ones are aggregated with momentum correction for later transmission.

Communication Frequency. Local (Parallel) SGD [226, 316, 375, 383, 399] entails performing local updates several times before parameter aggregation. Motivated by reducing the inter-node communication, this approach is also called model averaging. One-shot averaging [226, 399] considers only one aggregation during the whole training process. While [383] argues that one-shot averaging can cause inaccuracy and proposes more frequent communications, many works [206, 263, 374, 386] prove the applicability of the model averaging approach in various deep learning applications. In an asynchronous setting, the communication frequency can also be maneuvered through the push and pull operations in the worker nodes. DistBelief [84] has adopted this approach with a larger push interval compared to the pull interval.

2.2.4 Federated Learning. In traditional distributed machine learning, the training usually happens on the cloud data center with aggregated training data generated by collecting, labelling and shuffling raw data. The training data is thus considered *identical and independent distributed (IID)* and balanced. This facilitates the training process as one only needs to consider distributing the training task across various computation units and updating the model by aggregating all local gradient updates. However, this is not the case when IoT comes into play. The ML-based IoT applications differ from the traditional ML applications in that they usually generate data from heterogeneous geo-distributed devices (e.g., user behavior data from mobile phones). These data can be privacy-sensitive as users usually prefer not to leak personal information, making conventional distributed ML algorithms infeasible for solving such problems. Thus novel optimization techniques are required to enable training in such scenarios.

Federated learning (FL) [180] is a type of distributed machine learning research that moves the training close to the distributed IoT devices. It learns a global model by aggregating local gradient updates and does not require the movement of the raw data to the cloud center. FederatedAveraging (FedAvg) [227] is a decentralized learning algorithm specifically designed for the FL. It implements synchronous local SGD [54] on each device with a global server averaging over a fraction of all the model updates per iteration. FedAvg is capable of training high-accuracy models on various datasets with many fewer communication rounds. Following this work, [180] proposed two approaches: *Structured updates* and *sketched updates* for reducing the communication cost, achieving higher communication efficiency. Further

research addresses the privacy limitation of FL by Differential Privacy [228] and Secure Aggregation [34]. Finally, [32] delivers system-level implementation of FL based on previously mentioned techniques. It is able to train deep learning models with local data stored on mobile phones.

FL is still developing rapidly with many challenges remaining to be solved. On the one hand, FL shares similar challenges as in conventional distributed machine learning methods in terms of more efficient communication protocol, synchronization strategy as well as parallel optimization algorithms. On the other hand, the distinct setting of FL requires more research preserving the privacy of training data, ensuring the fairness and addressing bias in the data. For a more thorough survey on details of FL, one can refer to [170].

2.2.5 Knowledge Transfer Learning. The knowledge learnt from trained models can be transferred and adapt to new tasks, and it is especially helpful when limited data/labels are available. In this section we introduce four types of knowledge transfer learning (KTL) approaches: Transfer learning, Meta learning, Online learning and Continual learning.

Transfer Learning [327] – transferring knowledge across datasets– is the most popular KTL approach. It trains a model in the source domain (with adequate data/labels, e.g., on ImageNet [87] for general visual recognition tasks), and fine-tunes the model parameters in the target domain to accommodate the new tasks (e.g., medical imaging analysis on rare diseases). The rationale behind is that low-level and mid-level features can be representative enough and thus shared across different domains. In this case, only the parameters related to high-level feature extraction need to be updated. This mechanism does not require a large amount of data annotation for learning reliable representation in the new tasks, which could be useful in cases when annotations are expensive (e.g., medical applications). **Meta Learning** [340] is another popular KTL approach; instead of transferring knowledge across datasets, it focuses on knowledge transfer across tasks. Meta learning means learning knowledge or patterns from a large number of tasks, then transfer this knowledge for more efficient learning of new tasks. When with continuous data streaming, it is also desirable to update the model with incoming data, and in this case **Online Learning** [149] can be used. However, it is difficult to model when the incoming data is from a different distribution or a different task. Most recently, **Continual Learning**[255] was proposed to address this issue. Not only can it accommodate the new tasks or data with unknown distribution, it can also maintain the performance on the old/historical tasks (i.e., no forgetting [172]), making it a practical tool for real-world IoT applications. These four KTL approaches are similar in concept yet have different use cases. Transfer/Meta learning are focused on knowledge transfer across datasets or tasks (irrespective of data types), while online/continuous learning are more suitable for data streaming and can transfer the knowledge continuously to the new incoming data or tasks.

2.2.6 Discussion. Effort has been devoted to implementing the distributed machine learning on top of modern deep learning frameworks. Remarkable results have been achieved where with proper implementation [124] with Tensorflow, the training time of the state-of-art ImageNet can be reduced from days to one hour. Compared with Tensorflow, more efficient implementation such as Horovod can increase the GPU utilization for even more acceleration. Horovod has already been incorporated in various deep learning framework ecosystems (e.g., Pytorch, Mxnet).

Deep learning ecosystems free the researcher from heavy implementation effort. There are however, challenges for model generation in a distributed setting: (1) *The choice of hardware.* The same implementation can have different performance on different devices. One would have to be aware of the device features for efficient acceleration. (2) *Parallel hyperparameter tuning strategy.* Compared with single machine training, the distributed system is more complex and it is thus more difficult to find an optimal structure. (3) *Effective work of DL frameworks with other big data application like*

Hadoop/Spark. Existing big data frameworks (e.g., Spark/Hadoop) can also be applied for effectively distributing the DL training pipeline, and a deeper integration of both frameworks is urgently required.

2.3 Model Optimization

We have discussed the model selection and model generation where a model is generated catering to the specific needs of IoT applications. There are, however, still things to be considered before model deployment. The IoT application differs significantly from other areas in terms of deployment devices and data sources. The limited computational budget of edge devices requires smaller models for small-scale computational workload to ensure low inference latency. Also, heterogeneous data sources in IoT environments usually contain redundant information that can even mislead the decision of the ML models. It is important to select only relevant and informative features or to perform model compression for performance optimization. In this section, we discuss these two topics.

2.3.1 Feature Selection. The amount of high-dimensional data in IoT environments poses challenges on the training of the ML algorithms. Noisy and redundant signals exist and they may consume substantial computational power. For unstructured data (e.g., images, audio, etc.), a common way to reduce the data redundancy is to perform feature extraction, which can be realised in many ways, e.g., via feature engineering, supervised/unsupervised learning, etc. The extracted features can be deemed as structured data, and its redundancy can be further reduced by feature selection. Feature selection can help reducing the computational complexity, which may improve the performance in terms of both effectiveness and efficiency—crucial factors in the limited-resource IoT environments. Briefly, feature selection is the process of preserving relevant features while discarding irrelevant/redundant features. There are generally three categories of feature selection approach, namely the *Filter* approach [18, 81, 135], the *Wrapper* approach [119, 132, 169, 223] and the *Embedded* approach [42, 270].

The application of feature selection in the IoT literature is emerging especially when data quality plays a significant role in the performance of the learning application. In *Smart Farming*, early detection and treatment of plant disease or nutrient deficiency would increase the health of the crops. An IoT based *Smart Farming* system can give real-time decisions with optimal fertilizer volume. [171] [232] show that when there are many sensors with a large volume of IoT streaming data, feature selection can remove the redundant information and maximize the prediction accuracy. Another application is *Intrusion Detection*, which aims to detect online behaviour that compromises the confidentiality of the network. In this case, feature selection can be used to extract the most informative patterns [188, 193], serving as the basis of an effective intrusion detection system.

For future research, one could extend the single-object optimization to multi-object optimization. For example in IoT systems, optimal feature selection algorithms assist the machine learning models to optimize the execution time. We can explore the modifications of the feature selection algorithm to minimize the energy consumption of routing decisions as well [98]. One can also study how to detect the dynamics within the data flow and then adaptively apply the search algorithms accordingly to further improve the performance of the feature selection algorithms.

2.3.2 Model efficiency. The state-of-the-art DL models often require high computational resources beyond the capabilities of IoT devices. Those models that perform well on large CPU and GPU clusters may suffer from unacceptable inference latency or even be unable to run on edge devices (e.g. Raspberry Pi). Tuning the deep neural network architectures to increase the efficiency without sacrificing much accuracy has been an active research area. In this section, we will cover three main optimization directions: *Efficient architecture design*, *Neural architecture search* and *Model compression*.

Efficient architecture design. There exist neural networks that can specifically match the resource and application constraints. They aim to explore highly efficient basic architecture specially designed for platforms such as mobiles, robots as well as other IoT devices. MobileNets [151] is among the most famous works and proposed to use depth-wise separable convolutions [307] to build CNN models. By controlling the network hyper-parameters, MobileNets can strike an optimal balance between the accuracy and the constraints (e.g., computing resources). Later in MobileNetv2 [288], the inverted residual with linear bottleneck architecture was introduced to significantly reduce the operations and memory usage. Other important works include Xception [64] ShuffleNet [389], ShuffleNetv2 [217], CondenseNet [156]. These neural networks optimize on-device inference performance via efficient design of building blocks, achieving much less computational complexity while keeping or even raising accuracy on various computer vision datasets. Some work even outperforms the neural architectures generated through exhaustive automatic model search. Also, different building blocks can be combined together for even lighter models.

Neural architecture search (NAS). Another research direction named neural architecture search aims at searching an optimal network structure in a predefined search space. There are usually three types of algorithms: reinforcement learning approach [209, 261], Genetic Algorithm (GA) based [210, 276], and other algorithms [15, 44].

The models generated by these methods are normally constrained to smaller model sizes. Model size and operation quality are the two most common metrics to be optimized, over other metrics such as inference time or power consumption. Representative works including MONAS [153], DPP-Net [91], RENA [398], Pareto-NASH [102] and MnasNet [328] are interested in finding the best model architectures to meet these constraints. These approaches are more straightforward as they optimize directly over real world performance. However, one drawback of NAS is the extensive computing power required for finding the optimal neural architectures. Thus, the already generated architectures can be utilized as guidance for future design for more efficient neural network architecture.

Model Compression. As modern state-of-art DL models can be very large, reducing the model computation cost is crucial for deploying the models on IoT devices, especially for those latency-sensitive real-time applications. Model compression methods can be divided into four categories, *Parameter pruning and sharing*, *Low-rank factorization*, *Transferred/compact convolutional filters* and *Knowledge distillation*. Hereby we briefly summaries the categories of model compression techniques and list several important works.

Parameter pruning and sharing method aims to find and remove the redundant parameters (of DL models) for higher efficiency and generalization. One direction is to apply quantization techniques. The DL's parameters/weights are usually stored in memory with 32-bits, and quantization techniques can compress them into 16-bits [131], 8-bits [339] or even 1-bit [74, 75, 275]. On the other hand, weight pruning and sharing (in pre-trained DL models) has also attracted interest among the community. Some popular methods imposed L1 or L2 regularization constraints [191, 355], which can penalize models with more parameters, yielding the effect of pruning unnecessary parameters.

Low-rank factorization method decomposes the CNN or DNN tensors to lower ranks. The tensor matrix decomposition is implemented for each layer of the DL model. That is, once the decomposition for a certain layer is completed, the parameter size will be fixed (for this layer) and the decomposition will proceed to the next layer. Interesting work can be found in [190]. However, there are two major drawbacks. For very large DL models, it may be very expensive to perform decomposition owing to large parameter matrices. On the other hand, its layer-wise nature may also yield cumulative error, diverting the compression results to be far from optimal.

Transferred/compact convolutional filters method reduces the memory consumption by implementing special structural convolutional filters. Motivated by the equivariant group theory [69], the transferred convolutional filter transforms the

network layers to a more compact structure, thus reducing the overall parameter space. The family of transformation functions [195, 298, 378] operates in the spatial domain of the convolutional filters to compress the whole network parameters. Compared to other model compression methods, transferred convolutional filters methods are less stable due to the strong transfer assumptions. However, when the assumption holds, the compact convolutional filter can have very good performance. In [326, 360], the filters were decomposed from 3×3 or bigger to 1×1 convolutions—ideal operations for IoT devices.

Knowledge distillation method learns a new, more compact model that mimics the function presented by the original complex DL model. The idea came from the work in [45], where a neural network was applied to mimic the behavior of a large classifier ensemble system. Later this idea has been extended to the complex DL methods [145], more details can be found in [16, 282, 377]. However, currently the knowledge distillation methods are limited to classification tasks and further development is required.

Discussion. Types of model compression techniques have their own strengths and weaknesses and thus optimal choice is based on specific user requirements. *Parameter pruning and sharing* methods are the most commonly applied techniques for compression models from original models. It is stable as with proper tuning, this approach usually delivers no or few accuracy losses. On the other hand, *Transferred/compact convolutional filters* methods address the compression from scratch. This end-to-end efficient design for improving the CNN performance approach shares similar insights to the efficient *neural architecture design* approach as we discussed earlier. *Knowledge distillation* methods are promising when working with relatively small datasets as the student model can benefit from the teacher model with less data. All these methods are not mutually exclusive, we can make combinations based on specific use cases to optimize the models that are more suitable for low-resource IoT devices.

2.4 Model Evaluation

After the models have been trained, based on suitable metrics their performance should be evaluated before deployment. Accuracy is one of the most popular evaluation metrics in classification tasks, yet it faces several problems in different scenarios. For example, it is an overall measure without indicating the recognition capability for each class, which may be heavily biased if there exists a significant class imbalance problem. There are various evaluation metrics and it is key to select the most appropriate one. For the rest of this section we investigate several widely used metrics for classification and regression tasks. For classification/regression tasks, one aims to construct a model (i.e. $f(\cdot)$) that can predict the value of dependent variable Y from independent variable X . The difference between these two tasks is the fact that the dependent variable Y is numerical for regression and categorical for classification.

2.4.1 Classification Problem based metric. In classification tasks, one of the most effective evaluation metrics is a confusion matrix [332]. As demonstrated in Table 2 for a binary classification task, in a confusion matrix the row represents the predicted class and the column represents the ground truth (actual class). The entries True Positive (tp) and True Negative (tn) represent the correctly classified positive and negative samples, while the entries False Negative (fn) and False Positive (fp) denote the misclassified positive and negative samples, respectively.

	Actual Positive Class	Actual Negative Class
Predicted Positive Class	True Positive (tp)	False Negative (fn)
Predicted Negative Class	False Positive (fp)	True Negative (tn)

Table 2. Confusion Matrix for Classification

Based on the confusion matrix, several evaluation metrics can be derived. The *accuracy* (i.e., $\frac{tp+tn}{tp+tn+fp+fn}$) and *error rate* (i.e., $\frac{fp+fn}{tp+tn+fp+fn}$ or $1 - accuracy$) are the most commonly used metrics because it is more understandable and intuitive for humans. However, these two metrics are powerless in terms of class-wise informativeness [218], which may neglect the minority class [51] (if there is a class imbalance problem).

The metrics *precision* and *recall* can be used to measure the performance irrespective of the class imbalance problem (more definitions of classification evaluation metrics can be found in **Table 4 in Appendix C**). In binary classification problems, as mentioned earlier tp, fn, fp are defined as the the number of “positives correctly classified as positives”, “positives incorrectly classified as negatives”, “negatives incorrectly classified as positives” respectively. Then we can see *recall* (i.e., $\frac{tp}{tp+fn}$) indicates the ability of a classifier to detect (true) positives out of all positive instances, while *precision* (i.e., $\frac{tp}{tp+fp}$) is the percentage of detected (true) positives out of all the detected ones. Since the binary classification’s decision may highly depend on the threshold, there is a trade-off between precision and recall. For example, if a high threshold has been chosen—the similarity scores (the model outputs) have to be higher to give positive decisions—the classifier tends to have high fn and low fp , yielding low recall and high precision. Similarly, reducing the value of the threshold may increase recall and decrease precision accordingly. For different applications, one needs to consider the optimal threshold for their requirements. For example, forensic applications may prioritize high precision (i.e., low in fp) while a medical diagnosis may prioritize high recall (i.e., low in fn).

In some tasks when both recall and precision are important, the *F1-score* (i.e., $2 \frac{precision \cdot recall}{precision+recall}$)—a measure that can balance the precision/recall trade-off—is normally used. It is worth noting that for multi-class cases, the multi-class confusion matrix can be calculated, and the aforementioned precision/recall/F1-score can be extended to measure the class-wise performance. Depending on the data/applications, the overall performance can be measured by aggregating all the class-wise metrics. Two popular aggregation operations are averaging, and weighted averaging, e.g., mean F1-score, or weighted F1-score (over all the class-wise F1-scores).

2.4.2 Regression Problem based metric. For regression problems, the evaluation metrics are different from the classification ones. Popular evaluation metrics include Mean Squared Error (MSE), Mean Absolute Error (MAE), Mean Percentage Error (MPE), etc. Details and formulas of these metrics can be found in **Table 5 Appendix C**.

Mean Absolute Error (MAE) and *Mean Squared Error (MSE)* are the simplest metrics for regression evaluation. They denote the expected model errors defined in terms of absolute difference and squared difference (between the predicted value and the ground truth), respectively. Alternatively, the *Mean Absolute Percentage Error (MAPE)* and *Mean Percentage Error (MPE)* can also be applied to regression problems. The MAPE is similar to MAE but more intuitive as it shows percentage. The MPE lacks the absolute term on MAPE, which means the positive and negative errors will cancel out. In this case, the MPE can not be directly used to measure the performance of a model. However, it can be used to check whether the model systematically underestimates (more negative errors) or overestimates (more positive errors).

All of the above metrics can be applied to the regression problem, but it is important to consider the property of the dataset beforehand. For example, some fields may (or may not) be more prone to outliers, and the corresponding (effective) evaluation metrics may be different.

3 MODEL DEPLOYMENT

When the ML model development process is finished, the developed models are to be deployed and composed as an application in the complex IoT environments. To simplify the deployment, the ML models and underlying infrastructure need to be specified (§3.1). Next, the optimization techniques can be applied to generate the deployment plans that select

the suitable ML models for the deployment, optimizes the resource utilization of the model deployment and improve the reusability of the deployed models (§3.2). Once the deployment plans are generated, the models will be deployed over the specified infrastructure and the deployed models will be composed as defined in the plan (§3.3).

3.1 Declarative Machine Learning and Deployment

Declarative ML. Declarative ML aims to use high-level language to specify ML tasks by separating the applications from the underlying data representation, model training and computing resources. There are *three* general properties of declarative ML. First, the high-level specification only considers data types of input, intermediate results and output. They are exposed as abstract data types without considering the physical representation of the data or how the data is processed by the underlying ML models. Second, the ML tasks are specified as high-level operations through well-defined semantics. The basic operation primitives and their expected accuracy levels (or confidence interval) are defined accordingly. Based on the operation semantics, declarative ML systems select the features and underlying ML models automatically or semi-automatically, optimize the model performance and accuracy for varying data characteristics and runtime environments. Notably, the selection is based on the available models, provided as services. Finally, the correctness of the ML models must be satisfied when a given model produces the equivalent results in any computing resources with the same input data and configurations. As a result, the declarative ML enables execution of the ML models over various hardware and computation platforms (such as Apache Spark) without any changes. Besides, these specification languages also bring flexibility and usability in the ML model deployment stage.

SystemML [31] is an implementation of declarative ML on Apache Spark. Through domain-specific languages, it specifies the ML models as abstract data types and operations, independent of their implementation. The system is able to specify the majority of ML models: matrix factorizations, dimension reduction, survival models for training and scoring, classification, descriptive statistics, clustering and regression. There is also other state-of-the-art research on declarative ML, including TUPAQ [313] and Columbus [380]. They utilize language specification and modelling technologies to describe the ML models for automatic model and feature selection, performance and resource optimization, model and data reuse.

Declarative Deployment. Hardware in the IoT environment consists of three basic types of device: *data generating* devices, *data processing* devices and *data transferring* devices. *Data generating* devices are also called “Things” (e.g., sensors, CCTV) and are used to collect environmental data. *Data transferring* devices such as router, IoT gateway, base station are used to transfer the generated data to the *data processing* devices. *Data processing* devices are used to run the analytic jobs. They can be GPU, CPU and TPU servers running in cloud or ARM based edge device such as Raspberry Pi and Arduino. An ML-based IoT application is usually running across a fully distributed environment, such that it requires correct specification of the component devices as well as the precise interoperation between these devices. [309] lists fundamental aspects that may simplify the hardware specification, i.e., processor, clock rate, general purpose input/output (GPIO), connectivity methods (Wi-Fi, Bluetooth, wired connection) and communication protocols (serial peripheral interface), universal asynchronous receiver-transmitter (UART).

Regarding the software, it is often categorized into three groups based on operation levels: *operating system (OS)*, *programming language* and *platform*. IoT OS allows users to achieve the basic behavior of a computer within internet-connected devices. The choice of OS in different layers of the IoT environment depends on the hardware properties such as memory and CPU. The *programming language* helps the developers to build various applications in different working environments with diverse constraints. The choice depends on the capability of devices and the purpose of

the application [47]. The IoT software *platform* is a system that simplifies the development and deployment of the ML-based IoT application. It is an essential element of a huge IoT ecosystem which can be leveraged to connect new elements to the system. For more details of the most popular OSs, programming languages and platforms in IoT domain, one can refer to **Appendix A**. The ML development *platforms* have been discussed in § 2.2.1.

The heterogeneity of IoT infrastructures makes the deployment very complicated and difficult to automate. To overcome this issue, the infrastructure must be described and specified by machine understandable languages. Then, the declarative deployment systems are able to automatically map the ML models to the infrastructures and generate the deployment plans that optimize the performance and the accuracy.

The declarative TOSCA model [80] is able to specify the common infrastructures such as Raspberry Pis and cloud VM (hardware), MQTT and XMPP (communication protocol). The deployment logic can be defined through *TOSCA Lifecycle Interface* that allows users to customize the deployment steps. However, this declarative model is still very basic and can not handle complex deployments such as specifying the details of ML based application. Moreover, the IoT applications consist of installing devices and sensors which require *human tasks*. These tasks are not natively supported by any available declarative deployment [43]. The imperative tool (e.g., kubectl commands) allows the technical experts with diverse knowledge of different deployment systems and APIs to interact with a deployment system and decide what actions should be taken. However, current imperative frameworks such as Juju, Kubernetes still do not support interactions such as sensor installation. In future, declarative deployment systems should interact with declarative ML systems to deploy a complex application over the heterogeneity of IoT infrastructure while supporting the *human tasks* through a more human centered imperative deployment model.

3.2 Deployment Optimization

When the infrastructures and deployment workflow of the ML models are specified, the deployment optimization problem can be formed as a mathematical expression subject to a set of system constraints. Then, resource allocation algorithms can be used to efficiently and precisely find the best solution for the given mathematical expressions. Moreover, the optimization objectives are a set of QoS parameters including storage and memory space, budget, task execution time and communication delay etc.,. These algorithms can be divided into *two* classes based on whether an optimal solution can be guaranteed: *meta-heuristic method* and *iterative method* (or *mathematical optimization*). Nowadays, ML methods are becoming popular and being applied to solve these resource allocation problems by learning “good” solutions from the data. We investigate the representative works in resource allocation based on these *three* classes.

Iterative-based method. This class of algorithm generates a sequence of improved approximate solutions where each solution is driven by previous solutions. Eventually, the solutions will converge to an optimal point proved by a rigorous mathematical analysis. The heuristic-based iterative methods are also very common, but we categorize this type of algorithms into the *meta-heuristic based method*. The most popular algorithms of this class include newton’s method [219, 220], gradient method [19] and ellipsoid method [214]. To apply and adapt iterative-based algorithms to optimize resource allocation requires strong mathematical background, which can be an obstruction for software developers to utilize these algorithms to optimize their deployment. Furthermore, the algorithms have the variety of performance for different problems in terms of efficiency and accuracy. As a result, more algorithms from iterative-based methods need to be studied and simplified by the system researchers, providing toolkits (or solvers) to tackle different optimization problems in IoT application deployment.

Meta-heuristic based method. The optimization problems in IoT applications can have large search spaces or be time-sensitive. The *meta-heuristic* based method is faster than *iterative-based method* in finding a near-optimal solution. This type of method consists of two subclasses: *trajectory*-based method and *population*-based method. The *trajectory*-based method finds a suitable solution with a trajectory defined in the search space. First, the resource allocation problems are mapped into a set of search problems such as variable neighborhood search, iterated local search, simulated annealing and tabu search. Then, the *meta-heuristic* algorithms are used to find the solutions. Many survey papers [137, 211, 310] have reviewed the algorithms applied for resource allocation in IoT, cloud computing, mobile computing. Additionally, *population*-based methods aim to find a suitable solution in the search space that is described as the evolution of a population of solutions. This method is also called evolutionary computation and the most well-known algorithm is the genetic algorithm. [379] investigates the resource allocation problems solved by evolutionary approaches in cloud computing.

Machine learning based method. ML based method is inspired by the ability of data to represent the performance and utilization of the contemporary systems. The ML based methods are used to build data-driven models that allows the target systems to learn and generate an optimized deployment plan. The proposed algorithms have been used to optimize various QoS parameters such as latency [222, 369], resource utilization [235], energy consumption [23] and many others. Zhang et al. [381] have given a comprehensive survey of the ML based methods used for resource allocation in mobile and wireless networking.

Deployment (or resource allocation) optimization problems have been studied for decades, and remain a huge legacy for overcoming the optimization problems in deploying ML-based IoT applications. Instead of developing new optimization algorithms, more efforts are required to model the complex optimization problems, in which the system scale, conditions and diversity have been amplified significantly.

3.3 Action and Model Composition

Deployment of ML models in a pipeline requires proper model composition to maximize the user QoS. As shown in §1.1, a smart car navigator system comprises multiple ML models, including speech recognition, text classification, text generation and text-to-speech (TTS) model.

Action composition is defined by composing a set of basic actions for complex decisions. In a self-driving car operating system, actions can be accelerating, braking, turning left and right, etc. The combination of various action spaces increases the difficulties of learning optimal decisions in such complex systems. Hierarchical abstract machines (HAM) [311] are well studied in the context of reinforcement learning [256, 331] by allowing agents to select from a constrained list of action spaces, speeding up the learning and adaptation to the new environment.

Model composition aims to create a ML-based IoT application by using reusable, portable, self-contained modules via inserting new components or removing existing components. Apache Airflow¹¹ is an open-source platform for creating, scheduling and monitoring workflows in Python. The Valohai¹² operator is an extension of Airflow that utilizes the docker container to build self-contained modules for each model while providing the flexibility for users to define the steps to execute. [173] reported the following challenges for chaining the ML models:

¹¹<https://airflow.apache.org/>

¹²<https://valohai.com/>

- How to allocate computing resources automatically for different models. An application is chained by various ML models require different computing resources across heterogeneous infrastructures. It is challenging to provision the computing resources efficiently for the chained ML models while meeting their performance requirements.
- How to chain the dependent models. Each individual ML model has its own specification and data format of the inputs and outputs. The challenge is to design a data messaging system to orchestrate the data flow across different models while considering their required specification and data format.
- How to meet the security requirements. The inferences are performed through various components, with each deployed across different computing resources. This introduces a set of challenges including privacy, verification of outputs of each model, changing the security policies of components, etc.
- How to monitor failure. The composed application consists of a set of ML models that needs to be monitored, ensuring that everything is streamlined and executed as anticipated.

Apart from challenges mentioned above, the literature discusses the techniques to improve the performance of individual models via system configuration, including *model batch size*, *model replica* and *system buffering*.

Per-Model Batch Size. Batching the received user queries optimizes throughput by fully utilizing the features of the pre-trained models, which is faster than processing one query at a time. However, batching query can potentially increase latency because the the model will wait for a whole batch of queries to come before it starts to proceed. The first query is not returned until the final query is processed [77]. The choice of the per-model batch is challenging due to the sequential composition between the models.

Model Replica. In heavy or bursting loads, a system must quickly respond to the query fluctuations to meet the latency requirements. To alleviate the system congestion and achieve high throughput, it is critical to identify the bottleneck, which can be challenging due to the system dynamics. The bottleneck models can be resolved by replicating the model instances across multiple devices [78], therefore balancing the workload. However, distributing the queries across more model replica in a parallel setting [78] is also challenging since the optimal placement depends on the model performance and the device capacity.

System Buffering. Serving system as a stream processing system comprises components across multiple devices. These devices usually process at different speeds, making system buffering across nodes necessary. Message queues are usually implemented to ensure smooth running within the system. However, buffering mechanism would increase the latency based on various system configurations [78]. It is thus challenging to design proper strategies to balance the message queue overhead and the system latency.

4 MODEL AUDIT

Audit aims to evaluate whether the application is operating effectively, safely and reliably with the collected evidence. To this end, we must know what we should audit. Most work focuses on monitoring or debugging the issues caused by infrastructure failures [184], implementation bugs [92, 182] and deployment errors [322]. In this section, we investigate the security, reliability and performance issues caused by ML models, especially DL models.

4.1 Security

There are many surveys regarding IoT security issues and challenges. The security of IoT standardized communication protocols were evaluated in [125] based on their proposed model. [306] categorized the security issues of IoT into *eight* domains including authentication, access control, confidentiality, privacy, trust, secure middleware, mobile security and

policy enforcement. [281] studied the main challenges and solutions of designing and deploying security mechanisms in centralized and distributed IoT architectures. [205] discussed the security features of IoT and categorized the attacks into *four* layers. i.e., the perception layer, the network layer, and the application layer.

In this subsection, we discuss security issues for deep learning based IoT applications: *Model exploratory attack*, *Data poisoning attacks* and *Evasion attacks*. *Model exploratory attacks* do not happen during training, instead the attacker tries to discover information from the trained model including the model itself and training data. *Data poisoning attacks* happen during the training phase, where the attacker attempts to shift the boundary of DL models in their favor by polluting the training data. Finally, *evasion attacks* maliciously craft the inputs for the deep learning based IoT application to trigger abnormal model behavior. Interestingly, the development of the research on adversarial learning has started an arms race between adversaries and defenders.

The following subsection summarizes the most popular attacks and defenses of these attacks. We also propose research directions for development of robust IoT applications.

4.1.1 Model exploratory attack. This type of attack is usually performed on open-source frameworks such as PredictionIO and cloud-based machine learning service. This ML-as-a-service may allow users to input partial feature vectors while still being able to receive confidence values in addition to prediction results. Thus, the attacker can leverage this feature to either extract the model or the sensitive information underlying the model. *Model stealing* and *Membership leakage* are *two* main types of model exploratory attack. *Model stealing* attack aims to duplicate the functionality of the model that allows the attacker to evade detection by the stolen model [12, 249]. [334] proposed a method that learns the target models via a prediction API. Evaluations show that this method successfully extracts the models including logistic regression, SVM, neural network and decision tree from BigML and Amazon Web Services. More attack methods can be created based on the extensive literature on learning theory, e.g., PAC learning [337] and its variants [21]. *Membership leakage* attacks are interested in stealing the information from the training data which may not be publicly available and may contain some sensitive information such as trade secrets, medical records etc. In this type of attack, an attacker is able to infer the members of the population or the members of the training dataset. Attacking the members of the population means that the types of data are used to create the model. Therefore, the target model has not been generalized for the adversary, because he/she has the sample of the entire population of the training dataset. The attacks were successfully performed in the dataset including voice, handwritten images, network traffic, online shopping, record of hospital stays etc [12, 304]. The members of the training dataset attack aim to identify the individuals whose data are used for training the model, which causes a serious privacy issue. For example, if an attacker knows that a patient's medical record was used to train a disease detection model, it also reveals that the patient has this disease. The experiments in [334] show that the attacks are able to extract the training dataset when the model is based on kernel logistic regression.

Defending model exploratory attack. The most straightforward defense of this type of attack is to constrain the API, not returning confidences and not responding to probing queries. The minimization can be achieved in three approaches: *Rounding confidence*, *Differential privacy* and *Ensemble methods*

Rounding confidence is a type of defense that rounds confidence sources of an application to some fixed precisions [108]. Notably, some online ML service providers are already working on it. For instance, BigML and Amazon provide five decimal places and 16 significant digits for their confidence scores respectively when answering queries. Limiting the precision can decrease the success rate of attacks, this is because the outputs of the model are approximated. In

equation-solving attacks, for example, if the output of an equation-system is rounded, it will increase the difficulty for attackers to guess the target function.

Differential privacy is a class of mechanisms to protect, especially the privacy of training data [343]. A set of *Differential privacy* methods has been applied to protect regressions [50, 384], SVMs [283], decision trees [162] and neural networks [303]. The main idea of *Differential privacy* is to avoid a query that allows an adversary to distinguish closely neighboring model parameters.

Ensemble methods returns an aggregated output predicted by a set of models. This prevention method was mentioned in [334] and may have more resilience against the model exploratory attacks, compared with other methods. Bonawitz [33] designed a new protocol to compute the sum of a large subset of models supporting a secure federated learning setting.

4.1.2 Data poisoning attack. Unlike *model exploratory attack*, an adversary performs the attacks during the model training phase. These attacks insert carefully constructed poison instances into the training dataset to manipulate the performance of a system. We introduce types of data poisoning attacks both in traditional machine learning and deep learning.

Data poisoning attack in machine learning. (1) supervised learning. A causative attack was proposed by Xiao et al. against SVMs which utilizes *label flipping* to poison the training data [363]. A *label flipping* attack attempts to add a noise label to the training data. These flipping labels are able to cause some malicious samples to be labeled as legal, or make legal samples appear to be malicious. To improve the efficiency of the attack, Biggio and Laskov [25] utilized the gradient descent algorithm to find the best attack points to flip the labels.

(2) unsupervised learning. The poisoning attack has been demonstrated against various clustering algorithms. The idea is to introduce carefully crafted data points to the training dataset to cause clusters to merge. In [26, 279], the authors assumed that the attacker has full knowledge of the clustering algorithm and then reduced the attack to an optimization problem. The evaluations show that the clustering algorithms are compromised significantly with a very small percentage of poisoned input data.

Data poisoning attack in deep learning. There are very few data poisoning attacks in neural networks. [315] showed that a deep learning model lost 11% accuracy after modifying 3% training data. Moreover, if the attacks are focus on attacking the specific test instances, the successful rate, time consumption and required resources (the number of modified samples) can be reduced significantly [58, 128, 315]. [296, 321] targeted real-world scenarios where the labels are examined by human reviewers and malware detectors. The authors aimed to overfit the deep learning models by poisoning the training data. Thus, the target instances (trained models) would not perform well during inference time.

Defending data poisoning attack. Compared with other systems, defending data poisoning attacks is more critical in machine learning systems, because training data coming from the outside world is very easy to poison. Steinhardt et al. [315] developed an outlier detector for linear classifiers to generate approximate upper bounds of the training dataset. The data (poisoned) that exceeds the bounds will be removed from the training dataset, not changing the distribution of the clean data. The paper [284] aimed to develop a strategy to defend against attacks that attempts to poison the PCA based anomaly detection models. Since the PCA methods are less sensitive to outliers, an ideal approach against the variance injection caused by the perturbed data is to filter the poisoned data with predefined threshold. The paper also proposed two approaches [284] for selecting the threshold: the first uses covariance matrix, and the second is to find the maximal scale estimate of the data projection.

4.1.3 *Evasion attack.* With the explosive development of machine learning, evasion attacks are becoming the most prevalent type of attack in machine learning, attracting people’s attention from both academia and industry. Fig. 10 shows the arms race between the attacker and defender. It shows that the attacker attempts to confuse the defender with a crafted adversarial example, while the defender aims to strengthen its ability to filter out illegitimate input.

During both training and inference, the attacker can generate *adversarial examples* by modifying the samples. The training phase modification is similar to *data poisoning attacks* in that the decision boundary of the defender classifier is modified by insertion, modification or deletion of the training dataset. There are two approaches for generating adversarial samples, *white-box* or *black-box*. In the *white-box* setup, the adversarial samples are crafted based on the attacker who has access to both the training data and the targeted model. Therefore, an adversary is able to obtain the boundaries of the targeted model by carefully modifying the training data. To be more explicit, as shown in Fig. 10, the defender aims to stop using the illegitimate input X to train the itself. In order to fool the defender, the attacker attempts to learn the boundaries of the defender by adding the perturbations to X and then performing the attack. This process is repeated until the adversarial samples break the boundaries. The most representative techniques [123, 186, 254] are based on the attacker who has knowledge of both target model and instance of data. In the *black-box* setup, the attack introduced in [253] is not aware of the training data and the targeted model. The only observation of the targeted model is the inputs and their labels given by the targeted model. Based on this, a local model is trained to replace the target DNN. 84% of the adversarial examples generated by the local substitute model are misclassified by the targeted DNN.




Fig. 10. Arms race game between the attacker and the defender. The attacker takes an illegitimate example (X) as the input of his/her neural network and generates an adversarial example (X'). This example attempts to fool the defender which is a classification neural network. If the classifier recognizes the adversarial example as Y which belongs to the legitimate input, it means the attacker wins the game. On the contrary, the defender wins the game if the adversarial example is classified as X . When the attacker fails, it will try to update its model to generate a stronger adversary example based on the feedback. Similarly, the defender will enhance its model based on the lesson learned from the successful attack.

Defending evasion attack. Crafting the adversarial samples is a complex optimization process as it is very hard to build a general tool for defending. Thus, *adversarial training* which adds the adversarial examples to the training set is a straightforward way to increase the model robustness [216, 297, 333]. Moreover, Defense-GAN [287] leverages a generative model to generate more samples similar to the training data, reducing the adversarial perturbation significantly. Apart from these general defending methods, some defenses use model hardening techniques. [366] proposed a feature squeezing method to reduce the complexity of representing the data. Less important adversarial perturbations are filtered out afterwards. In the following section, we survey the attacks and the defenses in real IoT applications and then discuss the challenges of building a secure IoT application.

4.1.4 *The system challenges of building a secure ML-based IoT application.* Most of the attacks and defenses reviewed in previous sections focus on developing the algorithms for *functional tasks* such as computer version, natural language

processing, audio speech processing etc. To build a secure IoT application, we must consider the security issues from the system perspective, as the *functional tasks* cannot perform well when the system is under attack. We discuss two system challenges to improve the security of ML-based IoT application.

Developing new attacking and defending models. ML has been widely used in IoT system developments including network engineering [202, 222], resource allocation [196, 368], system debugging [49, 308], network intrusion detection [178, 212, 345] and network operations[155]. These systems can be exposed to the aforementioned attacks as well. The literature [3, 104, 136] has revealed successful examples of attacks and defenses. Further to the efforts on *functional tasks*, more research and development (R&D) work is required to improve the security of ML-based IoT applications.

Developing new security platforms/frameworks. We have discussed the arms race game between the attacker and defender (see Fig. 10). This can be utilized to ensure the resilience of IoT applications to various attacks. At a high level, an ideal platform would be able to launch various attacks via a predefined deployment pipeline to attack the experimental group. Meanwhile, the attack behaviors and system performance will be monitored to reinforce the capacity of the defender. To this end, three research questions need to be answered. 1) *How to automate the attacks.* Unlike the traditional software deployment problem, deploying attacks is much more complicated. For example in an evasion attack, the proposed platform must be able to use various ML models to craft the adversarial examples. It is very difficult to automate this process. Due to the difference between the model inputs and outputs, the models may need to be retrained based on the observation of the real world to generate better adversarial examples. 2) *How to monitor the attacks.* As discussed in previous sections, attacks can happen in data collection, model training and model inference. Therefore, the traditional log system is not able to handle this complexity. In a data poisoning attack, for instance, the traditional log system is unable to capture the impact caused by fake data points injection into the system, thus the training of a defender is unfeasible. 3) *How to coordinate the attacker with the defender.* At the high level, the arms race game between the attacker and the defender is very logical. The challenge here is to continuously select the suitable attacks and thereby improve the defender's performance. This can be formalized as an optimization problem where one of the objective is to maximize the ability of a system in defending against certain types of attack.

4.2 Fault Tolerance

Distributed system fault tolerance has been studied for decades, many representative works have been proposed to handle the failures including system architecture [273] and algorithm design [48]. In the IoT environment, the probability of failure increases significantly, and many faults are very hard to detect. Our previous papers [115, 359] reviewed the state-of-the-art research, and then discussed the key research directions. In this subsection, we will introduce some of the most common faults in ML applications.

4.2.1 Faults in ML. Generalisation is crucial for ML models, which measure the prediction capacity on unseen test data [38]. Generally, ML training can be regarded as an optimisation process. For example, the model can be trained by minimising a certain loss function. However, overfitting may occur when ML models are trained on less representative, noisy or small data, and in this case, trivial error patterns may be learned, causing lack of generalisation (i.e., faults) at the test stage. There are many ways to reduce the overfitting effect, such as regularization (e.g., with regularization terms such as L1/L2 norm), *Stochastic gradient descent (SGD)* [35] or dropout [314] (for DL models), early stopping [22] (stop training when validation error starts to increase), etc.

Data imbalance, on the other hand, is very common in real-world scenarios and it may also cause overfitting. The model may mainly learn patterns from the majority classes while it may easily ignore the contributions from the

minority classes (with limited training samples), yielding severe faults at the inference stage. Various approaches have been proposed for mitigation including data augmentation (e.g., [346]), data upsampling (e.g., GAN-based data generation [329], [344]), cost-sensitive learning (which will impose a larger penalty on training errors with minority classes), transfer learning, etc.

In addition to overfitting effect, faults can also be attributed to the optimisation process. For example, with very deep DL models or with RNN, gradient vanishing/explosion may occur during the optimisation process, causing representation learning to be challenging or even infeasible. There are also several approaches to address this issue, e.g., *Batch Normalization (BN)* [160] (through normalizing the gradients in each layer), residual connection structure in DL (to preserve the gradient across many layers). For federated learning or distributed learning, RSA (i.e., Byzantine-Robust Stochastic Aggregation) [199] has also been proposed to prevent the incorrect gradient aggregation.

4.2.2 Fault tolerance in neural networks. At a high level of abstraction, the neural network can be viewed as a distributed system. Therefore, the failure can happen in *neuron* or *synapse*. In [233], Mhanmadi and Guerraoui proposed a general model to describe the fault model of neural networks. The *neuron* may stop computing (Crash) or generate some abnormal outputs (Byzantine). Similarly, the failures of *synapse* can be abstracted as Crash and Byzantine. Crash represents that the transmission has not succeeded, and Byzantine is that the incorrect messages are sent from the source *neuron* to the destination *neuron*. Thus, we assume that a given neural network \mathcal{N} performs an expected output $F_{\mathcal{N}}(X)$, and $F_{\mathcal{N}_{fault}}(X)$ is the output of the faulty network obtained from \mathcal{N} . The distance ϵ between $F_{\mathcal{N}}(X)$ and $F_{\mathcal{N}_{fault}}(X)$ represents the fault tolerance of \mathcal{N} , when there are at most n faulty components (including *neuron* and *synapse*):

$$\| F_{\mathcal{N}}(X) - F_{\mathcal{N}_{fault}}(X) \| \leq \epsilon \quad (1)$$

where X is the training dataset, applied to both \mathcal{N} and \mathcal{N}_{fault} . In order to guarantee the robustness of the neural model, the designer needs to ensure that the error (left hand in Equation 1) is below a predefined threshold (right hand in Equation 1). The threshold depends on the performance of the network and its intended application [233, 264].

Like traditional fault tolerance in distributed system, the fault tolerance in neural networks also has two types: *Passive* and *Active*.

In **passive** fault tolerance, no diagnostics, relearning, or reconfiguration is required thereby avoiding fault detection and location. The most common *passive* fault tolerance approach which is also one of the important features of neural network is inserting redundancy. Such methods learn a small network from the given input/output, and then add the replicated hidden neurons to share the load of the critical nodes, after the model has been trained. Representative works [63, 67, 103] addressed the fault tolerance by adding extra links or nodes to the well trained neural network. The authors in [62] proposed a solution that adds artificial faults to the network during the training time. Therefore, the network can tolerate the specific faults. However, this approach requires that the neural network designers are aware of all the faulty scenarios while building the network. Also, adding redundancies makes the models very complex and huge, which brings the challenges of deploying them over lightweight and low-power IoT devices.

Active fault tolerance aims to recover the neural model from faults by resetting the neural network into a fault-free state. However, it does not attract too much attention from research, a common strategy is to utilize high-performance computation resources to re-compute the lost work when the hardware fails [1, 354]. Notably, Qiao et al. proposed a checkpoint-based fault tolerance for deep learning in [266]. This new method partially recovers the model from the checkpoints based on the priority of the checkpoints thereby significantly reducing the cost of recomputing.

4.3 Performance Evaluation

In this section, we consider several performance criteria that need to be considered for evaluating the efficiency of the obtained ML models. The criteria is identified as *two* main dimensions: *model precision* and *execution latency*.

Model Precision. In a typical IoT application, the software performance is assumed stable after deployment. However, this is not the case for ML application where precision degradation is always expected after deployment. Precision degradation can happen as various unexpected external changes lead to shift in data distribution. Device location change, time and the weather are all important factors that may decrease the model performance. Therefore, it is critical that the model performance is monitored and new data is introduced continuously for retraining of the model. In ML, we define lifelong learning [255] as continually acquiring data and extracting new information without catastrophic forgetting of past knowledge. Lifelong learning keeps the model precision at a steady level.

Execution Latency. Many IoT applications are latency-sensitive depending on their tasks. For example, in the aforementioned smart transportation system (see §1.1) where sensors monitor and detect car accidents, instant decisions have to be made to warn the drivers of potential hazards. Various factors, listed below, have to be evaluated to ensure seamless communication among the distributed components of a smart IoT application.

Bandwidth Usage. In distributed IoT networks, large scale IoT sensors are generating a huge amount of data all the time. It is not possible to send all the data to the cloud for data analysis. Fog computing proposed to move the computing close to the sensors to reduce the data transmission over the IoT network. However, the bandwidth of sensor network and edge network are still limited, some nodes may experiences high latency due to the network congestion. This may cause huge latency for the whole system as well. We need to monitor and evaluate this network dynamic [221] in order to provide solutions to alleviate the congestion in the networks.

Resource Consumption. Hardware in IoT applications varies in computing power, memory and storage capacity. For any resource-intensive tasks, for example those computation-heavy or memory-heavy ones, resource exhaustion in one node may lead to unacceptable latency for the whole application. It is thus necessary to design efficient resource management systems [241, 372] to monitor and optimize task allocation for these physical devices,

System Throughput. The ML-based IoT applications may be developed to serve millions of people, for example, the *smart traffic routing application* mentioned in 1.1. This massive number of users may send the requests simultaneously. Responding to these requests quickly without losing user satisfaction is still an unsolved problem in cloud computing. However, this issue is amplified in ML-based IoT applications, in which the queries may be performed on various devices and models. Some database optimization techniques such as caching frequent queries, batching queries and approximate computing are applied [78, 259]. There are remaining gaps in optimizing the query plans by considering heterogeneousness of the computing resources, uncertainty of the network, and diversity of ML models.

5 DATA ACQUISITION

Apart from ML models' actions and operations, acquisition of data plays a crucial role in ML-based IoT applications. Data is one of the most important constituents in developing a ML model as the prediction accuracy of the model has a strong positive correlation with the quality of the input data [236]. To provide high quality data for a ML-based IoT application, we orchestrate the data acquisition process into several steps. In the whole data acquisition pipeline, we collect raw data from various data sources (§5.1). With proper preprocessing techniques (§5.2) to remove redundant information and annotate the data, we are capable of performing several different ML tasks. While we have more data sources during the development process, we can also fuse (§5.3) them to provide more consistent and useful information.

The following subsections will focus on the mentioned steps and discuss how data acquisition can support development of a robust ML-based IoT application.

5.1 Data Collection

The IoT data can be broadly categorized into *Structured data* and *Unstructured data* based on its representation. *Structured data* can be represented in a pre-defined format (rows and columns). The meaning of each field is explicit which eases the analysis and storage of the data. Examples of structured data include employee register information, visiting logs, etc. On the other hand, *Unstructured data* lacks any specific structure or format. Varying from text, audio, video to mails and messages, it accounts for a large proportion of IoT data. These two types of data are generated in *three* formats: *signal data*, *log data* and *packet data*. The *signal data* collects the daily life signal through various hardware such as sensors, sound recorders, cctv cameras, etc. The *log data* is usually used to capture the system status. Finally, the *packet data* is the data sent over the network and each unit transmitted consists of a header and the actual data. To collect these data, *three* important factors need to be considered: 1) *Data exchange*, 2) *Resource consumption* and 3) *Concept Drift*.

Data exchange. The data is generated from an IoT device and sent to an edge (sink) node or other IoT devices and eventually the data will be collected and stored in the cloud. The computation power of gateways and edge nodes is improving, which brings an opportunity to remove data redundancy while saving the energy and bandwidth required for transferring data to downstream nodes [357]. This aggregation requires application of various data summarization techniques [71] including sampling, sketching, histograms, wavelets and adaptation of these techniques to meet the constraints of the hardware and the time-varying channel conditions. Henriette et al. [280] investigated the state-of-the-art stream processing systems that can be used to implement these data summarization techniques and execute them in a parallel and elastic manner. However, it still requires a lot of effort to develop new data summarization techniques and stream processing systems to handle the difficulty of processing high volumes data from various sources with multi-modality.

Resource consumption. As mentioned earlier, IoT devices are very limited by resources such as processing capability, storage capacity, wireless bandwidth and battery power. Thus, it is very critical to optimize the resource utilization while processing, storing or transferring data to the edge device or cloud. To this end, we need to consider *three* issues: *resource allocation*, *energy control* and *task allocation*. *Resource allocation* in the context of data collection is to assign computing, storage or bandwidth resources to the data generated by IoT devices before transferring to edge or cloud. Sending streaming data drains the battery at a faster rate while limited storage capacity does not enable large data storage. *Energy control* focuses on optimizing the energy consumption when the IoT data is processed and transferred over the devices. *Task allocation* aims to balance the resources consumption in IoT devices while minimizing the overall latency. These *three* factors are sometimes considered together and most of the available algorithms are based on market-enabled pricing schemes, which dynamically exchange the resources among the devices in IoT infrastructure by creating an artificial market [105, 165]. In ML-based IoT application, the ML models should be considered as the special tasks that are running on extremely heterogeneous computing resources in a distributed manner, and these tasks are usually compute-intensive, data-intensive and network-intensive. As a result, it is crucial to develop new market models to describe these special resource consumption problems and new algorithms to solve the problems.

Concept Drift. Due to the dynamicity of the IoT environment, data distribution becomes very uncertain and changes frequently over time leading to *concept drift* [114]. Changes can occur abruptly or gradually correlated with the occurrence of other events. Additionally, the change can be in different forms i.e., input data characteristics or relation

between input data and target variables with single or multiple occurrence (constant or variable recurrence). For the successful execution of IoT applications, these drifts need to be predicted, distinguished from noise and handled properly. Numerous algorithms are proposed for managing concept drift. [82, 154] review the generic algorithms to handle the concept drift. There are two main detection methods, *performance-based* and *data distribution-based*. The former can work well if the data is labeled which may not be possible for all cases while the latter is able to detect only a subset of available drifts. Since IoT-based ML application data are not always labeled and high accuracy is desired, it is essential to develop new algorithms which are able to detect and manage the concept drift.

5.2 Data Preprocessing

The real-world data collected from heterogeneous IoT devices usually contains outliers or is incomplete in nature, which makes it difficult to feed it into ML models directly. Data preprocessing deals with these anomalies and improves the data quality and practicality. There are several things that need to be considered, namely *data cleaning*, *data annotation* and *feature engineering*. We have discussed the details of *feature engineering* in §2.3.1, will not consider it in this section.

5.2.1 Data cleaning. Much data contains noise that is bound to confuse the ML models and reduce the accuracy of the prediction results. Data cleaning resolves this problem by completing several routine tasks such as *filling missing values*, *smoothing noise data* and *removing outliers* [6, 29]. Empty records in the data set can be replaced manually by a specific value, for example the attribute mean or the most common attribute in the set. It can also be marked with "unknown" or just ignored if the dataset is large enough. Noisy data, though, can be smoothed by grouping first and then averaging over each group. Data outliers can also be detected during this process if the value exceeds a predefined threshold. There are other common practices such as data normalization [267], which is used to scale all dimensions of data to a specific range. This is a very efficient method for the case where there is high variation for different dimensions of the data.

5.2.2 Data Annotation. As discussed in §2.1, data annotation is necessary for supervised learning based ML models, in which both the data and the corresponding target act as the input sample. The model is trained with the labeled data which is used to predict the target for new unseen data. This is usually costly and complex due to the requirement for a large volume of labeled data needed for the training. The following investigates different annotation methods that can be applied according to the size of the data to be annotated and the cost of annotation per data.

Manual Annotation. At the initial stage of a ML project, quick prototyping of a workable model requires only few labeled data. In this sense, the developers can manually annotate the collected data to create a small dataset. This is usually done by reviewing the data samples and attaching labels following the annotation guidelines. Manual annotation by the engineers is quick and precise without any professional training, and the data quality is usually great. the problem with this approach is the lack of scalability.

Crowdsourcing Annotation. Crowdsourcing annotation is a scalable and cost-effective method. It is usually orchestrated by an online platform that provides access to a workforce of people to complete the annotation tasks. Famous crowdsourcing platforms include Amazon Mechanical Turk (MTurk). Compared to manual annotation, this approach can be scaled to large dataset labeling. However, the crowdsourcing method requires delicate design on quality control mechanisms to ensure the annotation quality, and the incentives or rewards for the crowds.

Active Learning. Active learning [118, 295] aims to design a system capable of choosing and learning from less training data while still achieving the same or even higher accuracy. An active learning system consists of two components: a

learning module that trains a model with the current training sample and a *sample selection module* that selects the most informative samples from the unlabeled samples. The selected samples will then be annotated manually and added to the training set. The iterative process continues until the training converges. The key here is the sample selection module which can be approximately subdivided into *five* categories, *risk reduction*, *uncertainty*, *diversity*, *density* and *relevance* according to the selection criteria [347]. These criteria can be used either single-handedly (e.g., risk reduction [107], uncertainty [168], relevance [13]) or in a combination. In [348], *uncertainty*, *diversity*, *density* and *relevance* are combined for multi-modality video annotation. Similarly, work in [148] combines *uncertainty*, *diversity* and *density* metrics and the evaluation proves the combination performs well on medical image classification tasks.

All the above explained methods work well for the case of static machine learning scenarios with batch data available beforehand. However, this may not be suitable for IoT-specific streaming data imminent with high concept drifts. In this case, the model needs to learn continuously with the upcoming data. Since the new data does not have any label, multiple delayed learning concepts [121, 127, 262] are proposed to handle the non-negligible delay in data labeling. These methods are adequate for the scenario where labeling takes a constant time and latency is not a determining factor. For IoT data with variable constant drifts, cleaning and labeling may not take uniform time. Also, latency is one of the deciding factors for IoT-based ML applications. Thus, new sets of methods are essential for data labeling which considers the fluctuating IoT data with minimum possible delay.

5.3 Data Fusion

Data fusion aims to combine the data from multiple sources to provide more accurate and useful information. It offers numerous advantages for ML-based application by enhancing the data quality (finding the missing values), detecting any anomalies, conducting the prediction and finding any correlations among the distributed dataset [27, 189, 370]. However, there are multiple challenges in combining heterogeneous IoT data [5] such as data frequency, data imperfection, data correlation, data alignment and dynamic iterative process. To handle these challenges effectively, numerous data fusion methods are available in the literature. It is mainly categorized into *three* groups as given below.

5.3.1 Probabilistic Data Fusion Algorithms. This group consists of the algorithms that use density function or probability distributions as a core method for data fusion. It includes Bayesian techniques [28], Markov models [185], evidential reasoning [367] and other methods. These methods are simple and widely used in different applications to express the co-relation and dependency between numerous datasets. However, there are certain drawbacks with probabilistic data fusion methods highlighted in [5]. First, it can not scale with the size and modality of the data. Second, uncertain and noisy data can not be handled properly. Finally, prior probabilities and density functions are difficult to obtain.

5.3.2 Knowledge-based Data Fusion Algorithms. To overcome the uncertainty of data and to increase the accuracy of fusion methods, knowledge-based data fusion methods are proposed. This method accumulates knowledge from the imprecise big data and apply over the fusion process. Different aggregation techniques and ML methods are used for the data fusion process. For example, [24, 231] (supervised learning method) and [112, 388] (unsupervised learning method) are used to discover the distribution of the complex datasets. However, the complexity of this type of method is higher than the probabilistic methods. This class of method, thus, may consume more computing resources and cost more time to process.

5.3.3 Evidence-based Data Fusion Algorithms. This group of methods is based on Demster-Shafer Theorem (DST) and recursive operations. As compared to probabilistic methods, where there are only two states (happening or non-happening) of an event, DST includes an unknown state to capture real-world uncertainty. [163, 286] are the applications of DST for data fusion. However, increasing the data evidence also increases the complexity of this method. Therefore, this method is not suitable for the applications running on less powerful computing resources.

5.4 Discussion

In this section, we reviewed core components in the data acquisition process and discussed how they can contribute to generation of high quality, ready-to-use data for IoT-ML application. There are *multiple* research directions which can be considered to leverage others' efforts, thereby improving the performance of training a ML model. First, *Data reuse*, with the scaling of the data volume, past data is stored and usually ignored after use. However, it can be reused and mined for more values. For example it can be used for boosting semi-supervised data annotation [396], or it can be integrated with newly collected data for model training. Second, *Data re-organization*, there exist datasets for different tasks in similar areas. They may not be the same, but can be re-organized to extract the common distributions. Proper identification and extraction can be explored to save effort on data collection. *Feature evolution* is also an important trait of streaming data as a feature may appear and disappear over time. If a feature appears and is found to be relevant, it is necessary to incorporate that for the learning process. In this case, disappearance of a feature can be considered as a drift and the unavailability is treated as missing values. Ignoring this feature may lead to inaccurate prediction. Taking the relevancy of feature evolution for different problem domains. Other challenges that related to the unbalanced data have been discussed in §4.2.1 as well.

6 CONCLUSION

Growing numbers of internet-connected things (IoT) produce vast amounts of data, build applications and provide various services in domains such as smart cities, energy, mobility, and smart transportation. ML is becoming a preliminary technique for analyzing IoT data. It produces high-level abstraction and insight that is fed to the IoT systems for fine-tuning and improvement of the services. In this survey, we reviewed the characteristics of the IoT development lifecycle and the role of ML for individual steps. Specifically, we divided the development lifecycles into different modules and presented a novel taxonomy to characterize and analyze various techniques used to build an ML-based IoT application. In summary, this survey seeks to provide systematic and insightful information for researchers. It assists the development of future orchestration solutions by providing a holistic view on the current status of ML-based IoT application development, deriving key open research issues that were identified based on our critical review.

REFERENCES

- [1] M. Abadi, P. Barham, et al. 2016. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on OSDI*. 265–283.
- [2] A. Abdiansah and R. Wardoyo. 2015. Time complexity analysis of support vector machines (SVM) in LibSVM. *International journal computer and application* (2015).
- [3] S. Achleitner, T. La Porta, et al. 2017. Adversarial network forensics in software defined networking. In *Proceedings of the Symposium on SDN Research*. ACM, 8–20.
- [4] A. F. Aji and K. Heafield. 2017. Sparse Communication for Distributed Gradient Descent. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. 440–445.
- [5] F. Alam, R. Mehmood, et al. 2017. Data fusion and IoT for smart ubiquitous environments: A survey. *IEEE Access* 5 (2017), 9533–9554.
- [6] S. A. Alasadi and W. S. Bhaya. 2017. Review of data preprocessing techniques in data mining. *Journal of Engineering and Applied Sciences* 12, 16 (2017), 4102–4107.
- [7] D. Alistarh, D. Grubic, et al. 2017. QSGD: Communication-efficient SGD via gradient quantization and encoding. In *Advances in Neural Information Processing Systems*. 1709–1720.
- [8] Y. Alsouda, S. Pillana, et al. 2019. IoT-based Urban Noise Identification Using Machine Learning: Performance of SVM, KNN, Bagging, and Random Forest. In *Proceedings of the International Conference on Omni-Layer Intelligent Systems*. ACM, 62–67.

- [9] F. A. Aoudia, M. Gautier, et al. 2018. RLMan: an energy manager based on reinforcement learning for energy harvesting wireless sensor networks. *IEEE Transactions on Green Communications and Networking* 2, 2 (2018), 408–417.
- [10] J. Appleyard, T. Kocisky, et al. 2016. Optimizing performance of recurrent neural networks on gpus. *arXiv preprint arXiv:1604.01946* (2016).
- [11] K. Arulkumaran, M. P. Deisenroth, et al. 2017. A brief survey of deep reinforcement learning. *arXiv preprint arXiv:1708.05866* (2017).
- [12] G. Ateniese, G. Felici, et al. 2013. Hacking smart machines with smarter ones: How to extract meaningful data from machine learning classifiers. *arXiv* (2013).
- [13] S. Ayache and G. Quénot. 2007. Evaluation of active learning strategies for video indexing. *Signal Processing: Image Communication* 22, 7-8 (2007), 692–704.
- [14] J. Ba, R. Grosse, et al. 2016. Distributed second-order optimization using Kronecker-factored approximations. (2016).
- [15] B. Baker, O. Gupta, et al. 2017. Accelerating neural architecture search using performance prediction. *arXiv preprint arXiv:1705.10823* (2017).
- [16] A. K. Balan, V. Rathod, et al. 2015. Bayesian dark knowledge. In *Advances in Neural Information Processing Systems*. 3438–3446.
- [17] P. Baldi. 2012. Autoencoders, unsupervised learning, and deep architectures. In *Proceedings of ICML workshop on unsupervised and transfer learning*. 37–49.
- [18] R. Battiti. 1994. Using mutual information for selecting features in supervised neural net learning. *IEEE Transactions on neural networks* 5, 4 (1994), 537–550.
- [19] A. Beck, A. Nedić, et al. 2014. An $o(1/k)$ gradient method for network resource allocation problems. *IEEE Transactions on Control of Network Systems* 1, 1 (2014), 64–73.
- [20] T. Ben-Nun, E. Levy, et al. 2015. Memory access patterns: the missing piece of the multi-GPU puzzle. In *SC'15: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 1–12.
- [21] G. M. Benedek and A. Itai. 1991. Learnability with respect to fixed distributions. *Theoretical Computer Science* 86, 2 (1991), 377–389.
- [22] Y. Bengio. 2012. Practical recommendations for gradient-based training of deep architectures. In *Neural networks: Tricks of the trade*. Springer, 437–478.
- [23] J. L. Berral, Í. Goiri, et al. 2010. Towards energy-aware scheduling in data centers using machine learning. In *Proceedings of the 1st International Conference on energy-Efficient Computing and Networking*. ACM, 215–224.
- [24] B. Bigdeli, F. Samadzadegan, et al. 2015. Fusion of hyperspectral and LIDAR data using decision template-based fuzzy multiple classifier system. *International Journal of Applied Earth Observation and Geoinformation* 38 (2015), 309–320.
- [25] B. Biggio, B. Nelson, et al. 2012. Poisoning attacks against support vector machines. *arXiv preprint arXiv:1206.6389* (2012).
- [26] B. Biggio, K. Rieck, et al. 2014. Poisoning behavioral malware clustering. In *Proceedings of the 2014 workshop on artificial intelligent and security workshop*. ACM, 27–36.
- [27] F. H. Bijarbooneh, W. Du, et al. 2016. Cloud-assisted data fusion and sensor selection for internet of things. *IEEE Internet of Things Journal* 3, 3 (2016), 257–268.
- [28] T. A. Biresaw, A. Cavallaro, et al. 2015. Tracker-level fusion for robust Bayesian visual tracking. *IEEE Transactions on Circuits and Systems for Video Technology* 25, 5 (2015).
- [29] A. Blázquez-García, A. Conde, et al. 2020. A review on outlier/anomaly detection in time series data. *arXiv preprint arXiv:2002.04236* (2020).
- [30] J. Boedecker, J. T. Springenberg, et al. 2014. Approximate real-time optimal control based on sparse gaussian process models. In *2014 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*. IEEE, 1–8.
- [31] M. Boehm, M. W. Dusenberry, et al. 2016. Systemml: Declarative machine learning on spark. *Proceedings of the VLDB Endowment* 9, 13 (2016), 1425–1436.
- [32] K. Bonawitz, H. Eichner, et al. 2019. Towards federated learning at scale: System design. *arXiv preprint arXiv:1902.01046* (2019).
- [33] K. Bonawitz, V. Ivanov, et al. 2017. Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1175–1191.
- [34] K. Bonawitz, F. Salehi, et al. 2019. Federated learning with autotuned communication-efficient secure aggregation. *arXiv preprint arXiv:1912.00131* (2019).
- [35] L. Bottou. 2010. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*. Springer, 177–186.
- [36] L. Bottou. 2012. Stochastic gradient descent tricks. In *Neural networks: Tricks of the trade*. Springer, 421–436.
- [37] L. Bottou, F. E. Curtis, et al. 2018. Optimization methods for large-scale machine learning. *Siam Review* 60, 2 (2018), 223–311.
- [38] O. Bousquet, S. Boucheron, et al. 2003. Introduction to statistical learning theory. In *Summer School on Machine Learning*. Springer, 169–207.
- [39] J. K. Bradley, A. Kyröla, et al. 2011. Parallel coordinate descent for l1-regularized loss minimization. *arXiv preprint arXiv:1105.5379* (2011).
- [40] D. Brauckhoff, K. Salamati, et al. 2009. Applying PCA for traffic anomaly detection: Problems and solutions. In *IEEE INFOCOM 2009*. IEEE, 2866–2870.
- [41] L. Breiman. 2001. Random forests. *Machine learning* 45, 1 (2001), 5–32.
- [42] L. Breiman. 2017. *Classification and regression trees*. Routledge.
- [43] U. Breitenbücher, K. Kepes, et al. 2017. Declarative vs. Imperative: How to Model the Automated Deployment of IoT Applications? *Proceedings of the 11th Advanced Summer School on Service Oriented Computing* (2017), 18–27.
- [44] A. Brock, T. Lim, et al. 2017. SMASH: one-shot model architecture search through hypernetworks. *arXiv preprint arXiv:1708.05344* (2017).
- [45] C. Bucilua, R. Caruana, et al. 2006. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM.
- [46] R. H. Byrd, S. L. Hansen, et al. 2016. A stochastic quasi-Newton method for large-scale optimization. *SIAM Journal on Optimization* 26, 2 (2016), 1008–1031.
- [47] B. Cabé, E. I. W. Group, et al. 2018. IoT Developer Survey 2018. *SlideShare, April 13* (2018).
- [48] M. Castro, B. Liskov, et al. 1999. Practical Byzantine fault tolerance. In *OSDI*, Vol. 99. 173–186.
- [49] Z. B. Celik, L. Babun, et al. 2018. Sensitive information tracking in commodity IoT. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*. 1687–1704.
- [50] K. Chaudhuri and C. Monteleoni. 2009. Privacy-preserving logistic regression. In *Advances in neural information processing systems*. 289–296.
- [51] N. V. Chawla, N. Japkowicz, et al. 2004. Special issue on learning from imbalanced data sets. *ACM Sigkdd Explorations Newsletter* 6, 1 (2004), 1–6.
- [52] K. Chellapilla, S. Puri, et al. 2006. High performance convolutional neural networks for document processing. In *10th International Workshop on Frontiers in Handwriting Recognition*. Suvisoft.
- [53] C.-C. Chen, C.-L. Yang, et al. 2018. Efficient and robust parallel dnn training through model parallelism on multi-gpu platform. *arXiv preprint arXiv:1809.02839* (2018).
- [54] J. Chen, X. Pan, et al. 2016. Revisiting distributed synchronous SGD. *arXiv preprint arXiv:1604.00981* (2016).
- [55] L. Chen, J. Lu, et al. 2018. Part-activated deep reinforcement learning for action prediction. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 421–436.
- [56] L. Chen, Y. Zhang, et al. 2018. Detection of DNS DDoS Attacks with Random Forest Algorithm on Spark. *Procedia computer science* 134 (2018), 310–315.
- [57] T. Chen, T. Moreau, et al. 2018. {TVM}: An Automated End-to-End Optimizing Compiler for Deep Learning. In *13th USENIX Symposium on OSDI*. 578–594.
- [58] X. Chen, C. Liu, et al. 2017. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526* (2017).
- [59] Y. Cheng, D. Wang, et al. 2017. A survey of model compression and acceleration for deep neural networks. *arXiv preprint arXiv:1710.09282* (2017).
- [60] S. Chetlur, C. Woolley, et al. 2014. cudnn: Efficient primitives for deep learning. *arXiv preprint arXiv:1410.0759* (2014).
- [61] T. Chilimbi, Y. Suzue, et al. 2014. Project adam: Building an efficient and scalable deep learning training system. In *11th USENIX Symposium on OSDI*. 571–582.
- [62] C.-T. Chin, K. Mehrotra, et al. 1994. Training techniques to obtain fault-tolerant neural networks. In *Proceedings of IEEE 24th International Symposium on Fault-Tolerant Computing*. IEEE, 360–369.

- [63] C.-T. Chiu, K. Mehrotra, et al. 1993. Robustness of feedforward neural networks. In *Neural Networks, 1993., IEEE International Conference on*. IEEE, 783–788.
- [64] F. Chollet. 2017. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1251–1258.
- [65] P. Chopra and S. K. Yadav. 2015. Fault detection and classification by unsupervised feature extraction and dimensionality reduction. *Complex & Intelligent Systems* (2015).
- [66] C.-T. Chu, S. K. Kim, et al. 2007. Map-reduce for machine learning on multicore. In *Advances in neural information processing systems*. 281–288.
- [67] L.-C. Chu and B. W. Wah. 1990. Fault tolerant neural networks with hybrid redundancy. In *Neural Networks, 1990., 1990 IJCNN International Joint Conference on*. IEEE.
- [68] K. Chua, R. Calandra, et al. 2018. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems*. 4754–4765.
- [69] T. Cohen and M. Welling. 2016. Group equivariant convolutional networks. In *International conference on machine learning*. 2990–2999.
- [70] I. Colin, L. Dos Santos, et al. 2019. Theoretical Limits of Pipeline Parallel Optimization and Application to Distributed Deep Learning. In *Advances in Neural Information Processing Systems*. 12350–12359.
- [71] G. Cormode, M. Garofalakis, et al. 2011. Synopses for massive data: Samples, histograms, wavelets, sketches. *Foundations and Trends® in Databases* 4, 1–3 (2011), 1–294.
- [72] N. Corporation. 2015. NVIDIA Collective Communications Library (NCCL). Retrieved March 2, 2020 from <https://developer.nvidia.com/nccl>
- [73] C. Cortes and V. Vapnik. 1995. Support-vector networks. *Machine learning* 20, 3 (1995), 273–297.
- [74] M. Courbariaux, Y. Bengio, et al. 2015. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Advances in NIPS*.
- [75] M. Courbariaux, I. Hubara, et al. 2016. Binarized neural networks: Training deep neural networks with weights and activations constrained to+ 1 or-1. *arXiv* (2016).
- [76] T. M. Cover, P. Hart, et al. 1967. Nearest neighbor pattern classification. (1967).
- [77] D. Crankshaw, G.-E. Sela, et al. 2018. InferLine: ML Inference Pipeline Composition Framework. *arXiv preprint arXiv:1812.01776* (2018).
- [78] D. Crankshaw, X. Wang, et al. 2017. Clipper: A low-latency online prediction serving system. In *14th USENIX Symposium on NSDI*. 613–627.
- [79] H. Cui, H. Zhang, et al. 2016. Geeps: Scalable deep learning on distributed gpus with a gpu-specialized parameter server. In *Proceedings of the Eleventh European Conference on Computer Systems*. 1–16.
- [80] A. C. F. da Silva, U. Breitenbücher, et al. 2017. Internet of Things Out of the Box: Using TOSCA for Automating the Deployment of IoT Environments.. In *CLOSER*. 330–339.
- [81] M. Dash and H. Liu. 2003. Consistency-based search in feature selection. *Artificial intelligence* 151, 1-2 (2003), 155–176.
- [82] R. S. M. de Barros and S. G. T. de Carvalho Santos. 2019. An overview and comprehensive comparison of ensembles for concept drift. *Information Fusion* 52 (2019), 213–244.
- [83] C. M. De Sa, C. Zhang, et al. 2015. Taming the wild: A unified analysis of hogwild-style algorithms. In *Advances in neural information processing systems*. 2674–2682.
- [84] J. Dean, G. Corrado, et al. 2012. Large scale distributed deep networks. In *Advances in neural information processing systems*. 1223–1231.
- [85] M. Deisenroth and C. E. Rasmussen. 2011. PILCO: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*. 465–472.
- [86] O. Dekel, R. Gilad-Bachrach, et al. 2012. Optimal distributed online prediction using mini-batches. *Journal of Machine Learning Research* 13, Jan (2012), 165–202.
- [87] J. Deng, W. Dong, et al. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 248–255.
- [88] S. Devi and T. Neetha. 2017. Machine Learning based traffic congestion prediction in a IoT based Smart City. (2017).
- [89] G. Diamos, S. Sengupta, et al. 2016. Persistent mns: Stashing recurrent weights on-chip. In *International Conference on Machine Learning*. 2024–2033.
- [90] M. Ding and H. Tian. 2016. PCA-based network traffic anomaly detection. *Tsinghua Science and Technology* 21, 5 (2016), 500–509.
- [91] J.-D. Dong, A.-C. Cheng, et al. 2018. Dpp-net: Device-aware progressive search for pareto-optimal neural architectures. In *Proceedings of the ECCV*. 517–531.
- [92] W. Dong, C. Chen, et al. 2013. D2: Anomaly detection and diagnosis in networked embedded systems by program profiling and symptom mining. In *RTSS. IEEE*, 202–211.
- [93] N. Dryden, N. Maruyama, et al. 2019. Improving strong-scaling of CNN training by exploiting finer-grained parallelism. In *2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 210–220.
- [94] N. Dryden, N. Maruyama, et al. 2019. Channel and filter parallelism for large-scale CNN training. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–20.
- [95] N. Dryden, T. Moon, et al. 2016. Communication quantization for data-parallel training of deep neural networks. In *2016 2nd Workshop on Machine Learning in HPC Environments (MLHPC)*. IEEE, 1–8.
- [96] W. Du and Z. Zhan. 2002. Building decision tree classifier on private data. In *Proceedings of the IEEE international conference on Privacy, security and data mining-Volume 14*. Australian Computer Society, Inc., 1–8.
- [97] J. Duchi, E. Hazan, et al. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research* 12, Jul (2011), 2121–2159.
- [98] S. Egea, A. R. Mañez, et al. 2017. Intelligent IoT traffic classification using novel search strategy for fast-based-correlation feature selection in industrial environments. *IEEE Internet of Things Journal* 5, 3 (2017), 1616–1624.
- [99] A. Elk. 2019. Distributed Machine Learning Toolkit: BIG DATA, BIG MODEL, FLEXIBILITY, EFFICIENCY. Retrieved March 7, 2019 from <http://www.dmtk.io>
- [100] Y. Elovici, A. Shabtai, et al. 2007. Applying machine learning techniques for detection of malicious code in network traffic. In *Annual Conference on Artificial Intelligence*. Springer, 44–50.
- [101] M. F. Elrawy, A. I. Awad, et al. 2018. Intrusion detection systems for IoT-based smart environments: a survey. *Journal of Cloud Computing* 7, 1 (2018), 21.
- [102] T. Elsken, J. H. Metzen, et al. 2018. Multi-objective architecture search for cnns. *arXiv preprint arXiv:1804.09081* 2 (2018).
- [103] M. D. Emmerson and R. I. Damper. 1993. Determining and improving the fault tolerance of multilayer perceptrons in a pattern-recognition application. *IEEE transactions on neural networks* 4, 5 (1993), 788–793.
- [104] T. Erpek, Y. E. Sagduyu, et al. 2018. Deep learning for launching and mitigating wireless jamming attacks. *TCCN* (2018).
- [105] S. Eswaran, A. Misra, et al. 2012. Utility-based bandwidth adaptation in mission-oriented wireless sensor networks. *TOSN* (2012).
- [106] M. Everingham, L. Van Gool, et al. 2010. The Pascal Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision* 88, 2 (June 2010), 303–338.
- [107] H. I. Fawaz, G. Forestier, et al. 2019. Deep learning for time series classification: a review. *Data Mining and Knowledge Discovery* 33, 4 (2019), 917–963.
- [108] M. Fredrikson, S. Jha, et al. 2015. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1322–1333.
- [109] G. Frewat, C. Baroud, et al. 2016. Android voice recognition application with multi speaker feature. In *2016 18th MELECON*. IEEE, 1–5.
- [110] N. Friedman, D. Geiger, et al. 1997. Bayesian network classifiers. *Machine learning* 29, 2-3 (1997), 131–163.
- [111] S. Fujimoto, H. van Hoof, et al. 2018. Addressing function approximation error in actor-critic methods. *arXiv preprint arXiv:1802.09477* (2018).
- [112] C. E. Fuss, A. A. Berg, et al. 2016. DEM Fusion using a modified k-means clustering algorithm. *International journal of digital earth* 9, 12 (2016), 1242–1255.
- [113] Y. Gal, R. McAllister, et al. 2016. Improving PILCO with Bayesian neural network dynamics models. In *Data-Efficient Machine Learning workshop, ICML, Vol. 4*.

- [114] J. a. Gama, I. Žliobaitundefined, et al. 2014. A Survey on Concept Drift Adaptation. *ACM Comput. Surv.* 46, 4, Article Article 44 (March 2014), 37 pages.
- [115] P. Garraghan, R. Yang, et al. 2018. Emergent failures: Rethinking cloud reliability at scale. *IEEE Cloud Computing* 5, 5 (2018), 12–21.
- [116] A. L. Gaunt, M. A. Johnson, et al. 2017. AMPNet: Asynchronous model-parallel training for dynamic neural networks. *arXiv preprint arXiv:1705.09786* (2017).
- [117] A. Gholami, A. Azad, et al. 2018. Integrated model, batch, and domain parallelism in training neural networks. In *Proceedings of the 30th on Symposium on Parallelism in Algorithms and Architectures*. 77–86.
- [118] R. Gilyazev and D. Y. Turdakov. 2018. Active Learning and Crowdsourcing: A Survey of Optimization Methods for Data Labeling. *Programming and Computer Software* 44, 6 (2018), 476–491.
- [119] D. E. Goldberg. 1989. *Genetic Algorithms in Search, Optimization and Machine Learning* (1st ed.). Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA.
- [120] M. Goldstein and S. Uchida. 2016. A comparative evaluation of unsupervised anomaly detection algorithms for multivariate data. *PLoS one* 11, 4 (2016), e0152173.
- [121] H. M. Gomes, A. Bifet, et al. 2017. Adaptive random forests for evolving data stream classification. *Machine Learning* 106, 9-10 (2017), 1469–1495.
- [122] I. Goodfellow, J. Pouget-Abadie, et al. 2014. Generative adversarial nets. In *Advances in neural information processing systems*. 2672–2680.
- [123] I. J. Goodfellow, J. Shlens, et al. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014).
- [124] P. Goyal, P. Dollár, et al. 2017. Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677* (2017).
- [125] J. Granjal, E. Monteiro, et al. 2010. A secure interconnection model for IPv6 enabled wireless sensor networks. In *2010 IFIP Wireless Days*. IEEE, 1–6.
- [126] A. Graves, A. rahman Mohamed, et al. 2013. Speech recognition with deep recurrent neural networks. *ICASSP* (2013), 6645–6649.
- [127] M. Grzenda, H. M. Gomes, et al. 2019. Delayed labelling evaluation for data streams. *Data Mining and Knowledge Discovery* (2019), 1–30.
- [128] T. Gu, B. Dolan-Gavitt, et al. 2019. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733* (2017).
- [129] L. Guan, W. Yin, et al. 2019. XPipe: Efficient Pipeline Model Parallelism for Multi-GPU DNN Training. *arXiv preprint arXiv:1911.04610* (2019).
- [130] Y. Guan and T. Plötz. 2017. Ensembles of deep lstm learners for activity recognition using wearables. *IMWUT* 1, 2 (2017), 11.
- [131] S. Gupta, A. Agrawal, et al. 2015. Deep learning with limited numerical precision. In *International Conference on Machine Learning*. 1737–1746.
- [132] R. Gutierrez-Osuna. 2002. Pattern analysis for machine olfaction: a review. *IEEE Sensors journal* (2002).
- [133] T. Haarnoja, A. Zhou, et al. 2018. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905* (2018).
- [134] H. HaddadPajouh, A. Dehghantanha, et al. 2018. A deep Recurrent Neural Network based approach for Internet of Things malware threat hunting. *Future Generation Computer Systems* 85 (2018), 88–96.
- [135] M. A. Hall. 2000. Correlation-based feature selection of discrete and numeric class machine learning. (2000).
- [136] Y. Han, D. Hubczenko, et al. 2019. Adversarial Reinforcement Learning under Partial Observability in Software-Defined Networking. *arXiv:1902.09062* (2019).
- [137] P. Hansen, N. Mladenović, et al. 2010. Variable neighbourhood search: methods and applications. *Annals of Operations Research* 175, 1 (2010), 367–407.
- [138] A. Harlap, D. Narayanan, et al. 2018. PipeDream: Fast and efficient pipeline parallel dnn training. *arXiv preprint arXiv:1806.03377* (2018).
- [139] J. A. Hartigan and M. A. Wong. 1979. Algorithm AS 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 28, 1 (1979).
- [140] K. He, G. Gkioxari, et al. 2017. Mask R-CNN. *2017 IEEE International Conference on Computer Vision (ICCV)* (2017), 2980–2988.
- [141] X. He, D. Mudigere, et al. 2017. Distributed hessian-free optimization for deep neural network. In *Workshops at the Thirty-First AAAI Conference on Artificial Intelligence*.
- [142] N. Heess, S. Sriram, et al. 2017. Emergence of locomotion behaviours in rich environments. *arXiv preprint arXiv:1707.02286* (2017).
- [143] V. Hema. 2015. DoS Attack Detection Based on Naive Bayes Classifier. (2015).
- [144] G. Hinton. 2019. RMSprop. Retrieved June 7, 2019 from http://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf
- [145] G. Hinton, O. Vinyals, et al. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015).
- [146] Q. Ho, J. Cipar, et al. 2013. More effective distributed ml via a stale synchronous parallel parameter server. In *Advances in neural information processing systems*. 1223–1231.
- [147] S. Hochreiter and J. Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [148] S. C. Hoi, R. Jin, et al. 2006. Batch mode active learning and its application to medical image classification. In *Proceedings of the 23rd ICML*. ACM, 417–424.
- [149] S. C. Hoi, D. Sahoo, et al. 2018. Online learning: A comprehensive survey. *arXiv preprint arXiv:1802.02871* (2018).
- [150] L. Hou, R. Zhang, et al. 2019. Analysis of Quantized Models. In *International Conference on Learning Representations*. https://openreview.net/forum?id=ryM_loAqYX
- [151] A. G. Howard, M. Zhu, et al. 2017. MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017).
- [152] K. Hsieh, A. Harlap, et al. 2017. Gaia: Geo-distributed machine learning approaching {LAN} speeds. In *14th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 17)*. 629–647.
- [153] C.-H. Hsu, S.-H. Chang, et al. 2018. Monas: Multi-objective neural architecture search using reinforcement learning. *arXiv preprint arXiv:1806.10332* (2018).
- [154] H. Hu, M. Kantardzic, et al. 2019. No Free Lunch Theorem for concept drift detection in streaming data classification: A review. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* (2019), e1327.
- [155] T. Hu and Y. Fei. 2010. QELAR: A machine-learning-based adaptive routing protocol for energy-efficient and lifetime-extended underwater sensor networks. *TMC* (2010).
- [156] G. Huang, S. Liu, et al. 2018. Condensenet: An efficient densenet using learned group convolutions. In *Proceedings of the IEEE Conference on CVPR*. 2752–2761.
- [157] Y. Huang, Y. Cheng, et al. 2019. Gpipe: Efficient training of giant neural networks using pipeline parallelism. In *Advances in Neural Information Processing Systems*. 103–112.
- [158] I. Hubara, M. Courbariaux, et al. 2017. Quantized neural networks: Training neural networks with low precision weights and activations. *The Journal of Machine Learning Research* 18, 1 (2017), 6869–6898.
- [159] Z. Huo, B. Gu, et al. 2018. Decoupled parallel backpropagation with convergence guarantee. *arXiv preprint arXiv:1804.10574* (2018).
- [160] S. Ioffe and C. Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167* (2015).
- [161] J. Ivanecky and S. Mehlhase. 2012. An in-car speech recognition system for disabled drivers. In *International Conference on Text, Speech and Dialogue*. Springer, 505–512.
- [162] G. Jagannathan, K. Pillaiipakkamatt, et al. 2009. A practical differentially private random decision tree classifier. In *2009 IEEE International Conference on Data Mining Workshops*. IEEE, 114–121.
- [163] H. Jamshidi, T. Lukaszewicz, et al. 2011. Fusion of digital map traffic signs and camera-detected signs. In *ICSPCS*. IEEE, 1–7.
- [164] S. Jeaugey. 2017. NCCL 2.0. Retrieved March 2, 2020 from <http://on-demand.gputechconf.com/gtc/2017/presentation/s7155-jeaugey-nccl.pdf>
- [165] D. N. Jha, P. Michalak, et al. 2019. Multi-objective Deployment of Data Analysis Operations in Heterogeneous IoT Infrastructure. *IEEE Transactions on Industrial Informatics* (2019).
- [166] Z. Jia, M. Zaharia, et al. 2018. Beyond data and model parallelism for deep neural networks. *arXiv preprint arXiv:1807.05358* (2018).
- [167] P. Jin, B. Ginsburg, et al. 2018. Spatially parallel convolutions. (2018).
- [168] A. J. Joshi, F. Porikli, et al. 2009. Multi-class active learning for image classification. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2372–2379.

- [169] D. Jouan-Rimbaud, D.-L. Massart, et al. 1995. Genetic algorithms as a tool for wavelength selection in multivariate calibration. *Analytical Chemistry* 67, 23 (1995), 4295–4301.
- [170] P. Kairouz, H. B. McMahan, et al. 2019. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977* (2019).
- [171] A. P. Kale and S. P. Sonavane. 2019. IoT based Smart Farming: Feature subset selection for optimized high-dimensional data using improved GA based approach for ELM. *Computers and Electronics in Agriculture* 161 (2019), 225–232.
- [172] R. Kemker, M. McClure, et al. 2018. Measuring catastrophic forgetting in neural networks. In *Thirty-second AAAI conference on artificial intelligence*.
- [173] H. Kim, M. Kim, et al. 2018. Nsm: Meet the mlaas platform with a real-world case study. *arXiv preprint arXiv:1810.09957* (2018).
- [174] J. K. Kim, Q. Ho, et al. 2016. STRADS: a distributed framework for scheduled model parallel machine learning. In *Proceedings of the Eleventh European Conference on Computer Systems*. 1–16.
- [175] D. P. Kingma and J. Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [176] M. M. Kirmani and S. I. Ansarullah. [n. d.]. Prediction of Heart Disease using Decision Tree a Data Mining Technique. ([n. d.]).
- [177] J. Ko and D. Fox. 2009. GP-BayesFilters: Bayesian filtering using Gaussian process prediction and observation models. *Autonomous Robots* 27, 1 (2009), 75–90.
- [178] C. Koliás, V. Koliás, et al. 2017. TermID: A distributed swarm intelligence-based approach for wireless intrusion detection. *ISECURE* (2017).
- [179] V. R. Konda and J. N. Tsitsiklis. 2000. Actor-critic algorithms. In *Advances in neural information processing systems*. 1008–1014.
- [180] J. Konečný, H. B. McMahan, et al. 2016. Federated optimization: Distributed machine learning for on-device intelligence. *arXiv preprint arXiv:1610.02527* (2016).
- [181] W. Kong, Z. Y. Dong, et al. 2017. Short-term residential load forecasting based on LSTM recurrent neural network. *IEEE Transactions on Smart Grid* 10, 1 (2017), 841–851.
- [182] N. Kothari, T. Millstein, et al. 2008. Deriving state machines from TinyOS programs using symbolic execution. In *IPSN*. IEEE Computer Society.
- [183] A. Krizhevsky. 2014. One weird trick for parallelizing convolutional neural networks. *arXiv preprint arXiv:1404.5997* (2014).
- [184] V. Krunić, E. Trumpler, et al. 2007. NodeMD: Diagnosing node-level faults in remote wireless sensor systems. In *MobiSys*. ACM, 43–56.
- [185] P. Kumar, H. Gauba, et al. 2017. Coupled HMM-based multi-sensor data fusion for sign language recognition. *Pattern Recognition Letters* 86 (2017), 1–8.
- [186] A. Kurakin, I. Goodfellow, et al. 2016. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236* (2016).
- [187] S. Lakshminarasimman, S. Ruswin, et al. 2017. Detecting DDoS attacks using decision tree algorithm. In *2017 Fourth International Conference on Signal Processing, Communication and Networking (ICSCN)*. IEEE, 1–6.
- [188] T. Lappas and K. Pelechrinis. 2007. Data mining techniques for (network) intrusion detection systems. *CSE at UC Riverside* (2007).
- [189] B. P. L. Lau, S. H. Marakkalage, et al. 2019. A survey of data fusion in smart city applications. *Information Fusion* 52 (2019), 357–374.
- [190] V. Lebedev, Y. Ganin, et al. 2014. Speeding-up convolutional neural networks using fine-tuned cp-decomposition. *arXiv preprint arXiv:1412.6553* (2014).
- [191] V. Lebedev and V. Lempitsky. 2016. Fast convnets using group-wise brain damage. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- [192] S. Lee, J. K. Kim, et al. 2014. On model parallelization and scheduling strategies for distributed machine learning. In *Advances in neural information processing systems*. 2834–2842.
- [193] W. Lee, S. J. Stolfo, et al. 2000. Adaptive intrusion detection: A data mining approach. *Artificial Intelligence Review* 14, 6 (2000), 533–567.
- [194] S. Levine and P. Abbeel. 2014. Learning neural network policies with guided policy search under unknown dynamics. In *Advances in Neural Information Processing Systems*. 1071–1079.
- [195] H. Li, W. Ouyang, et al. 2016. Multi-bias non-linear activation in deep neural networks. In *International conference on machine learning*. 221–229.
- [196] J. Li, H. Gao, et al. 2018. Deep reinforcement learning based computation offloading and resource allocation for MEC. In *WCNC*. IEEE, 1–6.
- [197] J. Li, W. Monroe, et al. 2016. Deep reinforcement learning for dialogue generation. *arXiv preprint arXiv:1606.01541* (2016).
- [198] L. Li, Y. Lv, et al. 2016. Traffic signal timing via deep reinforcement learning. *IEEE/CAA Journal of Automatica Sinica* 3, 3 (2016), 247–254.
- [199] L. Li, W. Xu, et al. 2019. Rsa: Byzantine-robust stochastic aggregation methods for distributed learning from heterogeneous datasets. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 1544–1551.
- [200] M. Li, D. G. Andersen, et al. 2014. Scaling distributed machine learning with the parameter server. In *11th USENIX Symposium on OSDI*. 583–598.
- [201] M. Li, D. G. Andersen, et al. 2014. Communication efficient distributed machine learning with the parameter server. In *Advances in Neural Information Processing Systems*.
- [202] R. Li, C. Zhang, et al. 2019. Learning Driven Mobility Control of Airborne Base Stations in Emergency Networks. *ACM SIGMETRICS Performance Evaluation Review* (2019).
- [203] W. Li, P. Yi, et al. 2014. A new intrusion detection system based on KNN classification algorithm in wireless sensor network. *Journal of Electrical and Computer Engineering* 2014 (2014).
- [204] T. P. Lillicrap, J. J. Hunt, et al. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971* (2015).
- [205] J. Lin, W. Yu, et al. 2017. A survey on internet of things: Architecture, enabling technologies, security and privacy, and applications. *IEEE Internet of Things Journal* 4, 5 (2017), 1125–1142.
- [206] T. Lin, S. U. Stich, et al. 2018. Don't Use Large Mini-Batches, Use Local SGD. *arXiv preprint arXiv:1808.07217* (2018).
- [207] Y. Lin, S. Han, et al. 2017. Deep gradient compression: Reducing the communication bandwidth for distributed training. *arXiv preprint arXiv:1712.01887* (2017).
- [208] R. Lioutikov, A. Paraschos, et al. 2014. Sample-based information-theoretic stochastic optimal control. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 3896–3902.
- [209] C. Liu, B. Zoph, et al. 2018. Progressive neural architecture search. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 19–34.
- [210] H. Liu, K. Simonyan, et al. 2017. Hierarchical representations for efficient architecture search. *arXiv preprint arXiv:1711.00436* (2017).
- [211] L. Liu, M. Zhang, et al. 2014. A survey on workflow management and scheduling in cloud computing. In *2014 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*. IEEE, 837–846.
- [212] M. Lopez-Martin, B. Carro, et al. 2017. Conditional variational autoencoder for prediction and feature recovery applied to intrusion detection in iot. *Sensors* (2017).
- [213] G. Louppe. 2014. Understanding random forests: From theory to practice. *arXiv preprint arXiv:1407.7502* (2014).
- [214] Q. Lu, T. Peng, et al. 2010. Utility-based resource allocation in uplink of OFDMA-based cognitive radio networks. *International Journal of Communication Systems* 23, 2 (2010), 252–274.
- [215] M.-T. Luong, H. Pham, et al. 2015. Effective approaches to attention-based neural machine translation. *EMNLP* (2015).
- [216] C. Lyu, K. Huang, et al. 2015. A unified gradient regularization family for adversarial examples. In *2015 IEEE International Conference on Data Mining*. IEEE, 301–309.
- [217] N. Ma, X. Zhang, et al. 2018. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- [218] D. J. MacKay and D. J. Mac Kay. 2003. *Information theory, inference and learning algorithms*. Cambridge university press.
- [219] R. Madan, J. Borran, et al. 2010. Cell association and interference coordination in heterogeneous LTE-A cellular networks. *IEEE Journal on selected areas in communications* 28, 9 (2010), 1479–1489.

- [220] R. Madan, S. P. Boyd, et al. 2010. Fast algorithms for resource allocation in wireless cellular networks. *IEEE/ACM Transactions on Networking (TON)* 18, 3 (2010), 973–984.
- [221] H. Mao, M. Alizadeh, et al. 2016. Resource management with deep reinforcement learning. In *Proceedings of the 15th ACM Workshop on Hot Topics in Networks*. ACM.
- [222] H. Mao, R. Netravali, et al. 2017. Neural adaptive video streaming with pensieve. In *SIGCOMM*. ACM.
- [223] A. Marciano-Cedeno, J. Quintanilla-Domínguez, et al. 2010. Feature selection using sequential forward selection and classification applying artificial metaplasticity neural network. In *IECON*. IEEE.
- [224] M. Mathieu, M. Henaff, et al. 2013. Fast training of convolutional networks through fits. *arXiv preprint arXiv:1312.5851* (2013).
- [225] W. McColl. 1995. Bulk synchronous parallel computing. *Abstract Machine Models for Highly Parallel Computers*, Oxford University Press, Oxford (1995).
- [226] R. McDonald, K. Hall, et al. 2010. Distributed training strategies for the structured perceptron. In *Human language technologies: The 2010 annual conference of the North American chapter of the association for computational linguistics*. Association for Computational Linguistics, 456–464.
- [227] H. B. McMahan, E. Moore, et al. 2016. Communication-efficient learning of deep networks from decentralized data. *arXiv preprint arXiv:1602.05629* (2016).
- [228] H. B. McMahan, D. Ramage, et al. 2017. Learning differentially private recurrent language models. *ICLR 2018* (2017).
- [229] Y. Meidan, M. Bohadana, et al. 2017. Detection of unauthorized iot devices using machine learning techniques. *arXiv preprint arXiv:1709.04647* (2017).
- [230] X. Meng, J. Bradley, et al. 2016. Mlib: Machine learning in apache spark. *The Journal of Machine Learning Research* 17, 1 (2016), 1235–1241.
- [231] A. Merentitis and C. Debes. 2015. Automatic fusion and classification using random forests and features extracted with deep learning. In *IGARSS*. IEEE, 2943–2946.
- [232] A. Meunkaewjinda, P. Kumsawat, et al. 2008. Grape leaf disease detection from color imagery using hybrid intelligent system. In *ECTI-CON*, Vol. 1. IEEE, 513–516.
- [233] E. M. E. Mhamdi and R. Guerraoui. 2017. When Neurons Fail. In *2017 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. 1028–1037. <https://doi.org/10.1109/IPDPS.2017.66>
- [234] A. Mirhoseini, A. Goldie, et al. 2018. A hierarchical model for device placement. (2018).
- [235] A. Mirhoseini, H. Pham, et al. 2017. Device placement optimization with reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning—Volume 70*. JMLR.org, 2430–2439.
- [236] T. M. Mitchell. 1999. Machine learning and data mining. *Commun. ACM* 42, 11 (1999), 30–36.
- [237] V. Mnih, A. P. Badia, et al. 2016. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*. 1928–1937.
- [238] V. Mnih, K. Kavukcuoglu, et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529.
- [239] M. Mohammadi, A. Al-Fuqaha, et al. 2018. Deep learning for IoT big data and streaming analytics: A survey. *IEEE Communications Surveys & Tutorials* 20, 4 (2018).
- [240] P. Moritz, R. Nishihara, et al. 2016. A linearly-convergent stochastic L-BFGS algorithm. In *Artificial Intelligence and Statistics*. 249–258.
- [241] P. Moritz, R. Nishihara, et al. 2019. Ray: A distributed framework for emerging {AI} applications. In *13th USENIX Symposium on OSDI*. 561–577.
- [242] S. Mukherjee and N. Sharma. 2012. Intrusion detection using naive Bayes classifier with feature reduction. *Procedia Technology* 4 (2012), 119–128.
- [243] M. S. Munir, S. F. Abedin, et al. 2017. RNN based Energy Demand Prediction for Smart-Home in Smart-Grid Framework.
- [244] R. Munos, T. Stepleton, et al. 2016. Safe and efficient off-policy reinforcement learning. In *Advances in Neural Information Processing Systems*. 1054–1062.
- [245] A. Murad, F. A. Kraemer, et al. 2019. Autonomous Management of Energy-Harvesting IoT Nodes Using Deep Reinforcement Learning. *arXiv:1905.04181* (2019).
- [246] A. Nagabandi, G. Kahn, et al. 2018. Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 7559–7566.
- [247] P. Nakkiran, G. Kaplan, et al. 2019. Deep double descent: Where bigger models and more data hurt. *arXiv preprint arXiv:1912.02292* (2019).
- [248] P. Naraei, A. Abhari, et al. 2016. Application of multilayer perceptron neural networks and support vector machines in classification of healthcare data. In *FTC*. 848–852.
- [249] B. Nelson, B. I. Rubinstein, et al. 2012. Query strategies for evading convex-inducing classifiers. *Journal of Machine Learning Research* 13, May (2012), 1293–1332.
- [250] D. Oh and I. Yun. 2018. Residual error based anomaly detection using auto-encoder in smd machine sound. *Sensors* 18, 5 (2018), 1308.
- [251] K. Osawa, Y. Tsuji, et al. 2018. Second-order optimization method for large mini-batch: Training resnet-50 on imagenet in 35 epochs. *arXiv preprint arXiv:1811.12019* (2018).
- [252] X. Pan, M. Lam, et al. 2016. Cyclades: Conflict-free asynchronous machine learning. In *Advances in Neural Information Processing Systems*. 2568–2576.
- [253] N. Papernot, P. McDaniel, et al. 2017. Practical black-box attacks against machine learning. In *ASIACCS 2017*. ACM, 506–519.
- [254] N. Papernot, P. McDaniel, et al. 2016. The limitations of deep learning in adversarial settings. In *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE.
- [255] G. I. Parisi, R. Kenker, et al. 2019. Continual lifelong learning with neural networks: A review. *Neural Networks* (2019).
- [256] R. Parr and S. J. Russell. 1998. Reinforcement learning with hierarchies of machines. In *NIPS*.
- [257] S. Peddabachigari, A. Abraham, et al. [n. d.]. Intrusion detection systems using decision trees and support vector machines. ([n. d.]).
- [258] K. Peng, V. Leung, et al. 2018. Intrusion detection system based on decision tree over big data in fog environment. *Wireless Communications and Mobile Computing* 2018 (2018).
- [259] Z. Peng, X. Chen, et al. 2018. AXNet: Approximate computing using an end-to-end trainable neural network. In *ICCAD*. ACM, 11.
- [260] A. Petrowski, G. Dreyfus, et al. 1993. Performance analysis of a pipelined backpropagation parallel algorithm. *IEEE Transactions on Neural Networks* 4, 6 (1993), 970–981.
- [261] H. Pham, M. Y. Guan, et al. 2018. Efficient neural architecture search via parameter sharing. *arXiv preprint arXiv:1802.03268* (2018).
- [262] J. Plasse and N. Adams. 2016. Handling delayed labels in temporally evolving data streams. In *2016 IEEE International Conference on Big Data (Big Data)*. IEEE, 2416–2424.
- [263] D. Povey, X. zhang, et al. 2017. Parallel training of DNNs with Natural Gradient and Parameter Averaging. *ICLR 2015 - (Aug 2017)*.
- [264] P. W. Protzel, D. L. Palumbo, et al. 1993. Performance and fault-tolerance of neural networks for optimization. *IEEE transactions on Neural Networks* 4, 4 (1993), 600–614.
- [265] N. Qian. 1999. On the momentum term in gradient descent learning algorithms. *Neural networks* 12, 1 (1999), 145–151.
- [266] A. Qiao, B. Aragam, et al. 2018. Fault Tolerance in Iterative-Convergent Machine Learning. *arXiv preprint arXiv:1810.07354* (2018).
- [267] J. Quackenbush. 2002. Microarray data normalization and transformation. *Nature genetics* 32, 4s (2002), 496.
- [268] D. Quillen, E. Jang, et al. 2018. Deep reinforcement learning for vision-based robotic grasping: A simulated comparative evaluation of off-policy methods. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 6284–6291.
- [269] J. R. Quinlan. 1986. Induction of decision trees. *Machine learning* 1, 1 (1986), 81–106.
- [270] J. R. Quinlan. 2014. *C4.5: programs for machine learning*. Elsevier.
- [271] R. Raina, A. Madhavan, et al. 2009. Large-scale deep unsupervised learning using graphics processors. In *Proceedings of the 26th annual ICML*. ACM.
- [272] N. Ramakrishnan and T. Soni. 2018. Network Traffic Prediction Using Recurrent Neural Networks. In *2018 17th IEEE ICMLA*. IEEE.
- [273] B. Randell. 1975. System structure for software fault tolerance. *Ieee transactions on software engineering* 2 (1975), 220–232.
- [274] T. Rao, N. Rajasekhar, et al. [n. d.]. An efficient approach for weather forecasting using support vector machines.

- [275] M. Rastegari, V. Ordonez, et al. 2016. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European Conference on Computer Vision*. Springer.
- [276] E. Real, A. Aggarwal, et al. 2019. Regularized evolution for image classifier architecture search. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [277] B. Recht, C. Re, et al. 2011. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in neural information processing systems*. 693–701.
- [278] L. Ren, X. Yuan, et al. 2018. Deep reinforcement learning with iterative shift for visual tracking. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 684–700.
- [279] K. Rieck, P. Trinius, et al. 2011. Automatic analysis of malware behavior using machine learning. *Journal of Computer Security* 19, 4 (2011), 639–668.
- [280] H. Röger and R. Mayer. 2019. A Comprehensive Survey on Parallelization and Elasticity in Stream Processing. *ACM Computing Surveys (CSUR)* 52, 2 (2019), 36.
- [281] R. Roman, J. Zhou, et al. 2013. On the features and challenges of security and privacy in distributed internet of things. *Computer Networks* 57, 10 (2013), 2266–2279.
- [282] A. Romero, N. Ballas, et al. 2014. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550* (2014).
- [283] B. I. Rubinstein, P. L. Bartlett, et al. 2009. Learning in a large function space: Privacy-preserving mechanisms for SVM learning. *arXiv preprint arXiv:0911.5708* (2009).
- [284] B. I. Rubinstein, B. Nelson, et al. 2009. Antidote: understanding and defending against poisoning of anomaly detectors. In *SIGCOMM on Internet measurement*. ACM.
- [285] F. Ruelens, B. J. Claessens, et al. 2016. Residential demand response of thermostatically controlled loads using batch reinforcement learning. *IEEE Transactions on Smart Grid* 8, 5 (2016), 2149–2159.
- [286] V. Saeidi, B. Pradhan, et al. 2014. Fusion of airborne lidar with multispectral spot 5 image for enhancement of feature extraction using dempster–shafer theory. *TGRS* (2014).
- [287] P. Samangouei, M. Kabkab, et al. 2018. Defense-gan: Protecting classifiers against adversarial attacks using generative models. *arXiv preprint arXiv:1805.06605* (2018).
- [288] M. Sandler, A. Howard, et al. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- [289] B. Sanjaan and E. Chuluun. 2013. Malware detection using linear SVM. In *Ifost*, Vol. 2. IEEE, 136–138.
- [290] T. Schaul, J. Quan, et al. 2015. Prioritized experience replay. *arXiv preprint arXiv:1511.05952* (2015).
- [291] J. Schulman, S. Levine, et al. 2015. Trust region policy optimization. In *International Conference on Machine Learning*. 1889–1897.
- [292] J. Schulman, F. Wolski, et al. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [293] F. Seide, H. Fu, et al. 2014. 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech dnns. In *Fifteenth Annual Conference of the International Speech Communication Association*.
- [294] A. Sergeev and M. Del Balso. 2018. Horovod: fast and easy distributed deep learning in TensorFlow. *arXiv preprint arXiv:1802.05799* (2018).
- [295] B. Settles. 2012. Active learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 6, 1 (2012), 1–114.
- [296] A. Shafahi, W. R. Huang, et al. 2018. Poison frogs! targeted clean-label poisoning attacks on neural networks. In *Advances in Neural Information Processing Systems*.
- [297] U. Shaham, Y. Yamada, et al. 2018. Understanding adversarial training: Increasing local stability of supervised models through robust optimization. *Neurocomputing* (2018).
- [298] W. Shang, K. Sohn, et al. 2016. Understanding and improving convolutional neural networks via concatenated rectified linear units. In *ICML*. 2217–2225.
- [299] N. Shazeer, Y. Cheng, et al. 2018. Mesh-TensorFlow: Deep Learning for Supercomputers. In *Neural Information Processing Systems*.
- [300] W. Shi, J. Cao, et al. 2016. Edge computing: Vision and challenges. *IEEE internet of things journal* 3, 5 (2016), 637–646.
- [301] T. A. Shinde and J. R. Prasad. 2017. IoT based animal health monitoring with naive Bayes classification. *IJETT* 1, 2 (2017).
- [302] J. Shlens. 2014. A tutorial on principal component analysis. *arXiv preprint arXiv:1404.1100* (2014).
- [303] R. Shokri and V. Shmatikov. 2015. Privacy-preserving deep learning. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*. ACM, 1310–1321.
- [304] R. Shokri, M. Stronati, et al. 2017. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 3–18.
- [305] A. Shustanov and P. Yakimov. 2017. CNN design for real-time traffic sign recognition. *Procedia engineering* 201 (2017), 718–725.
- [306] S. Sicari, A. Rizzardi, et al. 2015. Security, privacy and trust in Internet of Things: The road ahead. *Computer networks* 76 (2015), 146–164.
- [307] L. Sifre and S. Mallat. 2014. Rigid-motion scattering for image classification. *Ph. D. dissertation* (2014).
- [308] A. K. Sikder, H. Aksu, et al. 2017. 6thsense: A context-aware sensor-based attack detector for smart devices. In *26th {USENIX} Security Symposium*. 397–414.
- [309] K. J. Singh and D. S. Kapoor. 2017. Create Your Own Internet of Things: A survey of IoT platforms. *IEEE Consumer Electronics Magazine* 6, 2 (2017), 57–68.
- [310] S. Singh and I. Chana. 2016. A survey on resource scheduling in cloud computing: Issues and challenges. *Journal of grid computing* 14, 2 (2016), 217–264.
- [311] S. P. Singh. 1992. Reinforcement learning with a hierarchy of abstract models. In *Proceedings of the National Conference on Artificial Intelligence*. JOHN WILEY & SONS LTD.
- [312] S. L. Smith, P.-J. Kindermans, et al. 2017. Don’t decay the learning rate, increase the batch size. *arXiv preprint arXiv:1711.00489* (2017).
- [313] E. R. Sparks, A. Talwalkar, et al. 2015. Automating model search for large scale machine learning. In *Proceedings of the Sixth ACM Symposium on Cloud Computing*. ACM, 368–380.
- [314] N. Srivastava, G. Hinton, et al. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1 (2014), 1929–1958.
- [315] J. Steinhardt, P. W. Koh, et al. 2017. Certified defenses for data poisoning attacks. In *Advances in neural information processing systems*. 3517–3529.
- [316] S. U. Stich. 2019. Local SGD Converges Fast and Communicates Little. In *ICLR 2019 ICLR 2019 International Conference on Learning Representations*.
- [317] I. Stoica, D. Song, et al. 2017. A berkeley view of systems challenges for ai. *arXiv preprint arXiv:1712.05855* (2017).
- [318] N. Strom. 2015. Scalable distributed DNN training using commodity GPU cloud computing. In *Sixteenth Annual Conference of the International Speech Communication Association*.
- [319] J. Su and H. Zhang. 2006. A fast decision tree learning algorithm.
- [320] M.-Y. Su. 2011. Real-time anomaly detection systems for Denial-of-Service attacks by weighted k-nearest-neighbor classifiers. *Expert Systems with Applications* 38, 4 (2011), 3492–3498.
- [321] O. Suci, R. Marginean, et al. 2018. When does machine learning FAIL? generalized transferability for evasion and poisoning attacks. In *USENIX Security 18*). 1299–1316.
- [322] S. Sultana, D. Midi, et al. 2014. Kinesis: a security incident response and prevention system for wireless sensor networks. In *SenSys*. ACM.
- [323] J. Supancic III and D. Ramanan. 2017. Tracking as online decision-making: Learning a policy from streaming videos with reinforcement learning. In *Proceedings of the IEEE International Conference on Computer Vision*. 322–331.
- [324] R. S. Sutton and A. G. Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- [325] R. B. S. Swetha and K. G. Meena. [n. d.]. Smart Grid-A Network based Intrusion Detection System. *International Journal of Computer Applications* 975 ([n. d.]), 8887.
- [326] C. Szegedy, S. Ioffe, et al. 2017. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- [327] C. Tan, F. Sun, et al. 2018. A survey on deep transfer learning. In *International conference on artificial neural networks*. Springer, 270–279.

- [328] M. Tan, B. Chen, et al. 2019. Mnasnet: Platform-aware neural architecture search for mobile. In *CVPR*.
- [329] H. Tang, W. Wang, et al. 2019. Expression conditional gan for facial expression-to-expression translation. In *2019 IEEE International Conference on Image Processing (ICIP)*. IEEE, 4449–4453.
- [330] TensorFlow. 2019. High performance inference with TensorRT Integration. <https://medium.com/tensorflow/high-performance-inference-with-tensortrt-integration-c4d78795fbfc>
- [331] S. Thrun and A. Schwartz. 1995. Finding structure in reinforcement learning. In *NIPS*.
- [332] J. T. Townsend. 1971. Theoretical analysis of an alphabetic confusion matrix. *Perception & Psychophysics* 9, 1 (1971), 40–50.
- [333] F. Tramer, A. Kurakin, et al. 2017. Ensemble adversarial training: Attacks and defenses. *arXiv preprint arXiv:1705.07204* (2017).
- [334] F. Tramer, F. Zhang, et al. 2016. Stealing machine learning models via prediction apis. In *25th USENIX Security Symposium (USENIX Security 16)*. 601–618.
- [335] A. Ullah, J. Ahmad, et al. 2018. Action Recognition in Video Sequences using Deep Bi-Directional LSTM With CNN Features. *IEEE Access* 6 (2018), 1155–1166.
- [336] M. Vacher, B. Lecouteux, et al. 2015. Speech and speaker recognition for home automation: Preliminary results. In *SpED*. IEEE.
- [337] L. G. Valiant. 1984. A theory of the learnable. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing*. ACM, 436–445.
- [338] H. Van Hasselt, A. Guez, et al. 2016. Deep reinforcement learning with double q-learning. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- [339] V. Vanhoucke, A. Senior, et al. 2011. Improving the speed of neural networks on CPUs. (2011).
- [340] J. Vanschoren. 2018. Meta-learning: A survey. *arXiv preprint arXiv:1810.03548* (2018).
- [341] V. Vapnik. 2013. *The nature of statistical learning theory*. Springer science & business media.
- [342] S. Venugopalan, H. Xu, et al. 2014. Translating videos to natural language using deep recurrent neural networks. *arXiv preprint arXiv:1412.4729* (2014).
- [343] S. A. Vinterbo. 2012. Differentially private projected histograms: Construction and use for prediction. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 19–34.
- [344] R. Volpi, H. Namkoong, et al. 2018. Generalizing to unseen domains via adversarial data augmentation. In *Advances in Neural Information Processing Systems*. 5334–5344.
- [345] G. Wang, J. Hao, et al. 2010. A new approach to intrusion detection using Artificial Neural Networks and fuzzy clustering. *Expert systems with applications* (2010).
- [346] J. Wang and L. Perez. 2017. The effectiveness of data augmentation in image classification using deep learning. *Convolutional Neural Networks Vis. Recognit* (2017).
- [347] M. Wang and X.-S. Hua. 2011. Active learning in multimedia annotation and retrieval: A survey. *ACM Transactions on Intelligent Systems and Technology (TIST)* (2011).
- [348] M. Wang, X. S. Hua, et al. 2007. Interactive video annotation by multi concept multi modality active learning. *ICSC* (2007).
- [349] M. Wang, C.-c. Huang, et al. 2019. Supporting very large models using automatic dataflow graph partitioning. In *Proceedings of the Fourteenth EuroSys Conference 2019*. 1–17.
- [350] S. Wang, T. Tuor, et al. 2018. When edge meets learning: Adaptive control for resource-constrained distributed machine learning. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. IEEE, 63–71.
- [351] Z. Wang, T. Schaul, et al. 2015. Dueling network architectures for deep reinforcement learning. *arXiv preprint arXiv:1511.06581* (2015).
- [352] C. J. Watkins and P. Dayan. 1992. Q-learning. *Machine learning* 8, 3-4 (1992), 279–292.
- [353] D. Weerasiri, M. C. Barukh, et al. 2017. A taxonomy and survey of cloud resource orchestration techniques. *ACM Computing Surveys (CSUR)* 50, 2 (2017), 26.
- [354] J. Wei, W. Dai, et al. 2015. Managed communication and consistency for fast data-parallel iterative analytics. In *Proceedings of the Sixth ACM Symposium on Cloud Computing*. ACM, 381–394.
- [355] W. Wen, C. Wu, et al. 2016. Learning structured sparsity in deep neural networks. In *Advances in neural information processing systems*. 2074–2082.
- [356] W. Wen, C. Xu, et al. 2017. Terngrad: Ternary gradients to reduce communication in distributed deep learning. In *Advances in neural information processing systems*. 1509–1519.
- [357] Z. Wen, P. Bhatotia, et al. 2018. Approxiot: Approximate analytics for edge computing. In *ICDCS*. IEEE.
- [358] Z. Wen, D. O’Neill, et al. 2015. Optimal demand response using device-based reinforcement learning. *IEEE Transactions on Smart Grid* 6, 5 (2015), 2312–2324.
- [359] Z. Wen, R. Yang, et al. 2017. Fog orchestration for internet of things services. *IEEE Internet Computing* 21, 2 (2017), 16–24.
- [360] B. Wu, F. Iandola, et al. 2017. SqueezeDet: Unified, small, low power fully convolutional neural networks for real-time object detection for autonomous driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 129–137.
- [361] S. Wu, G. Li, et al. 2018. Training and Inference with Integers in Deep Neural Networks. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=HJGXzmspb>
- [362] Y. Wu, M. Schuster, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144* (2016).
- [363] H. Xiao, H. Xiao, et al. 2012. Adversarial Label Flips Attack on Support Vector Machines. In *ECAI*. 870–875.
- [364] E. P. Xing, Q. Ho, et al. 2015. Petuum: A new platform for distributed machine learning on big data. *IEEE Transactions on Big Data* 1, 2 (2015), 49–67.
- [365] C. Xiong, V. Zhong, et al. 2017. Dcn+: Mixed objective and deep residual coattention for question answering. *arXiv preprint arXiv:1711.00106* (2017).
- [366] W. Xu, D. Evans, et al. 2017. Feature squeezing: Detecting adversarial examples in deep neural networks. *arXiv preprint arXiv:1704.01155* (2017).
- [367] X. Xu, J. Zheng, et al. 2017. Data classification using evidence reasoning rule. *Knowledge-Based Systems* 116 (2017), 144–151.
- [368] Z. Xu, Y. Wang, et al. 2017. A deep reinforcement learning based framework for power-efficient resource allocation in cloud RANs. In *ICC*. IEEE.
- [369] N. J. Yadwadkar, B. Hariharan, et al. 2016. Multi-task learning for straggler avoiding predictive job scheduling. *The Journal of Machine Learning Research* 17, 1 (2016), 3692–3728.
- [370] C. Yang, L. Feng, et al. 2018. A novel data fusion algorithm to combat false data injection attacks in networked radar systems. *TSIPN* (2018).
- [371] J. Yang, Y. Chen, et al. 2008. Detecting sybil attacks in wireless and sensor networks using cluster analysis. In *2008 5th IEEE International Conference on Mobile Ad Hoc and Sensor Systems*. IEEE, 834–839.
- [372] E. Yigitoglu, M. Mohamed, et al. 2017. Foggy: a framework for continuous automated IoT application deployment in fog computing. In *AIMS*. IEEE, 38–45.
- [373] Y. You, I. Gitman, et al. 2017. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888* (2017).
- [374] H. Yu, S. Yang, et al. 2018. Parallel restarted SGD for non-convex optimization with faster convergence and less communication. *arXiv preprint arXiv:1807.06629* 2, 4 (2018), 7.
- [375] H. Yu, S. Yang, et al. 2019. Parallel restarted SGD with faster convergence and less communication: Demystifying why model averaging works for deep learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 5693–5700.
- [376] S. Yun, J. Choi, et al. 2017. Action-decision networks for visual tracking with deep reinforcement learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2711–2720.

- [377] S. Zagoruyko and N. Komodakis. 2016. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. *arXiv preprint arXiv:1612.03928* (2016).
- [378] S. Zhai, Y. Cheng, et al. 2016. Doubly convolutional neural networks. In *Advances in neural information processing systems*. 1082–1090.
- [379] Z.-H. Zhan, X.-F. Liu, et al. 2015. Cloud computing resource scheduling and a survey of its evolutionary approaches. *ACM Computing Surveys (CSUR)* 47, 4 (2015), 63.
- [380] C. Zhang, A. Kumar, et al. 2016. Materialization optimizations for feature selection workloads. *ACM Transactions on Database Systems (TODS)* 41, 1 (2016), 2.
- [381] C. Zhang, P. Patras, et al. 2019. Deep learning in mobile and wireless networking: A survey. *IEEE Communications Surveys & Tutorials* (2019).
- [382] H. Zhang, J. Li, et al. 2017. Zipml: Training linear models with end-to-end low precision, and a little bit of deep learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 4035–4043.
- [383] J. Zhang, C. De Sa, et al. 2016. Parallel SGD: When does averaging help? *arXiv preprint arXiv:1606.07365* (2016).
- [384] J. Zhang, Z. Zhang, et al. 2012. Functional mechanism: regression analysis under differential privacy. *Proceedings of the VLDB Endowment* 5, 11 (2012), 1364–1375.
- [385] Q.-s. Zhang and S.-C. Zhu. 2018. Visual interpretability for deep learning: a survey. *Frontiers of Information Technology & Electronic Engineering* 19, 1 (2018), 27–39.
- [386] S. Zhang, A. E. Choromanska, et al. 2015. Deep learning with elastic averaging SGD. In *Advances in neural information processing systems*. 685–693.
- [387] W. Zhang, W. Guo, et al. 2018. LSTM-based analysis of industrial IoT equipment. *IEEE Access* 6 (2018), 23551–23560.
- [388] W.-A. Zhang, B. Chen, et al. 2016. Hierarchical fusion estimation for clustered asynchronous sensor networks. *IEEE Trans. Automat. Control* 61, 10 (2016), 3064–3069.
- [389] X. Zhang, X. Zhou, et al. 2018. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *CVPR*. 6848–6856.
- [390] C. Zhao, C. Chen, et al. 2018. Application of Auxiliary Classifier Wasserstein Generative Adversarial Networks in Wireless Signal Classification of Illegal Unmanned Aerial Vehicles. *Applied Sciences* 8, 12 (2018), 2664.
- [391] C. Zhao, M. Shi, et al. 2018. Research on the open-categorical classification of the Internet-of-things based on generative adversarial networks. *Applied Sciences* (2018).
- [392] D. Zhao, Y. Chen, et al. 2016. Deep reinforcement learning with visual attention for vehicle classification. *IEEE Transactions on Cognitive and Developmental Systems* (2016).
- [393] W. Zhao, S. Tang, et al. [n. d.]. An improved kNN algorithm based on essential vector. *Elektronika ir Elektrotechnika* 123, 7 ([n. d.]), 119–122.
- [394] X. Zhao, Z. Gao, et al. 2011. A fault detection algorithm based on cluster analysis in wireless sensor networks. In *2011 Seventh International Conference on Mobile Ad-hoc and Sensor Networks*. IEEE, 354–355.
- [395] B. Zhou, J. Cao, et al. 2010. Adaptive traffic light control in wireless sensor network-based intelligent transportation system. In *VTC. IEEE*, 1–5.
- [396] S. Zhou, Q. Chen, et al. 2013. Active deep learning method for semi-supervised sentiment classification. *Neurocomputing* 120 (2013), 536–546.
- [397] S. Zhou, Y. Wu, et al. 2016. Dorefa-net: Training low bandwidth convolutional neural networks with low bandwidth gradients. *arXiv preprint arXiv:1606.06160* (2016).
- [398] Y. Zhou, S. Ebrahimi, et al. 2018. Resource-efficient neural architect. *arXiv preprint arXiv:1806.07912* (2018).
- [399] M. Zinkevich, M. Weimer, et al. 2010. Parallelized stochastic gradient descent. In *Advances in neural information processing systems*. 2595–2603.

A APPENDIX A

Software Specification	Cloud	Edge	IoT devices
Main OS	<ul style="list-style-type: none"> - Ubuntu - CentOS - Debian - RHEL - Windows Server - Amazon Linux 	<ul style="list-style-type: none"> - Raspbian - NOOBS - Amazon FreeRTOS - RIOT - Google Fuchsia OS - Windows 10 IoT 	<ul style="list-style-type: none"> - Amazon FreeRTOS - Contiki - TinyOS - RIOT - Ubuntu Core - Mbed OS
Programming Language	<ul style="list-style-type: none"> - Java - ASP.NET - Python - PHP - Ruby 	<ul style="list-style-type: none"> - Java - Python - C - C++ - JavaScript 	<ul style="list-style-type: none"> - C - C++ - Java - JavaScript - Python
Platforms	<ul style="list-style-type: none"> - AWS - Azure - Google Cloud Platform - IBM Cloud - Oracle Cloud 	<ul style="list-style-type: none"> - Amazon Greengrass - EdgeX - Cisco IOx - Akraino Edge Stack - Eclipse ioFog 	<ul style="list-style-type: none"> - AWS IoT - Azure IoT - GCP IoT - IBM Watson - Cisco IoT cloud connect

Table 3. List of OS, programming language and platform in IoT layers

B APPENDIX B

B.1 Traditional Machine Learning (TML) methods

In this subsection, we give the details of several TML algorithms, as well as their IoT applications.

Logistic Regression (LR). Logistic regression is a linear classifier capable of performing binary or multi-class classification. It is among the simplest classification algorithms. In a binary classification setting, prediction target y is usually formulated as $y \in \{0, 1\}$ and the prediction probability of the positive class given a d -dimensional input $\mathbf{x} = [x_1, x_2, \dots, x_d] \in \mathbb{R}^d$ can be calculated via:

$$p(y = 1|\mathbf{x}; \boldsymbol{\theta}) = \sigma(\mathbf{w}^T \mathbf{x} + b) = \sigma\left(\sum_{i=1}^d w_i x_i + b\right), \quad (2)$$

In the equation $\boldsymbol{\theta} = \{\mathbf{w} = [w_1, w_2, \dots, w_d] \in \mathbb{R}^d, b \in \mathbb{R}\}$ denotes the model parameters, and $\sigma(\cdot)$ is an activation function used to squash the linear output within the range $[0, 1]$ for probabilistic interpretation. The training of the classifier aims to learn suitable values for the parameters $\boldsymbol{\theta} = \{\mathbf{w}, b\}$, starting from some random initialization, through minimizing of a loss or cost function $J(\boldsymbol{\theta})$. For LR, log (Cross-Entropy) loss is used, i.e., $J(\boldsymbol{\theta}) = -\ln p(y|\mathbf{x}; \boldsymbol{\theta})$ for data point $\{\mathbf{x}, y\}$ to facilitate the calculation of the model gradient and to minimize the model loss. With the trained model, for any query data the classification decision can be made via thresholding the predicted probability.

Extend the binary LR to support c -class prediction scenarios, the target $y \in \{1, 2, \dots, c\}$. Softmax function is applied to the output layer to normalize the c outputs into probabilities. Different from binary LR, there are c sets of parameters $\boldsymbol{\theta} = \{\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_c]^T, \mathbf{b} = [b_1, \dots, b_c]^T\}$ to be estimated by minimizing the cross-entropy loss (i.e., log loss). With the trained model parameters, at the inference stage the class label will be assigned to the one with the largest classification probability.

Decision Trees (DT). Decision Tree is a tree-like model for classification or regression tasks. A decision tree is made up of nodes and edges where a node can be seen as a feature, and the edge represents a condition for classification. The learning process of DT is to select the optimal feature that can best split the training examples in a recursive manner. Based on certain criteria such as information gain [269] or Gini impurity [96], the root node can be selected from the features, which will divide the whole training population into two or more homogeneous sets. For each sub-population, sub-nodes will be selected in a similar manner and this process will repeat until all the subsets are pure (i.e., with the same class label for each subset). These nodes and edges constitute the trained model and can be used for inference.

Different from most of the other classifiers (such as LR, SVM, ANN, DL) which require the input to be normalised to numerical values, the unique tree-structure of DT make it possible to take both numerical and nominal values, making it a highly interpretable tool for various classification tasks (e.g., medical records diagnosis). However, DT suffers from the “curse of dimensionality”, and it faces an overfitting problem when the input dimensionality is too high. With unstructured high-dimensional data, feature engineering/extraction is one of the necessary steps to take (for dimensionality reduction), before DT is applied.

DT can be used as a main classifier (e.g., for low-dimensional structured data) or collaborative classifier with other machine learning algorithms on various IoT applications, such as Intrusion detection system [258] in fog environment.

Random Forest (RF). As its name implies, a RF is an ensemble model with many DTs as base classifiers. The individual DTs are constructed by random sampling the features and the training examples for diversity, boosting the performance of multiple classifier systems. The random sampling process makes the individual DTs less correlated—with different prediction errors—and the aggregating function can smooth the large prediction variance, making RF a robust classifier with high generalization capabilities. There are several key hyper-parameters for RF, and two main ones are: feature number for individual DTs, and number of DT classifiers. For individual DTs, there is a trade-off between generalization and discrimination capabilities with respect to feature dimension, and one popular heuristic




Fig. 11. Basic concepts in SVM

value is to use the square root of original dimension number (e.g., the default setting in scikit-learn). For the RF though, the performance of the model usually correlates with the number of DTs. Yet the performance gain tends to become less significant since the diversity (among the DT classifiers, which are correlated to some extent) will decrease accordingly. Since both the efficiency and storage are proportional (in a linear manner) to the classifier number, to find a number configuration that can result in optimal effectiveness, storage and efficiency performance is of vital importance.

RF is one of the most popular classifiers due to its great generalization capability. For example, it has been used for intrusion detection [56, 229] and anomaly detection. In a previous study [229], the authors collected and manually labeled data from 17 distinct IoT devices and used RF algorithms to recognize IoT device categories from the white list.

Support Vector Machine (SVM). Support Vector Machine (SVM), also referred to as large margin classifier, aims to find a decision boundary (separating hyperplane) that can best separate (i.e., with the largest margin) positive/negative classes. Fig. 11 shows some basic concepts in SVM including support vectors, margin and slack variable ξ . ξ is a non-negative variable that is used to measure the misclassified instances and those within the margin (i.e., margin violation, as shown in Fig. 11). The objective function of (soft-margin) SVM can be constructed to maximize the margin while penalizing these instances:

$$\operatorname{argmin}_{w, b, \xi_i} \frac{\|w\|^2}{2} + C \sum_i^N \xi_i, \quad \text{subject to } y_i(w^T x_i + b) \geq 1 - \xi_i, \text{ where } \xi_i \geq 0, \quad i = 1, 2, \dots, N. \quad (3)$$

In Eq (3), $\{x_i, y_i\}_{i=1}^N$ are the N training sets with $y_i \in \{-1, 1\}$, and C is a regularization hyper-parameter that trades-off between the margin and errors (i.e., measured by $\sum_i^N \xi_i$). It is obvious to see that if we set $C = 0$, we can get a classifier with the large margin at a cost of potential high training errors. On the other hand, if we set C to a very large number (or ∞ in theory), errors are less likely to be tolerant, and we may end up with a classifier with narrow or hard margin. It is worth noting that Eq.(3) can be further simplified into $\operatorname{argmin}_{w, b} \frac{\|w\|^2}{2} + C \sum_i^N \max(0, 1 - y_i(w^T x_i + b))$, where the term $\max(0, 1 - y_i(w^T x_i + b))$ is also referred to as hinge loss.




Fig. 12. Margins for linear (left) and kernel (right) SVM




Fig. 13. Majority voting process in KNN

On the other hand, for highly non-linearly separable data, instead of employing feature engineering/extraction processes, an elegant alternative—kernel SVM can be applied, and Radial Basis Function (RBF) is one of the most popular kernel functions. Kernel SVM also has the aforementioned characteristics such as soft margin, and it tends to have great performance on small non-linearly separable data. Fig. 12 shows the margins on linearly separable (by linear SVM) and non-linearly separable (by kernel SVM) data, respectively.

SVM is notable for its generalization capability and is suitable for those small datasets with high-dimensional features [341], and there are many IoT applications such as Android malware detection system [289], smart weather prediction [274], etc.

K-Nearest Neighbour (KNN). KNN is a nonparametric, instance-based, non-linear classifier. Under the classification setting, given any query data, KNN essentially performs majority voting among the K most similar training samples, as shown in Fig. 13. The similarity can be measured by some distance metrics such as Minkowski distance $D(\mathbf{x}, \mathbf{y}) = (\sum_{j=1}^d |x_j - y_j|^p)^{\frac{1}{p}}$ (for d -dimensional vectors \mathbf{x}, \mathbf{y}). It is worth noting that when $p = 1$ and $p = 2$, the Minkowski distance can be seen as Manhattan Distance and Euclidean distance respectively, yet for different applications the optimal distance metric may vary from case to case.

One of the key properties is that KNN does not require any training process (i.e., lazy learning), and for any query data, the distance calculation has to be performed for each sample in the whole training set, which makes KNN a less-scalable approach for large datasets. Another issue is the selection of hyper-parameter K . A small K may make KNN sensitive to outliers in the training set while a large one may make KNN less discriminate. Nevertheless, KNN is a powerful non-linear classifier with low-dimensional small datasets, and there are many IoT applications such as network intrusion detection [203], anomaly detection [320] and Urban noise identification [8], etc.

Naive Bayes (NB). NB is a probabilistic classifier which takes the class prior distribution into account, and assumes the features are conditionally independent. Based on the Bayesian theory and the posterior probability, it can be presented as $p(y|\mathbf{x}) = \frac{p(\mathbf{x}|y)p(y)}{p(\mathbf{x})}$, where $p(\mathbf{x}|y)$ is the likelihood; $p(y)$ is the prior; $p(\mathbf{x})$ is the evidence. The classification process is to assign the label with the largest posterior probability, and in this case the term evidence $p(\mathbf{x})$ remains a constant which can be cancelled out, i.e., $p(y|\mathbf{x}) \propto p(\mathbf{x}|y)p(y)$, which can be further written into $p(y|\mathbf{x}) \propto \prod_{j=1}^d p(x_j|y)p(y)$ for d -dimensional input \mathbf{x} due to the feature independence assumption. In practice, logarithm is often used to avoid the

Table 4. Summary of Traditional Machine Learning Models (Note: In time complexity, m represents the number of training sample, n represents the feature dimension, k represents the selected K value, c represents the class number and t indicates the tree number)

Method	Learning model	Category	Typical input data	Time Complexity	Characteristics	IoT Application
DT	Supervised	Discriminative	Various	$O(m \cdot n^2)$ [319]	<ul style="list-style-type: none"> Dividing training samples to branches and leaves High interpretability method Require large memory space due to the construction nature. 	<ul style="list-style-type: none"> Intrusion detection [257] Suspicious detection on traffic sources [100] Future Heart Attack Quantity prediction [176]
SVM		Discriminative	Various	$O(m^2 \cdot n)$ [66]	<ul style="list-style-type: none"> Good Generalization capability and suitability for small dataset with large feature Difficult on selection optimal kernel High computation complexity on large dataset with complex kernel 	<ul style="list-style-type: none"> Malware detection [289] Attack detection in smart grids [325] Smart weather prediction [274]
NB		Discriminative	Various	$O(m \cdot n + n \cdot c)$ [66]	<ul style="list-style-type: none"> Ease to implement Generalizes well to multi-class problem Low dependence on large dataset and robustness Hard to capture relation information 	<ul style="list-style-type: none"> Detection of network intrusion [242] Animal health monitoring [301]
KNN		Discriminative	Various	$O(m \cdot n \cdot k)$ [393]	<ul style="list-style-type: none"> Nonparametric, instance-based No training process, no model construction Sensitive to the outlier Difficult on selection optimal K 	<ul style="list-style-type: none"> Anomalies detection [120] Urban noise identification [8]
RF		Discriminative	Various	$O(t \cdot n^2 \cdot \log n)$ [213]	<ul style="list-style-type: none"> Robust to over-fitting Bypasses feature selection Impractical in specific real-time application 	<ul style="list-style-type: none"> DDoS attack detection [187] Unauthorized IoT devices detection [229]
K-Means	Unsupervised	Clustering	Various	$O(m \cdot n \cdot k)$ [66]	<ul style="list-style-type: none"> Ease to use on Unlabelled data Produces tighter cluster than hierarchical clustering Less effective than supervised learning method Difficult on selection optimal K 	<ul style="list-style-type: none"> Sensor fault detection [394] Sybil detection in industrial WSNs [371]
PCA		Dimension Reduction	Various	$O(m \cdot n^2 + n^3)$ [66]	<ul style="list-style-type: none"> Used for dimensional reduction Consequently reduce the complexity of the model Should be used with other ML methods 	<ul style="list-style-type: none"> Real-time detection systems in IoT environments [101] Traffic anomaly detection [90]

problem of floating point underflow, and the class label \hat{y} can be assigned via:

$$\hat{y} = \underset{y \in \{1,2,\dots,c\}}{\operatorname{argmax}} \ln p(y) + \sum_{j=1}^d \ln p(x_j|y), \quad (4)$$

NB is good at modelling both continuous and discrete data. For example the likelihood of a discrete feature can be calculated by frequency while the likelihood of a continuous feature can be calculated by density estimation (e.g., Gaussian). It takes the prior of class distribution into consideration, which is helpful in data imbalanced problems. However, it also suffers from the ‘‘curse of dimensionality’’ like DT, and normally can not be used directly on unstructured data before feature extraction/engineering approaches have been applied. When the feature independence assumption is not significantly violated, it is normally served as an efficient and effective classifier. There are many IoT applications, such as network traffic analysis for DoS attach detection [143], animal health monitoring [301], etc.

K-Means. Different from the supervised classification models above, K-means is an unsupervised clustering algorithm without using class label information for training. Given a number of data points, K-Means aims to find K centroids (i.e., means), and the corresponding nearest samples to form the clusters. Various distance metrics can be used for K-means algorithm, and the most common one is based on Euclidean distance, whose objective function is:

$$\operatorname{argmin} \sum_{i=1}^N \sum_{k=1}^K \gamma_{ik} \|x_i - \mu_k\|^2, \quad (5)$$

where μ_k is the centroid of the k th cluster and $\gamma_{ik} \in \{0, 1\}$ denotes whether sample x_i belongs to the k th cluster (1) or not (0). K-means clustering is an heuristic process—starting from random values (of the centroids $\{\mu_k\}_{k=1}^K$), it will 1) assign each example to the nearest cluster and 2) update the K centroids (by re-calculating the means of the

corresponding samples for each cluster). It is an iterative process and the updating will stop until Eq.(5) is minimized (e.g., lower than a pre-defined threshold).

K-means is among the most popular clustering algorithms due to its simplicity. However, using a Euclidean distance based method is limited to spherical datasets, and finding the most suitable distance metrics (for different applications/datasets) is one of the key issues for improving the performance of K-means algorithm. K-means is widely used for IoT applications such as sensor fault detection [394], Sybil detection [371], etc.

Principal component analysis (PCA). PCA is another unsupervised learning approach and it is normally used for dimensionality reduction or feature decorrelation. Covariance matrix \mathbf{S} which reflects the correlation of the features can be calculated via $\mathbf{S} = \frac{1}{N} \sum_{i=1}^N (\mathbf{x}_i - \mu)(\mathbf{x}_i - \mu)^T$ for N d -dimensional training data points $\{\mathbf{x}_i\}_{i=1}^N$ with $\mu = \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i$. Eigenvalue decomposition can be performed on \mathbf{S} such that the leading $s (< d)$ eigenvectors can be used as a transformation matrix for feature decorrelation or dimensionality reduction. PCA is one of the most popular feature extraction tools due to its simplicity, and there are many IoT applications such as traffic anomaly detection [40].

Table 4 summarizes the aforementioned TML, including their advantages, disadvantages and applications in IoT systems.

B.2 Deep Learning (DL) methods

In this section, we detail several deep learning algorithms: Deep Neural Networks (DNN), Convolutional Neural Network (CNN), Recurrent Neural Network (RNN), etc.

Deep Neural Networks (DNN). As previously mentioned, LR is one of the most popular linear classifiers, which means the decision boundary is a line or hyperplane in the feature space. However, most of the data are more complex and are not linearly separable in real-world scenarios. In this case, directly applying linear classifiers such as LR on the raw data would yield unsatisfied classification performance. One could empirically use hand-crafted features followed by LR, yet this trial-and-error process can be time-consuming, and it may be challenging to define the high-order discriminant features where the data can be linearly separable. One alternative is to use data-driven methods, and it is straightforward to extend LR to DNN (i.e., Multilayer Perceptron MLP).

DNN is built upon LR with a stack of hidden layers. By applying activation functions like sigmoid, tanh, or ReLU on the hidden layers, raw features can be transformed in a non-linear manner, yielding discriminant (i.e., linearly separable) features before the output layer. Similar to LR, the training process is to minimise the log loss. However, the parameters of DNN are the layer-wise connections, which can be learned using the back-propagation approach. DNN performs end-to-end learning, which means it learns the feature extractors and classifiers simultaneously.

DNN can be deemed as the simplest form of deep learning model which comprises one input layer, one output layer as well as multiple hidden layers for more complex feature extraction. It is worth noting that Artificial Neural Network (ANN) is a special case of DNN with only one hidden layer. The layers are organised in a hierarchical manner, with each layer being a function of the layer that preceded it. For example, the second layer $\mathbf{h}^{(2)}$ (i.e., also the first hidden layer) can be expressed as a function of the first (i.e., input) layer (e.g., with input vector $\mathbf{x} \in \mathbb{R}^d$):

$$\mathbf{h}^{(2)} = g^{(2)}(\mathbf{W}^{(2)T} \mathbf{x} + \mathbf{b}^{(2)}), \quad (6)$$




Fig. 14. An architecture of CNN

where $\{\mathbf{W}^{(2)}, \mathbf{b}^{(2)}\}$ are the parameters and $g^{(2)}$ is the activation function (e.g., ReLU) for the second layer. We also call these layers fully connected or dense layers. Similarly, the l^{th} ($l > 2$) hidden layer can be written as

$$\mathbf{h}^{(l)} = g^{(l)}(\mathbf{W}^{(l)T}\mathbf{h}^{(l-1)} + \mathbf{b}^{(l)}). \quad (7)$$

Assuming there are a total number of L layers, then an output layer linearly transforms the previous hidden units $\mathbf{h}^{(L-1)}$, followed by a softmax function for probability scaling:

$$p(\mathbf{y}|\mathbf{h}^{(L-1)}) = \text{softmax}(\mathbf{W}^{(L)T}\mathbf{h}^{(L-1)} + \mathbf{b}^{(L)}), \quad (8)$$

Note the model parameters include the layer-wise weight matrices and bias vectors, i.e., $\boldsymbol{\theta} = \{\mathbf{W}^{(l)}, \mathbf{b}^{(l)}\}_{l=2}^L$, which needed to be estimated by minimising the loss function (e.g., log loss). DNN has been widely used for various IoT applications, including wearable-based activity recognition [336], traffic congestion prediction[88], healthcare[248], etc.

Convolutional Neural Networks (CNNs). Similar to the DNN, CNN also comprises multiple hidden layers. It learns to map from 2D data (e.g., images) at the input, output the class probabilities at the output. DNN usually contains three types of hidden layers: convolution layer and pooling layer at each stage and FC/dense layer right before the output layer. In a convolution layer, local patterns can be extracted by learning several convolutional kernels/filters (with a predefined size, such as 3×3), each of which has the same shared weights to be estimated (via back-propagation). This weight-sharing scheme makes CNN effective in dealing with high-dimensional data (such as high-resolution images). After convolution operation with a number of kernels/filters, the corresponding feature maps can be formed, and then a pooling layer can be employed for down-sampling these feature maps for more compact representation. In CNN, normally there are multiple convolution/pooling layers, before FC/dense layers can be applied; Fig. 14 shows an example architecture of CNN. Since cameras are one of the major parts of the IoT environment, CNN can be a very useful tool and has been widely used such as in traffic sign detection on autonomous driving [305].

Recurrent Neural Networks (RNNs). RNN is a type of deep model and it is designed to model and recognise sequential data (e.g., time-series data). Compared with DNN, a hidden unit of RNN can feed-forward on itself in the next timestamp and thus can memorise the temporal information for sequential inference. Fig. 15 shows how to extend an ANN (DNN with one hidden layer) to a three-layer vanilla RNN by modelling the temporal information, we use an arrow to represent FC/dense layer for better visual effects. Compared with ANN, RNN has an additional set of parameters, i.e., the hidden-to-hidden transformation matrix \mathbf{W}_{hh} to be estimated, and the full parameters can be expressed as:




Fig. 15. From ANN to RNN

$\Theta = \{\mathbf{W}_{xh}, \mathbf{W}_{hh}, \mathbf{W}_{ho}, \mathbf{b}_o, \mathbf{b}_h\}$. For \mathbf{x}_t , the input vector at the t^{th} timestamp, the feed-forward pass can be written as:

$$p(y|\mathbf{x}_t) = \text{softmax}(\mathbf{W}_{ho}^T \mathbf{h}_t + \mathbf{b}_o)$$

$$\text{where } \mathbf{h}_t = \tanh\left(\begin{bmatrix} \mathbf{W}_{xh} \\ \mathbf{W}_{hh} \end{bmatrix}^T [\mathbf{x}_t, \mathbf{h}_{t-1}] + \mathbf{b}_h\right), \quad (9)$$

It is clear that the current hidden state \mathbf{h}_t is calculated based on the previous hidden state \mathbf{h}_{t-1} and the current signal \mathbf{x}_t , while the output layer remains the same as DNN (see Eq.(8)). We can see RNN is an end-to-end method where both the sample-wise features \mathbf{h}_t and predictions $p(y|\mathbf{x}_t)$ can be learned simultaneously.

For complex time-series data analysis, we need a deeper RNN (i.e., larger window, multiple layers, etc.) to capture the high-level temporal/contextual information. Yet vanilla RNN has a gradient explosion/vanishing problem owing to its numerical properties. In the 1990s, a complex hidden unit named long short term memory (LSTM) [147] was proposed for large-scale recurrent neural network construction, which can preserve the error/gradient that can be back-propagated effectively through time and layers. LSTM includes four different gates organised in a special internal structure, in contrast to tanh function for activation in vanilla RNN. Accordingly, it also has four sets of specific gate parameters to be estimated.

RNNs achieve great performance in time-series applications, for example machine translation and speech recognition. In IoT applications, it has shown its great performance in sensor-based power station condition prediction [387], wearable-based activity recognition [130], etc.

Auto-encoder (AE). An Auto-encoder(AE) is a type of unsupervised neural network, and it can transform the data into latent code/representation/feature (e.g., in lower-dimensionality), from which the original data can be reconstructed. AE is widely used for data compression and feature extraction, and a typical AE includes two parts: encoder, and decoder. Fig. 16 shows a three-layer encoder-decoder structure for AE, and the d -dimensional input vector \mathbf{x} can be transformed into the latent representation \mathbf{h} via a linear transformation followed by an activation function $g(\cdot)$, i.e., $\mathbf{h} = g(\mathbf{W}_{\text{encoder}}^T \mathbf{x} + \mathbf{b}_0)$. Similarly, decoder can be used to reconstruct the data such that $\mathbf{x}' = g(\mathbf{W}_{\text{decoder}}^T \mathbf{h} + \mathbf{b}_1)$. Note $\Theta = \{\mathbf{W}_{\text{encoder}}, \mathbf{b}_0, \mathbf{W}_{\text{decoder}}, \mathbf{b}_1\}$ are the model parameters that can be estimated via minimising a reconstruction loss

$$\underset{\Theta}{\text{argmin}} L(\mathbf{x} - \mathbf{x}'). \quad (10)$$




Fig. 16. The encoder-decoder structure of a three-layer Auto-encoder

Note in Eq. (10), $L(\cdot)$ can be very flexible, and some popular ones include mean squared error, log loss (for binary input), etc. It is worth noting that with small dataset, a popular constraint is to set $W_{\text{encoder}} = W_{\text{decoder}}^T$, which can limit the degree of freedom of this model with better generalisation capabilities. AE can be used as a powerful tool for dimensionality reduction. Compared with PCA, the non-linear transformation may extract more discriminant information for classification. It can also be used for anomaly detection, and the intuition is that the trained AE (based on normal class) cannot reconstruct the abnormal class well, yielding large reconstruction errors. There are many IoT applications for AE, such as fault diagnosis in hardware devices [65], and anomaly detection in the performance of assembly lines [250].

Generative Adversarial Network (GAN). Generative Adversarial Networks (GAN) is a generative model with adversarial architecture. Fig. 17 shows the basic network architecture of a GAN, which contains a generator and a discriminator.

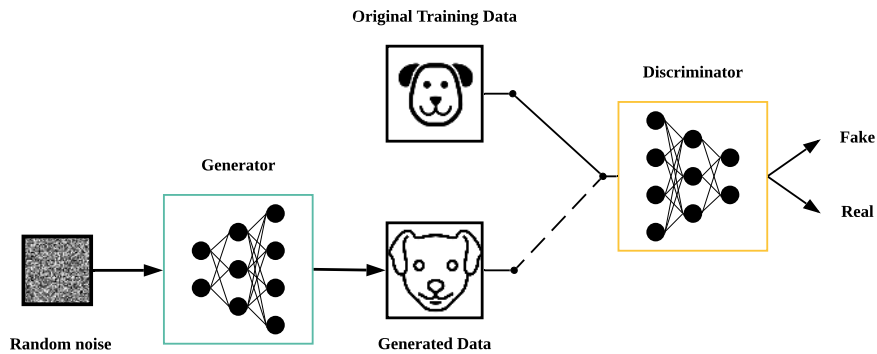


Fig. 17. Generative Adversarial Network Architecture

Table 5. Summary of Deep Machine Learning Models

Method	Learning Type	Category	Input data Type	Characteristics	IoT Application
CNNs	Supervised	Discriminative	2-D (image, sound, etc.)	<ul style="list-style-type: none"> • Mainly used on image processing • Less connection compared to DNNs. • Require large training samples. 	<ul style="list-style-type: none"> • Traffic sign detection • Plant disease detection • Bridge crack detection
RNNs		Discriminative	Sequential data	<ul style="list-style-type: none"> • Mainly used to analyze sequential data • Useful in IoT applications with time-dependent data 	<ul style="list-style-type: none"> • Identify movement pattern • Behavior detection • Human activity recognition • Mobility prediction
AEs	Unsupervised	Generative	Various	<ul style="list-style-type: none"> • Mainly used for feature extraction, and dimensionality reduction • Optimized by reconstructs input data • Can be used on unlabeled data 	<ul style="list-style-type: none"> • Emotion recognition • Machinery fault diagnosis • Intrusion detection • Failure detection
GANs		Generative	Various	<ul style="list-style-type: none"> • Learn data distribution • Can be used as a data generation tool • Two part networks: a generator and a discriminator 	<ul style="list-style-type: none"> • Localization and way finding • Image to text

The generator aims at generating indistinguishable samples compared to the real data while the discriminator works adversarially to distinguish the generated fake samples from the real data. It is an iterative competition process that will eventually lead to a state where the generated samples are indistinguishable from the real data. The generative model aims to learn the data distribution and generate data samples for those domains that lack data samples.

The decent idea of applying adversarial structure on deep neural networks brought impressive results related to content generated, and realistic visual content. Moreover, it can be applied to many applications such as image restoration, style transfer, sample generation. However, there are three main drawbacks on GAN frameworks. The first one is that the network is hard to train and the training process is destabilized (i.e., “Non-convergence”). The second drawback is that the generator will tend to produce limited varieties of samples at the end of a training period (i.e., “Mode collapse”). The last drawback of GAN is the gradient diminishing problem that occurs when the balance between generator and discriminator were broken.

GANs have been recently implemented in the IoT environment, mainly on IoT security. For example, a previous study [391] proposed a GAN based framework for improving open-categorical classification on individual identity authentication application. Moreover, GAN has been used as a tool to generate large datasets that do not need manual annotation. For example, a study[390] explored the feasibility of using GAN to generate illegal Unmanned Aerial Vehicles (UAVs) dataset and obtain a better classification model with better accuracy. As the data generated from sensors may be unlabelled, GANs may have more potential applications in the IoT environment.

B.3 Reinforcement Learning (RL) methods

The goal of a reinforcement learning agent is to find an optimal policy to maximize the expected sum of future rewards $J(\theta)$ parameterized by θ . At each time step t , reward $r_t = r(a_t, s_t)$ is given when an agent takes an action $a_t \in A$ at state $s_t \in S$.

$$\arg \max_{\theta} J(\theta) = \mathbb{E}_{t \sim p_{\theta}(t)} \left[\sum_t r(s_t, a_t) \right] \quad (11)$$

In the equation, $p_{\theta}(t)$ represents the interactions between the RL agent and outer environments. The trajectory depends on two factors: the agent policy and the environment dynamic. Agent takes actions based on its policy $\pi_{\theta}(a_t|s_t)$. The dynamic transition for the environment $s_t \times a_t \rightarrow s_{t+1}$ can be expressed as $p(s_{t+1}|s_t, a_t)$ and is usually unknown.

Overall, the whole trajectory $p_\theta(t)$ can be represented as:

$$P_\theta(s_1, a_1 \dots s_T, a_T) = p(s_1) \prod_{t=1}^T \pi_\theta(a_t|s_t) p(s_{t+1}|s_t, a_t) \quad (12)$$

It is clear from the above equation that, if we are smart on the agent policy $\pi_\theta(a_t|s_t)$ or the transition dynamic $p(s_{t+1}|s_t, a_t)$, we can find the best trajectory $p_\theta(t)$ which maximizes reward at each time step. The RL community has formalized these two approaches as *Model-Free* and *Model-Based* approaches based on the fact that the later one directly models the system dynamic transitions.

Fig. 8 gives a broad view of the reinforcement learning world. The reinforcement learning method can be divided into *Model-free* methods and *Model-based* methods with correlations combining the features of both methods. The model-free approach can be subdivided into *Policy based*, *Value based* and *Actor-Critic* approaches according to the ways that the best policy is generated. In this section, we will detail both approaches with simple mathematical expressions and list the most famous algorithms. We will also cover some works that take advantage of both methods.

Model-Free Methods. They learn a policy and decide the best action to take given a certain state. It can be categorized as *policy-gradient*, *value-based* and *actor-critic* methods, based on how the policy is generated.

Value-Based methods learn a value function to estimate the “goodness” $V(s_t)$ for reaching a certain state s_t , or the “goodness” $Q(s_t, a_t)$ for taking certain action a_t given the state s_t . Hereby the “goodness” function estimates the sum of future rewards from current state s_t till the end s_T (given a finite trajectory). At each step, the agent chooses the action with the highest score based on the estimated value function. *Value-based* approach is deterministic, and may not be sufficient to solve complex problems. A list of the most prevalent algorithms includes Q-Learning [352], DQN [238], Prioritised Experience Replay [290], Dueling DQN [351], Double DQN [338] and Retrace [244].

Policy-Gradient methods provide an attractive paradigm by directly maximising $J(\theta)$ (Equa 11) with respect to the parameters θ of the policy $\pi_\theta(a|s)$. The gradient with respect to the parameters θ can be derived as

$$\nabla_\theta J = \mathbb{E}_\theta \left[\sum_t \nabla_\theta \log \pi_\theta(a_t|s_t) (R_t - b_t) \right] \quad (13)$$

$R_t = \sum_{t'=t} \gamma^{t'-t} r(s_{t'}, a_{t'})$ with a discounting factor γ emphasizing the agent more on recent rewards. To reduce the policy variance given different datasets, a baseline b_t that does not depend on future states or actions is subtracted. In practice, the expected future return is sampled and aggregated within a trajectory. Several works have been proposed based on this paradigm to either reduce the policy variance [291], increase scalability [142] or reduce sample complexity [292].

Actor-Critic algorithms [179] are similar to the policy-gradient approach in updating the policy, while an estimated value function $V(s_t)$ is leveraged in place of the constant baseline b in the original equation 13. The term $(R_t - b_t)$ is thus an estimate of the *Advantage* defined as $A(a_t, s_t) = Q(a_t, s_t) - V(s_t)$ where a_t is the action and s_t denotes the current state, with R_t estimating $Q(a_t, s_t)$. Actor-critic methods experience much lower variance with the policy π and value function V seen as actors and critics respectively. Works utilising this architecture benefit from both the policy-gradient and value-based methods, with DDPG [204] combining deep Q-learning for continuous action space, A3C [237] for concurrent training. Other prevalent works include TD3 [111], Soft-Critic SAC [133].

Model-based Methods. The model-based algorithms differ from the model-free methods in that the latter cares less about the environment’s inner working and the rewards are estimated through sampling. On the contrary, a model-based approach focuses on the model to predict the next state at each time step. Model-based RL achieves good sample

efficiency by learning the transition dynamics of the environments directly. During learning, sample trajectories are collected and trained with supervised learning. Several different approaches have been applied to study the dynamics, Gaussian Process approaches [85][177][30], Time-varying linear models approaches [194][208] or Deep networks [246][113][68].

C APPENDIX C

Metrics	Formula	Type	Evaluation Focus
Accuracy(acc)	$\frac{tp+tn}{tp+fp+tn+fn}$	Binary	This metric measures the correct percentage of the total samples
	$\frac{\sum_{i=1}^l \frac{tp_i+tn_i}{tp_i+tn_i+fp_i+fn_i}}{l}$	Multi class	Average accuracy for all classes
Error Rate(err)	$\frac{fp+fn}{tp+fp+tn+fn}$	Binary	This metric measures the miss-classification percentage over evaluated samples
	$\frac{\sum_{i=1}^l \frac{fp_i+fn_i}{tp_i+tn_i+fp_i+fn_i}}{l}$	Multi class	Average Error Rate of all classes
Precision(p)	$\frac{tp}{tp+fp}$	Binary	Precision measures correct classified positive samples in a positive class
	$\frac{\sum_{i=1}^l \frac{tp_i}{tp_i+fp_i}}{l}$	Multi class	Average of precision on each class
Recall(r)	$\frac{tp}{tp+fn}$	Binary	Recall measures the fraction of correct classified positive samples
	$\frac{\sum_{i=1}^l \frac{tp_i}{tp_i+tn_i}}{l}$	Multi class	The average of recall for each class
F1-Score(FS)	$\frac{2*p*r}{p+r}$	Binary	This metric measures the harmonic mean of recall and precision
	$\frac{2*p_M*r_M}{p_M+r_M}$	Multi class	The average F1-Score
Geometric Mean(GM)	$\sqrt{tp * tn}$	Binary	This metric is similar to the F1-Score but aims to maximize the tp rate and tn rate

Table 6. Evaluation Metrics for Classification Problem.
Note: i indicates the class C_i , and M donates the macro-averaging.

Metric	Formula	Evaluation Focus
Mean Squared Error	$\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y})^2$	MSE calculates the square residual for every data point, contribute more to the outlier
Mean Absolute Error	$\frac{1}{N} \sum_{i=1}^N y_i - \hat{y} $	MAE calculates the absolute residual for every data point, so that negative and positive residuals do not cancel out
Mean Absolute Percentage Error	$\frac{1}{N} \sum_{i=1}^N \left \frac{y_i - \hat{y}}{y} \right $	Percentage equivalent of MAE
Mean Percentage Error	$\frac{1}{N} \sum_{i=1}^N \left(\frac{y_i - \hat{y}}{y} \right)$	MPE can help to check whether the model underestimates or overestimates

Table 7. Evaluation Metrics for Classification Problem. Note: y_i indicates the predicted value for i th sample, and \hat{y} denotes the ground truth value. N indicates the total number of samples.