

Machine Learning for Wireless Communications in the Internet of Things: A Comprehensive Survey

Jithin Jagannath^{† ‡}, Nicholas Polosky[‡], Anu Jagannath[‡],
Francesco Restuccia[†], and Tommaso Melodia[†]

[‡]ANDRO Advanced Applied Technology, ANDRO Computational Solutions, LLC, Rome, NY, 13440

[†]Department of Electrical and Computer Engineering, Northeastern University, Boston, MA, 02115

E-mail: {jjagannath, npolosky, ajagannath}@androcs.com

{jagannath.j, melodia, frestuc}@northeastern.edu

Abstract

The Internet of Things (IoT) is expected to require more effective and efficient wireless communications than ever before. For this reason, techniques such as spectrum sharing and dynamic spectrum access will soon become essential components of the IoT wireless communication process. In this vision, IoT devices must be able to not only learn to autonomously extract spectrum knowledge on-the-fly from the network but also leverage such knowledge to dynamically change appropriate wireless parameters (*e.g.*, frequency band, symbol modulation, coding rate, route selection etc) to reach the network's optimal operating point. Given that the majority of the IoT will be composed of tiny, mobile, and energy-constrained devices, traditional techniques based on *a priori* network optimization may not be suitable, since (i) an accurate model of the environment may not be readily available in practical scenarios; (ii) the computational requirements of traditional optimization techniques may prove unbearable for IoT devices. To address the above challenges, much research has been devoted to exploring the use of machine learning to address problems in the IoT wireless communications domain. The reason behind machine learning's popularity is that it provides a general framework to solve very complex problems where a model of the phenomenon being learned is too complex to derive or too dynamic to be summarized in mathematical terms.

This work provides a comprehensive survey of the state of the art in the application of machine learning techniques to address key problems in IoT wireless communications with an emphasis on its ad hoc networking aspect. First, we present extensive background notions on machine learning techniques. Then, by adopting a bottom-up approach, we examine existing work on machine learning for the IoT at the physical, data-link and network layer of the protocol stack. Thereafter, we discuss directions taken by the community towards hardware implementation to ensure the feasibility of these techniques. Finally, we provide a series of research challenges associated with the applications of machine learning techniques for IoT wireless communications.

Keywords: Machine learning, internet of things, wireless ad hoc network, spectrum sensing and classification, medium access control, and routing protocol.

1. Introduction

The Internet of Things (IoT) will soon become the most pervasive technology worldwide. In the next few years, cars, kitchen appliances, televisions, smartphones, utility meters, intra-body sensors, thermostats, and almost anything we can imagine will be accessible from anywhere on the planet [1]. The revolution brought by the IoT has been compared to the building of roads and railroads during the Industrial Revolution of the 18th to 19th centuries [2] – and is expected to radically transform the education, health-care, smart home, manufacturing, mining, commerce, transportation, and surveillance fields, just to mention a few [3].

As the IoT gains momentum in every aspect of our lives, the demand for wireless resources will accordingly increase in an unprecedented way. According to the latest Ericsson’s mobility report, there are now 5.2 billion mobile broadband subscriptions worldwide, generating more than 130 exabytes per month of wireless traffic [4]. Moreover, over 50 billion devices are expected to be in the IoT by 2020, which will generate a global network of “things” of dimensions never seen before [5]. Given that only a few radio spectrum bands are available to wireless carriers [6], technologies such as radio frequency (RF) spectrum sharing through beamforming [7, 8, 9], dynamic spectrum access (DSA) [10, 11, 12, 13, 14] and anti-jamming technologies [15, 16, 17] will become essential in the near future. These technologies usually require coordination among wireless devices to optimize spectrum usage – often, they need to be implemented in a distributed manner to ensure scalability, reduce overhead and energy consumption. To address this challenge, machine learning (ML) has been widely recognized as the technology of choice for solving classification or regression problems for which no well-defined mathematical model exists.

The recent introduction of ML to wireless communications in the IoT has in part to do with the new-found pervasiveness of ML throughout the scientific community at large, and in part to do with the nature of the problems that arise in IoT wireless communications. With the advent of advances in computing power and ability to collect and store massive amounts of data, ML techniques have found their way into many different scientific domains in an attempt to put both of the aforementioned to good use. This concept is equally true in wireless communications. Additionally, problems that arise in wireless communication systems are frequently formulated as classification, detection, estimation, and optimization problems; all of which ML techniques can provide elegant and practical solutions to. In this context, the application of ML to wireless communications seems almost natural and presents a clear motivation [18, 19, 20].

The objective of this paper is to provide a detailed insight on the influence ML has had on the IoT and the broader context of wireless ad hoc networks (WANETs). Our hope is to elicit more research in the field to solve some of the key challenges of modern IoT communication systems. To begin, we provide an overview of the ML techniques in Section 2. In Sections 3 and 4, we discuss the applications of ML to physical layer to improve the communication and acquire signal intelligence respectively. Next, in Section 5, we discuss how ML has been exploited to advance protocol design at the data-link and network layers of the protocol stack. In Section 6, we discuss the implications of hardware implementations in context of ML. Finally, in Section 7, we discuss the open problems that may be currently hindering the application of ML in IoT wireless systems.

2. Overview of Machine Learning Techniques

Before we begin, we would like to introduce some standard notations that will be used throughout this paper. We use boldface upper and lower-case letters to denote matrices and column vectors, respectively. For a vector \mathbf{x} , x_i denotes the i -th element, $\|\mathbf{x}\|$ indicates the Euclidean norm, \mathbf{x}^\top its transpose, and $\mathbf{x} \cdot \mathbf{y}$ the Euclidean inner product of \mathbf{x} and \mathbf{y} . For a matrix \mathbf{H} , H_{ij} will indicate the (i,j) -th element of \mathbf{H} . The notation \mathcal{R} and \mathcal{C} will indicate the set of real and complex numbers, respectively. The notation $\mathbb{E}_{x \sim p(x)} [f(x)]$ is used to denote the expected value, or average of the function $f(x)$ where the random variable x is drawn from the distribution $p(x)$. When a probability distribution of a random variable, x , is conditioned on a set of parameters, θ , we write $p(x; \theta)$ to emphasize the fact that θ parameterizes the distribution and reserve the typical conditional distribution notation, $p(x|y)$, for the distribution of the random variable x conditioned on the random variable y . We use the standard notation for operations on sets where \cup and \cap are the infix operators denoting the union and intersection of two sets, respectively. We use $S_k \subseteq S$ to say that S_k is either a strict subset of or equal to the set S and $x \in S$ to denote that x is an element of the set S . \emptyset is used to denote the empty set and $|S|$ the cardinality of a set S . Lastly, the convolution operator is denoted as $*$.

2.1. Introduction to Machine Learning

The primary purpose of this section is to provide a brief overview of the field of ML itself as well as provide a fundamental description of the algorithms and techniques presented as solutions to the wireless communications problems introduced in subsequent sections. This section aims to be as rigorous as necessary to allow the reader to understand how the presented algorithms are applied to wireless communications problems but does not aim to give an all-encompassing, comprehensive survey of the field of ML. Interested readers are urged to refer to [21], [22] and [23] for a comprehensive understanding of ML. The material presented in this section is given from a probabilistic perspective, as many of the concepts of ML are rooted in probability and information theory. The rest of this Section 2.1 provides a kind of road map for the Section 2 as a whole.

2.1.1. Taxonomy

Most introductory texts in ML split the field into two subdivisions: supervised learning and unsupervised learning. We follow suit and will make the distinction of which subdivision each presented algorithm falls under. As will be shown in later sections of this paper, many problems in WANET can be solved using an approach called reinforcement learning (RL). RL in its most fundamental form can be viewed as a third and separate subdivision of ML, thus we will denote representative algorithms as such. It is important to note that many advanced RL algorithms incorporate techniques from supervised and unsupervised learning yet we will still denote these as RL algorithms.

Another common type of learning discussed in ML literature is that of deep learning (DL). We view DL techniques not as a separate subdivision of ML but as a means to achieve the ends associated with each of the three subdivisions stated above. DL typically refers to the use of a deep neural network (DNN), which we present with more rigor later in Section 2.2.4. Thus the “Deep”

qualifier denotes an algorithm that employs a deep neural network to achieve the task. (ex: A deep reinforcement learning (DRL) algorithm would use a DNN in a RL framework)

2.1.2. A Note on Modularity

The concept of modularity is pervasive throughout engineering disciplines and is certainly prevalent in communications. We adopt this precedent throughout this text and present each of the algorithms using a common learning algorithm framework. This framework is primarily composed of the model, the optimization algorithm, the loss function, and a data set.

At its core, a learning algorithm is any algorithm that learns to accomplish some goal given some data to learn from. A common formalism of this definition is given [24]: “A Computer program is said to learn from experience E with respect to some class of tasks T and performance measure P , if its performance at tasks T , as measured by P , improves with experience E .” While this definition of a learning algorithm is commonly agreed upon, formal definitions of a task, experience, and performance measure are less endemic within the ML community, thus we provide examples of each.

In the context of ML, tasks usually define some way of processing an object or data structure. A classification task is the process of assigning a class label to an input object or data structure. While different **examples**(*objects*) within the data set will give rise to different class labels, the task of assigning a given example a label is the same for the entire data set. Other examples of tasks addressed in this text include regression (assigning a real value to an example) and structured output (assigning a separate data structure with a pre-defined form to an example).

The performance measure, P , essentially defines the criteria by which we evaluate a given learning algorithm. In the case of classification, the performance is typically the accuracy of the algorithm, or how many examples the algorithm assigns the correct class label to divided by the total number of examples. It is common practice to divide the entire available data set into two separate data sets, one used for training the algorithm and one used to test the algorithm. The latter, called the test set, is kept entirely separate from the algorithm while training and is used to evaluate the trained algorithm. The performance measure is often a very important aspect of the learning algorithm as it will define the behavior of the system.

The experience, E , that a learning algorithm has while learning essentially characterizes the algorithm into one of the three subdivisions defined earlier. Supervised learning algorithms are provided with a data set that contains examples and their associated labels or targets. An unsupervised learning algorithm experiences data sets containing only examples and attempts to learn the properties of the data set. RL algorithms experience examples produced by the environment with which they interact. The environment often provides feedback to the RL algorithm along with examples.

2.2. Supervised Learning

2.2.1. Overview

Recall from the previous discussion that in a supervised learning setting the learning algorithm experiences a data set containing examples and their respective labels or targets. An example will typically be denoted as x and its label or target as y . Together, we have training examples $(x, y) \in D$ existing in our data set D . In supervised learning problems, we attempt to learn to

predict the label y from the example x , or equivalently, estimate the conditional distribution $p(y|x)$. Taking this approach, we will want to obtain a model of the conditional distribution and we will denote the parameters of such a model as θ . Assuming a set of i.i.d data $D = \{x_1, x_2, \dots, x_n\}$ drawn from the data generating distribution $p_{data}(x)$, the maximum likelihood estimator of the parameters, θ , of a model of the data generating distribution is given as,

$$\theta_{ML} = \underset{\theta}{\operatorname{argmax}} p_{model}(D; \theta) = \underset{\theta}{\operatorname{argmax}} \prod_{i=0}^n p_{model}(x_i; \theta) \quad (1)$$

where p_{model} is a function space of probability distributions over the parameters θ . To make the above more computationally appealing, we can take the logarithm on both sides, as this does not change the optimization problem, which gives us,

$$\theta_{ML} = \underset{\theta}{\operatorname{argmax}} \sum_{i=0}^n \log(p_{model}(x_i; \theta)) \quad (2)$$

Additionally, we can divide the right hand side of the equation by n , as this does not change the optimization problem either, and we obtain the expectation of the log-probability of the model over the empirical data generating distribution,

$$\theta_{ML} = \underset{\theta}{\operatorname{argmax}} \mathbb{E}_{x \sim \hat{p}_{data}} \log(p_{model}(x; \theta)) \quad (3)$$

Alternatively, we could formulate the maximum likelihood estimation as the minimization of the KL divergence between the empirical data generating distribution and the model distribution given as,

$$D_{KL}(\hat{p}_{data} || p_{model}) = \mathbb{E}_{x \sim \hat{p}_{data}} [\log(\hat{p}_{data}(x)) - \log(p_{model}(x))] \quad (4)$$

Since the data generating distribution is not a function of the model, we can solve the same minimization problem by minimizing

$$- \mathbb{E}_{x \sim \hat{p}_{data}} \log(p_{model}(x)) \quad (5)$$

which is exactly equivalent to the maximization problem stated in the maximum likelihood formulation. The above is referred to as the negative log-likelihood of the model distribution and minimizing results in the minimization of the cross-entropy between the data generating distribution and the model distribution. The significance of this is two-fold. Firstly, the terms cross entropy and negative log-likelihood are often used in literature to describe the loss functions that are being used to evaluate a given ML model, and the above minimization problem is what is being referred to. Secondly, this gives rise to the narrative that the model associated with the maximum likelihood estimate is in fact the same model that most closely resembles the empirical data distribution. This is important considering what it is we want our model to do, namely, produce correct labels or targets for data drawn from the data generating distribution that the model has not seen before.

For completeness, the maximum likelihood estimator for the conditional distribution, which is what we want in order to predict a label from a given example, is given as,

$$\theta_{ML} = \underset{\theta}{\operatorname{argmax}} \sum_{i=0}^n \log(p_{\text{model}}(y_i|x_i; \theta)) \quad (6)$$

for i.i.d examples, x_i .

Often times some form of regularization on the parameters of the model is desirable, as regularization can lead to better generalization of the model. This is most frequently seen in the different types of neural network models that will be described later in this section. Building on the maximum likelihood perspective of the loss function, we can show that adding a regularization function to our optimization function can be seen as inducing a prior over the model parameters and subsequently changing our estimator to the maximum a posteriori (MAP) point estimate. Inducing a prior probability on the model parameter results in the following optimization problem,

$$\theta_{MAP} = \underset{\theta}{\operatorname{argmax}} p(\theta|D) = \underset{\theta}{\operatorname{argmax}} \log(p(D|\theta)) + \log(p(\theta)) \quad (7)$$

Here, we have made use of Bayes' Rule, the properties of logarithm, and the fact that the optimization problem does not depend on the data generating distribution. If we wish to put a Gaussian prior on the parameters, $p(\theta) \sim \mathcal{N}(0, \frac{1}{\lambda}I^2)$ we obtain a log prior proportional to $\lambda\theta^T\theta$, which yields the popular L2-Regularization scheme. Again, we have made use of the fact that the Gaussian prior does not depend on the data distribution and contains constants that do not affect the optimization problem. Thus, the L2-Regularizer can be seen as a cost associated with the magnitude of the model's parameters as well as the placement of a Gaussian prior on the model parameters.

2.2.2. Support Vector Machines

The support vector machines (SVMs) were initially developed to perform the task of binary classification. Since their introduction into the ML community, they have been successfully extended to perform regression and multi-class classification tasks as well. SVMs are non-parametric models, meaning that the number of parameters that compose the model is not fixed whilst constructing the model. In contrast, a parametric model would have a fixed number of tunable parameters defined before constructing the model. We will first define the SVM in the context of linear regression and then later in the section extensions to the algorithm will be expanded upon.

Linear regression is perhaps one of the most well known and prevalent linear classifiers known throughout the ML and statistical community. It is typically formulated as follows,

$$y_i = \mathbf{w}^T \mathbf{x}_i + w_0 \quad (8)$$

where y_i are the target values, \mathbf{x}_i are individual training examples and weights, \mathbf{w} , are the model parameters. A common approach to solving such a problem is to vectorize the output and input variables and solve the normal equations, giving a closed form solution for the minimum mean square error (MMSE). A typical approach to adapt this problem formulation for classification is the well known logistic regression given as,

$$p(y = 1|\mathbf{x}; \theta) = \sigma(\theta^T \mathbf{x}) \quad (9)$$

where σ is the logistic *sigmoid* function given as,

$$\sigma(x) = \frac{1}{1 + e^{-x}} \quad (10)$$

Here we have denoted the model parameters as θ . One favorable property of logistic regression is that it has a well defined probabilistic interpretation that can be viewed as maximizing the likelihood of the conditional distribution $p(y|x)$. An alternative formulation for a linear classifier is given in what is known as the perceptron algorithm [25]. The perceptron algorithm aims to find a hyperplane in the input space that linearly separates inputs that correspond to different classes. It does this using a zero-one loss function, meaning that the model is penalized equally for every point in the training data that it classifies incorrectly. The obvious problem with this formulation is that the algorithm converges to any hyperplane that separates the data; it need not be the *optimal* hyperplane.

The linear SVM [26], attempts to find the hyperplane that best separates the data, where the optimal hyperplane maximizes the margin between the nearest points in each class on either side of the plane. While this solution is better, the true power of SVMs come from the kernelization of the linear SVM, which allows the model to find nonlinear boundaries between different classes by representing the input data in a higher dimensional space. Kernelization of an algorithm is a process by which the parameters of the model are written in terms of a linear combination of the input vectors, which allows the computation of the inner product between a new input vector and the parameter vector of the model to be written as an inner product of the new input and the training inputs. A kernel function can then be substituted for the inner products between training vectors, which can be intuitively interpreted as a function that returns a real value which represents the similarity between two vectors. The kernelization of the SVM leads to the kernel SVM [27]. The most common kernels used to kernelize SVMs are the linear, polynomial, and radial basis function (RBF) kernels, given as,

$$k(x_i, x_j) = x_i^T x_j, \quad (11)$$

$$k(x_i, x_j) = (x_i^T x_j + 1)^d, \text{ and} \quad (12)$$

$$k(x_i, x_j) = e^{-\frac{(x_i - x_j)^2}{\sigma^2}} \quad (13)$$

respectively, where σ is a user defined parameter.

2.2.3. Decision Trees

Decision trees can be employed for both the tasks of classification and regression. ML decision tree algorithms are similar to nearest neighbor type algorithms in a sense that labels for examples lying near each other in input space should be similar; however they offer a much lighter weight solution to these problems.

A decision tree is essentially nothing more than an aggregation of if conditions that allow a new example to traverse the tree. The tree is traversed until happening upon a leaf node, which would specify the output label. Decision trees can be constructed in a number of different ways, but a common approach is to create trees that minimize some measure of impurity while splitting

the data. There are many such impurity measures but each of them essentially conveys how diverse the data in either child node would be if a given split of the data were to occur. If all of the training examples in either child node have the same label, that child node is called a pure leaf. Decision trees are often constructed to have pure leaves.

We now discuss two of the most popular impurity functions used in decision tree construction. We first define the training data as $S = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$, $y_i \in \{1, \dots, c\}$ where c is the number of classes. Additionally, we have $S_k \subseteq S$ where $S_k = \{(\mathbf{x}, y) \in S : y = k\}$ and $S = S_1 \cup \dots \cup S_c$. We then define the fraction of inputs in S with label k as,

$$p_k = \frac{|S_k|}{|S|} \quad (14)$$

and the Gini Impurity of a leaf node and a tree, respectively as,

$$G(S) = \sum_{k=1}^c p_k(1 - p_k), \text{ and} \quad (15)$$

$$G^T(S) = \frac{|S_L|}{|S|} G^T(S_L) + \frac{|S_R|}{|S|} G^T(S_R) \quad (16)$$

where $S = S_L \cup S_R$, $S_L \cap S_R = \emptyset$. The idea is then to choose splits in the tree that minimize this measure of impurity. Another popular impurity function is the entropy function. The entropy of the tree has its derivation in using the KL-divergence between the tree label distribution and the uniform distribution to determine how impure it is. Leaving the derivation to the interested reader, we define,

$$H(S) = - \sum_k p_k \log(p_k), H(S) = \frac{|S_L|}{|S|} H(S_L) + \frac{|S_R|}{|S|} H(S_R) \quad (17)$$

as the entropy of a leaf and the tree respectively. While decision trees can be strong classifiers on their own, they often benefit from a technique called bagging. We omit the statistical derivation of the benefits of bagging and simply state the essence of bagging: by training many classifiers and considering the average output of the ensemble we can greatly reduce the variance of the overall ensemble classifier. Bagging is often done with decision trees as decision trees are not very robust to errors due to variance in the input data.

Perhaps the most popular bagged algorithm is that of the Random Forest. Random forests are bagged decision trees generated by the following procedure,

- Sample m datasets D_1, \dots, D_m from D with replacement.
- For each D_i train a decision tree classifier $h_j()$ to the maximum depth and when splitting the tree only consider a subset of features k .
- The ensemble classifier is then the mean output decision i.e.
$$h(x) = \frac{1}{m} \sum_{i=1}^m h_i(x)$$

The number of trees m can be set to any number that can be afforded by the designer or application of the algorithm. If d is the number of features in each training example, the parameter $k \leq d$ is typically set to $k = \sqrt{d}$.

2.2.4. Feedforward Neural Networks

The original formulation of feedforward neural networks was proposed in [28]. It can be seen as an extension to the previously mentioned perceptron algorithm with an element-wise nonlinear transition function applied to the linear classifier. This nonlinear transition function allows the hyperplane decision boundary to take a nonlinear form, allowing the model to separate training data that is not linearly separable. The formulation for a single layer is as follows,

$$y = \sigma(\mathbf{w}^T \mathbf{x} + b) \quad (18)$$

where \mathbf{x} is the training example input, y is the layer output, \mathbf{w} are the layer weights, b is the bias. One common approach to handling the bias is to add an additional parameter to the weight vector and append a 1 to the input vector. When a bias term is omitted this formulation can be assumed unless otherwise stated throughout the section.

The nonlinear transition function, σ , is also referred to the activation function throughout literature. This is often chosen from a handful of commonly used nonlinear functions for different applications. The most widely used activation functions are the following,

$$\sigma(z) = \frac{1}{1 + e^{-z}}, \quad (19)$$

$$ReLU = \max(0, z), \text{ and} \quad (20)$$

$$\tanh(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}} \quad (21)$$

Additionally, the RBF kernel function described earlier in Section 2.2.2 can be used as an activation function and doing so give rise to the radial basis function neural network (RBFNN) [29]. To increase the complexity of the model, and thus its ability to learn more complex relationships between the input features, network layers can be subsequently added to the model that accept the previous layer's output as input. Doing so results in a DNN. The function of the network as a whole $\phi(x)$ thus becomes,

$$\phi(\mathbf{x}) = \mathbf{W}^{(3)} \sigma(\mathbf{W}^{(2)} \sigma(\mathbf{W}^{(1)} \mathbf{x})) \quad (22)$$

where the weight matrices $\mathbf{W}^{(i)}$ are indexed according to the layer they belong to. Intuitively, this allows the first layer to learn linear functions between the input features, the second layer to learn nonlinear combinations of these functions, and the third layer to learn increasingly more complex nonlinear combinations of these functions. This formulation additionally gives rise to a nice graphical interpretation of the model, which is widely used in literature and given in Figure 1.

This graphical interpretation is also where the feedforward neural network gets its loose biological interpretation. Each solid line in Figure 1 denotes a weighted connection in the graph. The input, output, and hidden layers are denoted as such in the graph and a close up of one node in the graph is provided. This close up calls the single node a neuron, but it can equivalently be referred to simply as a unit in this text and throughout literature. The close up also shows the inputs to the neuron, the weighted connections from the previous layer, the weighted sum of inputs, and the activation value, denoted as a_i^{l-1} , w_{ik}^l , z_k^l , and a_k^l , respectively. Occasionally, a neuron employing a given activation function may be referred to as that type of unit in this text and throughout literature, i.e. a units with a *ReLU* activation function may be called a “*ReLU* unit”.

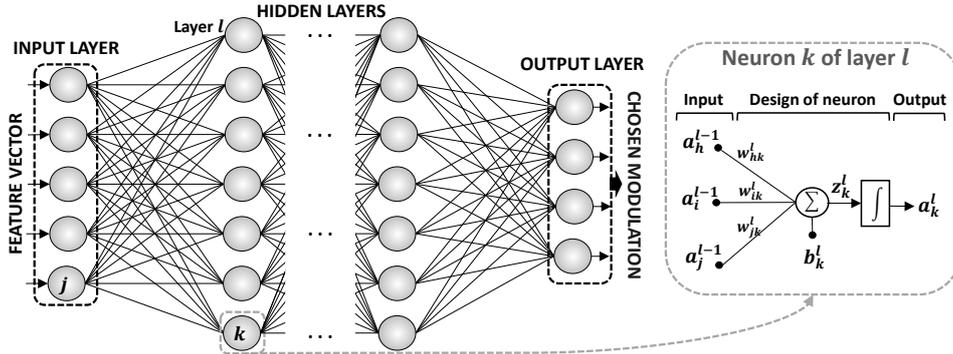


Figure 1: Standard Framework of Feed Forward Neural Network

The most common way to train most kinds of neural networks is the optimization method called stochastic gradient descent (SGD). SGD is similar to well known gradient descent methods with the exception that the true gradient of the loss function with respect to the model parameters is not used to update the parameters. Usually the gradient is computed using the loss with respect to a single training example or some subset of the entire training set, which is typically referred to as a mini-batch, resulting in mini-batch SGD. This results in the updates of the network following a noisy gradient, which in fact, often helps the learning process of the network by being able to avoid convergence on local minima which are prevalent in the non-convex loss landscapes of neural networks. The standard approach to applying SGD to the model parameters is through the repeated application of the chain rule of derivation using the famous back-propagation algorithm [30].

The last layer in a given neural network is called the output layer. The output layer differs from the inner layers in that the choice of the activation function used in the output layer is tightly coupled with the selection of the loss function and the desired structure of the output of the network. Generally, the following discussion of output layers and loss functions applies to all neural networks, including the ones introduced later in this section.

Perhaps the simplest of output unit activation functions is that of the linear output function. It takes the following form,

$$\hat{y} = \mathbf{W}^T \mathbf{h} + b \quad (23)$$

where \mathbf{W} is the output layer weight matrix, \mathbf{h} are the latent features output from the previous layer, and \hat{y} are the estimated output targets. Coupling a linear output activation function with a mean squared error loss function results in the maximizing the log-likelihood of the following conditional distribution,

$$p(y|\mathbf{x}) = N(y; \hat{y}, I) \quad (24)$$

Another task that we have already touched upon in our discussion of SVMs and perceptrons is that of binary classification. In a binary classification task, the output target assumes one of two values and thus can be characterized by a Bernoulli distribution, $p(y = 1|x)$. Since the output of a purely linear layer has a range over the entire real line, we motivate the use of a function that “squashes” the output to lie in the interval $[0, 1]$, thus obtaining a proper probability. We have seen that the logistic *sigmoid* does exactly this and it is in fact the preferred method to obtain a

Bernoulli output distribution. Accordingly, the output layer becomes,

$$\hat{y} = \sigma(\mathbf{W}^T \mathbf{x} + b) \quad (25)$$

The negative log-likelihood loss function, used for maximum likelihood estimation, of the above output layer is given as,

$$L(\theta) = -\log(p(y|x)) = f((1 - 2y)z) \quad (26)$$

where $f(x) = \log(1 + e^x)$ is called the *softplus* function and $z = \mathbf{W}^T \mathbf{x} + b$ is called the activation value. The derivation of (26) is not provided here but can be found in [21] for the interested reader.

In the case when the task calls for a multi-class classification, we want a Multinoulli output distribution rather than a Bernoulli output distribution. The Multinoulli distribution assigns a probability that a particular example belongs to a particular class. Obviously, the sum over class probabilities for a single example should be equal to 1. The Multinoulli distribution is given as the conditional distribution: $\hat{y}_i = p(y = i|\mathbf{x})$. It is important to note that the output, $\hat{\mathbf{y}}$, is now an n -dimensional vector containing the probability that \mathbf{x} belongs to class $i \in [0, n]$ at each index i in the output vector. The targets for such a classification task are often encoded as an n -dimensional vector containing $(n - 1)$ 0's and one 1, located at an index j which denotes that the associated training example belongs to the class j . This type of target vector is commonly referred to as a one-hot vector. The output function that achieves the Multinoulli distribution in the maximum likelihood setting is called the *softmax* function and is given as,

$$\text{softmax}(\mathbf{z})_i = \frac{e^{z_i}}{\sum_j e^{z_j}} \quad (27)$$

where z_j is the linear activation at an output unit j . *Softmax* output units are almost exclusively coupled with a negative log-likelihood loss function. Not only does this give rise to the maximum likelihood estimate for the Multinoulli output distribution but the log in the loss function is able to undo the exponential in the *softmax* which keeps the output units from saturating and allows the gradient to be well-behaved, allowing learning to proceed [31].

2.2.5. Convolutional Neural Networks

The convolutional neural networks (CNN) was originally introduced in [32] as a means to handle grid-like input data more efficiently. Input of this type could be in the form of a time-series but is more typically found as image-based input. The formulation of CNNs additionally has biological underpinnings related to the human visual cortex.

CNNs are very similar to the feedforward networks introduced previously with the exception that they use a convolution operation in place of a matrix multiplication in the computation of a unit's activation value. In this section, we assume the reader is familiar with the concept of the convolution operation on two continuous functions, where one function, the input function, is convolved with the convolution kernel. The primary differences from the aforementioned notion of convolution and convolution in the CNN setting are that the convolution operation is discretized (for practical implementation purposes) and that it is often truly the cross-correlation operation that is performed in CNNs rather than true convolution. This means that the kernel is not typically

flipped before convolving it with the input function. This is also primarily done for practical implementation purposes and does not typically affect the efficacy of the CNN in practice.

Convolution in the context of CNNs is thus defined as the following, for an input image I ,

$$S(i, j) = (K * I)(i, j) = \sum_m \sum_n I(m, n)K(i - m, j - n) \quad (28)$$

where K is the convolution kernel and the output, S , is often referred to as the feature map throughout literature. It is important to note that the above formulation is for two dimensional convolution but can be extended to input data of different dimensions. The entries of K can be seen as analogous to the weight parameters described previously (Section 2.2.4) and can be learned in a similar manner using SGD and the back-propagation (BP) algorithm. Intuitively, one can imagine having multiple K kernels in a single CNN layer being analogous to having multiple neurons in a single feedforward neural network layer. The output feature maps will be grid-like and subsequent convolutional layers can be applied to these feature maps after the element-wise application of one of the aforementioned nonlinear activation functions.

In addition to convolutional layers, CNNs often employ a separate kind of layer called pooling layers. The primary purpose of a pooling layer is to replace the output of the network at a certain location with a type of summarization of the outputs within a local neighborhood. Examples of pooling layers include max pooling [33], average pooling, L^2 norm pooling, and distance weighted average pooling. A max pooling layer would summarize some rectangular region of the input image by selecting only the maximum activation value present in the region as output from the pooling layer. Pooling layers improve the efficacy of CNNs in a few different ways. First, they help make the learned representation of the input invariant to small translations, which is useful when aiming to determine the presence of a feature in the input rather than its location. Second, pooling layers help condense the size of the network since convolutional layers don't inherently do so. A binary classification task taking image data with size $256 \times 256 \times 3$ will need to reduce the size of the net to a single output neuron to make use of the output layer and cost function pairs described previously in Section 2.2.4. Lastly, pooling layers lead to infinitely strong prior distributions making the CNN more statistically efficient [21].

Some common adaptations applied to CNNs come in the form of allowing information flow to skip certain layers within the network. While the following adaptations were demonstrated on CNNs and long short term memories (LSTMs) (a type of recurrent neural network (RNN)), the concepts can be applied to any of the networks presented in this paper. A residual network (RN), or ResNet [34], is a neural network which contains a connection from the output of layer, say L_{i-2} , to the input of the layer L_i . This connection allows the activation of the L_{i-2} to skip over the layer L_{i-1} such that a "residual function" is learned from layer L_{i-2} to layer L_i . A highway neural network [35], is similar in that it allows a skip connection over layers but additionally apply weights and activation functions to these connections. Lastly, a dense neural network [36] is a network that employs such weighted connections between each layer and all of its subsequent layers. The motivation behind each of these techniques is similar in that they attempt to mitigate learning problems associated with vanishing gradients [37]. For each of these networks, the BP algorithm used must be augmented to incorporate the flow of error over these connections.

2.2.6. Recurrent Neural Networks

The RNN was first introduced in [30] as a way to handle the processing of sequential data. These types of neural networks are similar to CNNs in the sense that they make use of parameter sharing. In RNNs, parameters are shared across time steps or indices in the sequential input. Recurrent nets get their name from the fact that they have recurrent connections between hidden units. We can denote this mathematically as follows,

$$\mathbf{h}^{(t)} = f(\mathbf{h}^{(t-1)}, \mathbf{x}^{(t)}; \theta) \quad (29)$$

where the function f could be considered the activation output of a single unit, $\mathbf{h}^{(i)}$ are called the state of the hidden units at a time i , $\mathbf{x}^{(i)}$ is the input from the sequence at the index i , and θ are the weight parameters of the network. Note, θ is not indexed by i , signifying that the same network parameters are used to compute the activation at all indices in the input sequence. Output layers and loss functions appropriate for the desired task are then applied to the hidden unit state \mathbf{h} .

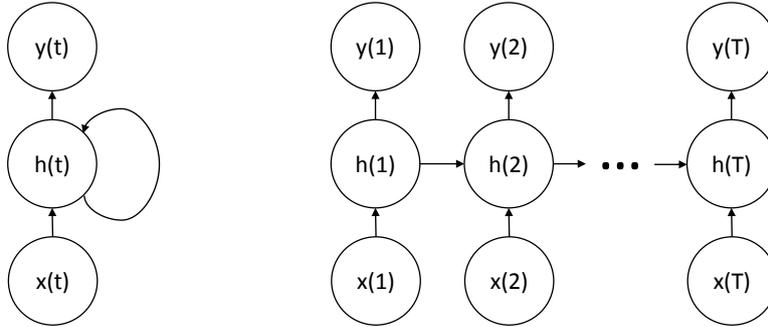


Figure 2: Equivalent graphical formulations for Recurrent Neural Networks

Two equivalent graphical representations of RNNs are provided as reference in Figure 2. The left representation shows the network “rolled up” with a recurrent connection onto itself. The right representation shows the network “unrolled” with the recurrent connections now propagating information forward in time. We now provide the forward propagation equations for the hidden unit and use the *softmax* output layer as an example of how the hidden state would be used as input to the output layer. The loss function can then be applied to the *softmax* output as previously discussed in the paper.

$$\begin{aligned} \mathbf{a}^{(t)} &= \mathbf{W}\mathbf{h}^{(t-1)} + \mathbf{U}\mathbf{x}^{(t)} + b \\ \mathbf{h}^{(t)} &= \tanh(\mathbf{a}^{(t)}) \\ \mathbf{o}^{(t)} &= \mathbf{V}\mathbf{h}^{(t)} + c \\ \hat{y}^{(t)} &= \text{softmax}(\mathbf{o}^{(t)}) \end{aligned} \quad (30)$$

The matrices \mathbf{W} , \mathbf{U} , and \mathbf{V} are the weight matrices shared across hidden units. They are used to weight the connections between hidden units from one time step to the next, between the input and hidden state at the current time step, and the hidden state and output at the current time step. The parameters b and c are biases that are also shared across time steps.

The loss for a single sequential training example is accumulated over the entire sequence, thus using a negative log-likelihood loss for a sequence $x^{(t)}$ with output targets $y^{(t)}$ the loss would be,

$$L(\{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(\tau)}\}, \{y^{(1)}, \dots, y^{(\tau)}\}) = - \sum_t \log(p_{model}(y^{(t)}|\{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(t)}\})) \quad (31)$$

The computation for the gradient of the loss with respect to the model parameters is an involved computations and is out of the scope of this paper. For the interested reader, SGD is used to train RNNs and the algorithm for the computation of the gradients is back-propagation through time (BPTT) [38].

Many extensions to above RNN model exist and are worth mentioning. Perhaps the most obvious extension is to add more recurrent layers following the single recurrent layer that was described above, resulting in Deep RNNs [39]. This provides similar advantages that were discussed in the motivation for extending feedforward networks to multiple layers. Another form of extension for RNNs is adding more recurrent connections. These additional recurrent connections could skip over time steps, skip over layers, or even move information backward in time giving bidirectional RNNs [40]. These additional recurrent connections would be weighted and a nonlinear activation function could be applied in the same manner that the basic recurrent connection operates.

The most prevalent extensions to the original RNNs are those of the LSTM and gated recurrent unit (GRU), developed originally in [41] and [42], respectively. LSTMs augment the traditional RNN framework by adding a self loop on the state of the network. This self loop is coupled with input, output, and forget gates which control whether input values are written to the state, the state values are forgotten within the state, or the state values are written to the output of the network, respectively. These adaptations allow the network to better “remember” relevant information over longer periods in time. Each of the gates are weighted and have a logistic *sigmoid* activation applied to them, allowing the network to learn how to best use these gates with respect to the task. GRUs operate in a similar fashion but use what are called update and reset gates. The update gate controls to what degree the state of the network at the given time step is written back to the state variable as well as what parts of the new state to write to the current state. The reset gates control what parts of the current state to use in the next computation of the new state. Both of the LSTM and GRU have the ability to retain information over longer time periods and mitigate negative learning effects associated with vanishing gradients.

Recurrent networks can also take forms that are significantly different from the models described above. In particular, a hopfield neural network (HNN) [43] is a special type of recurrent network formulated to recover corrupted patterns. More specifically it is a recurrent network where each unit is connected to all other units in the graph except for itself. Additionally, the weight between units is shared. Each unit in the network encodes a binary state value, typically either 1 or -1 . This formulation aims to mimic the forms of associative memory present in human cognition models is often trained using a form of Hebbian Learning [44]. The famous summarization of Hebbian learning “cells that fire together wire together” drives the idea that when part of the pattern that the HNN is trained to recognize is present, all of the units associated with that pattern will “fire” and the entire pattern will be represented by the network. Another interesting difference from the previously described RNN structures is that the HNN does not make use of any type of

training targets y . This makes the HNN a type of unsupervised learning algorithm, more of which we discuss in further detail in the next section.

2.3. Unsupervised Learning

Unsupervised learning, a separate learning paradigm from the previous described supervised learning, attempts to learn useful properties of the training data rather than learning to map inputs to specific outputs. Examples of unsupervised learning tasks include probability density estimation, denoising, and clustering. Unsupervised learning algorithms only experience the training data examples and are given no target outputs, which are obviously preferable in scenarios when data sets are produced without targets and it would be impractical for a human to go through and label the data set with a target value. Thus, without targets, unsupervised learning algorithms usually try to present the data set in a simpler or easier to understand representation. This simpler representation most commonly manifests itself in the form of lower dimensional representations of data, sparse representations of data, and independent representations of the data.

While some unsupervised learning algorithms draw techniques from previously mentioned supervised learning algorithms, they employ different types of loss functions. Usually the best types of loss functions to use in unsupervised learning settings will reward the algorithm for preserving information about the input data but penalize the algorithm for not representing the data in one of the three ways discussed in the previous paragraph. The reader may be familiar with the Principal component analysis (PCA) algorithm, which is a great example of a linear unsupervised learning algorithm that aims to decorrelate the input data.

2.3.1. Self Organizing Maps

The self-organizing maps (SOMs) [45] were originally introduced as a type of unsupervised learning algorithm with the goal of performing dimensionality reduction and data clustering. The reader may be familiar with the simple clustering algorithm referred to as k -means clustering, in which each example in the training data is required to belong to one of k different clusters. The obvious pitfall of this algorithm is that the designer of the algorithm must choose the parameter k prior to constructing the model, hence the models usefulness is contingent on the users estimate of the appropriate number of clusters. The SOM algorithms avoid this by learning the appropriate number of clusters. Additionally, the SOM algorithm typically aims to represent the training data as a two dimensional grid, where examples that are near each other in the input topological space are embedded near each other in the two dimensional latent representation.

The canonical SOM formulation can be viewed as a fully connected single layer feedforward neural network, with units arranged in a two dimensional grid. As the network sees each input, it computes the similarity between the input vector and each unit in the grid using some discriminant function similar to,

$$d_j(x) = \sum_{i=1}^D (\mathbf{x}_i - \mathbf{w}_{ji})^2 \quad (32)$$

where $d_j(\mathbf{x})$ is the value of the discriminant function at unit j , \mathbf{w}_j is the weight vector associated with unit j , and $i \in [1, D]$ indexes the D dimensional input and weight vectors. This is often called

the competitive process of SOMs as it is representative of a type of learning called competitive learning.

Once the discriminant function is computed at each unit for a training example the unit with the least value for the discriminant function is selected for what is called the cooperative process of SOM. The cooperative process attempts to update the neurons in some local neighborhood around the neuron that provides the closest representation of the input vector (i.e. the neuron with the minimal discriminant function). This creates neighborhoods in the map that will activate similarly for similar input values, thus creating clusters within the map. The topological neighborhood is usually defined as,

$$T_{j,I(\mathbf{x})} = \exp\left(\frac{-S_{j,I(\mathbf{x})}^2}{2\sigma^2}\right) \quad (33)$$

where $I(\mathbf{x})$ represents the index in the map where the minimal discriminant function occurred and $S_{j,i}$ denotes the distance from a neuron j to a neuron i . σ is a parameter chosen by the designer and is typically decayed over time using the following schedule for time-dependence,

$$\sigma(t) = \sigma_0 \exp\left(\frac{-t}{\tau_\sigma}\right) \quad (34)$$

Once the topological neighborhood is computed, the weight vectors associated with the units in the neighborhood are updated. This is usually referred to as the adaptive process in the context of SOMs. The change applied to the weight vectors is given as,

$$\delta \mathbf{w}_{ji} = \eta(t) T_{j,I(\mathbf{x})}(t) (\mathbf{x}_i - \mathbf{w}_{ji}) \quad (35)$$

where $\eta(t)$ is the learning rate parameter and is also decayed over time using a similar schedule to that of the σ parameter,

$$\eta(t) = \eta_0 \exp\left(\frac{-t}{\tau_\eta}\right) \quad (36)$$

This process is repeated many times for each training example in the training data set, resulting in the SOM.

2.4. Reinforcement Learning

RL is a learning paradigm that can be considered separate from supervised and unsupervised learning. That being said, RL techniques often use ideas and algorithms from both unsupervised and supervised learning. We first describe the problem formulation for RL and then present a solution and how it can be extended to include concepts from other learning paradigms [46].

RL is built on the idea of an agent performing actions within an environment, based on observations of the environment. The agent generally carries out actions according to a policy, which defines how the agent behaves at a given time. The agent receives reward signals, which define the ultimate goal of the algorithm, from the environment which indicates how well off the agent is at the time step the reward is given. The agent then aims to maximize its cumulative reward by observing its environment and the reward signal received, and then performing actions based on these inputs. The maximization of the cumulative reward is typically defined in terms of a value function. The value function differs from the reward signal in that the reward represents what is

a desirable immediate setting and the value function represents how much reward the agent can obtain in the future given the agents current state. Additionally, RL problems typically define a model of the environment. The model is estimated by the agent to determine the dynamics of the environment and is then subsequently used by the agent to devise some sort of plan about how to act.

RL problems, as described above, are usually formalized mathematically using finite markov decision process (MDP). The tuple $(S, A, P_a(\cdot), R_a(\cdot))$ defines the dynamics of the MDP as well as the state and action spaces, S and A . At a given time step, an agent observes a state s , chooses an action a , receives a reward r , and transitions to a new state s' . The functions $P_a(\cdot)$ and $R_a(\cdot)$ define the transition probabilities between states and reward received from the environment when transitioning to a new state. The transition probability function takes the current state, s , and a possible new state, s' and outputs the probability of transitioning to that new state, conditioned on an action, a . i.e.,

$$P_a(s, s') = Pr(S_{t+1} = s' | S_t = s, A_t = a) \quad (37)$$

R is reward function such that it gives the reward obtained directly after transitioning to state s' from state s via action a and is defined as,

$$R_a(s, s') = E [R_{t+1} | S_t = s, A_t = a] \quad (38)$$

A policy is a function which defines how the agent will act given the state it is currently in. The policy is usually denoted as $\pi(a|s)$. Using such a policy, the agent moves about the environment and can start to construct value functions and action-value function based on the return they observe. The action-value function, q , for a policy, π is given as,

$$q_\pi(s, a) = E_\pi \left[\sum_{k=0}^{\infty} \gamma^k R_{t+k+1} | S_t = s, A_t = a \right] \quad (39)$$

where R_t are the observed returns over time and γ is a scaling parameter that is used to weight future returns less heavily than immediate returns. The action-value function can be plainly stated as the expected return starting in a state s , taking the action a , and subsequently following the policy π . Obtaining values for state action pairs allows for the agent to plan how to act in its environment. Equipped with the optimal action-value function, the solution to the MDP is merely choosing the action with the greatest action value.

2.4.1. Q-Learning

The method of Q-Learning was introduced in [47] and is what is called an off-policy control algorithm. The off-policy qualifier simply denotes that the algorithm does not depend on the policy the agent uses to navigate the environment. The Q-Learning is defined by the following,

$$Q(S_t, A_t) \leftarrow Q(S_t, A_t) + \alpha [R_{t+1} + \gamma \max_a Q(S_{t+1}, a) - Q(S_t, A_t)] \quad (40)$$

Using such an update scheme for action-value pairs will lead to the approximation of the optimal action-value function independent of the policy being followed.

As one would imagine, the state and action spaces of some RL problem become extremely vast, making the storage of the action-value function for all state-action pairs impractical. One way to overcome this issue is to introduce a function approximation method which learns to provide values for state-action pairs. This function approximator could be one of the different types of ML algorithms discussed previously in this section. Such is the case in the famous deep Q-network (DQN) [48], where a deep convolutional neural network was used to approximate the action-value function when learning to play Atari games.

3. Machine Learning For Physical Layer

The signal processing techniques that enable the physical layer functionalities have a direct impact on the data rate and sensitivity of the radio. With the increasing number of IoT devices that communicate over networks, some of which stream multimedia data, there is a growing need for high speed, low latency, and higher capacity systems. IoT devices are often deployed densely with several devices interconnected and communicating in the same spectrum posing severe constraints on bandwidth. To enable communication in such dense IoT networks, several challenges such as interference, power and bandwidth constraints come into play. Adaptive signal processing is a well researched topic aimed around suppressing interference and noise from received attenuated signal samples by estimating the interference plus noise covariance from the received samples and suppressing their effect to improve the spectral efficiency of the system [49, 50, 51]. Another well-known approach to increase spectral efficiency is to adjust the modulation and coding scheme *on-the-fly* based on instantaneous channel conditions. The promising capabilities of multiple input multiple output (MIMO) systems to increase channel capacity has led to their adoption in wireless communication standards. Significant performance gain can be achieved by learning and estimating the varying channel dynamics and nullifying the channel' effect from the received signal samples to estimate the actual transmitted bits, in what is commonly known as adaptive channel equalization. Research surrounding the physical layer has historically been aimed at pushing the boundaries against the norm to provide increased agility to the radios, subsequently enhancing their performance. Enabling the radios with cognitive skills at the physical layer can revolutionize their wireless communication capability of the IoT devices. The ML based solutions can transform the radios into “intelligent radios” that can adaptively decide which actions are necessary to achieve a certain objective based on parameters learned by the system. This section will explore the various aspects of signal processing at the physical layer and how ML based solutions can offer a better alternative.

3.1. Adaptive Rate and Power control

RL based solutions have been extensively used in wireless communications to estimate the dynamic system model on the fly [52, 53, 54, 55, 56, 57, 58, 59, 60]. In the context of the physical layer, RL based solutions can extensively improve the system data rate, bit error rate, goodput (*i.e.*, the amount of useful information that successfully arrived at the destination over the time-varying channel) and energy efficiency [61, 62, 63, 64]. Adaptive rate control can serve as a useful tool to selectively adapt the data rate depending on the instantaneous channel conditions. Such flexibility aids the system in leveraging the channel statistics to its benefit, essentially maximizing

the channel utilization. IoT devices are often battery powered and hence constrained in power. Each layer of the protocol stack must be designed to reduce the energy consumption and prolong the device' lifetime. Therefore, adaptive power control at the physical layer is imperative to the longevity of the device and consequently the IoT network lifetime.

In [65], an adaptive rate control strategy based on RL is proposed to learn the dynamically varying channel conditions. The time-varying fading channel is modeled as a finite state Markov chain, whose channel state transition probabilities are unknown but the instantaneous channel gains can be estimated. Now the optimization problem forms a MDP which can be solved in dynamic programming (DP). However, the DP approach is suited best for static systems and hence would not be suitable for a dynamic system where the channel statistics vary with time.

In this work, the authors propose to use $Q(\lambda)$ -learning [46] to track the varying environmental changes in pursuit of the optimal control policy online. $Q(\lambda)$ -learning is a popular RL based algorithm used to solve MDP when the system's state transition probabilities are unknown. The $Q(\lambda)$ -learning algorithm is similar to the standard Q-learning except that it updates the learning rate based on the Q value of the state-action pair. The incremental learning process involves the learning agent transitioning from system state of one block to another at the next block by choosing an action. For each chosen action, the agent observes the reward and modifies its control policy with an aim to maximize the expected reward for future actions. This foresighted iterative learning process will repeat at each block and the agent will eventually converge at the optimal policy.

In this context, the objective is to find a rate-control scheme that maximizes the system throughput subject to a fixed bit error rate (BER) constraint and long-term average power constraint. The system state is characterized by the instantaneous channel gain and buffer occupancy as $s_n = \{g_n, b_n\}$. The receiver estimates the channel gain and feeds back to the transmitter. In a practical system, this could be accomplished by having the receiver and transmitter exchange estimated statistics via control packets using a separate control channel. Consider the transmission buffer is a first-in first-out (FIFO) queue that can hold a maximum of N packets each of size B bits. The packet arrival process to the buffer follows a Poisson distribution $P_a = \frac{\nu^a e^{-\nu}}{a!}$, where a is the number of packets that arrived at the buffer and ν is the expected number of packets that will arrive in one block. The number of packets dropped from buffer in the n^{th} block can be expressed as $d_n = \max[b_{n-1} - p_{n-1} + a_n - N, 0]$, where p_n is the number of packets leaving the buffer in the n^{th} block.

Consider an M-ary quadrature amplitude modulation (M-QAM) system which, based on the learning agent's rate-control policy, can change the number of bits per symbol ($\log_2(M)$). There are numerous ways to change a system's transmission rate; (i) vary coding rate, (ii) vary modulation scheme, i.e., constellation size, and (iii) careful combination of both. Let us denote the bits per symbol in the n^{th} block as $m_n = \{1, 2, 3, \dots, K\}$ and the number of symbols in a block as S . Then, the number of packets that can be transmitted in the n^{th} block is $p_n = \frac{m_n S}{B}$, referred to herein as rate. For a W bandlimited system operating in an additive white gaussian noise environment with noise spectral density N_0 , the minimum transmission power required to maintain an acceptable BER (ϵ_*) in the n^{th} transmission block is,

$$P_n \geq \frac{W N_0 (-\log 5\epsilon_*) (2^{p_n B/S} - 1)}{g_n \cdot 1.5} \quad (41)$$

Now, the long-term average power consumption can be expressed as,

$$\bar{P} = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^n P_i. \quad (42)$$

The rate control scheme must now aim to maximize the system throughput ($T = \nu(1 - P_d)$) subject to the BER and average power constraints. Here, P_d is the packet drop probability. This escalates to a dual objective optimization; maximizing system throughput and minimizing average power. This multi-objective optimization will be solved to arrive at a Pareto-optimal solution (rate control policy). Q(λ)-learning aims to find the optimal control policy by estimating an action-value function for each state-action pair. The action-value function is the long-term discounted reward if the system starts at state s_n taking an action p_n . The reward per block for taking an action/transmission rate p_n at a state s_n has a Lagrangian form which essentially implies the system gets a larger reward if the packet drops and transmission power is lower. The negative cost (reward) per block can be expressed as,

$$r_{n+1} = - [E(d_{n+1}) + \lambda P_n] \quad (43)$$

The Q(λ)-learning can be solved in a way similar to the standard Q-learning except here the learning rate (ρ) is updated based on the state-action pair which is kept a constant in the standard Q-learning.

$$\rho = r_{n+1} + \gamma Q(s_{n+1}, p_{s_{n+1}}^*) - Q(s_n, p_n) \quad (44)$$

where γ is the discount factor and $p_{s_{n+1}}^*$ is the action which maximizes the action-value function $Q(s_{n+1}, p_{s_{n+1}}^*)$. The Q(λ)-learning demonstrates faster convergence compared to the standard Q-learning. The authors demonstrated the ability of learning agent to acclimate to the varying wireless channel to learn and adapt the rate control policy best suited for the channel conditions.

The authors of [61] attempt to solve the link adaptation problem of single carrier frequency domain equalization (SC-FDE) systems. SC-FDE systems use cyclically prefixed M-QAM to allow frequency domain equalization at the receiver. The authors approach the problem from a classification perspective such that the optimum modulation and coding scheme that would deliver the highest goodput for the current channel conditions would correspond to the best classification of the multidimensional data. The feature vectors considered include estimated post-processing signal-to-noise-ratio (SNR), estimated channel coefficients, and noise variance. PCA is used for dimensionality reduction such that an orthogonal transformation maps the features from a higher dimensional space to lower dimension. The kNN algorithm is used to classify the reduced dimensional feature vectors. A significant drawback of using kNN algorithm is that it requires storing the previously observed values which is memory intensive and computationally expensive. For a low power wireless device, such an algorithm is a poor choice for real-time operations. [62] tackle this problem to perform real-time link adaptation of MIMO-orthogonal frequency-division multiplexing (OFDM) systems by using online kernelized support vector regression (SVR). SVR attempts to minimize the generalization error bound to achieve generalized performance rather than minimizing training error like SVM. SVR requires minimal memory and computational power and was demonstrated to adapt quickly to varying channel conditions in their simulations. For every

packet, the receiver observes the packet failure/success, channel measurements and the modulation and coding scheme corresponding to that packet. To prevent memory explosion, the authors use a sparsification algorithm [66] such that only linearly independent samples are preserved in the dictionary. The SVR algorithm finds the linear regression function that corresponds to the minimum mean squared loss function. The authors compared the performance of online kNN versus online SVR to demonstrate the monotonically increasing memory and time consumption with online kNN while it remained constant for online SVR.

A RL based solution is proposed in [63] to achieve adaptive rate and power control for point-to-point communication and extend it to a multi-node scenario. The receiver is assumed to feedback channel gain and packet success/fail status (acknowledgement (ACK)/negative acknowledgement (NACK)) to the transmitter allowing it to choose the modulation and transmitter power based on the obtained information. Accordingly, the authors formulate the objective to maximize the throughput per total consumed energy considering the channel conditions, queue backlog, modulation and transmit power. The authors incorporate buffer processing cost/energy into the total energy consumption cost such that there is a cost incurred for buffer overflows. Imposing buffer processing cost can be viewed as a quality of service (QoS) factor. The formulated MDP is solved using the Actor-Critic (AC) algorithm [46] which involves two parts: actor and critic. The actor decides the action and the critic estimates the state-value function and the error which criticizes the actor's action. The actor selects the action based on Gibbs *softmax* method [46] such that the action corresponding to the highest conditional probability of state-action is chosen. The authors demonstrated the throughput achieved with AC algorithm is twice that of a simple policy where the highest modulation order that maintains a predefined link SNR is chosen.

Another notable application of ML in improving real-time video streaming is presented in [67]. The authors propose Video Quality Aware Rate Control (QARC), a DL based adaptive rate control scheme to achieve high video quality and low latency. The complex challenge posed by the varying video quality and dynamic channel conditions is solved by two RL based models; a video quality prediction network (VQPN) and a video quality reinforcement learning (VQRL). The VQPN predicts future video quality based on previous video frames whereas the VQRL algorithm adopts the asynchronous advantage actor critic (A3C) RL [68] method to train the neural network. VQRL accepts historic network status and video quality predictions from VQPN as inputs.

The authors use a neural network in this video streaming application motivated by the effectiveness of the neural network in the prediction of time sequence data. The VQPN model adopts CNN layers to perform feature extraction of the input video frames to obtain spatial information. The CNN layers are followed by a two layered RNN which extracts temporal characteristics of the video frames in the past k sequences. The output of the VQPN is the prediction of video quality for the next time slot $t + 1$ denoted by V_{t+1} . The weights are updated based on the mean squared error loss function between the actual video quality score and the estimated video quality score. Specifically, the VQPN has 5 layers to perform feature extraction; a convolution layer with 64 filters each of size 5 with stride 1, a 3×3 average pooling layer, a second convolution layer with 64 filters of size 3 with stride 1, a 2×2 max-pooling layer and a hidden layer with 64 neurons. The output of the feature maps represents time series data which is fed into the RNN. The RNN comprises a GRU layer with 64 hidden units which then connects to another GRU layer of 64 hidden units. The hidden layer connects to the hidden output of the last GRU layer resulting in a 5-dimensional vector

output corresponding to the video quality scores for the bit rates [300, 500, 800, 1100, 1400] kbps. The authors use Adam gradient optimizer to train the VQPN with a learning rate of 10^{-4} . The VQPN was realized using the open source ML library, TensorFlow [69].

In modeling the VQRL, the neural network must be trained to learn the relationship between the video quality and bit rate. The sender serves as a learning agent who observes the future video quality and previous network status in the state space. The network status in the state space is comprised of sender's video transmission bit rate, received bit rate of past k sequences, delay gradient, and packet loss ratio of previous k sequences. The action taken refers to the video bit rate selected for the next time slot. Since in this case the states will be represented by continuous numbers which leads to a fairly large state space, it is unable to store them in a tabular form. As Q-learning cannot be effective in solving large state space problems, the authors have combined RL with a neural network. The authors solve this RL problem using A3C RL algorithm [68] whereby the policy training is achieved by means of a policy gradient algorithm. The authors further propose a multiple-training agent version to accelerate the training process. The multiple agents comprise of a central agent and several forward propagation agents. The forward propagation agent only decides with policy and critic via state inputs and neural network model received by the central agent for each step. The central agent uses the actor-critic algorithm to compute gradients and then updates its neural network model which is then pushed to the forward propagation agent. This can happen asynchronously among all agents with no thread locking between agents. The VQPN is trained and tested on two video datasets; VideoSet (a large scale compressed video quality dataset) and self-collected video sets (live concerts, music videos, short movies). To train VQRL, the authors use packet-level, chunk-level, and synthetic network traces. The quality of experience (QoE) metric is defined as a weighted function of video quality, sender's bit rate and delay gradient measured by the receiver at a time instant n . QARC algorithm was tested for its efficacy in real-world operating conditions by video-streaming on three different networks (public WiFi, Verizon cellular network and a wide area network between Shanghai and Boston) at a local coffee shop. The client was running on a MacBook Pro laptop connected to a server running on a Desktop machine located in Boston. The authors demonstrated the QARC algorithm outperform Skype and WebRTC in terms of the QoE metric.

3.2. Adaptive Channel Equalization

Another key area where ML algorithms, and more specifically neural network models, have been successfully employed to enhance the physical layer is adaptive channel equalization [70, 71, 72, 73, 74, 75, 76, 77, 78, 79]. IoT networks are usually dense comprising of several devices attempting to communicate simultaneously. Such dense deployment with multiple transmissions results in a harsh communication environment. Channel equalization techniques must be employed at the receiver for efficient signal demodulation. [78] employs multi-layer perceptrons (MLPs) to perform non-linear channel equalization of a 16-QAM system. The use of non-linear power amplifiers result in non-linear amplitude and phase distortion resulting in a non-linear channel model as expressed by the following relation,

$$r(t) = \mathcal{A}(a(t))e^{j[\phi(t)+\mathcal{P}(a(t))]} + g(t) \quad (45)$$

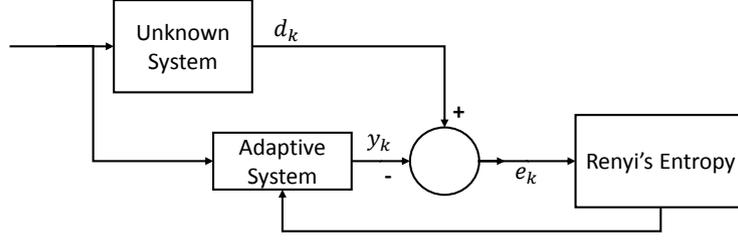


Figure 3: Adaptive System

such that $\mathcal{A}(a(t)) = \frac{\alpha_a x}{1 + \beta_a x^2}$ and $\mathcal{P}(a(t)) = \frac{\alpha_\phi x}{1 + \beta_\phi x^2}$ are the non-linear amplitude and phase distortions and $g(t)$ is the additive white Gaussian noise (AWGN).

The goal of non-linear channel equalization is to estimate the transmitted symbol from the received distorted symbols. The MLP is trained following a minimum error entropy criterion [80]. The adaptive system training aims to minimize/maximize the information potential based on the Renyi's entropy order $\gamma < 1 / \gamma > 1$. Figure 3 shows an adaptive system learning to update its weights such that the difference (e_i) between the output (y_i) and desired response (d_i) is minimized.

The weights (w) are trained based on the gradient of the information potential (I_ρ) as

$$\frac{\partial I_\rho}{\partial w} = \frac{\rho - 1}{N^\rho} \sum_j \left(\left[\sum_i \mathfrak{G}_\sigma(e_j - e_i) \right]^{\rho-2} \sum_i \mathfrak{G}'_\sigma(e_j - e_i) \frac{\partial y_i}{\partial w} \frac{\partial y_j}{\partial w} \right) \quad (46)$$

where $\mathfrak{G}_\sigma(\cdot)$ denotes the Gaussian kernel with standard deviation σ . The gradients of the outputs with respect to the weights can be computed using the standard BP algorithm. The proposed adaptive equalizer is composed of two MLPs operating in parallel, say MLP1 and MLP2. MLP1 is trained to learn the mapping of the transmitted signal amplitude $a_i, i = 1, 2, \dots, N$ to the received signal amplitude $|r_i|$. Since for a 16-QAM, the transmitted signal amplitude can only have three different amplitude levels, the output of the MLP1 corresponding to the transmitted signal amplitude is compared to the measured $|r_i|$. The output that gives the closest estimate to the possible values is chosen as the estimate for transmitted signal amplitude. The MLP2 is trained to learn the mapping from received signal amplitude to the non-linear phase distortion. From the estimated amplitude and phase from MLP1 and MLP2, the in-phase and quadrature components of the transmitted symbol are determined. In this work, authors train the system for an entropy order ($\rho = 3$), steepest ascent for information potential, a Gaussian kernel with $\sigma = 1$ and a dynamic step size. The training initially starts with unitary step size which then updates depending on the weight update such that the value increases when the update yields better performance and vice-versa. The authors demonstrated in simulations that the information potential maximization approach converges in fewer iterations than the mean squared error technique.

The authors of [79] explore the capabilities of DL for joint channel equalization and decoding. The DL model comprises of increased number of hidden layers to improve the representation capability of the neural network. Similar to the [78], the channel is assumed to introduce non-linear distortion to the transmitted symbols.

The authors train the network to minimize the mean squared error loss ($L = \frac{1}{N} \sum_i (y_i - m_i)^2$) between the transmitted symbol (m_i) and received symbol (y_i) as presented to the network in the training phase. The neuron weights in each layer can be updated using any gradient descent algorithms such as to minimize the loss function. The activation functions could be *ReLU* or *sigmoid* functions. The authors demonstrated the performance of the proposed DL for joint channel equalization and decoding with a six layer neural network comprising of 16, 256, 128, 64, 32 and 8 neurons in each layer. The modulation used is binary phase shift keying (BPSK) for a (16, 8) polar code and *ReLU* activation function.

In [81], three DL models for channel decoding are proposed. In this book, we will cite the CNN model for the channel decoder. The CNN employs a convolution operation which significantly reduces the number of parameters allowing the network to be deeper with fewer parameters. The hidden layers use either convolution or pooling. The input to the CNN is batch-normalized [82] such that any layer that previously received input x will receive $BatchNorm(x)$ which is a normalized, scaled and shifted version of original input with respect to a mini-batch.

$$BatchNorm(x) = \alpha \hat{x} + \beta \quad (47)$$

where $\hat{x} = \frac{x - \mu_\chi}{\sqrt{var_\chi + \epsilon}}$ is the normalized x over the mini-batch χ with mean (μ_χ) and variance (var_χ), α and β are the parameters to be learned.

The batch-normalized CNN is trained with mini-batch SGD to minimize the mean-squared error loss function. The authors observed the CNN decoder offered better performance compared to a MLP but at the cost of increased computational time.

The applicability of DL in channel estimation and signal detection in OFDM systems is demonstrated in [77]. The DL model is trained offline with simulated data to learn the channel distortions and reconstruct the transmitted symbols. Let $x(n)$ be the baseband OFDM modulated symbols transmitted over an N -path multipath channel $\{h(n)\}_{n=0}^{N-1}$ with AWGN $g(n)$ as shown by,

$$y(n) = x(n)h(n) + g(n) \quad (48)$$

After removing cyclic prefix and converting back to frequency domain the signal representation translates to

$$Y(k) = X(k)H(k) + G(k). \quad (49)$$

The pilot symbols are transmitted in the first OFDM block followed by user data in the subsequent blocks. The received pilot block and one data block are fed as input to the DL model. During the offline training stage, the model is trained with various received OFDM symbols generated with varying channel conditions under certain statistical profiles. The trained model when deployed for online signal detection, would estimate the signals without explicit channel estimation. The received signal and original transmitted symbols are supplied to the model to train it such that the difference between the model output and the original transmitted data are minimized. The model consists of five layers, three of which are hidden. Each layer comprises 256, 500, 250, 120 and 16 neurons respectively. The *ReLU* function is used as the activation function in all layers to map the input to the outputs of each layer except the last layer where *sigmoid* function is used to map to the interval $[0, 1]$.

A DL based method to improve the belief propagation (BEP) algorithm for decoding linear codes is proposed in [83]. BEP also known as Sum-Product algorithm is a message passing algorithm to derive statistical inferences from graphical models such as Bayesian networks. BEP is a form of Maximum A Posteriori (MAP) decoding of linear codes. BEP was first used in information theory by Gallager's iterative decoder for LDPC [84] which was a generalized case of belief propagation. Tanner graph forms a Bayesian network on which BEP operates. The DL model is trained with a single codeword. Conventional BEP decoder is constructed from the Tanner graph which is a graphical representation of the parity check matrix that describes the code [85]. The messages are transmitted over edges such that each edge calculates the outgoing message based on messages received over all its edges except for the transmitting edge.

To enable DL for a BEP decoder, the authors propose an alternative trellis representation where nodes in the hidden layer represent edges in the Tanner graph. If N denote the code block length, the number of neurons in the input layer is a vector of size N . The subsequent layers except for the final output layer i.e., the hidden layers have size E implying the number of edges in the Tanner graph. Each neuron in the hidden element corresponds to the message transmitted over some edge in the Tanner graph. The output layer has a size N that outputs the final decoded codeword. Let $e = (v, c)$ denote the neuron in the hidden layer $i, i \in 1, 2, \dots, 2L$ and $x_{i,e}$ represent the output message from the neuron after $\lfloor \frac{i-1}{2} \rfloor$ iterations. To allow the DL model, to learn based on the inputs, they are assigned weights which will be updated using the SGD method. The output of a neuron in the hidden layer, for an odd i is expressed as,

$$x_{i,e=(v,c)} = \tanh\left(\frac{1}{2}(w_{i,v}l_v + \sum_{e'=(v,c'),c' \neq c} w_{i,e,e'}x_{i-1,e'})\right). \quad (50)$$

and for an even i ,

$$x_{i,e=(v,c)} = 2 \tanh^{-1}\left(\prod_{e'=(v',c),v' \neq v} x_{i-1,e'}\right) \quad (51)$$

and the final v th output of the network is expressed as,

$$o_v = \sigma \left[w_{2L+1,v}l_v + \sum_{e'=(v,c')} w_{2L+1,v,e'}x_{2L,e'} \right] \quad (52)$$

where $\sigma(x) = (1 + e^{-x})^{-1}$ is the *sigmoid* function to map the final output codeword in the range $[0, 1]$. The goal is to train the weights $\{w_{i,v}, w_{i,e,e'}, w_{i,v,e'}\}$ to achieve an N -dimensional output codeword.

3.3. Adaptive Array Processing

MIMO systems are a trending physical layer solution to meet the increasing demand for high speed-high multiuser capacity communication systems. MIMO systems, due to their antenna arrays, can exploit spatial and temporal diversity to increase the communication data rate and spectral efficiency. Systems with adaptive antenna arrays can perform smart signal processing to combine the signals received at each array and nullify the interference and/or transmit the signals to

steer the beam in an intended direction. Multi-user MIMO [86] is already adopted in the developed and evolving communication standards like 3GPP LTE, 4G, and LTE-A. Another emerging MIMO technology, Massive MIMO, is the physical layer technology of choice for the latest 5G technology [87]. Massive-MIMO can revolutionize the 5G communication by providing reliable faster communication to more number of users simultaneously. Emerging 5G wireless networks promise ubiquitous connectivity, high data rates, energy efficiency, and spectrum availability. The dense, diverse and heterogeneous nature of IoT networks can be fulfilled by the disruptive 5G technologies such as massive MIMO, Non-Orthogonal Multiple Access (NOMA), machine-to-machine (M2M) etc. Beamforming is a prominent MIMO solution to enable communication to the desired device allowing to coexist with the other devices in the dense network. An essential step that enables beamforming is direction of arrival (DoA) estimation that allows the transmitter/receiver to learn the direction to/from which the signal should be directed/arrived. In this section, we will discuss a few prominent adaptive array techniques and how ML solutions can improvise them.

Several works [88, 89, 90, 91, 92, 93] address the problem of DoA estimation in array signal processing using artificial neural networks (ANNs). Let us look at each of these solutions. In [89], authors propose the use of a three-layer RBFNN that can learn multiple source-direction findings of a six-element linear antenna array. The RBFNN does not require training with all possible combinations of training sets. The network will generalize when trained with an expected range of input data. In this case, authors trained the network with input data whose DoA is uniformly distributed in the range -90° to 90° . The performance is compared to the conventional multiple signal classification (MUSIC) algorithm for DoA estimation of correlated and uncorrelated signals. The linear antenna array performs the mapping from the angle space to the sensor output space such that

$$s_i = \sum_{k=1}^K a_k e^{j2\pi f_0 d \sin \theta_k + \alpha_k} \quad (53)$$

where $i = 1, 2, \dots, M$ and k denote the respective antenna element and incident signal respectively, f_0 is the frequency of incident signal, d is the inter-element spacing, θ_k is the angle of arrival of k -th signal and α_k is the initial phase of k -th incident signal. The RBFNN is trained with N patterns to perform reverse mapping of received array data (s_i) to the angle space (θ_k). The incident array vectors are preprocessed prior to feeding them to the RBFNN. To train the neural network, the antenna array output vectors are generated ($\mathbf{s}(n)$, $n = 1, 2, \dots, N$). Each of the array output vector is further transformed to the spatial correlation matrix $\mathbf{R}(n)$. Since the diagonal elements of the correlation matrix does not carry any angle information, i.e. $R_{mm'} = \sum_{k=1}^K a_k$, only the cross-correlation terms are considered. These cross-correlated terms are arranged into an input vector $\mathbf{v}(n)$.

The output node subsequently computes the weighted sum of the hidden layer outputs.

$$\theta_k(j) = \sum_{i=1}^N w_i(k) f(\|\mathbf{s}(j) - \mathbf{s}(i)\|^2), k = 1, 2, \dots, K, j = 1, 2, \dots, N \quad (54)$$

where $w_i(k)$ represents the i -th weight of the network for the k -th incident signal and $f(\cdot)$ is the

Gaussian function performed by the hidden layer. Now, the equation 54 changes to

$$\theta_k(j) = \sum_{i=1}^N w_i(k) e^{-\|s(j)-s(i)\|^2/\sigma_g^2}, \quad (55)$$

where σ_g controls the influence of each basis function. The above equation can be rewritten in matrix form as,

$$\Theta = \mathbf{W}\mathbf{F}. \quad (56)$$

Here, Θ and \mathbf{W} are the $K \times L$ angle and weight matrices and \mathbf{F} is the $L \times N$ hidden layer matrix. L is chosen to be less than N to prevent ill-conditioning arising from large matrix. The input vectors $\mathbf{v}(n)$ are normalized according to equation 56. Using least-squares (LS) approach, the weights can be obtained as,

$$\hat{\mathbf{W}} = \Theta \mathbf{F}^\dagger \quad (57)$$

where is the pseudo-inverse given by

$$\mathbf{F}^\dagger = \mathbf{F}^T (\mathbf{F}\mathbf{F}^T)^{-1} \quad (58)$$

Now, the DoA estimate can be obtained as

$$\hat{\Theta} = \hat{\mathbf{W}}\mathbf{F} = \Theta^T (\mathbf{F}\mathbf{F}^T)^{-1} \mathbf{F} \quad (59)$$

The RBFNN is trained with the Normalized cumulative delta rule [94] such that the weight changes are accumulated over several training presentations as specified by the Epoch. The trained RBFNN will give the DoA estimates when presented with the normalized input vector. The authors demonstrated the computational advantage gained by adopting the RBFNN based DoA estimator as opposed to the conventional MUSIC algorithm. The estimation accuracy of the proposed DoA estimator in addition to the computational efficiency are the key merits of the proposed solution and presents itself as a computationally efficient alternative.

In a recent work [88], the DoA estimation is performed using a ANN with three layers; input, hidden and output layers. The authors study the estimation accuracy in terms of the number of neurons in the hidden layer. Unlike the RBF approach adopted by authors of [89], here the input activation function is the *hyperbolic tangent sigmoid transfer function* and the output activation function is the *logarithmic sigmoid*.

In yet another work [95], the authors employ RBFNN to perform adaptive beamforming. Adaptive beamforming is a method of updating weights of an adaptive antenna array such that the antenna radiation pattern will form beams such that strong beam is sent towards intended user's direction and nulls to sources of interference. The authors adopt a two-step approach to tackle this problem. First, the DoA of desired users are determined as in [89] and secondly, the beamformer weights are estimated to direct the beams. Similar to the DoA estimation problem, the authors approach this using RBFNN. For K incident signals, let the signal received at an M element linear antenna array at p^{th} time instant be

$$\mathbf{x}(p) = \sum_{k=1}^K \mathbf{b}_k s_k(p) + \mathbf{g}(p) = \mathbf{B}\mathbf{s}(p) + \mathbf{g}(p) \quad (60)$$

such that, $\mathbf{B} = [\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_M]^T$ is the array steering matrix,

$$\mathbf{b}_m = [1, e^{-j2\pi d \sin \theta_k / \lambda}, \dots, e^{-j2\pi(M-1)d \sin \theta_k / \lambda}] \quad (61)$$

that holds the spatial signal of the k^{th} source, $\mathbf{s} = [s_1, s_2, \dots, s_K]$ is the signal vector and $\mathbf{g}(p)$ is the noise vector. RBFNN is trained to compute the minimum variance distortionless response (MVDR) beamformer weights such that

$$\hat{\mathbf{w}} = \frac{\mathbf{R}^{-1}\mathbf{b}_m}{\mathbf{b}_m^H \mathbf{R}^{-1} \mathbf{b}_m} \quad (62)$$

where $\mathbf{R} = \frac{1}{P} \sum_p \mathbf{x}(p)\mathbf{x}^H(p)$ is the sample averaged covariance matrix computed from P snapshots of the received signal vector. The beamformer output can be denoted as $\mathbf{y}(p) = \hat{\mathbf{w}}^H \mathbf{x}(p)$. The beamformer vector estimation can be extended to any adaptive antenna array, the model considered in this example is a linear array for notational simplicity. The input and output layer of the RBFNN consists of $2M$ nodes to accommodate the in-phase and quadrature components of the input vector $\mathbf{x}(p)$ and the hidden layer' outputs. Much alike the DoA estimation problem in [89], the RBFNN is trained to perform an input output mapping from the received vector space to the beamformer weight space. The weights from the input to the hidden layer are identified using unsupervised k -means clustering and those from hidden to output layer follows the supervised *delta* learning rule. During the training phase, the RBFNN is trained with N_t training array output vectors $\mathbf{x}_n(p)$ and their corresponding $\hat{\mathbf{w}}_n \forall n \in 1, 2, \dots, N_t$. The array output vector is normalized prior to computing the corresponding covariance matrices \mathbf{R}_n . The beamformer weights $\hat{\mathbf{w}}_n$ are then computed according to equation 62. The trained RBFNN can be used to estimate the optimum MVDR beamformer weights for a presented normalized array output vector in a computationally inexpensive manner.

In this paper, we intended to give a glimpse of the various ML techniques to enhance the radio physical layer. The unification of such ML solutions with the radio hardware would be a prominent step in developing intelligent radios that can learn, adapt and behave under varying system and environmental dynamics. Table 1 enlists the various ML algorithms and their corresponding physical layer objective.

4. Machine Learning For Signal Intelligence

As IoT devices become more pervasive throughout society the available operational RF environment will contain more non-cooperative signals than seen before. Subsequently, the ability to garner information about signals within the spectrum of interest will become ever more important and complex, motivating the use of ML for signal intelligence. ML techniques for signal intelligence typically manifest themselves as solutions to discriminative tasks, and many applications specifically focus on multi-class or binary classification tasks. Perhaps the most prevalent signal intelligence task solved using ML techniques is that of automatic modulation classification (AMC). In short, this task involves determining what scheme was used to modulate the transmitted signal, given the raw signal observed at the receiver. Other signal intelligence tasks that employ ML solutions include wireless interference classification and signal detection. In this section, different state of the art ML solutions to signal intelligence tasks are discussed in further detail.

Table 1: Summary of Applications of ML in Physical layer

Physical layer solution	ML Algorithm	Objective
Li et al. [65]	RL	Adaptive rate control
Puljiz et al. [61]	kNN	Adaptive rate control
Yun and Caramanis [62]	SVR	Adaptive rate control
Li [96]	RL	Adaptive rate and power control
Huang et al. [67]	RL wt CNN and RNN	Adaptive rate control
Erdogmus et al. [78]	MLP	Non-linear channel equalization
Ye and Li [79]	DL	Non-linear channel equalization
Lyu et al. [81]	DL	Channel decoder
Ye et al. [77]	DL	Channel equalization in OFDM systems
Nachmani et al. [83]	DL	Improve BEP algorithm for decoding linear codes
Zooghby et al. [89]	RBFNN	DoA estimation
Nleren and Yaldiz [88]	ANN	DoA estimation
Zooghby et al. [95]	RBFNN	Adaptive beamforming

4.1. Modulation Classification

DL solutions to modulation classification tasks have received significant attention in the last couple of years [97, 98, 99, 100, 101, 102]. Several DL models are presented in [97] to address the modulation recognition problem. Hierarchical DNNs used to identify data type, modulation class and modulation order are discussed in detail in [102]. A conceptual framework for end-to-end wireless DL is presented in [101], followed by a comprehensive overview of the methodology for collecting spectrum data, designing wireless signal representations, forming training data and training deep neural networks for wireless signal classification tasks.

The task of AMC is pertinent in signal intelligence applications as the modulation scheme of the received signal can provide insight to what type of communication frameworks and emitters are present in the local RF environment. The problem at large can be formulated as estimating the conditional distribution, $p(y|x)$, where y represents the modulation structure of the signal and x is the received signal.

Traditionally, AMC techniques are broadly classified as maximum likelihood based approaches [103, 104, 105, 106, 107], feature-based approaches [108, 109, 110] and hybrid techniques [111]. Prior to the introduction of ML, AMC tasks were often solved using complex hand engineered features computed from the raw signal. While these features alone can provide insight about the modulation structure of the received signal, ML algorithms can often provide better generalization to new unseen data sets, making their employment preferable over solely feature based approaches. The logical remedy to the use of complex hand engineered feature based classifiers are models that aim to learn directly from received data. Recent work [112] show that deep convolutional neural

networks (DCNNs) trained directly on complex time domain signal data outperform traditional models using cyclic moment feature based classifiers. In [113], the authors propose a DCNN model trained on the two dimensional constellation plots generated from the received signal data and show that their approach outperforms other approaches using cumulant based classifiers and SVMs.

While strictly feature based approaches may become antiquated with the advent of the application of ML to signal intelligence, expert feature analysis can provide useful input to ML algorithms. In [114], we compute hand engineered features directly from the raw received signal and apply a feedforward neural network classifier to the features to provide a AMC. The discrete time complex valued received signal can be represented as,

$$y(n) = h(n)x(n) + w(n), \quad n = 1, \dots, N \quad (63)$$

where $x(n)$ is the discrete-time transmitted signal, $h(n)$ is the complex valued channel gain that follows a Gaussian distribution and $w(n)$ is the additive complex zero-mean white Gaussian noise process at the receiver with two-sided power spectral density (PSD) $N_0/2$. The received signal is passed through an Automatic Gain Control prior to the computation of feature values.

The first feature value computed from the received signal is the variance of the amplitude of the signal and is given by,

$$Var(|y(n)|) = \frac{\sum_{N_s} (|y(n)| - \mathbb{E}(|y(n)|))^2}{N_s} \quad (64)$$

where $|y(n)|$ is the absolute value of the over-sampled signal and $\mathbb{E}(|y(n)|)$ represents the mean computed from N_s samples. This feature provides information which help distinguish frequency shift keying (FSK) modulations from the phase shift keying (PSK) and quadrature amplitude modulation (QAM) modulation structures also considered in the classification task. The second and third features considered are the mean and variance of the maximum value of the power spectral density of the normalized centered-instantaneous amplitude, which is given as,

$$\gamma_{max} = \frac{\max |FFT(a_{cn}(n))|^2}{N_s}, \quad (65)$$

where $a_{cn}(n) \triangleq \frac{a(n)}{m_a} - 1$, $m_a = \frac{1}{N_s} \sum_{n=1}^{N_s} a(n)$, and $a(n)$ is the absolute value of the complex-valued received signal. This feature provides a measure of the deviation of the PSD from its average value. The mean and variance of this feature computed over subsets of a given training example are used as two separate entries in the feature vector input into the classification algorithm, corresponding to the second and third features, respectively.

The fourth feature used in our work was computed using higher order statistics of the received signal. Namely, cumulants, are known to be invariant to the various distortions commonly seen in random signals and are computed as follows,

$$C_{lk} = \sum_p^{\text{No. of partitions in } l} (-1)^{p-1} (p-1)! \prod_{j=1}^p \mathbb{E}\{y^{l_j-k_j} y^{*k_j}\}, \quad (66)$$

where l denotes the order and k denotes the number of conjugations involved in the computation of the statistic. We use the ratio, C_{40}/C_{42} as the fourth feature which is computed using,

$$C_{42} = \mathbb{E}(|y|^4) - |\mathbb{E}(y^2)|^2 - 2\mathbb{E}(|y|^2)^2, \quad (67)$$

$$C_{40} = \mathbb{E}(y^4) - 3\mathbb{E}(y^2)^2. \quad (68)$$

The fifth feature used in our work is called the in-band spectral variation as it allows discrimination between the FSK modulations considered in the task. We define $Var(f)$ as,

$$Var(f) = Var\left(\mathcal{F}(y(t))\right), \quad (69)$$

where $\mathcal{F}(y(t)) = \{Y(f) - Y(f - F_0)\}_{f=-f_i}^{+f_i}/F_0$, F_0 is the step size, $Y(f) = FFT(y(t))$, and $[-f_i, +f_i]$ is the frequency band of interest.

The final feature used in the classifier is the variance of the deviation of the normalized signal from the unit circle, which is denoted as $Var(\Delta_o)$. It is given as,

$$\Delta_o = \frac{|y(t)|}{\mathbb{E}(|y|)} - 1. \quad (70)$$

This feature helps the classifier discriminate between PSK and QAM modulation schemes.

The modulations considered in the work are the following: BPSK, quadrature phase shift keying (QPSK), 8PSK, 16QAM, continuous phase frequency shift keying (CPFSK), Gaussian frequency shift keying (GFSK), and Gaussian minimum shift keying (GMSK). This characterizes a seven class classification task using the aforementioned six features computed from each training example. To generate the data set, a total of 35,000 examples were collected: 1,000 examples for each modulation at each of the five SNR scenarios considered in the work. Three different feedforward neural network structures were trained at each SNR scenario using a training set consisting of 80% of the data collected at the given SNR and a test set consisting of the remaining 20%. The three feedforward nets differed in the number of hidden layers, ranging from one to three. Qualitatively, the feedforward network with one hidden layer outperformed the other models in all but the least favorable SNR scenario, achieving the highest classification accuracy of 98% in the most favorable SNR scenario. The seemingly paradoxical behavior is attributed to the over-fitting of the training data when using the higher complexity models, leading to poorer generalization in the test set.

This work has been further extended to evaluate other ML techniques using the same features. Accordingly, we found that training a random forest classifier for the same AMC task yielded similar results to the feedforward network classifier. Additionally, the random forest classifier was found to outperform the DNN approach in scenarios when the exact center frequency of the transmitter was not known, which was assumed to be given in the previous work. The random forest classifier was comprised of 20 classification and regression trees (CART) constructed using the gini impurity function. At each split a subset of the feature vectors with cardinality equal to 3 was considered.

An alternative approach to the previously described method is to learn the modulation of the received signal from different representations of the raw signal. [101] train DCNNs to learn the

modulation of various signals using three separate representations of the raw received signal. The authors denote the raw complex valued received signal training examples as $r_k \in C^N$, where k indexes the procured training data set and N is the number of complex valued samples in each training example. We inherit this notation for presentation of their findings. The data set in the work was collected by sampling a continuous transmission for a period of time and subsequently segmenting the received samples into N dimensional data vectors.

The authors train separate DCNNs on three different representations of the raw received signal and compare their results to evaluate which representation provides the best classification accuracy. The first of the three signal representations are given as a $2 \times N$ dimensional in-phase/quadrature (I/Q) matrix containing real valued data vectors carrying the I/Q information of the raw signal, denoted x_i and x_q , respectively. Mathematically,

$$x_k^{IQ} = \begin{bmatrix} x_i^T \\ x_q^T \end{bmatrix} \quad (71)$$

where $x_k^{IQ} \in R^{2 \times N}$. The second representation used is a mapping from the complex values of the raw received signal into two real valued vectors representing the phase, Φ and the magnitude, A ,

$$x_k^{A/\Phi} = \begin{bmatrix} x_A^T \\ x_\Phi^T \end{bmatrix} \quad (72)$$

where $x_k^{A/\Phi} \in R^{2 \times N}$ and the phase vector $x_\Phi^T \in R^N$ and magnitude vector $x_A^T \in R^N$ have elements,

$$x_{\Phi_n} = \arctan\left(\frac{r_{q_n}}{r_{i_n}}\right), x_{A_n} = (r_{q_n}^2 + r_{i_n}^2)^{\frac{1}{2}} \quad (73)$$

respectively. The third representation is a frequency domain representation of the raw time domain complex signal. It is characterized by two real valued data vectors, one containing the real components of the complex FFT, $\Re(X_k)$, and the other containing the imaginary components of the complex FFT, $\Im(X_k)$, giving,

$$x_k^F = \begin{bmatrix} \Re(X_k)^T \\ \Im(X_k)^T \end{bmatrix} \quad (74)$$

Using these three representations of the raw signal, the authors train three DCNNs with identical structure and compare the accuracy of the resultant models to determine which representation allows for learning the best mapping from raw signal to modulation structure.

The authors use training examples comprised of $N = 128$ samples of the raw signal sampled at 1 MS/s and consider the following 11 modulation formats: BPSK, QPSK, 8-PSK, 16-QAM, 64-QAM, CPFSK, GFSK, 4-pulse-amplitude modulation (PAM), wideband Frequency Modulation (WBFM), amplitude modulation (AM)-double-sideband modulation (DSB), and AM-single-sideband modulation (SSB). Thus, the training targets $y_k \in R^{11}$ are encoded as one-hot vectors where the index holding a 1 corresponds to the modulation of the signal. The authors use a total of 220,000 training examples $x_k \in R^{2 \times 128}$. Additionally, samples were acquired uniformly over different SNR scenarios ranging from $-20dB$ to $+20dB$.

The CNN structure used for each signal representation is the same, and consists of two convolutional layers, a fully connected layer, and a *softmax* output layer trained using the negative log-likelihood loss function. The activation function used in each of the convolutional layers and the fully connected layer is the *ReLU* function. The CNNs were trained using a training set comprised of 67% of the total data set, with the rest of the data set used as test and validation sets. An Adam optimizer [115] was used to optimize the training processes for a total of 70 epochs. The metrics used to evaluate each of the models include the precision, recall, and F1 score of each model. The authors give a range of values for the three aforementioned metrics for the CNN models trained on different data representations for three different SNR scenarios: high, medium, and low, corresponding to $18dB$, $0dB$, and $-8dB$, respectively. In the high SNR scenario, the authors report that the precision, recall, and F1 score of each of the three CNN models fall in the range of $0.67 - 0.86$. For the medium and low SNR scenarios, the same metrics are reported in the ranges of $0.59 - 0.75$ and $0.22 - 0.36$, respectively. The authors attribute the relatively low performance to the choice of the channel model used when generating the data, namely, a time-varying multipath fading channel.

The authors go on to evaluate the classification accuracy of each of the three models trained using different data representations under similar SNR conditions. Qualitatively, each of the three CNN models performs similarly at low SNR, while the CNN trained on the I/Q representation of data yields a better accuracy at medium SNR and the CNN trained on the amplitude and phase representation yields a better accuracy at high SNR. Interestingly, the CNN trained on the frequency domain representation of the data performs significantly worse than the I/Q and A/ϕ CNNs at high SNR. The authors mention that this could potentially be due to the similar characteristics exhibited in the frequency domain representation of the PSK and QAM modulations used in the classification problem. The primary takeaway from this work is that learning to classify modulation directly from different representations of the raw signal can be an effective means of developing a solution to the AMC task; however the efficacy of the classifier is dependent on how the raw signal is represented to the learning algorithm.

The following table provides the summary of the methods for AMC discussed in this section.

4.2. Wireless Interference Classification

The wireless interference classification (WIC) is a classification task that essentially refers to identifying what type of wireless emitters exist in the local RF environment. The motivation behind such a task is that it can be immensely helpful to know what type of emitters are present (WiFi, Zigbee, Bluetooth, etc.) in the environment so that effective interference avoidance and coexistence with other emitters can be achieved. Solutions to WIC tasks are often similar in nature to AMC techniques. For example, [116] employ DCNNs to classify IEEE 802.11 b/g, IEEE 802.15.4, and IEEE 802.15.1 emitters using a frequency domain representation of the captured signal. WIC tasks may also consider emitters in the environment that are not used in communication systems. In [117], an SVM solution is developed to classify interference in wireless sensor networks (WSNs) from IEEE 802.11 signals and microwave ovens. A recent work [118] shows the use of DCNNs to classify radar signals using both spectrogram and amplitude-phase representations of the received signal. In [119], DCNN models are proposed to accomplish interference classification on two dimensional time-frequency representations of the received signal to mitigate

Table 2: Summary of ML Solutions for Automatic Modulation Classification

Classifiers	Model	Representation	Objective
Jagannath et al. [114]	DNN	Feature-Based	7-Class task considering PSKs, FSKs, QAMs
Kulin et al. [101]	DCNN	I/Q, A/ Φ , FFT	11-Class task considering PSKs, FSK, QAMs, PAM, DSB, SSB
O’Shea and Corgan [112]	DCNN	I/Q	11-Class task considering PSKs, FSK, QAMs, PAM, DSB, SSB
Shengliang Peng and Yao [113]	DCNN	Constellation	4-Class task considering PSKs and QAMs
West and O’Shea [100]	DCNN, LSTM, RN	I/Q	11-Class task considering PSKs, FSK, QAMs, PAM, DSB, SSB
Karra et al. [102]	DCNN, DNN	I/Q, FFT	11-Class task considering PSKs, FSK, QAMs, PAM, DSB, SSB

the effects of radio interference in cosmological data. Additionally, the authors of [120] employ DCNN and LSTM models to achieve a similar end.

In [101], DCNNs are employed for the purpose of the wireless interference classification of three different wireless communication systems based on the WiFi, Zigbee, and Bluetooth standards. They look at five different channels for each of the three standards and construct a fifteen class classification task for which they obtain 225, 225 training vectors consisting of 128 samples each, where samples were collected at 10 MS/s. A flat fading channel with additive white Gaussian noise is assumed for this classification task.

Three DCNNs were trained and evaluated using the wireless interference classification data set described above. Each of the three CNNs was trained on one of the representations of the data that were presented in the previous section that discussed AMC. The CNN architectures were also the same as presented previously in Section 4.1.

Each of the three CNNs trained using different data representations was evaluated in a similar fashion to the evaluation method described in Section 4.1, namely, using precision, recall, and F1 score under different SNR scenarios. For the wireless interference classification task, the precision, recall, and F1 score of each of the three CNNs all fell in the interval from 0.98 – 0.99 under the high SNR scenario. For the medium and low SNR scenarios, the analogous intervals were from 0.94 – 0.99 and 0.81 – 0.90, respectively.

Additionally, the authors provide an analysis of classification accuracy for each of the three CNN models at varying SNRs. For the task of wireless interference classification, the CNN model trained on the frequency domain representation of the data outperforms the other models at all SNRs and especially so in lower SNR scenarios. The authors claim that these findings are due to the fact that the wireless signals that were considered have more expressive features in the frequency domain as they have different bandwidth, modulation, and spreading characteristics.

The authors of [121] take a different approach to the wireless interference classification task and primarily compare different types of learning models rather than different types of data representation. The models the authors propose include deep feedforward networks, deep convolutional networks, support vector machines using two different kernels, and a multi-stage training (MST) algorithm using two different learning algorithms. The authors consider 12 different transmitters and collect 1,000 packets from each transmitter for a total of 12,000 packets which comprise the entire data set. Each transmitter transmitted the same exact 1,000 packets, which were generated using pseudo-random values injected into the modem. All of the transmitters used a proprietary OFDM protocol with a QPSK modulation scheme and a baseband transmitter sample rate of 1.92 MS/s. At the receiver, each packet is represented by 10,000 time domain I/Q samples. Each of the models were trained on data sets consisting of training examples made up of 32, 64, 128, 256, 512, and 1024 samples from each packet, and their performance is compared across data sets. Given the complex valued received signal,

$$f = (f_1, f_2, \dots, f_N) \quad (75)$$

N samples were selected by skipping the first N_0 samples of a packet where $|\mathcal{R}(f_i)| < \tau$ for some $\tau > 0$ yielding the signal vector g ,

$$g = (f_{N_0}, f_{N_0+1}, \dots, f_{N_0+N-1}) \quad (76)$$

For the DNN, SVM, and MST models each training example was constructed by concatenating the real and imaginary parts of the signal vector, yielding a vector of dimension $2N$. For the CNN model the real and imaginary parts of the signal vector were stacked to generate $2 \times N$ dimensional training vectors.

The DNN architecture considered in the work consisted of two fully connected hidden layers, comprised of 128 *ReLU* units each and an output layer consisting of logistic *sigmoid* units. The network was trained using the Adam optimizer [115] and mini-batch size of 32.

The CNN model used by the authors was composed of two convolutional layers using 64 (8×2) and 32 (16×1) filters, respectively. Each convolutional layer was input into a max-pool layer with a pool size of 2×2 and 2×1 , respectively. The output of the second max-pool layer was fed into a fully-connected layer consisting of 128 *ReLU* units. An output layer employed logistic *sigmoid* units was used on top of the fully-connected layer.

The two SVM architectures analyzed in the work differ only in the kernel function used. The first architecture employed the polynomial kernel and the second employed the Pearson VII Universal Kernel [122]. Both architectures used Platt's Minimization Optimization algorithm to compute the maximum-margin hyperplanes.

The authors also analyze the performance of MST MLPs trained using first order and second order methods. A high level description of MST MLP is presented here and we refer the interested reader to [123] for a more rigorous derivation. The MST method to training neural networks, as presented in the work, is essentially a hierarchical way to solve an optimization problem by solving smaller constituent optimization problems. To this end, in what is called the first stage, a number of separate MLPs would be trained on different subsets of the training data set. This can be seen in the lowest layer of the hierarchical representation adapted from [121], and provided here in Figure 4.

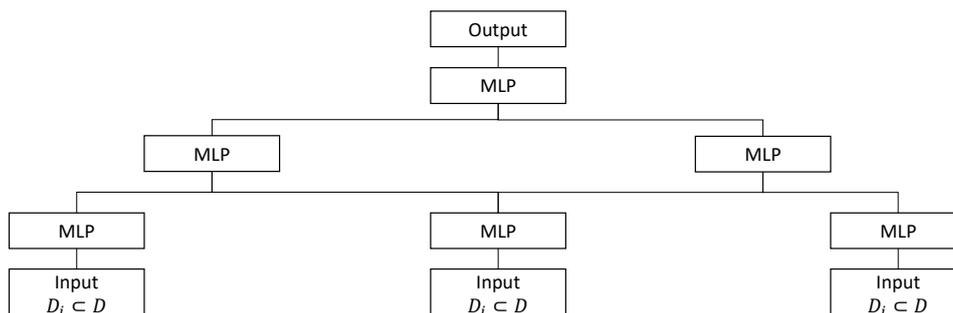


Figure 4: Adaptation of MST MLP used in [121].

Once the first stage is trained, a second stage is trained by taking the concatenation of the network outputs from the first stage as input. Training can continue in this fashion for subsequent stages. One of the advantages of training networks in this way is that the many smaller MLPs comprising the larger classifier can be efficiently trained using second order optimization methods. Second order optimization methods such as Newton, Gauss-Newton, or Levenberg-Marquardt methods are usually intractable due to the size of typical networks but can provide better convergence when applicable. The authors train two 3-stage MST systems, one using the first order method of SGD, and one using the second order Accelerated Levenberg-Marquardt method [124]. Each MST system had the identical structure where stage 1 consisted of 60 MLPs with 2 hidden layers and 10 units in each layer. Stage 2 and 3 had the same architecture and were comprised of 30 MLPs with each MLP consisting of 2 hidden layers made up of 15 units each. All hidden units employed the *tanh* activation function and all output layers contained linear units.

All of the models described above were trained on 10 different iterations of the collected data set and their performance was compared. Five data sets were constructed using training examples made up of 32, 64, 128, 256, and 512 samples and then each model was trained twice, using a training set comprised of 90% and 10% of the total data set, for a total of 10 different data sets for each model. In general, the MST system trained using second order methods on 90% of the training data performed best across all sizes of training examples, yielding a classification accuracy of 100% for each data set. All of the models performed better when trained using 90% of the data set as opposed to 10% of the training data set. Generally, each model performed better when provided with training examples that contained more samples, with the exception of the deep feedforward network model, which the authors attribute to the fact that longer sequences of samples may contain an increasing number of artifacts which the DNN may not be robust to. A summarization of the different models presented in this section is provided in Table 3.

5. Machine Learning For Higher Layers

The advancements in the higher layers, especially the data-link and the network layers have played a significant role in enabling IoT devices. The necessity to provide fair and efficient spectrum access has been a key motivating factor for researchers to design medium access control (MAC) protocols for IoT [125, 126]. In contrast to centralized designs where entities like base stations control and distribute resources, nodes in ad hoc IoT network have to coordinate resource

Table 3: Summary of ML Solutions for Wireless Interference Classification

Classifiers	Model	Representation	Objective
Kulin et al. [101]	DCNN	I/Q, A/ Φ , FFT	Classification of 15 WiFi, ZigBee, and Bluetooth Transmitters
Selim et al. [118]	DCNN	2D time-frequency, A/ Φ	Classification of Radar Signals
Akeret et al. [119]	DCNN	2D time-frequency	Classification of Cosmological Interference
Czech et al. [120]	DCNN, LSTM	2D time-frequency	Classification of Cosmological Interference
Youssef et al. [121]	DNN, DCNN, SVM, MST	I/Q	Classification of 12 OFDM Transmitters
Schmidt et al. [116]	DCNN	FFT	Classification of IEEE 802.11 b/g, IEEE 802.15.4, IEEE 802.15.1 signals
Grimaldi et al. [117]	SVM	Feature Based	Classification of IEEE 802.11 and Microwave Oven signals

allocation in a distributed manner. Similarly, to ensure scalability and reduce overhead, distributed designs are usually favored while designing routing algorithms for such networks. Recently, ML has made a significant impact on the design of these layers specifically to enhance scheduling and resource allocation, mitigating attacks like denial of service (DoS) in hostile environments, and efficient routing among others. In this section, we discuss in detail some of the advances made on this front.

5.1. Data Link Layer

A key functionality of the data link layer is to negotiate the access to the medium by sharing the limited spectrum resources in an ad hoc manner. Traditional MAC protocols designed for WANETs (including IoT networks) include carrier sense multiple access/collision avoidance (CSMA/CA) [127, 128], time division multiple access (TDMA) [129, 130], code division multiple access (CDMA) [131, 132] and hybrid approaches [133, 134, 135]. Here, we discuss some of the recent efforts to employ ML to enhance the data link layer.

The broadcast scheduling problem (BSP) is a key problem studied in a TDMA-based network to find an optimal TDMA schedule that provides transmission time slots to of all nodes while minimizing the TDMA frame size [136]. Several ML-based approaches have been proposed to solve this combinatorial optimization of BSP using variations of neural networks. This includes the work of [137] proposing a combination of HNN and genetic algorithm (GA) and [138] using sequential vertex coloring (SVC) and noisy chaotic neural network (NCNN). Subsequently, these

solutions were shown to be outperformed by fuzzy hopfield neural network (FHNN) proposed in [139]. Here, we describe how [139] tackles BSP.

Consider N nodes in a network with M time slots to share among these nodes. The slot assignment matrix \mathbf{A} , in which each element is defined as $A_{ij} = 1$, if time slot j is assigned to node i , otherwise $A_{ij} = 0$. The set of time slots to be assigned is given by set $T = \{t_1, t_2, \dots, t_M\}$. The fuzzy state, a degree that time slot t_x is assigned to node i is represented by μ_{xi} and matrix of all the fuzzy states, \mathbf{U} is called a fuzzy c-partition matrix. Next, the channel utilization of node i is defined as the fraction of total time slots assigned to node j from the total TDMA frame given as $\rho_j = (\sum_{j=1}^M A_{ij}/M)$. Accordingly, the total channel utilization for the network can be given as [140],

$$\rho = \frac{1}{MN} \sum_{j=1}^M \sum_{i=1}^N A_{ij} \quad (77)$$

The lower bound for the frame length is given by $\max_{i \in N} \deg(i) + 1$ where $\deg(i)$ is the number of edges incident to it. In case of FHNN, an energy function is considered as the distance between the current state of the HNN and its solution state. The objective is to minimize the energy function by solving the optimization problem. In this case, the energy function that considers all the constraints is defined as follows [139],

$$E = \frac{\alpha}{2} \sum_{x=1}^M \left(\sum_{i=1}^N \mu_{xi} - 1 \right)^2 + \beta \sum_{x=1}^M \sum_{i=1}^N \left(\sum_{y=1, y \neq i}^M d_{iy} \mu_{yi} \right. \\ \left. + \sum_{y=1, y \neq i}^M d_{iy} \sum_{y=1, y \neq i, k \neq y}^M d_{yk} (\mu_{xi})^f \left[t_x - \sum_{y=1}^M \frac{t_y}{\sum_{k=1}^M (\mu_{ki})^f} (\mu_{yi})^f \right]^2 \right) \quad (78)$$

where α and β are assumed to be positive coefficients, f is the fuzzification parameter, and $d_{iy} = 1$, if there is a connectivity between i and y . The first term in equation (78) ensures that M slots can only be distributed among the N classes (nodes). The second term minimizes the inter-class euclidean distance from a sample to the cluster center of all clusters. Accordingly, FHNN aims to classify M time slots into N nodes by minimizing E . In simulations, the proposed FHNN based BSP approach outperforms both [138] and [137] in terms of average time delay. Additionally, authors also show that performance improves with larger f at the expense of increased convergence time.

There have also been efforts to advance the current MAC protocols to react to different kinds of attack like DoS. In one such case [141], a MLP is used to modify carrier sense multiple access (CSMA)-based network to identify DoS attack and stay inactive for a duration to preserve the energy of the wireless sensor nodes. As shown in Figure 5, the MAC layer of each node consists of a MLP that has been trained prior to deployment. The parameters used by MLP include collision rate (c_r), packet request rate (P_{req}) and average packet wait time (t_w). The proposed solution is evaluated using both BP and the particle swarm optimization (PSO) [142] algorithm for training.

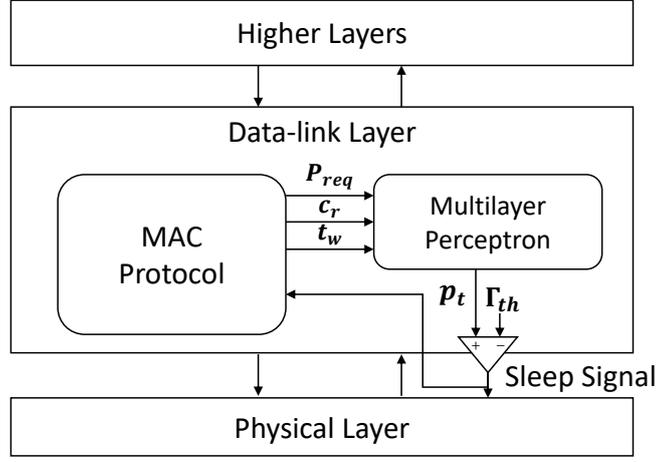


Figure 5: Block Diagram of MLP in MAC against DoS

The authors show that BP has lower computational cost compared to PSO but provides inferior convergence point in terms of quality of the weights. The output of MLP represents the probability that there is an active DoS attack (p_t). Based on the chosen threshold Γ_{th} , the nodes decide to sleep for a predetermined period of time when $p_t > \Gamma_{th}$. The work does not discuss the optimal value for Γ_{th} or the sleep time but provides an example of applying ML to mitigate the effects of such attacks.

Another interesting application where ML, specifically RL, has been successfully applied is in the domain of DSA for cognitive radio (CR) which could be instrumental in enabling modern IoT given the constrained availability of spectrum. An ALOHA-like scheme is developed for CRs by applying a multiagent RL framework [96]. In this work, a secondary user that is able to transmit successfully over an idle channel without collision receives a positive reward and zero otherwise. ACK packets received after transmission is used to ensure collision-free transmissions. Since the secondary user does not have control over the channel state, the Q-function is defined as the expected reward over a given time slot t . This in turn depends on the state of the overall system, $S(t) = s$ and the node i 's action ($a_i(t)$) at time slot t to transmit on channel h . The expectation is taken over the randomness of other secondary user's action and the primary user's activity which can be represented as,

$$Q_{ih}^s = E[R_i | a_i(t) = h, S(t) = s] \quad (79)$$

where R_i is the rewards for action. To ensure good channels are not neglected, the authors propose the use of a Boltzmann distribution for random exploration during the learning phase. Considering temperature \mathcal{T} , the exploration probability is given as,

$$P(i \text{ chooses channel } h | \text{ state } s) = \frac{\exp(Q_{ih}^s / \mathcal{T})}{\sum_{k=1}^N \exp(Q_{ik}^s / \mathcal{T})} \quad (80)$$

To accomplish this, each secondary user considers both the channel and other secondary users

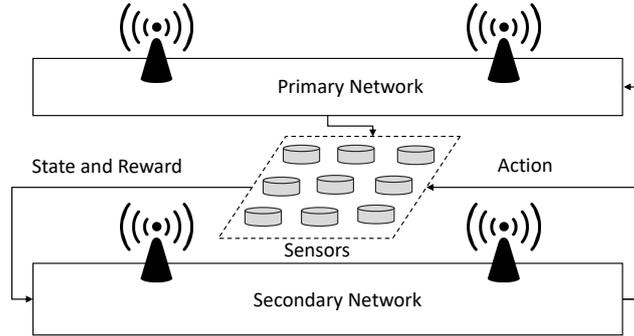


Figure 6: Framework for power control in cognitive network

to update its Q-values to choose the best action. It is important to remember that this is an extreme case where no control packets are exchanged between nodes similar to traditional ALOHA. Furthermore, the authors were able to show convergence in limited circumstances even when they extend the full observations to the case of partial observations. Simulations showed how secondary users can learn to avoid collision and outperform a scheme that uses Nash equilibrium.

A similar case is considered in [143] where authors propose a distributed DSA algorithm based on multi-agent reinforcement learning but this time employing DRL. We have seen how Q-learning provides adequate performance when the state-action space is relatively small. As the state-action space grows exponentially for larger problems, the direct application of Q-learning becomes inefficient and impractical as discussed previously. DQN which combines DNN with Q-learning can overcome this challenge. The goal is to enable users to learn a policy while dealing with the large state space without online coordination or message exchanges between users. In [143], the authors model their network state as partially observable for each user and the dynamics being non-Markovian and determined by the multi-user actions, they propose to use LSTM layer that maintains an internal state and aggregate observations over time.

To ensure feasibility, the training is set to happen offline where various training experiences with changing environment and topology are considered. This ensures that the algorithm can be deployed to operate in a distributed manner with the need to be updated only if the operating conditions are significantly different from the training set. After the training phase, each user determines which channel to select and associate “attempt probability” based on its observation. The proposed algorithm is compared against slotted-ALOHA that is assumed to have complete knowledge of the network and hence used optimal attempt probability. The proposed distributed algorithm that only used ACKs to learn outperforms slotted ALOHA by twice the channel throughput. They evaluated the network for two network utilities, (i) network rate maximization and (ii) individual rate maximization. In the case of users whose objective was to maximize the sum rate of the network, some learned to remain idle (sacrifice) incurring zero rate in order to maximize the overall network utility. In contrast, when each user aims to maximize its own rate they converged to a Pareto-optimal sharing policy.

The authors of [144] propose a DQN based framework that provides an Intelligent Power Control (IPC) algorithm for secondary users to coexist with the primary user while ensuring QoS for both. The overall architecture is depicted in Figure 6. The authors assume the presence of several

sensors that are deployed to monitor and convey the received signal strength (RSS) to the secondary users for decision making. The infinite states associated with the continuous RSS impose the need to employ DQNs. During the DQN’s training phase, secondary users assume complete knowledge of whether the QoS of every user (primary and secondary) are satisfied. The authors argue this can be achieved by overhearing the ACK packets. Once learning is complete, only the feedback from the sensors is required to determine the optimal power level for the secondary user to access the spectrum while satisfying the QoS constraint of both the networks. IPC is compared against the distributed constrained power control (DCPC) algorithm [145] which is an optimized solution. In contrast to IPC, an optimization-based technique like DCPC requires cooperation between both primary and secondary users. The simulation shows how IPC converges faster compared to DCPC while achieving a near optimal solution.

Realizing the role ML will play in the near future in maximizing the use of scarce spectrum, Defense Advanced Research Projects Agency (DARPA) initiated a three year competition known as Spectrum Collaboration Challenge (SC2). The goal was for teams to propose an ML based spectrum sharing strategy to allow peaceful coexistence between any unknown heterogeneous wireless networks. A solution inspired from this competition is presented in [146], which explores deep reinforcement learning multiple access (DLMA) for a heterogeneous wireless network consisting of various kind of networks (ALOHA, TDMA) that coexist. To accomplish this, authors use DRL to learn spectrum usage from a series of environmental observations and actions without actually being aware of the type of MAC protocols being operated. The goal is to maximize the total throughput of all the coexisting networks. They exploit neural networks to employ DRL as compared to traditional RL to enable fast convergence and ensure robustness to non-optimal parameters. Fast convergence is essential for wireless networks as convergence time is shorter than coherence time which will give nodes an opportunity to operate using an optimal strategy rather than trying to catch up with the changing efferent every time. Similarly, the lack of knowledge of existing networks makes it difficult to obtain optimal parameters.

The possible actions that can be taken by an agent is $a_t \in \{wait, transmit\}$. The observation after taking one of these actions can be $s_t \in \{success, collision, idleness\}$. Accordingly, the channel state at $t + 1$ is given as $c_{t+1} \triangleq (a_t, z_t)$. The reward is defined as,

$$r = \begin{cases} 1 & , \text{ if } z_t = success \\ 0 & , \text{ if } z_t = \{collision idleness\} \end{cases} \quad (81)$$

In this work, a DNN is used to approximate the state value function. Assuming ϕ is the parameter vector representing the weights of the DNN, the approximation can be represented as $q(s, a; \phi) \approx Q^*(s, a)$. The authors employ “experience replay” [147] which uses multiple experience samples (s, a, r_{t+1}, s_{t+1}) in each training step using the following loss equation,

$$L(\theta) = \sum_{(s,a,r,s') \in E_t} (y_{r,s'} - q(s, a; \phi))^2 \quad (82)$$

where,

$$y_{r,s'} = r + \gamma \max_{a'} q(s', a'; \phi_t) \quad (83)$$

where E_t is the set of experience samples used for training at time t . The authors argue the advantage of using DRL over RL by showing a faster convergence rate and a near-optimal strategy being achieved through simulations. They show how the network can learn and achieve near-optimal performance with respect to the sum throughput objective without the knowledge of co-existing MAC (TDMA, ALOHA). The work is further extended [148] by using a residual network [149] in place of the DNN. The authors show how a single RN architecture with fixed depth is suitable to ever changing wireless network scenarios as compared to the plain DNN which was shown to vary in performance based on the selected number of hidden layers.

These works provide a promising direction towards solving the spectrum crunch that will be experienced with the proliferation of IoT devices and 5G networks in the near future. We summarize the discussion of this section in Table 4.

Table 4: Summary of Application of ML in MAC protocols

MAC Protocol	ML Algorithm	Objective
Salcedo-Sanz et al. [137]	HNN wt GA	Proposed to solve BSP
Shi and Wang [138]	NCNN wt SVC	Proposed to solve BSP
Shen and Wang [139]	FHNN	Proposed to solve BSP
Kilkarni and Venayagamoorthy [141]	MLP	Tolerance against DoS
Li [96]	RL	ALOHA-like spectrum access
Naparstek and Cohen [143]	DQN wt LSTM	ALOHA-like spectrum access
Li et al. [144]	DQN	Intelligent power control
Yu et al. [146]	DQN	Non-cooperative heterogenous network
Yu et al. [148]	DQN wt RN	Non-cooperative heterogenous network

5.2. Network Layer

Routing protocols have evolved over the years to accommodate the needs of modern IoT WANETs. The design of the routing protocols primarily depends on the context and objective of the application and can be classified in several ways. Some of these classifications include geographical location based routing [150, 151, 152], hierarchical [153], QoS-based [154, 155], and recently cross-layer optimized routing [156, 157, 158]. Similar to earlier discussions, ML has elegantly found its way into this domain by providing a powerful tool to solve some of the problems associated with designing routing algorithms.

One of the earliest attempts to apply ML to routing algorithms is presented in [159] in the context of a traditional wired network including Local Access and Transport Area (LATA) telephone network. The proposed algorithm, referred to as Q-routing, uses a distributed approach which gathers estimated delay information from immediate neighbors to make the routing decision. The proposed Q-learning based routing algorithm can be represented as a variation of Bellman-Ford shortest path algorithm [160, 161] that replaces hop count by delivery time and performs the relaxation step online in an asynchronous manner. In [159], the authors clearly showed how Q-routing

is able to adapt to varying traffic loads after the initial inefficient learning period. When the load is low, Q-routing converges to using the shortest path and when the load increases, it is capable of handling the congestion more elegantly compared to the shortest path routing that is forced to use static routes.

In a recent effort [162], the need to reenvision router architectures and key routing strategies to meet the requirements of modern networks is highlighted. This was motivated by the advent of graphical processing unit accelerated software defined routers that are capable of massive parallel computing. Accordingly, authors propose to use DL, specifically, a deep belief network (DBN) based system that uses traffic patterns to determine the routes. The authors demonstrated with simulations the superiority of DBNs over open shortest path first (OSPF) in terms of throughput and average delay per hop. This can be attributed to the reduced overhead as DBNs does not use the traditional rule-based approach. Some of these ideas are extendable to WANETs after careful consideration of the challenges and characteristics of wireless networks.

One of the key challenges that will be faced by IoT devices operating in ad hoc mode is the reliability of routes that can get disconnected due to channel conditions or node failure. The authors of [163] study this problem in the context of multicast routing and apply cerebellar model articulation controller (CMAC) [164]. To ensure reliability, wireless mesh networks need to have the ability to recover from link disruption due to disrupted channel or node failure. The CMAC algorithm was first introduced around the same time that the perceptron algorithm was first introduced. While the CMAC framework can be considered a type of neural network, it is fundamentally different from the ones previously described in this paper. The CMAC architecture can be seen as an attempt to model human associative memory and employs a sort of look-up table technique. The CMAC framework is characterized by a mapping from input space to memory address space (look-up table) and a subsequent mapping from address space to output space. The mapping from input to address space is usually denoted as $S \rightarrow A$ where S is the input space and A is the address space. Typically, multiple mappings from input to address space are used such that a single input can “activate” multiple addresses in the address space. Each address in the address space contains a weight vector, $w \in A$, which is used in the subsequent mapping from address space to output space, usually denoted as $A \rightarrow P$. The function $f : A \rightarrow P$ is given to be the sum of the weight vectors contained in the activated memory regions. The training of the model can be conducted iteratively over training examples by updating the weight vectors used in the computation of the output by some proportion of the error observed for that training example.

In [163], authors use this concept to learn to estimate the route disconnection expectancy between itself and access points (APs) based on the following three parameters, (i) delay of packets in a node (i.e. sum of queuing delay and processing delay), (ii) number of node disconnections, and (iii) difference in delays between two packets that are separated by x number of packets. The proposed CMAC uses these three parameters to estimate the node disconnection probability (NDP). Then the NDP estimate enables nodes to predict possible node failure and react faster enabling better throughput, higher packet delivery ratio for multicast packets and provide minimum delay without prior knowledge of the topology.

Q-MAP is another multicast routing algorithm proposed to ensure reliable communication [165]. The algorithm is divided into two phases; in *join query forward* phase nodes use join query packets (JQPs) to explore all the possible routes to the multicast destination and *join reply*

backward phase uses join reply packets (JRPs) to establish the optimal route that maximizes the designed Q-value. The JQP can be considered as forwarding agents carrying the possible Q-values downstream, subsequently, JRP packets can be considered as backward agents carrying the optimal decision information upstream to the source.

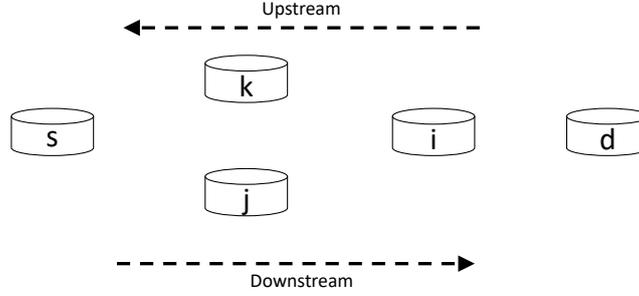


Figure 7: Framework for power control in cognitive network

In traditional unicast routing, the information used to make route decisions (such as resource reservation information, and Q value) are derived from downstream nodes. In contrast, Q-MAP gathers information from the upstream nodes that is used to make the route selection. A simple topology is depicted in Figure 7 where s is the source node and d is one of the destinations. In this example, node i needs to choose between j and k as the upstream node. Let us consider node i received a JQP from nodes j and k . Accordingly, node i computes its reinforcement Q-function and resource reservation data. Therefore, in this case, node i updates $Q(i, x)$ for any such JQP received from any upstream neighbor x (which in this example is j and k) as follows,

$$Q_t(i, x) \leftarrow (1 - \alpha)Q_{t+1}(i, x) + \alpha[r + \beta Q_t^*(x)] \quad (84)$$

Next, node i configures its own JQP and floods the packet downstream. Subsequently, when i receives a JRP from a downstream node (in this scenario d), node i will choose an upstream forwarding node that will eventually become part of the optimal route as follows,

$$Q_t^*(i) = \max_x Q_t(i, x) \quad \forall x \in (i, k) \quad (85)$$

Assuming that j is chosen as the forwarding node, in this case, node i creates its JRP and floods it. Node j receives this JRP and configures itself as the forwarding node for this multicast group by setting its forwarding flag. Each node maintains a forwarding table consisting of a source ID, group ID, forwarding flag indication and a timer field indicating the expiry of the forwarding group. In this manner, the multicast route is selected and maintained by source periodically initiating JQP. If any given receiver does not need to receive from a given source node, it just stops sending JRP for the corresponding multicast group. In this work, the authors do not discuss any experimental results, rather they keep the design general stating that the reward function is designed based on the objective of the network (maximize throughput, minimize energy consumption, minimize latency etc) and accordingly the corresponding resource reservation decision taken at each hop can include bandwidth, power or time slot allocation.

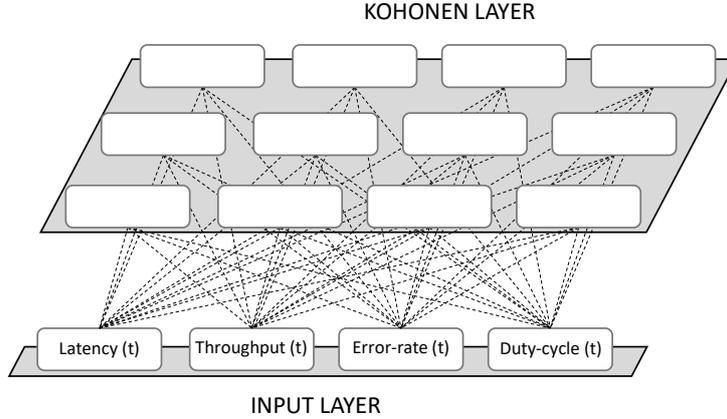


Figure 8: SOM Architecture used in SIR

An unsupervised learning based routing referred to as Sensor Intelligence Routing (SIR) is proposed in [166]. They modify the Dijkstra's algorithm utilizing SOM. Consider a directed connectivity graph $\mathcal{G}(k, E)$, where $k = \{k_0, k_1, \dots, k_N\}$ is a finite set of nodes, and $(i, j) \in E$ represents unidirectional wireless link from node k_i to node k_j (for simplicity, they are refer to them as node i and node j). Each edge (i, j) has a score associated with it denoted by γ_{ij} and it is assumed that $\gamma_{ij} = \gamma_{ji}$ which depends on QoS requirements of the network under consideration. In [166], the authors use latency, throughput, error-rate, and duty-cycle to represent a measure of QoS. Accordingly, the authors use these metrics for each link to represent their input of training vectors for a two-layer SOM architecture as shown in Fig 8. The input layer consists of $m = 4$ neurons, for each input vector of $\mathbf{x}(t) \in \mathcal{R}^m$. The second layer consists of a rectangular grid, where each neuron has m weight vectors connected from the input layer. During the learning phase, competitive learning is used such that the neuron whose vector most closely resembles the current input vector dominates. The SOM clusters these link by assigning each cluster a QoS rating. The learning phase is computationally intensive and hence needs to be performed offline. Meanwhile, execution can run on computational constrained sensor nodes and provide reliable performance as long as the topology and operational characteristics do not change.

The authors compare the proposed solution to Energy-Aware Routing (EAR) [167] and directed diffusion [168]. Directed diffusion is a data-centric routing protocol where the sink first broadcasts a request packet. This is used to set up a gradient (weighted reverse link) pointing to the sink (or the source of request). These gradients are used to find paths which are eventually pruned until the optimal path is determined. In the case of EAR, the source maintains a set of paths chosen by means of a certain probability that is inversely proportional to the energy consumption of that given route. The goal is to distribute traffic over multiple nodes to improve the network lifetime. Simulations show that the advantage of using SIR becomes evident only when nodes in the network start to fail. The parameters used to train SOM enable SIR to choose paths that are less prone to failure thereby providing better delay performance in scenarios where 40% nodes are prone to failure.

An example of RL in geographical routing can be seen in [169]. In Reinforcement Learning

based Geographic Routing (RLGR), they proposed a distributed algorithm that utilizes residual energy E_r and location of the neighbors. The MDP is characterized by the state of the packet which is defined by the current node where the packet resides and the action represents the choice of next-hop based on the Q-value ($Q(s, a)$). In this work, the reward function is given by,

$$r = \begin{cases} \alpha\tilde{\delta} + (1 - \alpha)\tilde{E} & , \text{ if next hop is not sink} \\ R_C & , \text{ if next hop is the sink} \\ -R_D & , \text{ if no next hop} \\ -R_E & , \text{ if next hop available but with low energy} \end{cases} \quad (86)$$

where $\tilde{\delta}$ represents the normalized advance towards the sink, \tilde{E} is the normalized residual energy. The authors consider a constant reward, R_C if the node is able to reach the sink directly. Finally, both R_D and R_E can be considered as the penalty suffered if no next-hop is found or if the existing next-hop has energy below the threshold. The proposed algorithm also uses ϵ to indicate the probability of exploration, i.e. how often the node will choose a random neighbor which may not be the next-hop that has the largest Q-value. For all other occasions (probability of $1 - \epsilon$), each node chooses a next-hop that provides the maximum Q-value. Their simulations showed significant improvement in network lifetime comparing RLGR to Greedy Perimeter Stateless Routing (GPSR) [150].

Next, we look at an example beyond RF terrestrial networks. In underwater acoustic networks (UANs), maximizing network lifetime is a key requirement. Accordingly, [170] propose a RL based approach that aims to distribute traffic among sensors to improve the lifetime of the network. In this work, the system state related to a packet is defined as the node that holds the packet. So s_i denotes the state of a packet held by node i . The action taken by node i to forward a packet to node j is denoted as a_j . If this action is successful, the state transitions from s_i to s_j with the transition probability of P_{ij}^j and stays in the same state s_i with transition probability of $P_{ii}^j = 1 - P_{ij}^j$ if it fails. Though these transition probabilities are unknown, authors argue that this can be estimated at runtime based on the history. Accordingly, the overall reward function at time instant t can be defined as follows,

$$r_t = P_{ij}^j R_{ij}^j + P_{ii}^j R_{ii}^j \quad (87)$$

where,

$$R_{ij}^j = -c - \alpha_1(E_i + E_j) + \alpha_2(D_i + D_j) \quad (88)$$

where α_1 and α_2 are tunable weights and c is the constant cost associated with consumption of resource (bandwidth, energy etc.) when a node chooses to transmit. E_i is the cost function associated with residual energy (E_i^{res}) and initial energy (E_i^{ini}). The energy cost function penalizes the system as residual energy decreases and is defined as,

$$E_i = 1 - \frac{E_i^{res}}{E_i^{ini}} \quad (89)$$

Similarly, D_i is defined to measure the energy distribution balance as follows,

$$D_i = \frac{2}{\pi} \arctan(E_i^{res} - \bar{E}_i) \quad (90)$$

where \bar{E}_i is the average residual energy of i and all its direct neighbors. This parameter increases the chance of neighbors with higher residual energy being preferred.

The reward function for the case where a packet forwarding attempt fails is defined as,

$$R_{ii}^j = -c - \beta_1 E_i + \beta_2 D_i \quad (91)$$

where β_1 and β_2 are again tunable weights. Authors use Q-learning at each node to enable them to learn about the environment using control packets and take action to improve network lifetime. The proposed solution is shown to outperform the vector-based forwarding protocol [151], a geographical routing protocol designed for UANs by achieving 20% longer lifetime. The authors claim the proposed solution can be applied for various UAN applications by tuning the trade-off between latency and energy efficiency for lifetime.

Feedback Routing for Optimizing Multiple Sinks (FROMS) [171] is proposed to achieve near-optimal routing from multiple source to multiple sink nodes. The goal of each node is to determine neighbor(s) for next-hop(s) towards the intended subset of sinks $S_p \subset S$. The state is defined as a tuple, $S = \{S_p, H_{S_p}^{NB}\}$, where $H_{S_p}^{NB}$ is the routing information through all neighboring nodes NB . The action is defined by a set $A_t = \{a_1, a_2, \dots, a_n\}$, such that $a_i = (nb_i, S_i)$, where $S_i \subset S_p$. The complete action set A must ensure that each sink $s \in S_p$ must be considered by exactly one element $a_i \in A$. The Q-value here is defined as follows,

$$Q_t(a) = \left(\sum_{i=1}^n Q_t(a_i) \right) - (n - 1) \quad (92)$$

where,

$$Q_t(a_i) = \left(\sum_{s \in S_i} H_s^{nb_i} \right) - 2(|D_i| - 1) \quad (93)$$

where $H_s^{nb_i}$ is the number of hops to the intended sink $s \in S_i$ through neighbor nb_i . $|S_i|$ denotes the number of sinks in S_i . The goal of the learning process is to decrease the Q-value as much as possible such that nodes pick the action that corresponds to the lowest Q-value. Accordingly, the reward function is defined as follows,

$$R_t(a_i) = C + \min_a Q(a) \quad (94)$$

where C is the cost of the action. In this manner, the Q-values propagate upstream facilitating the learning process. During the operational phase, it is assumed that nodes overhear neighbor's packets and use the information contained in the packets to update their Q-value. Eventually, the goal is to use routes that will deliver the packets to the desired subset of sinks through the least number of total hops. The authors also explore both greedy exploration and stochastic exploration techniques to avoid local minima. Simulation results validate the ability to learn shared routes to multiple sinks in an efficient manner to decrease the cost per packet compared to directed diffusion [172]. Additionally, they show how exploration can further reduce the cost per packet

albeit marginally. We summarize these routing algorithms and the ML techniques they apply in Table 5

Table 5: Application of ML in routing protocols

Routing Protocol	ML Algorithm	Objective/Comments
Boyan and Littman [159]	RL	Variation of Bellman-Ford proposed for wired network
Mao et al. [162] Sun et al. [165]	DBN RL	Outperform OSPF due to reduced overhead Multicast Routing Algorithm
Pourfakhar and Rahmani [163]	CMAC	Proposed to improve reliability by predicting disconnection probabilities
Barbancho et al. [166]	SOM	Modified version of Dijkstra’s Algorithm
Dong et al. [169]	RL	Energy efficient geographical routing
Hu and Fei [170]	RL	Lifetime-aware routing for UAN that aims to distribute traffic load among nodes
Forster and Murphy [171]	RF	Near-Optimal routing for multiple source to multiple sinks

6. Spectrum Sensing and Hardware Implementation

One of the key challenges in enabling real-time inference from spectrum data is how to *effectively* and *efficiently* extract *meaningful* and *actionable* knowledge out of the tens of millions of I/Q samples received every second by wireless devices. Indeed, a single 20 MHz-wide WiFi channel generates an I/Q stream rate of about 1.28 Gbit/s, if I/Q samples are each stored in a 4-byte word. Moreover, the RF channel is significantly time-varying (*i.e.*, in the order of milliseconds), which imposes strict timing constraints on the *validity* of the extracted RF knowledge. If (for example) the RF channel changes every 10ms, a knowledge extraction algorithm must run with latency (much) less than 10ms to both (i) offer an accurate RF prediction and (ii) drive an appropriate physical-layer response; for example, change in modulation/coding/beamforming vectors due to adverse channel conditions, local oscillator (LO) frequency due to spectrum reuse, and so on.

As discussed earlier, DL has been a prominent technology of choice for solving classification problems for which no well-defined mathematical model exists. It enables the analysis of unprocessed I/Q samples without the need of application-specific and computational-expensive feature extraction and selection algorithms [97], thus going far beyond traditional low-dimensional ML techniques. Furthermore, DL architectures are application-insensitive, meaning that the same architecture can be retrained for different learning problems.

Decision-making at the physical layer may leverage the spectrum knowledge provided by DL. On the other hand, RF DL algorithms must execute in *real-time* (*i.e.*, with static, known-a-priori latency) to achieve this goal. Traditional central processing unit (CPU)-based knowledge extraction

algorithms [173] are unable to meet strict time constraints, as general-purpose CPUs can be interrupted at-will by concurrent processes and thus introduce additional latency to the computation. Moreover, transferring data to the CPU from the radio interface introduces unacceptable latency for the RF domain. Finally, processing I/Q rates in the order of Gbit/s would require CPUs to run continuously at maximum speed, and thus consume enormous amounts of energy. For these reasons, RF DL algorithms must be closely integrated into the RF signal processing chain of the embedded device.

6.1. Existing work

Most of existing work is based on traditional low-dimensional machine learning [174, 175, 176, 177, 178], which requires (i) extraction and careful selection of complex features from the RF waveform (*i.e.*, average, median, kurtosis, skewness, high-order cyclic moments, etc.); and (ii) the establishment of tight decision bounds between classes based on the current application, which are derived either from mathematical analysis or by learning a carefully crafted dataset [179]. In other words, since feature-based machine learning is (a) significantly application-specific in nature; and (b) it introduces additional latency and computational burden due to feature extraction, its application to real-time hardware-based wireless spectrum analysis becomes impractical, as the wireless radio hardware should be changed according to the specific application under consideration.

Recent advances in DL [180] have prompted researchers to investigate whether similar techniques can be used to analyze the sheer complexity of the wireless spectrum. For a compendium of existing research on the topic, the reader can refer to [181]. Among other advantages, DL is significantly amenable to be used for real-time hardware-based spectrum analysis, since different model architectures can be reused to different problems as long as weights and hyper-parameters can be changed through software. Additionally, DL solutions to the physical layer modulation recognition task have been given much attention over recent years, as previously discussed in this work. The core issue with existing approaches is that they leverage DL to perform offline spectrum analysis only. On the other hand, the opportunity of real-time hardware-based spectrum knowledge inference remains substantially uninvestigated.

6.2. Background on System-on-Chip Computer Architecture

Due to its several advantages, we contend that one of the most appropriate computing platform for RF DL is a System on Chip (SoC). An SoC is an integrated circuit (also known as “IC” or “chip”) that integrates all the components of a computer, *i.e.*, CPU, random access memory (RAM), input/output (I/O) ports and secondary storage (*e.g.*, SD card) – all on a single substrate [182]. SoCs have low power consumption [183] and allow the design and implementation of *customized hardware* on the field-programmable gate array (FPGA) portion of the chip, also called programmable logic (PL). Furthermore, SoCs bring unparalleled flexibility, as the PL can be reprogrammed at-will according to the desired learning design. The PL portion of the SoC can be managed by the processing system (PS), *i.e.*, the CPU, RAM, and associated buses.

SoCs use the Advanced eXtensible Interface (AXI) bus specification [184] to exchange data (i) between functional blocks inside the PL; and (ii) between the PS and PL. There are three main AXI sub-specifications: *AXI-Lite*, *AXI-Stream* and *AXI-Full*. *AXI-Lite* is a lightweight, low-speed

AXI protocol for register access, and it is used to configure the circuits inside the PL. AXI-Stream is used to transport data between circuits inside the PL. AXI-Stream is widely used, since it provides (i) standard inter-block interfaces; and (ii) rate-insensitive design, since all the AXI-Stream interfaces share the same bus clock, the high-level synthesis (HLS) design tool will handle the handshake between DL layers and insert FIFOs for buffering incoming/outgoing samples. AXI-Full is used to enable burst-based data transfer from PL to PS (and *vice versa*). Along with AXI-Full, direct memory access (DMA) is usually used to allow PL circuits to read/write data obtained through AXI-Stream to the RAM residing in the PS. The use of DMA is crucial since the CPU would be fully occupied for the entire duration of the read/write operation, and thus unavailable to perform other work.

6.3. A Design Framework for Real-time RF Deep Learning

One of the fundamental challenges to be addressed is how to transition from a software-based DL implementation (*e.g.*, developed with the Tensorflow [173] engine) to a hardware-based implementation on an SoC. Basic notions of high-level synthesis and a hardware design framework are presented in Sections 6.3.1 and 6.3.2, respectively.

6.3.1. High-level Synthesis

HLS is an automated design process that interprets an algorithmic description of a desired behavior (*e.g.*, C/C++) and creates a model written in hardware description language (HDL) that can be executed by the FPGA and implements the desired behavior [185]. Designing digital circuits using HLS has several advantages over traditional approaches. First, HLS programming models can implement almost any algorithm written in C/C++. This allows the developer to spend less time on the HDL code and focus on the algorithmic portion of the design, and at the same time avoid bugs and increase efficiency, since HLS optimizes the circuit according to the system specifications. The clock speed of today's FPGAs is several orders of magnitude slower than CPUs (*i.e.*, up to 200-300 MHz in the very best FPGAs). Thus, parallelizing the circuit's operations is crucial. In traditional HDL, transforming the signal processing algorithms to fit FPGA's parallel architecture requires challenging programming efforts. On the other hand, an HLS toolchain can tell how many cycles are needed for a circuit to generate all the outputs for a given input size, given a target parallelization level. This helps to reach the best trade-off between hardware complexity and latency.

Loop Pipelining: In high-level languages (such as C/C++), the operations in a loop are executed sequentially and the next iteration of the loop can only begin when the last operation in the current loop iteration is complete. Loop pipelining allows the operations in a loop to be implemented in a concurrent manner.

Figure 9 shows an example of loop pipelining, where a simple loop of three operations, *i.e.*, read (RD), execute (EX), and write (WR), is executed twice. For simplicity, we assume that each operation takes one clock cycle to complete. Without loop pipelining, the loop would take 6 clock cycles to complete. Conversely, with loop pipelining, the next RD operation is executed concurrently to the EX operation in the first loop iteration. This brings the total loop latency to 4 clock cycles. If the loop length were to increase to 100, then the latency decrease would be even more evident: 300 versus 103 clock cycles, corresponding to a speedup of about 65%. An

```

for (int i=0; i<2;i++) {
    Op_Read;    /* RD */
    Op_Execute; /* EX */
    Op_Write;   /* WR */
}

```

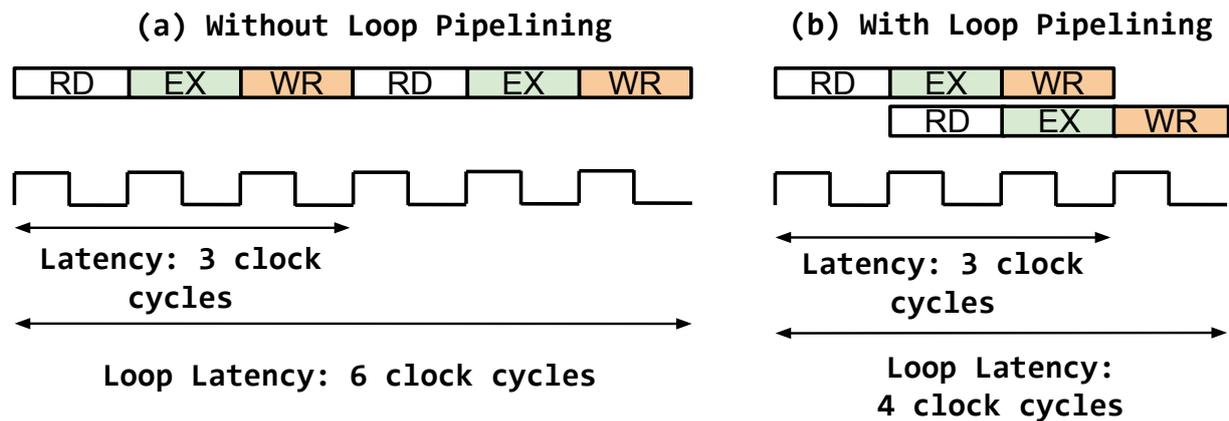


Figure 9: Loop pipelining.

important term for loop pipelining is called initiation interval (II), which is the number of clock cycles between the start times of consecutive loop iterations. In the example of Figure 9, the II is equal to one, because there is only one clock cycle between the start times of consecutive loop iterations.

Loop Unrolling: Loop unrolling creates multiple copies of the loop body and adjusts the loop iteration counter accordingly. For example, if a loop is processed with an unrolling factor (UF) equal to 2 (*i.e.*, two subsequent operations in the same clock cycle as shown in Figure 10), it may reduce a loop's latency by a factor of 50%, since a loop will execute in half the iterations usually needed. Higher UF and II may help achieve low latency, but at the cost of higher hardware resource consumption. Thus, the trade-off between latency and hardware consumption should be thoroughly explored.

```

for(int i = 0; i < 10; i++) {
    sum += a[i];
}

for(int i = 0; i < 10; i+=2) {
    sum += a[i];
    sum += a[i+1];
}

```

Figure 10: Loop unrolling.

6.3.2. Design Steps

Our framework presents several design and development steps, which are illustrated in Figure 11. Steps that involve hardware, middleware (*i.e.*, hardware description logic, or HDL), and software have been depicted with a blue, red, and green shade, respectively.

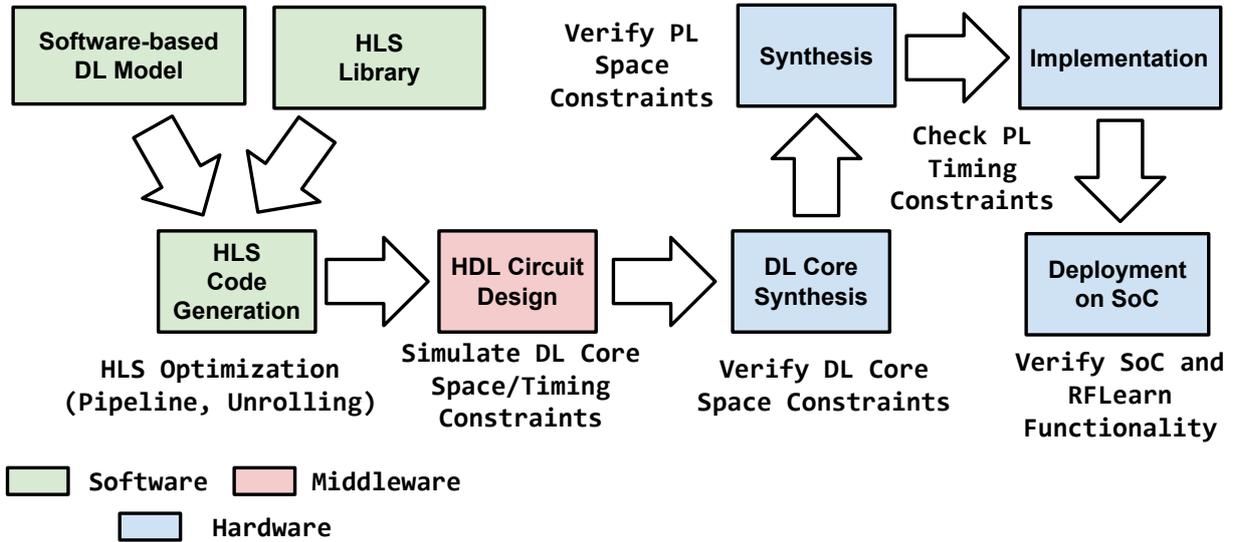


Figure 11: A Hardware Design Framework for RF Deep Learning.

The first major step of the framework is to take an existing DL model and convert the model in HLS language, so it can be optimized and later on synthesized in hardware. Another critical challenge is how to make the hardware implementation fully reconfigurable, *i.e.*, the weights of the DL model may need to be changed by the *Controller* according to the specific training. To address these issues, we distinguish between (i) the DL model architecture, which is the set of layers and hyper-parameters that compose the model itself, and (ii) the parameters of each layer, *i.e.*, the neurons' and filters' weights.

To generate the HLS code describing the software-based DL model, an *HLS Library*, which provides a set of HLS functions that parse the software-based DL model architecture and generates the HLS design corresponding to the desired architecture. The *HLS Library* supports the generation of convolutional, fully-connected, rectified linear unit, and pooling layers, and operated on fixed-point arithmetic for better latency and hardware resource consumption. The HLS code is subsequently translated to HDL code by an automated tool that takes into account optimization directives such as loop pipelining and loop unrolling. At this stage, the HDL describing the DL core can be simulated to (i) calculate the amount of PL resources consumed by the circuit (*i.e.*, flip-flops, BRAM blocks, etc); and (ii) estimate the circuit latency in terms of clock cycles. After a compromise between space and latency as dictated by the application has been found, the DC core can be synthesized and integrated with the other PL components, and thus total space constraints can be verified. After implementation (*i.e.*, placing/routing), the PL timing constraints can be verified, and finally the whole system can be deployed on the SoC and its functionality tested.

7. Open Problems

In this section, we discuss a set of open challenges overcoming which will accelerate the induction of ML techniques to the IoT devices.

7.1. Lack of Large-scale Wireless Signal Datasets

It is well known that learning algorithms require considerable amount of data to be able to effectively learn from a training dataset. Moreover, to compare the performance of different learning models and algorithms, it is imperative to use the same sets of data. More mature learning fields, such as computer vision and natural language processing (NLP) already have standardized datasets for these purposes [186, 187]. However, literature still lacks large-scale datasets for RF ML.

This is not without a reason. Although the wireless domain allows the synthetic generation of signals having the desired characteristics (*e.g.*, modulation, frequency content, and so on), problems such as RF fingerprinting and jamming detection require data that captures the unique characteristics of devices and wireless channels. Therefore, significant research effort must be put forth to build large-scale wireless signal datasets to be shared to the research community at large.

7.2. Choice of I/Q Data Representation Format

It is still subject of debate within the research community what is the best data representation for RF deep learning applications. For example, an I/Q sample can be represented as a tuple of real numbers or a single complex number, while a set of I/Q samples can be represented as a matrix or a single set of numbers represented as a string. It is common belief that there is no one-size-fits-all data representation solution for every learning problem, and that the right format might depend, among others, on the learning objective, choice of loss function, and the learning problem considered [97].

7.3. Choice of Learning Model and Architecture

While there is a direct connection between images and tensors, the same cannot be concluded for wireless signals. For example, while 3-D tensors have been proven to effectively model images (*i.e.*, red, green, and blue channels), and kernels in convolutional layers are demonstrably powerful tools to detect edges and contours in a given image, it is still unclear if and how these concepts can be applied to wireless signals. Another major difference is that, while images can be considered as stationary data, RF signals are inherently stochastic, non-stationary and time-varying. This peculiar aspect poses significant issues in determining the right learning strategy in the wireless RF domain. For example, while CNN seems to be able to effective at solving problems such as modulation recognition [100, 102, 97], it is still unclear if this is the case for complex problems such as RF fingerprinting. Moreover, DL has traditionally been used in static contexts [188, 189], where the model latency is usually not a concern. Another fundamental issue absent in traditional deep learning is the need to satisfy strict constraints on resource consumption. Indeed, models with high number of neurons/layers/parameters will necessarily require additional hardware and energy consumption, which are clearly scarce resources in embedded systems. Particular care must be devoted, therefore, when designing learning architectures to solve learning problems in the RF domain.

8. Conclusion

This paper provides a comprehensive account of advances in IoT wireless communication made possible by the application of ML. To accomplish this, we first provide readers with a detailed overview of some of the most prevalent ML techniques that are employed in wireless communication networks. Next, we discuss the impact ML has made on the design of the physical layer both in terms of enhancing the performance of communication links and in the domain of gathering signal intelligence. We then detail the increasing relevance of these techniques in the higher layers of the protocol stack enabling optimized utilization of limited resources which will be key to support the rapid growth of IoT devices. Realizing the importance of extending these techniques to hardware implementation, we discuss some steps that can be taken in those directions to ensure a rapid transition of these techniques to commercial hardware. Finally, we discuss some of the open problems that need to be tackled to further ease the adoption of ML for wireless networks. The overarching goal of this paper is to enable researchers with the fundamental tool to understand the application of ML in context of wireless communication in the IoT and apprise them of the latest advancements that will, in turn, motivate new and existing works.

References

References

- [1] A. Whitmore, A. Agarwal, and L. Da Xu, "The Internet of Things – A survey of topics and trends," *Information Systems Frontiers*, vol. 17, no. 2, pp. 261–274, 2015.
- [2] Glen Martin (Forbes), "How The Internet Of Things Is More Like The Industrial Revolution Than The Digital Revolution." <https://www.forbes.com/sites/oreillymedia/2014/02/10/more-1876-than-1995/#674c4e0b66d2>.
- [3] L. Da Xu, W. He, and S. Li, "Internet of things in industries: A survey," *IEEE Transactions on industrial informatics*, vol. 10, no. 4, pp. 2233–2243, 2014.
- [4] Ericsson Incorporated, "Ericsson Interim Mobility Report, February 2018." <https://www.ericsson.com/assets/local/mobility-report/documents/2018/emr-interim-feb-2018.pdf>, 2018.
- [5] Cisco Systems, "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2016-2021 White Paper." <http://tinyurl.com/zzo6766>, 2017.
- [6] Federal Communications Commission [2016], "Spectrum Crunch." <https://www.fcc.gov/general/spectrum-crunch>.
- [7] H. Shokri-Ghadikolaei, F. Boccardi, C. Fischione, G. Fodor, and M. Zorzi, "Spectrum sharing in mmwave cellular networks via cell association, coordination, and beamforming," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 11, pp. 2902–2917, 2016.
- [8] M. A. Vázquez, L. Blanco, and A. I. Pérez-Neira, "Hybrid analog–digital transmit beamforming for spectrum sharing backhaul networks," *IEEE transactions on signal processing*, vol. 66, no. 9, p. 2273, 2018.
- [9] L. Lv, J. Chen, Q. Ni, Z. Ding, and H. Jiang, "Cognitive non-orthogonal multiple access with cooperative relaying: A new wireless frontier for 5g spectrum sharing," *IEEE Communications Magazine*, vol. 56, no. 4, pp. 188–195, 2018.
- [10] X. Jin, J. Sun, R. Zhang, Y. Zhang, and C. Zhang, "SpecGuard: spectrum misuse detection in dynamic spectrum access systems," *to appear, IEEE Transactions on Mobile Computing*, 2018.
- [11] T. M. Chiwewe and G. P. Hancke, "Fast convergence cooperative dynamic spectrum access for cognitive radio networks," *IEEE Transactions on Industrial Informatics*, 2017.
- [12] J. Jagannath, S. Furman, T. Melodia, and A. Drozd, "Design and experimental evaluation of a cross-layer deadline-based joint routing and spectrum allocation algorithm," *IEEE Transactions on Mobile Computing*, pp. 1–1, 2018.
- [13] Federated Wireless, "Citizens Broadband Radio Service (CBRS) Shared Spectrum: An Overview." <https://www.federatedwireless.com/wp-content/uploads/2017/09/CBRS-Spectrum-Sharing-Overview.pdf>, 2018.

- [14] S. Agarwal and S. De, "eDSA: energy-efficient dynamic spectrum access protocols for cognitive radio networks," *IEEE Transactions on Mobile Computing*, vol. 15, no. 12, pp. 3057–3071, 2016.
- [15] L. Zhang, F. Restuccia, T. Melodia, and S. Pudlewski, "Learning to detect and mitigate cross-layer attacks in wireless networks: Framework and applications," in *Proc. of IEEE Conf. on Communications and Network Security*, (Las Vegas, NV, USA), October 2017.
- [16] J.-F. Huang, G.-Y. Chang, and J.-X. Huang, "Anti-jamming rendezvous scheme for cognitive radio networks," *IEEE Transactions on Mobile Computing*, vol. 16, no. 3, pp. 648–661, 2017.
- [17] G.-Y. Chang, S.-Y. Wang, and Y.-X. Liu, "A jamming-resistant channel hopping scheme for cognitive radio networks," *IEEE Transactions on Wireless Communications*, vol. 16, no. 10, pp. 6712–6725, 2017.
- [18] M. Bkassiny, Y. Li, and S. K. Jayaweera, "A survey on machine-learning techniques in cognitive radios," *IEEE Communications Surveys & Tutorials*, vol. 15, pp. 1136–1159, Third 2013.
- [19] C. Jiang, H. Zhang, Y. Ren, Z. Han, K. C. Chen, and L. Hanzo, "Machine learning paradigms for next-generation wireless networks," *IEEE Wireless Communications*, vol. 24, pp. 98–105, April 2017.
- [20] M. Chen, U. Challita, W. Saad, C. Yin, and M. Debbah, "Machine learning for wireless networks with artificial intelligence: A tutorial on neural networks," *CoRR*, vol. abs/1710.02913, 2017.
- [21] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [22] K. P. Murphy, *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012.
- [23] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Berlin, Heidelberg: Springer-Verlag, 2006.
- [24] T. M. Mitchell, *Machine learning*. McGraw Hill series in computer science, McGraw-Hill, 1997.
- [25] F. Rosenblatt, *The Perceptron, a Perceiving and Recognizing Automaton Project Para*. Report: Cornell Aeronautical Laboratory, Cornell Aeronautical Laboratory, 1957.
- [26] V. Vapnik and A. Lerner, "Pattern recognition using generalized portrait method," *Automation and Remote Control*, vol. 24, 1963.
- [27] C. Cortes and V. Vapnik, "Support-vector networks," *Mach. Learn.*, vol. 20, pp. 273–297, Sept. 1995.
- [28] F. Rosenblatt, *Principles of neurodynamics: perceptrons and the theory of brain mechanisms*. Report (Cornell Aeronautical Laboratory), Spartan Books, 1962.
- [29] D. Broomhead and D. Lowe, "Radial basis functions, multi-variable functional interpolation and adaptive networks," *ROYAL SIGNALS AND RADAR ESTABLISHMENT MALVERN (UNITED KINGDOM)*, vol. RSRE-MEMO-4148, 03 1988.
- [30] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, pp. 533–536, 1986. [;a href="https://www.nature.com/articles/323533a"](https://www.nature.com/articles/323533a)Commentary from News and Views section of Nature;[/a](https://www.nature.com/articles/323533a).
- [31] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio, *Deep learning*, vol. 1. MIT press Cambridge, 2016.
- [32] Y. LeCun *et al.*, "Generalization and network design strategies," *Connectionism in perspective*, pp. 143–155, 1989.
- [33] Y. T. Zhou and R. Chellappa, "Computation of optical flow using a neural network," *IEEE 1988 International Conference on Neural Networks*, pp. 71–78 vol.2, 1988.
- [34] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *CoRR*, vol. abs/1512.03385, 2015.
- [35] R. K. Srivastava, K. Greff, and J. Schmidhuber, "Highway networks," *CoRR*, vol. abs/1505.00387, 2015.
- [36] G. Huang, Z. Liu, and K. Q. Weinberger, "Densely connected convolutional networks," *CoRR*, vol. abs/1608.06993, 2016.
- [37] S. Hochreiter, Y. Bengio, P. Frasconi, and J. Schmidhuber, "Gradient flow in recurrent nets: the difficulty of learning long-term dependencies," 2001.
- [38] P. J. Werbos, "Backpropagation through time: what it does and how to do it," *Proceedings of the IEEE*, vol. 78, pp. 1550–1560, Oct 1990.
- [39] R. Pascanu, Ç. Gülçehre, K. Cho, and Y. Bengio, "How to construct deep recurrent neural networks," *CoRR*, vol. abs/1312.6026, 2013.
- [40] M. Schuster and K. K. Paliwal, "Bidirectional recurrent neural networks," *IEEE Transactions on Signal Pro-*

- cessing, vol. 45, pp. 2673–2681, Nov 1997.
- [41] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [42] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, “On the Properties of Neural Machine Translation: Encoder-Decoder Approaches,” *ArXiv e-prints*, Sept. 2014.
- [43] J. J. Hopfield, “Neural networks and physical systems with emergent collective computational abilities,” *Proceedings of the National Academy of Sciences*, vol. 79, no. 8, pp. 2554–2558, 1982.
- [44] D. Hebb, “The organisation of behaviour,” 1949.
- [45] T. Kohonen, “Self-organized formation of topologically correct feature maps. biological cybernetics,” *Biological Cybernetics*, pp. 59–69, 01 1982.
- [46] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press, 1998.
- [47] C. J. C. H. Watkins, *Learning from Delayed Rewards*. PhD thesis, King’s College, Oxford, 1989.
- [48] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. A. Riedmiller, “Playing atari with deep reinforcement learning,” *CoRR*, vol. abs/1312.5602, 2013.
- [49] P. P. Markopoulos, S. Kundu, and D. A. Pados, “Small-sample-support suppression of interference to pm-masked data,” *IEEE Transactions on Communications*, vol. 61, pp. 2979–2987, July 2013.
- [50] A. Jagannath and A. Amanna, “Realizing data driven and hampel preprocessor based adaptive filtering on a software defined radio testbed: A usrp case study,” in *Proc. of International Conference on Computing, Networking and Communications (ICNC)*, pp. 310–314, March 2018.
- [51] D. A. Pados and G. N. Karystinos, “An iterative algorithm for the computation of the mvdr filter,” *IEEE Transactions on Signal Processing*, vol. 49, pp. 290–300, Feb 2001.
- [52] J. P. Leite, P. H. P. de Carvalho, and R. D. Vieira, “A flexible framework based on reinforcement learning for adaptive modulation and coding in ofdm wireless systems,” in *Proc. of IEEE Wireless Communications and Networking Conference (WCNC)*, pp. 809–814, April 2012.
- [53] C. Wang, Z. Wang, W. Sun, and D. R. Fuhrmann, “Reinforcement learning-based adaptive transmission in time-varying underwater acoustic channels,” *IEEE Access*, vol. 6, pp. 2541–2558, 2018.
- [54] T. Ahmed, F. Ahmed, and Y. L. Moullec, “Optimization of channel allocation in wireless body area networks by means of reinforcement learning,” in *Proc. of IEEE Asia Pacific Conference on Wireless and Mobile (AP-WiMob)*, pp. 120–123, Sept 2016.
- [55] B. Peng, Q. Jiao, and T. Kürner, “Angle of arrival estimation in dynamic indoor thz channels with bayesian filter and reinforcement learning,” in *Proc. of European Signal Processing Conference (EUSIPCO)*, pp. 1975–1979, Aug 2016.
- [56] X. Liu, Y. Xu, Y. Cheng, Y. Li, L. Zhao, and X. Zhang, “A heterogeneous information fusion deep reinforcement learning for intelligent frequency selection of hf communication,” *China Communications*, vol. 15, pp. 73–84, Sept 2018.
- [57] A. R. Syed, K. L. A. Yau, H. Mohamad, N. Ramli, and W. Hashim, “Channel selection in multi-hop cognitive radio network using reinforcement learning: An experimental study,” in *Proce. of International Conference on Frontiers of Communications, Networks and Applications (ICFCNA)*, pp. 1–6, Nov 2014.
- [58] S. Tubachi, M. Venkatesan, and A. V. Kulkarni, “Predictive learning model in cognitive radio using reinforcement learning,” in *Proc. of IEEE International Conference on Power, Control, Signals and Instrumentation Engineering (ICPCSI)*, pp. 564–567, Sept 2017.
- [59] R. Ranjan and A. Phophalia, “Reinforcement learning for dynamic channel allocation in mobile cellular systems,” in *Proc. of International Conference on Recent Advances in Microwave Theory and Applications*, pp. 924–927, Nov 2008.
- [60] S. Singh and A. Trivedi, “Anti-jamming in cognitive radio networks using reinforcement learning algorithms,” in *Proc. of International Conference on Wireless and Optical Communications Networks (WOCN)*, pp. 1–5, Sept 2012.
- [61] Z. Puljiz, M. Park, and R. H. Jr., “A machine learning approach to link adaptation for sc-fde system,” in *Proc. of IEEE Global Communications Conference (GLOBECOM)*, pp. 1–5, Dec 2011.
- [62] S. Yun and C. Caramanis, “Reinforcement learning for link adaptation in mimo-ofdm wireless systems,” in *Proc. of IEEE Global Communications Conference (GLOBECOM)*, pp. 1–5, Dec 2010.

- [63] C. Pandana and K. J. R. Liu, "Near-optimal reinforcement learning framework for energy-aware sensor communications," *IEEE Journal on Selected Areas in Communications*, vol. 23, pp. 788–797, April 2005.
- [64] N. Mastrorarde and M. van der Schaar, "Joint physical-layer and system-level power management for delay-sensitive wireless communications," *IEEE Transactions on Mobile Computing*, vol. 12, pp. 694–709, April 2013.
- [65] X. Li, H. He, and Y. Yao, "Reinforcement learning based adaptive rate control for delay-constrained communications over fading channels," in *Proc. of International Joint Conference on Neural Networks (IJCNN)*, pp. 1–7, July 2010.
- [66] Y. Engel, S. Mannor, and R. Meir, "The kernel recursive least-squares algorithm," *IEEE Transactions on Signal Processing*, vol. 52, pp. 2275–2285, Aug 2004.
- [67] T. Huang, R.-X. Zhang, C. Zhou, and L. Sun, "Qarc: Video quality aware rate control for real-time video streaming based on deep reinforcement learning," 2018.
- [68] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. P. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, "Asynchronous methods for deep reinforcement learning," *CoRR*, vol. abs/1602.01783, 2016.
- [69] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015. Software available from tensorflow.org.
- [70] P. C. Kumar, P. Saratchandran, and N. Sundararajan, "Communication channel equalisation using minimal radial basis function neural networks," in *Neural Networks for Signal Processing VIII. Proceedings of the 1998 IEEE Signal Processing Society Workshop (Cat. No.98TH8378)*, pp. 477–485, Sept 1998.
- [71] L. Zhang and X. Zhang, "Mimo channel estimation and equalization using three-layer neural networks with feedback," *Tsinghua Science and Technology*, vol. 12, pp. 658–662, Dec 2007.
- [72] D. Jianping, N. Sundararajan, and P. Saratchandran, "Communication channel equalisation using complex-valued minimal radial basis function neural network," in *Proc. of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium*, vol. 5, pp. 372–377 vol.5, July 2000.
- [73] M. Birgmeier, "A neural network trained with the extended kalman algorithm used for the equalization of a binary communication channel," in *Proc. of IEEE Workshop on Neural Networks for Signal Processing*, pp. 527–534, Sept 1994.
- [74] W. R. Kirkland and D. P. Taylor, "On the application of feed forward neural networks to channel equalization," in *Proc. of IJCNN International Joint Conference on Neural Networks*, vol. 2, pp. 919–924 vol.2, June 1992.
- [75] S.-H. Heo, S.-K. Park, and S.-W. Nam, "Channel equalization for severe intersymbol interference and non-linearity with a radial basis function neural network," in *IJCNN'99. International Joint Conference on Neural Networks. Proceedings (Cat. No.99CH36339)*, vol. 6, pp. 3992–3995 vol.6, July 1999.
- [76] M. M. A. Moustafa and S. H. A. El-Ramly, "Channel estimation and equalization using backpropagation neural networks in ofdm systems," in *Proc. of IFIP International Conference on Wireless and Optical Communications Networks*, pp. 1–4, April 2009.
- [77] H. Ye, G. Y. Li, and B. Juang, "Power of deep learning for channel estimation and signal detection in ofdm systems," *IEEE Wireless Communications Letters*, vol. 7, Feb 2018.
- [78] D. Erdogmus, D. Rende, J. C. Principe, and T. F. Wong, "Nonlinear channel equalization using multilayer perceptrons with information-theoretic criterion," in *Neural Networks for Signal Processing XI: Proceedings of the 2001 IEEE Signal Processing Society Workshop (IEEE Cat. No.01TH8584)*, pp. 443–451, Sept 2001.
- [79] H. Ye and G. Y. Li, "Initial results on deep learning for joint channel equalization and decoding," in *Proc. of Vehicular Technology Conference (VTC-Fall)*, pp. 1–5, Sept 2017.
- [80] D. Erdogmus and J. C. Principe, "Generalized information potential criterion for adaptive system training," *IEEE Transactions on Neural Networks*, vol. 13, pp. 1035–1044, Sept 2002.
- [81] W. Lyu, Z. Zhang, C. Jiao, K. Qin, and H. Zhang, "Performance evaluation of channel decoding with deep neural networks," in *Proc. of IEEE International Conference on Communications (ICC)*, pp. 1–6, May 2018.

- [82] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *CoRR*, vol. abs/1502.03167, 2015.
- [83] E. Nachmani, Y. Be'ery, and D. Burshtein, "Learning to decode linear codes using deep learning," in *Proc. of Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Sept 2016.
- [84] R. G. Gallager, "Low-density parity-check codes," *IRE Transactions on Information Theory*, vol. IT-8, pp. 21–28, January 1962.
- [85] R. Tanner, "A recursive approach to low complexity codes," *IEEE Transactions on Information Theory*, vol. 27, pp. 533–547, September 1981.
- [86] J. Lee, J.-K. Han, and J. Zhang, "Mimo technologies in 3gpp lte and lte-advanced," *EURASIP J. Wirel. Commun. Netw.*, vol. 2009, pp. 3:1–3:10, Mar. 2009.
- [87] I. F. Akyildiz, S. Nie, S.-C. Lin, and M. Chandrasekaran, "5g roadmap," *Comput. Netw.*, vol. 106, pp. 17–48, Sept. 2016.
- [88] M. F. Ünlerşen and E. Yaldiz, "Direction of arrival estimation by using artificial neural networks," in *2016 European Modelling Symposium (EMS)*, pp. 242–245, Nov 2016.
- [89] A. H. E. Zooghyby, C. G. Christodoulou, and M. Georgiopoulos, "Performance of radial-basis function networks for direction of arrival estimation with antenna arrays," *IEEE Transactions on Antennas and Propagation*, vol. 45, pp. 1611–1617, Nov 1997.
- [90] S. H. Zainud-Deen, H. A. Malhat, K. H. Awadalla, and E. S. El-Hadad, "Direction of arrival and state of polarization estimation using radial basis function neural network (rbfn)," in *Proc. of National Radio Science Conference*, pp. 1–8, March 2008.
- [91] Y. L. Sit, M. Agatonovic, and T. Zwick, "Neural network based direction of arrival estimation for a mimo ofdm radar," in *Proc. of European Radar Conference*, pp. 298–301, Oct 2012.
- [92] E. Efimov, T. Shevgunov, and D. Filimonova, "Angle of arrival estimator based on artificial neural networks," in *2016 17th International Radar Symposium (IRS)*, pp. 1–3, May 2016.
- [93] M. Hirari, K. Gotoh, and M. Hayakawa, "Doa estimation of distributed sources using neural networks," in *Proc. of Fourth International Conference on Signal Processing (ICSP)*, pp. 335–338 vol.1, Oct 1998.
- [94] S. Haykin, *Neural Networks: A Comprehensive Foundation*. Prentice Hall, 1994.
- [95] A. Zooghyby, C. Christodoulou, and M. Georgeiopoulos, "Neural network based beamforming for interference cancellation," in *SPIE proceedings series*, pp. 420–429, Society of Photo-Optical Instrumentation Engineers, 1998.
- [96] H. Li, "Multiagent-learning for aloha-like spectrum access in cognitive radio systems," vol. 2010, p. 876216, May 2010.
- [97] T. J. O'Shea, T. Roy, and T. C. Clancy, "Over-the-air deep learning based radio signal classification," *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, pp. 168–179, Feb 2018.
- [98] T. J. O'Shea and J. Hoydis, "An introduction to deep learning for the physical layer," *IEEE Transactions on Cognitive Communications and Networking*, vol. 3, no. 4, pp. 563–575, 2017.
- [99] T. Wang, C.-K. Wen, H. Wang, F. Gao, T. Jiang, and S. Jin, "Deep learning for wireless physical layer: Opportunities and challenges," *China Communications*, vol. 14, no. 11, pp. 92–111, 2017.
- [100] N. E. West and T. O'Shea, "Deep architectures for modulation recognition," in *Proc. of IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, (Baltimore, MD, USA), pp. 1–6, March 2017.
- [101] M. Kulin, T. Kazaz, I. Moerman, and E. D. Poorter, "End-to-end learning from spectrum data: A deep learning approach for wireless signal identification in spectrum monitoring applications," *IEEE Access*, vol. 6, pp. 18484–18501, 2018.
- [102] K. Karra, S. Kuzdeba, and J. Petersen, "Modulation recognition using hierarchical deep neural networks," in *Proc. of IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, (Baltimore, MD, USA), pp. 1–3, March 2017.
- [103] O. Ozdemir, R. Li, and P. Varshney, "Hybrid Maximum Likelihood Modulation Classification Using Multiple Radios," *IEEE Communications Letters*, vol. 17, pp. 1889–1892, October 2013.
- [104] O. Ozdemir, T. Wimalajeewa, B. Dulek, P. K. Varshney, and W. Su, "Asynchronous Linear Modulation Classification with Multiple Sensors via Generalized EM Algorithm," *IEEE Transactions on Wireless Communications*, vol. 14, pp. 6389–6400, November 2015.

- [105] T. Wimalajeewa, J. Jagannath, P. K. Varshney, A. L. Drozd, and W. Su, "Distributed Asynchronous Modulation Classification Based on Hybrid Maximum Likelihood Approach," in *Proc. of IEEE Conf. on Military Communications (MILCOM)*, (Tampa, FL, USA), October 2015.
- [106] S. Foulke, J. Jagannath, A. L. Drozd, T. Wimalajeewa, P. K. Varshney, and W. Su, "Multisensor Modulation Classification (MMC) Implementation considerations – USRP case study," in *Proc. of IEEE Conf. on Military Communications (MILCOM)*, (Baltimore, MD, USA), October 2014.
- [107] J. Jagannath, H. M. Saarinen, and A. L. Drozd, "Framework for Automatic Signal Classification Techniques (FACT) for Software Defined Radios," in *Proc. of IEEE Symposium on Computational Intelligence in Security and Defense Applications (CISDA)*, (Verona, NY, USA), May 2015.
- [108] E. E. Azzouz and A. K. Nandi, *Automatic Modulation Recognition of Communication Signals*. Norwell, MA: Kluwer Academic Publishers, 1996.
- [109] A. Hazza, M. Shoaib, S. AlShebeili, and A. Fahd, "Automatic modulation classification of digital modulations in presence of HF noise.," *EURASIP Journal on Adv. in Signal Processing*, vol. 2012, p. 238, 2012.
- [110] A. Kubankova, J. Prinosil, and D. Kubanek, "Recognition of Digital Modulations Based on Mathematical Classifier," in *Proc. of the European Conference of Systems (ECCS)*, (Stevens Point, WI), 2010.
- [111] J. Jagannath, D. O'Connor, N. Polosky, B. Sheaffer, L. N. Theagarajan, S. Foulke, P. K. Varshney, and S. P. Reichhart, "Design and Evaluation of Hierarchical Hybrid Automatic Modulation Classifier using Software Defined Radios," in *Proc. of IEEE Annual Computing and Communication Workshop and Conference (CCWC)*, (Las Vegas, NV, USA), January 2017.
- [112] T. J. O'Shea and J. Corgan, "Convolutional radio modulation recognition networks," *CoRR*, vol. abs/1602.04105, 2016.
- [113] H. W. H. A. Shengliang Peng, Hanyu Jiang and Y.-D. Yao, "Modulation classification using convolutional neural network based deep learning model.," *WOCC*, 2017.
- [114] J. Jagannath, N. Polosky, D. O'Connor, L. Theagarajan, B. Sheaffer, S. Foulke, and P. Varshney, "Artificial Neural Network based Automatic Modulation Classifier for Software Defined Radios," in *Proc. of IEEE International Conference on Communications (ICC)*, (Kansas City, MO, USA), May 2018.
- [115] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014.
- [116] M. Schmidt, D. Block, and U. Meier, "Wireless interference identification with convolutional neural networks," *CoRR*, vol. abs/1703.00737, 2017.
- [117] S. Grimaldi, A. Mahmood, and M. Gidlund, "An svm-based method for classification of external interference in industrial wireless sensor and actuator networks," *Journal of Sensor and Actuator Networks*, vol. 6, 2017.
- [118] A. Selim, F. Paisana, J. A. Arokkiyam, Y. Zhang, L. Doyle, and L. A. DaSilva, "Spectrum monitoring for radar bands using deep convolutional neural networks," *CoRR*, vol. abs/1705.00462, 2017.
- [119] J. Akeret, C. Chang, A. Lucchi, and A. Refregier, "Radio frequency interference mitigation using deep convolutional neural networks," *Astronomy and Computing*, vol. 18, pp. 35–39, Jan. 2017.
- [120] D. Czech, A. Mishra, and M. Inggs, "A CNN and LSTM-based approach to classifying transient radio frequency interference," *Astronomy and Computing*, vol. 25, pp. 52–57, Oct. 2018.
- [121] K. Youssef, L.-S. Bouchard, K. Z. Haigh, H. Krovi, J. Silovsky, and C. P. Vander Valk, "Machine Learning Approach to RF Transmitter Identification," *ArXiv e-prints*, Nov. 2017.
- [122] B. Üstün, W. J. Melssen, and L. M. C. Buydens, "Facilitating the application of support vector regression by using a universal pearson vii function based kernel," 2005.
- [123] K. Youssef, N. N. Jarenwattananon, and L. Bouchard, "Feature-preserving noise removal," *IEEE Transactions on Medical Imaging*, vol. 34, pp. 1822–1829, Sept 2015.
- [124] L.-S. B. K. Youssef, "Training artificial neural networks with reduced computational complexity."
- [125] C. Cormio and K. R. Chowdhury, "A survey on mac protocols for cognitive radio networks," *Ad Hoc Networks*, vol. 7, pp. 1315–1329, Sept. 2009.
- [126] A. Al-Fuqaha, M. Guizani, M. Mohammadi, M. Aledhari, and M. Ayyash, "Internet of things: A survey on enabling technologies, protocols, and applications," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 4, pp. 2347–2376, 2015.
- [127] S. . Lien, C. . Tseng, and K. . Chen, "Carrier sensing based multiple access protocols for cognitive radio networks," in *Proc. of IEEE International Conference on Communications (ICC)*, pp. 3208–3214, May 2008.

- [128] N. Jain, S. R. Das, and A. Nasipuri, "A multichannel csma mac protocol with receiver-based channel selection for multihop wireless networks," in *Proc. of Tenth International Conference on Computer Communications and Networks (Cat. No.01EX495)*, pp. 432–439, Oct 2001.
- [129] C. Cordeiro and K. Challapali, "C-mac: A cognitive mac protocol for multi-channel wireless networks," in *Proc. of IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, pp. 147–157, April 2007.
- [130] M. Haddad, P. Muhlethaler, A. Laouiti, R. Zagrouba, and L. A. Saidane, "Tdma-based mac protocols for vehicular ad hoc networks: A survey, qualitative analysis, and open research issues," *IEEE Communications Surveys Tutorials*, vol. 17, pp. 2461–2492, Fourthquarter 2015.
- [131] A. Muqattash and M. Krunz, "CDMA-based MAC Protocol for Wireless Ad Hoc Networks," in *Proceedings of the 4th ACM International Symposium on Mobile Ad Hoc Networking & Computing, MobiHoc '03*, (New York, NY, USA), pp. 153–164, ACM, 2003.
- [132] S. Kumar, V. S. Raghavan, and J. Deng, "Medium Access Control Protocols for Ad Hoc Wireless Networks: A Survey," *Ad Hoc Network*, vol. 4, pp. 326–358, May 2006.
- [133] L. Sitanayah, C. J. Sreenan, and K. N. Brown, "Er-mac: A hybrid mac protocol for emergency response wireless sensor networks," in *Proc. of Fourth International Conference on Sensor Technologies and Applications*, pp. 244–249, July 2010.
- [134] H. Su and X. Zhang, "Opportunistic mac protocols for cognitive radio based wireless networks," in *Proc. of 41st Annual Conference on Information Sciences and Systems*, pp. 363–368, March 2007.
- [135] J. Jagannath, A. Saji, H. Kulhandjian, Y. Sun, E. Demirors, and T. Melodia, "A Hybrid MAC Protocol with Channel-dependent Optimized Scheduling for Clustered Underwater Acoustic Sensor Networks," in *Proc. of ACM Intl. Conf. on UnderWater Networks and Systems (WUWNet)*, (Kaohsiung, Taiwan), November 2013.
- [136] C. W. Commander, *Broadcast scheduling problem Broadcast Scheduling Problem*, pp. 339–345. Boston, MA: Springer US, 2009.
- [137] S. Salcedo-Sanz, C. Bousoño-Calzon, and A. R. Figueiras-Vidal, "A mixed neural-genetic algorithm for the broadcast scheduling problem," *IEEE Transactions on Wireless Communications*, vol. 2, pp. 277–283, March 2003.
- [138] H. Shi and L. Wang, "Broadcast scheduling in wireless multihop networks using a neural-network-based hybrid algorithm," *Neural Networks*, vol. 18, no. 5, pp. 765 – 771, 2005. IJCNN 2005.
- [139] Y.-J. Shen and M.-S. Wang, "Broadcast scheduling in wireless sensor networks using fuzzy Hopfield neural network," *Expert Systems with Applications*, vol. 34, no. 2, pp. 900 – 907, 2008.
- [140] G. Wang and N. Ansari, "Optimal broadcast scheduling in packet radio networks using mean field annealing," *IEEE Journal on Selected Areas in Communications*, vol. 15, pp. 250–260, Feb 1997.
- [141] R. V. Kulkarni and G. K. Venayagamoorthy, "Neural network based secure media access control protocol for wireless sensor networks," in *Proc. of International Joint Conference on Neural Networks*, pp. 1680–1687, June 2009.
- [142] J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proc. of International Conference on Neural Networks (ICNN)*, vol. 4, pp. 1942–1948 vol.4, Nov 1995.
- [143] O. Naparstek and K. Cohen, "Deep Multi-User Reinforcement Learning for Dynamic Spectrum Access in Multichannel Wireless Networks," in *Proc. of IEEE Global Communications Conference (GLOBECOM)*, pp. 1–7, Dec 2017.
- [144] X. Li, J. Fang, W. Cheng, H. Duan, Z. Chen, and H. Li, "Intelligent Power Control for Spectrum Sharing in Cognitive Radios: A Deep Reinforcement Learning Approach," *IEEE Access*, vol. 6, pp. 25463–25473, 2018.
- [145] S. A. Grandhi, J. Zander, and R. Yates, "Constrained power control," *Wireless Personal Communications*, vol. 1, pp. 257–270, Dec 1994.
- [146] Y. Yu, T. Wang, and S. C. Liew, "Deep-Reinforcement Learning Multiple Access for Heterogeneous Wireless Networks," in *Proc. of IEEE International Conference on Communications (ICC)*, pp. 1–7, May 2018.
- [147] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, pp. 529–533, Feb. 2015.

- [148] Y. Yu, T. Wang, and S. C. Liew, "Deep-Reinforcement Learning Multiple Access for Heterogeneous Wireless Networks," *CoRR*, vol. abs/1712.00162, 2017.
- [149] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, June 2016.
- [150] B. Karp and H. T. Kung, "GPSR: Greedy Perimeter Stateless Routing for Wireless Networks," in *Proc. of the 6th Annual International Conference on Mobile Computing and Networking, MobiCom '00*, (New York, NY, USA), 2000.
- [151] P. Xie, J.-H. Cui, and L. Lao, "Vbf: Vector-based forwarding protocol for underwater sensor networks," in *NETWORKING 2006. Networking Technologies, Services, and Protocols; Performance of Computer and Communication Networks; Mobile and Wireless Communications Systems* (F. Boavida, T. Plagemann, B. Stiller, C. Westphal, and E. Monteiro, eds.), (Berlin, Heidelberg), pp. 1216–1221, Springer Berlin Heidelberg, 2006.
- [152] J. Jagannath, S. Furman, T. Melodia, and A. Drozd, "design and experimental evaluation of a cross-layer deadline-based joint routing and spectrum allocation algorithm," *IEEE Transactions on Mobile Computing*, 2018.
- [153] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan, "Energy-efficient communication protocol for wireless microsensor networks," in *Proceedings of the 33rd Annual Hawaii International Conference on System Sciences*, pp. 10 pp. vol.2–, Jan 2000.
- [154] K. Akkaya and M. Younis, "An energy-aware qos routing protocol for wireless sensor networks," in *Proc. of International Conference on Distributed Computing Systems Workshops*, pp. 710–715, May 2003.
- [155] K. Akkaya and M. Younis, "A survey on routing protocols for wireless sensor networks," *Ad Hoc Networks*, vol. 3, pp. 325–349, 2005.
- [156] V. Srivastava and M. Motani, "Cross-layer design: a survey and the road ahead," *IEEE Communications Magazine*, vol. 43, pp. 112–119, 2005.
- [157] L. Kuo and T. Melodia, "Cross-layer Routing on MIMO-OFDM Underwater Acoustic Links," in *Proc. of IEEE Conf. on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, (Seoul, Korea), pp. 227–235, June 2012.
- [158] J. Jagannath, T. Melodia, and A. Drozd, "DRS: Distributed Deadline-Based Joint Routing and Spectrum Allocation for Tactical Ad-hoc Networks," in *Proc. of IEEE Global Communications Conference (GLOBECOM)*, (Washington D.C., USA), December 2016.
- [159] J. A. Boyan and M. L. Littman, "Packet Routing in Dynamically Changing Networks: A Reinforcement Learning Approach," in *Proceedings of the 6th International Conference on Neural Information Processing Systems, NIPS'93*, (San Francisco, CA, USA), pp. 671–678, Morgan Kaufmann Publishers Inc., 1993.
- [160] R. Bellman, "ON A ROUTING PROBLEM," *Quarterly of Applied Mathematics*, vol. 16, no. 1, pp. 87–90, 1958.
- [161] D. R. Ford and D. R. Fulkerson, *Flows in Networks*. Princeton, NJ, USA: Princeton University Press, 2010.
- [162] B. Mao, Z. M. Fadlullah, F. Tang, N. Kato, O. Akashi, T. Inoue, and K. Mizutani, "Routing or computing? the paradigm shift towards intelligent computer network packet transmission based on deep learning," *IEEE Transactions on Computers*, vol. 66, pp. 1946–1960, Nov 2017.
- [163] E. Pourfakhar and A. M. Rahmani, "A hybrid QoS multicast framework-based protocol for wireless mesh networks," *Computer Communications*, vol. 33, no. 17, pp. 2079 – 2092, 2010. Special Issue: Applied sciences in communication technologies.
- [164] J. S. Albus, "A new approach to manipulator control: The cerebellar model articulation controller (cmac)," *ASME Journal of Dynamic Systems, Measurement, and Control*, vol. 15, Sept 1975.
- [165] R. Sun, S. Tatsumi, and G. Zhao, "Q-MAP: a novel multicast routing method in wireless ad hoc networks with multiagent reinforcement learning," in *Proc. of IEEE Region 10 Conference on Computers, Communications, Control and Power Engineering. TENCOP*, vol. 1, pp. 667–670 vol.1, Oct 2002.
- [166] J. Barbancho, C. León, F. J. Molina, and A. Barbancho, "A new QoS routing algorithm based on self-organizing maps for wireless sensor networks," *Telecommunication Systems*, vol. 36, pp. 73–83, Nov 2007.
- [167] R. C. Shah and J. M. Rabaey, "Energy aware routing for low energy ad hoc sensor networks," in *2002 IEEE Wireless Communications and Networking Conference Record. WCNC 2002 (Cat. No.02TH8609)*, vol. 1, pp. 350–355 vol.1, March 2002.

- [168] C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed diffusion: A scalable and robust communication paradigm for sensor networks," in *Proc. of the Annual International Conference on Mobile Computing and Networking*, MobiCom '00, (New York, NY, USA), pp. 56–67, ACM, 2000.
- [169] S. Dong, P. Agrawal, and K. Sivalingam, "Reinforcement Learning Based Geographic Routing Protocol for UWB Wireless Sensor Network," in *Proc. of IEEE Global Telecommunications Conference (GLOBECOM)*, pp. 652–656, Nov 2007.
- [170] T. Hu and Y. Fei, "Qelar: A machine-learning-based adaptive routing protocol for energy-efficient and lifetime-extended underwater sensor networks," *IEEE Transactions on Mobile Computing*, vol. 9, pp. 796–809, June 2010.
- [171] A. Forster and A. L. Murphy, "FROMS: Feedback Routing for Optimizing Multiple Sinks in WSN with Reinforcement Learning," in *Proc. of International Conference on Intelligent Sensors, Sensor Networks and Information*, pp. 371–376, Dec 2007.
- [172] F. Silva, J. Heidemann, R. Govindan, and D. Estrin, *Frontiers in Distributed Sensor Networks*, ch. Directed Diffusion, p. to appear. Boca Raton, Florida, USA: CRC Press, Inc., Oct. 2004. refernece is superseded by [Silva05a].
- [173] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, *et al.*, "Tensorflow: A system for large-scale machine learning.," in *OSDI*, vol. 16, pp. 265–283, 2016.
- [174] M. L. D. Wong and A. K. Nandi, "Automatic digital modulation recognition using spectral and statistical features with multi-layer perceptrons," in *Proc. of the Sixth International Symposium on Signal Processing and its Applications (Cat.No.01EX467)*, vol. 2, pp. 390–393, 2001.
- [175] J. L. Xu, W. Su, and M. Zhou, "Software-defined radio equipped with rapid modulation recognition," *IEEE Transactions on Vehicular Technology*, vol. 59, pp. 1659–1667, May 2010.
- [176] S. U. Pawar and J. F. Doherty, "Modulation recognition in continuous phase modulation using approximate entropy," *IEEE Transactions on Information Forensics and Security*, vol. 6, pp. 843–852, Sept 2011.
- [177] Q. Shi and Y. Karasawa, "Automatic modulation identification based on the probability density function of signal phase," *IEEE Transactions on Communications*, vol. 60, pp. 1033–1044, April 2012.
- [178] S. Ghodeswar and P. G. Poonacha, "An SNR estimation based adaptive hierarchical modulation classification method to recognize M-ary QAM and M-ary PSK signals," in *Proc. of International Conference on Signal Processing, Communication and Networking (ICSCN)*, (Chennai, India), pp. 1–6, March 2015.
- [179] S. Shalev-Shwartz and S. Ben-David, *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- [180] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, p. 436, 2015.
- [181] Q. Mao, F. Hu, and Q. Hao, "Deep learning for intelligent wireless networks: A comprehensive survey," *to appear, IEEE Communications Surveys & Tutorials*, 2018.
- [182] R. F. Molanes, J. J. Rodríguez-Andina, and J. Fariña, "Performance characterization and design guidelines for efficient processor - FPGA communication in Cyclone V FPSoCs," *IEEE Transactions on Industrial Electronics*, vol. 65, pp. 4368–4377, May 2018.
- [183] Pete Bennett (EE Times), "The Why, Where and What of Low-Power SoC Design." https://www.eetimes.com/document.asp?doc_id=1276973, 2004.
- [184] Xilinx Inc., "AXI Reference Guide, UG761 (v13.1) March 7, 2011." https://www.xilinx.com/support/documentation/ip_documentation/ug761_axi_reference_guide.pdf, 2011.
- [185] F. Winterstein, S. Bayliss, and G. A. Constantinides, "High-level synthesis of dynamic data structures: A case study using vivado hls," in *Proc. of International Conference on Field-Programmable Technology (FPT)*, (Kyoto, Japan), pp. 362–365, 2013.
- [186] L. Deng, "The mnist database of handwritten digit images for machine learning research [best of the web]," *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 141–142, 2012.
- [187] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pp. 248–255, Ieee, 2009.
- [188] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, pp. 1097–1105, 2012.

- [189] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, *et al.*, “Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups,” *IEEE Signal processing magazine*, vol. 29, no. 6, pp. 82–97, 2012.