

## Two Applications of Deep Learning in the Physical Layer of Communication Systems

Deep learning has proven itself to be a powerful tool to develop data-driven signal processing algorithms for challenging engineering problems. By learning the key features and characteristics of the input signals instead of requiring a human to first identify and model them, learned algorithms can beat many human-made algorithms. In particular, deep neural networks are capable of learning the complicated features of nature-made signals, such as photos and audio recordings, and using them for classification and decision making.

The situation is rather different in communication systems, where the information signals are human made, the propagation channels are relatively easy to model, and we know how to operate close to the Shannon capacity limits. Does this mean that there is no role for deep learning in the development of future communication systems?

### Relevance

The answer to the question above is “no,” but, for the aforementioned reasons, we need to be careful not to reinvent the wheel. We must identify the right problems to tackle with deep learning and, even then, not start from a blank sheet of paper. There are many signal processing problems in the physi-

cal layer of communication systems that we already know how to solve optimally, for example, using well-established estimation, detection, and optimization theory. Nonetheless, there are also important practical problems where we lack acceptable solutions, for example, due to a lack of appropriate models or algorithms. In this lecture note, we first introduce the key properties of artificial neural networks and deep learning. The focus is not on technicalities around the training process or choice of network structure but on what we can practically achieve, assuming the training is carried out successfully. We then describe three application categories in communication engineering, whereof one exposes some fundamental weaknesses of deep learning and two illustrate important advances that can be made by utilizing deep learning.

### Prerequisites

This lecture note requires basic knowledge of linear algebra, digital communications, and probability.

### Problem statement and solution

We begin by briefly describing what artificial neural networks are and formulating the problem of using them as function approximators. Consider a system that takes an  $n_0$ -length input vector  $\mathbf{x}_0 \in \mathbb{R}^{n_0}$  and produces a  $k$ -length output vector  $\mathbf{y} \in \mathbb{R}^k$ , as illustrated in

Figure 1(a). The output is determined by the input via a deterministic function  $\hat{f}$ :

$$\mathbf{y} = \hat{f}(\mathbf{x}_0; \boldsymbol{\theta}). \quad (1)$$

The function is fixed but is characterized by an  $m$ -dimensional parameter vector  $\boldsymbol{\theta} \in \mathbb{R}^m$ . Many different input-output relations can be modeled in this way by changing the parameter vector  $\boldsymbol{\theta}$ , but they all share an underlying structure determined by the initial choice of  $\hat{f}$ . This is called a *gray-box model*.

When the function  $\hat{f}$  is selected to resemble the biological neural networks in human brains, the gray box is called an *artificial neural network*. The input vector  $\mathbf{x}_0$  is then viewed as the values in  $n_0$  neurons from which the function  $\hat{f}$  produces the values of  $\mathbf{y}$  in  $k$  other neurons. There are many different examples of this. The classical one is a fully connected feed-forward network, which is illustrated in Figure 1(b). In this case,  $\hat{f}$  is a composition of  $L$  functions,  $\hat{f}_1, \dots, \hat{f}_L$ , which describes transitions between neurons in an input layer to neurons in an output layer via  $L - 1$  intermediate hidden layers.  $L$  characterizes how deep the network is. The function  $\hat{f}_l$  is determined by the parameters  $\boldsymbol{\theta}_l = \{\mathbf{W}_l, \mathbf{b}_l\}$  and modeled as

$$\hat{f}_l(\mathbf{x}_{l-1}; \boldsymbol{\theta}_l) = \sigma_l(\mathbf{W}_l \mathbf{x}_{l-1} + \mathbf{b}_l), \quad (2)$$

where  $\mathbf{W}_l \in \mathbb{R}^{n_l \times n_{l-1}}$  is called a *weight matrix*,  $\mathbf{b}_l \in \mathbb{R}^{n_l}$  is called a *bias vector*,

and  $\sigma_l: \mathbb{R}^{n_l} \rightarrow \mathbb{R}^{n_l}$  is an element-wise nonlinear function that is called an *activation function*. With inspiration from the structure of the human brain, the function  $\hat{f}_l$  can be interpreted as taking the values  $\mathbf{x}_{l-1}$  in the  $n_{l-1}$  neurons of layer  $l-1$ , mixing the values together according to the affine transition relation  $\mathbf{W}_l \mathbf{x}_{l-1} + \mathbf{b}_l$ , and, finally, applying the activation function  $\sigma_l$  to determine the values of the  $n_l$  neurons of layer  $l$ .

If there are four layers, as in Figure 1(b), then  $L=3$ , and the complete input-output relation is


$$\mathbf{y} = \hat{f}_3(\hat{f}_2(\hat{f}_1(\mathbf{x}_0; \boldsymbol{\theta}_1); \boldsymbol{\theta}_2); \boldsymbol{\theta}_3). \quad (3)$$

Hence, the composite function  $\hat{f}$  is determined by the parameter vector  $\boldsymbol{\theta}$  containing the  $\sum_{l=1}^L n_l(n_{l-1} + 1)$  parameter values from  $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \boldsymbol{\theta}_3$  (i.e., the weights and biases from all layers).

### Problem statement

Artificial neural networks are generally used to approximate other functions by selecting the parameter vector  $\boldsymbol{\theta}$  to somehow minimize the approximation error. In particular, the category of fully connected feed-forward networks is capable of approximating any continuous function arbitrarily well by utilizing a (possibly) large but finite number of parameters (and neurons) [1]. This important result can be viewed as a generalization of Taylor polynomial approximations to functions with vector inputs and vector outputs. Two other categories are convolutional neural networks and recurrent neural networks [2]. Each category is believed to be better at approximating certain types of functions, in the sense of requiring fewer parameters to achieve a certain approximation error and/or it being easier to find appropriate parameter values in practice.

Selecting the right category of neural network is important but beyond the scope of this lecture note. Instead, our problem statement is as follows: what are the important use cases where the function approximation capability can be utilized in the physical layer of communication systems to achieve large improvements compared to conventional techniques?



**FIGURE 1.** (a) An arbitrary gray box taking  $\mathbf{x}_0$  as input and giving  $\mathbf{y}$  as output. (b) A fully connected feed-forward network with four layers ( $L=3$ ) that fits into the box in (a). The gray-box input-output model in (a) is characterized by  $\hat{f}$  and a parameter vector  $\boldsymbol{\theta}$ . It is called an *artificial neural network* if  $\hat{f}$  has a particular structure, such as the one illustrated in (b).

### Solution

To identify practically important use cases, we first need to understand how the function approximation is carried out. The parameter vector of an artificial neural network can be tuned/trained to approximate a (possibly unknown) function that we call  $f$ ; that is,  $\hat{f}$  should be trained to become a good estimate of  $f$ . This is, preferably, done by supervised learning using a set of  $T$  training examples consisting of input vectors  $\mathbf{x}_t^{\text{train}}$  and the corresponding output vectors  $\mathbf{y}_t^{\text{train}} = f(\mathbf{x}_t^{\text{train}})$  that we want the neural network to reproduce, for  $t=1, \dots, T$ . Let us represent these training examples as the columns of two matrices:

$$\mathbf{X}^{\text{train}} = [\mathbf{x}_1^{\text{train}} \ \dots \ \mathbf{x}_T^{\text{train}}], \quad (4)$$

$$\mathbf{Y}^{\text{train}} = [\mathbf{y}_1^{\text{train}} \ \dots \ \mathbf{y}_T^{\text{train}}].$$

The inputs should, ideally, be selected independently at random from the distribution of inputs that appears when using  $f$  in reality. The training basically consists of finding the parameter  $\boldsymbol{\theta}^*$  that minimizes a loss function  $\ell$  that measures the approximation mismatch:

$$\boldsymbol{\theta}^* = \underset{\boldsymbol{\theta}}{\text{argmin}} \ell(\boldsymbol{\theta}, \mathbf{X}^{\text{train}}, \mathbf{Y}^{\text{train}}). \quad (5)$$

For example, the loss can be measured in the mean-squared sense as

$$\ell(\boldsymbol{\theta}, \mathbf{X}^{\text{train}}, \mathbf{Y}^{\text{train}}) = \frac{1}{T} \sum_{t=1}^T \|\mathbf{y}_t^{\text{train}} - \hat{f}(\mathbf{x}_t^{\text{train}}, \boldsymbol{\theta})\|^2. \quad (6)$$

The goal is that the trained neural network  $\hat{f}(\mathbf{x}_0; \boldsymbol{\theta}^*)$  will provide approximately the right outputs, not only for the training examples, but for any input signal  $\mathbf{x}_0$  generated in the same way. This desired property is called *generalization*. Intuitively, if the unknown function  $f$  is continuous and has limited variability, we should be able to approximate it well from a large training set. We can, once again, make a parallel to polynomial approximations; any scalar polynomial of order  $T-1$  is uniquely determined by  $T$  samples (training examples) of the inputs and outputs. If the polynomial order is unknown, or if the function is only approximately polynomial, we need a larger number of samples to ensure a good approximation.

Since the training in (5) is a complicated, nonconvex optimization problem, huge efforts have been

dedicated to finding computationally and performance-wise acceptable suboptimal solutions. Moreover, the generalization to unseen inputs can be improved by various regularization, hyperparameter choices, and network designs [2]. These choices affect the model complexity. A simple model cannot capture complex dependencies. A too-complex model explains the training data only (this is called *overfitting*). A correct complexity tradeoff gives good generalization and is, typically, found using cross-validation. Such empirical craftsmanship is not the focus of this lecture note, but we conclude the following:

- Artificial neural networks can approximate any continuous function.
- The supervised training requires a large training set with inputs/outputs to achieve a low approximation error.

There are many functionalities in communication systems that can be described by a mathematical function  $f$  and, thus, can be approximated by a neural network. To identify the promising use cases, we first explain the basic methodology and its weaknesses by giving a concrete example.

A deep learning solution to signal detection

The physical layer of a communication system determines how an information-bearing signal is sent from the transmitter to the receiver over a physical channel. A critical task is the signal detection, where the receiver tries to identify what information was sent. To describe some key properties of deep learning, we exemplify how it can be used for signal detection.

We consider a classical additive white Gaussian noise (AWGN) channel, where a signal vector  $s \in \mathbb{R}^2$  of length 2 is sent. The received signal  $\mathbf{r} \in \mathbb{R}^2$  is given by

$$\mathbf{r} = \mathbf{s} + \mathbf{n}, \quad (7)$$

where  $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$  is an independent Gaussian noise vector where the entries have variance  $\sigma^2$ . We assume two bits of information are encoded into  $\mathbf{s}$  using a quadrature phase-shift keying (QPSK) constellation. Hence, there are four

possible signal points that are equally spaced on the unit circle:

$$s \in \left\{ \begin{bmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix}, \begin{bmatrix} 1/\sqrt{2} \\ -1/\sqrt{2} \end{bmatrix}, \begin{bmatrix} -1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix}, \begin{bmatrix} -1/\sqrt{2} \\ -1/\sqrt{2} \end{bmatrix} \right\}. \quad (8)$$

The mapping between information bits and signals is illustrated in Figure 2(a). Due to the additive noise, the received signal  $\mathbf{r}$  can take any value, but the Gaussian distribution makes values close to one of the signal points in (8) more likely than values far away. This can be seen from the red dots in Figure 2(a), which represent  $\mathbf{r}$  for 10,000 noise realizations with  $\sigma^2 = 0.2$  that are added to each of the four signal points.

Based on the received signal  $\mathbf{r}$ , the receiver needs to guess (detect) what signal  $\mathbf{s}$  was sent. We have trained a neural network for this task by taking the received signal  $\mathbf{x}_0 = \mathbf{r}$  as input and letting the output  $\mathbf{y}$  be a length-4 vector that is one for the detected signal and has zeroes elsewhere. We used the 40,000 red dots in Figure 2(a) and the signals  $\mathbf{s}$  that generated these  $\mathbf{r}$  to train a fully connected neural network using standard training methods. We then applied the neural network to a wide range of possible received signals to illustrate how it is making its detection.

The colored areas in Figure 2(b) show in which regions the received signals are mapped to the respective information signals. The regions are separated by lines, which is expected since each layer performs linear algebra operations; in particular, each activation function determines if the input is below/above a line that has been selected by training. Note that we have zoomed out, and the range of values that was shown in Figure 2(a) is indicated by the black square.

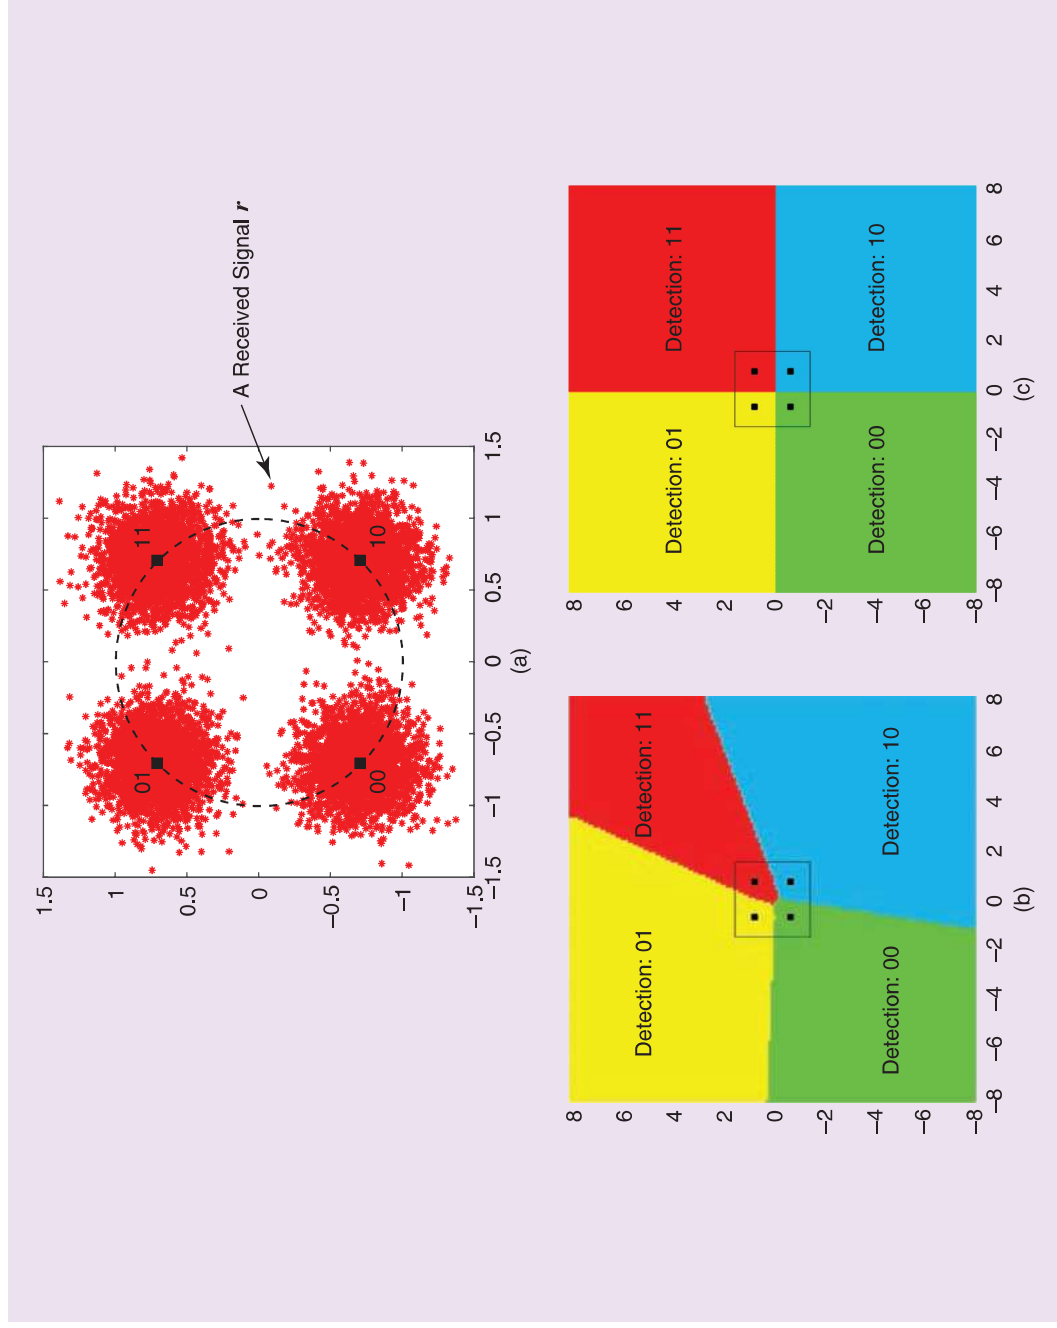
The colored detection regions produced by the neural network have peculiar asymmetric shapes, which are not optimal. In fact, the optimal detection regions for AWGN channels are well known [3, Ch. 6]: the received signal should be mapped to the closest signal point in terms of Euclidean distance. The optimal detection regions are shown in Figure 2(c). The regions are quite similar within the black square, but they greatly deviate farther away.

Several important observations can be made from this example.

- If there is a known optimal algorithm, a trained neural network cannot outperform it. The detection error probability is, however, almost the same in this particular example since most received signals appear within the black square where the neural network has a decent behavior.
- There are two reasons why the detection regions in Figure 2(b) are wrongly shaped. First, the shape inside the black square (around the signal points) is wrong due to overfitting; the training examples in Figure 2(a) can be approximately separated by many piecewise linear boundaries, including the ones shown in Figure 2(b). Second, since all training examples are inside the square, the behavior outside the square is somewhat random; the neural network has learned to interpolate between training examples but not to extrapolate outside the square. This is a practical issue since received signals far outside the square occasionally appear due to the long-tailed Gaussian distribution. This is a general phenomenon; neural networks are good at handling typical inputs but may generalize poorly to atypical inputs.
- We could have used prior domain knowledge (from digital communications) to preprocess the input signals. In this example, the neural network had to rediscover where the constellation points are, how the noise is distributed, and how to make the right detection. If we, instead, compute the Euclidean distance between the received signal and each of the four signal constellation points, we could use that as input to a neural network. This gives more accurate and reliable results since we have utilized our domain knowledge to ensure that the neural network has fewer characteristics to learn. However, it still cannot beat the optimal detection.

Is there a role of deep learning in communications?

Since signal detection in AWGN channels is easy to perform optimally, it



**FIGURE 2.** (a) The QPSK constellation used for information encoding and the corresponding received signals. We send QPSK signals over an AWGN channel, as shown, and try to detect the signals at the receiver. (b) The detection regions produced with a trained neural network. (c) The optimal detection regions obtained from detection theory.

makes little sense to utilize artificial neural networks for that purpose. There are many similar tasks in communications where deep learning cannot make any meaningful improvements. For example, the fundamental performance limits were derived by Shannon [4], and we can operate close to those limits using modern channel codes. Moreover, it is known how to perform optimal channel estimation, multiuser multiple-input multiple-output (MIMO) processing, and transmit power allocation in many wireless communication scenarios [5]. The fact that the information signals are human made gives us strong prior information that makes it easier to devise effective human-made algorithms than in many other fields, where the signals are created by nature.

There are, nevertheless, some important roles that deep learning can play in communications. First, there are many problems where a known algorithm finds the optimum but has prohibitively high complexity for real-time implementation. Second, there are cases where the standard system models are inadequate or incomplete. It is sufficient to replace the noise distribution in the previous example with an unknown one to find a case where learning can help. We elaborate on these two applications in the remainder of this lecture note. Before that, however, we stress that errors are unavoidable in the physical layer of communication systems and are conventionally dealt with using retransmissions. This built-in fault tolerance is positive when it comes to the utilization

of deep learning. It gives robustness to the strange behaviors that occasionally occur when an atypical signal is fed into a neural network that has been trained to work well for typical input signals. However, adversaries can also exploit atypical signals to perform jamming more efficiently [6].

### Application 1: Algorithmic approximation

The first important application of deep learning in communications is to approximate a known but computationally complicated algorithm. There are many examples of iterative algorithms that asymptotically find a global (or local) optimum to an optimization problem but require very many iterations for convergence and/or complicated operations in each iteration

[7]. Such algorithms might not be practically useful in communication systems where latency constraints require execution times below a millisecond.

The general procedure for training a neural network for algorithmic approximation is illustrated in Figure 3. Suppose we have a known algorithm, represented by the function  $y = f(\mathbf{x})$ , that cannot be implemented in real time. To address this problem using deep learning, we can first create a training set containing a large number  $T$  of input signals  $\mathbf{x}_t^{\text{train}}$  for  $t = 1, \dots, T$ . We then run the algorithm  $T$  times to compute the outputs:

$$\mathbf{y}_t^{\text{train}} = f(\mathbf{x}_t^{\text{train}}). \quad (9)$$

After having generated the training set, we can train an artificial neural network to provide approximately the same outputs for these inputs. More precisely, we should find an optimized parameter vector  $\theta^*$  in accordance with (5). If the training is performed well, the neural network will generalize well (i.e., provide good outputs) to previously unseen input signals that were generated in the same way as the inputs used for training. Simply speaking, this means that

$f(\mathbf{x}) \approx \hat{f}(\mathbf{x}; \theta^*)$  for all inputs  $\mathbf{x}$  of practical interest.

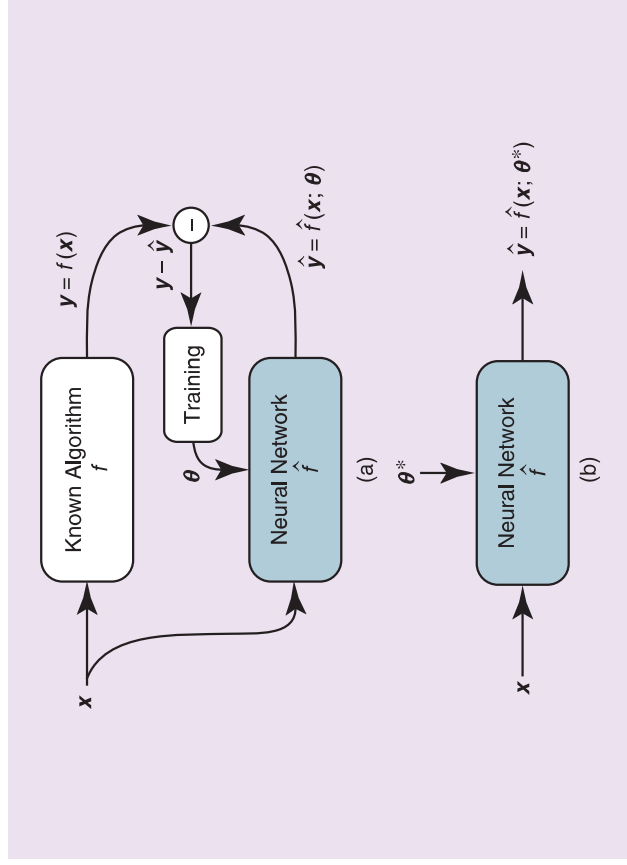
There are many optimization problems to be solved in communication systems. For example, at the transmitter, power allocation between concurrent transmissions is important to limit interference [5], [7]. At the receiver, nonlinear signal detection problems must be solved to deal with interference in MIMO systems [8]. Some of these problems are convex and can be solved by off-the-shelf optimization software. Other problems are nonconvex, but there exist iterative algorithms that converge to local or global optima. In both cases, the computational complexity is often prohibitive for real-time applications, where similar optimization problems with different input data are solved repeatedly. A neural network can, then, be trained to learn approximately how the solution depends on the input data. This approximate input–output map can be evaluated with substantially lower computational cost, as exemplified in [7] and [8]. Domain knowledge can be utilized to preprocess the input data to focus the learning on the problem that the algorithm is solving and not on rediscovering known properties (e.g.,

that the desired signal lies in a certain subspace).

There are two main approaches. One can learn the input–output mapping based on training data, as described previously, while ignoring how it was produced. Alternatively, the shape of the neural network can be selected so that each layer mimics one iteration of a known algorithm that converges asymptotically to an optimum. This is called *deep unfolding* and exploits the fact that many first-order iterative optimization methods have the same structure as a (recurrent) neural network [9]. The parameters of the neural network are, then, trained to give a nearly optimum solution after a predefined number of iterations, thereby speeding up the convergence. In [8], the authors unfold a gradient-descent-like algorithm for MIMO detection to create a neural network where each layer performs similar operations but with optimized parameters. When using this approach, (9) needs not to be determined in advance, which simplifies the training.

The practical benefit of this application is the complexity reduction it can provide; the neural network will, essentially, learn how to make algorithmic shortcuts to strike a good balance between accuracy and computational complexity. Another important benefit is related to hardware implementation. To solve a practical problem with real-time constraints, we conventionally would first need to design an algorithm and then develop a dedicated circuit based on it, which can be very time consuming. With the help of deep learning, we can, instead, predesign a general-purpose circuit that implements a neural network of a given maximum size (i.e., the numbers of layers and neurons) with a predetermined runtime. We can, then, train a neural network to perform the algorithmic task we need and, finally, load the corresponding trained parameters (i.e., weights and biases) onto the circuit. This new approach to hardware implementation can greatly reduce the time from when the algorithmic design begins to when a product can hit the market.

A main issue with this application is the highly computationally demanding



**FIGURE 3.** (a) A known algorithm  $f$  can be approximated by training a neural network  $\hat{f}$  to make  $f(\mathbf{x}) \approx \hat{f}(\mathbf{x}; \theta^*)$  for all possible inputs, as shown. The training procedure will iteratively update  $\theta$  to gradually reduce the approximation errors until it converges to some  $\theta^*$ . (b) This how the trained neural network is used at runtime. If the neural network is designed to have sufficiently low complexity, then the trained neural network shown can be used in real-time applications.

generation of desired outputs: the more complex the algorithm  $f$  is, the longer the time it takes to compute  $f(\mathbf{x}_t^{\text{train}})$  for  $t = 1, \dots, T$ . We are, basically, moving the complexity issue from the algorithmic runtime to the design process. There is a practical limit to which algorithms we can approximate in this way. If it takes 1 h to generate one training example, it will take 11.4 years or extreme parallelism to generate 100,000 examples.

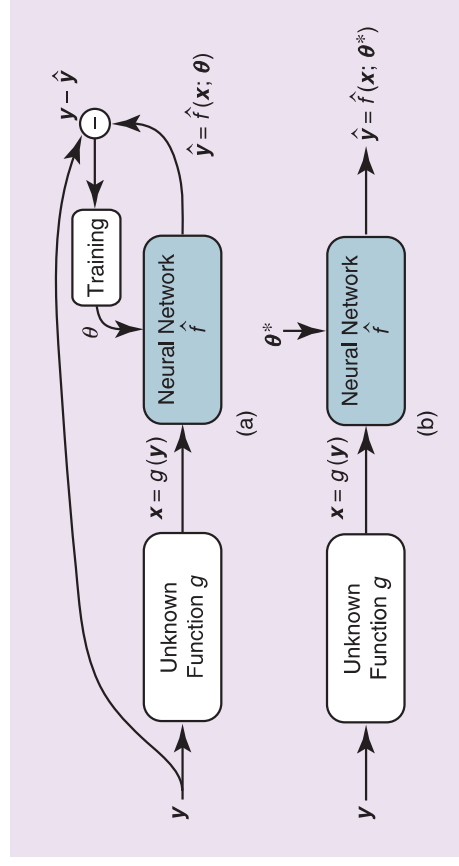
## Application 2: Inversion of an unknown function

The second important application is to invert an unknown function. In particular, nonlinear distortion can occur between the transmitter and receiver. Three prominent examples are finite-resolution quantization in the receiver hardware, nonlinear amplifiers in the transmitter hardware [10], and nonlinear fiber-optic channels [11]. While quantizers, typically, are designed with known properties, the latter two examples can be represented by an unknown function  $g$  that takes a signal  $\mathbf{y}$  as input and produces a distorted output  $\mathbf{x} = g(\mathbf{y})$ . The conventional way to undo the distortion is to identify an appropriate parameterized model of the function; then, estimate the parameters from measurements; and, finally, create an inverse function based on the estimates. This three-step approach is suboptimal and prone to error propagation. An alternative is to train a neural network to directly invert the function, without explicit modeling or parameter estimation. Whenever only suboptimal conventional algorithms exist, a learned algorithm can theoretically provide better performance and robustness, but only if the training is carried out successfully.

The general procedure for training a neural network for function inversion is illustrated in Figure 4. We need to generate a large number  $T$  of possible communication signals  $\mathbf{y}_t^{\text{train}}$  and send them through the unknown function to measure

$$\mathbf{x}_t^{\text{train}} = g(\mathbf{y}_t^{\text{train}}) \quad \text{for } t = 1, \dots, T. \quad (10)$$

It is, then,  $\mathbf{x}_t^{\text{train}}$  that is used as input to the neural network, while  $\mathbf{y}_t^{\text{train}}$  is the desired output.



**FIGURE 4.** (a) An unknown function  $g$  with input  $\mathbf{y}$  is inverted using a neural network  $\hat{f}$  by training it to achieve  $\hat{f}(g(\mathbf{y}); \theta) \approx \mathbf{y}$ , as shown. The training procedure will iteratively update  $\theta$  to gradually reduce the approximation errors until it converges to some  $\theta^*$ . (b) The trained neural network can be used to counteract the unknown function at runtime without having to explicitly model it and estimate model parameters.

Different from application 1, the creation of a training set can be very computationally efficient in application 2 because the outputs are human made. It is, typically, created to be statistically equivalent to the signals observed at runtime, but one can also create a biased training set to emphasize typical or atypical examples. Online learning when operating the communication system is possible by occasionally sending predefined reference signals to generate new training data. This is useful when the function  $g$  is time varying (e.g., due to temperature variations in the hardware). The key to successful utilization of deep learning is to identify tasks in communication systems that currently lack an optimal solution—there is, then, an opportunity to beat the state of the art. For example, a common way to deal with nonlinear communication hardware is to apply the Bussgang decomposition [12] to write the output of the nonlinear function  $g$  as  $g(\mathbf{y}) = \mathbf{D}\mathbf{y} + \mathbf{n}$ , where  $\mathbf{D}$  is a deterministic matrix, and  $\mathbf{n}$  is the distortion noise that is uncorrelated with  $\mathbf{y}$  but statistically dependent. By pretending as if  $\mathbf{n}$  were independent noise, one can often develop communication algorithms (e.g., for channel estimation or data detection) that partially mitigate distortion, but such algorithms are suboptimal since the distortion is, in fact, dependent on the input. As shown in [10], one can achieve substantially

better performance by training neural networks instead.

## What we have learned

Although many parts of communication systems can be designed optimally, there are important cases where deep learning can give large improvements. In particular, it can be used to reduce the computational complexity of known algorithms or deal with nonlinear hardware or channels in an efficient way.

## Acknowledgments

This work was supported by the Wallenberg AI, Autonomous Systems, and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

## Authors

**Emil Björnson** (emil.bjornson@liu.se) received his M.Sc. degree in engineering mathematics from Lund University, Sweden, in 2007 and his Ph.D. degree in telecommunications from the KTH Royal Institute of Technology, Sweden, in 2011. He is currently an associate professor at Linköping University, Sweden. He has authored *Optimal Resource Allocation in Coordinated Multi-Cell Systems* (2013) and *Massive MIMO Networks: Spectral, Energy, and Hardware Efficiency* (2017). He received the 2018 IEEE Marconi Prize Paper Award in Wireless Communications, the 2019 EURASIP Early CAREER Award,

the 2019 IEEE Communications Society Fred W. Ellersick Prize, and the 2019 IEEE Signal Processing Magazine Best Column Award.

**Pontus Giselsson** (pontus.giselsson@control.lth.se) received his M.Sc. and Ph.D. degrees from Lund University, Sweden, in 2006 and 2012, respectively. During 2013 and 2014, he held a postdoctoral position at Stanford University. He is an associate professor in the Department of Automatic Control at Lund University. In 2012, he received the Young Author Prize at the Advanced Control of Chemical Processes IFAC Symposium; in 2014, he received the Young Author Prize at the IFAC World Congress; and in 2015, he received the Ingvar Carlsson Award from the Swedish Foundation for Strategic Research. His research interests include

mathematical optimization and its wide range of applications, e.g., in machine learning, control, signal processing, and wireless communication.

## References

- [1] G. Cybenko, "Approximation by superpositions of a sigmoidal function," *Math. Control, Signals Syst.*, vol. 2, no. 4, pp. 303–314, Dec. 1989. doi: 10.1007/BF02551274.
- [2] J. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA: MIT Press, 2016.
- [3] U. Madhok, *Introduction to Communication Systems*. Cambridge, U.K.: Cambridge Univ. Press, 2014.
- [4] C. E. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.*, vol. 27, no. 3, pp. 379–423, 1948. doi: 10.1002/j.1538-7305.1948.tb00917.x.
- [5] E. Björnson, J. Hoydis, and L. Sanguinetti, "Massive MIMO networks: Spectral, energy, and hardware efficiency," *Found. Trends Signal Process.*, vol. 11, no. 3–4, pp. 154–655, 2017. doi: 10.1561/20000000093.
- [6] M. Sadeghi and E. G. Larsson, "Adversarial attacks on deep-learning based radio signal classification," *IEEE*

*Wireless Commun. Lett.*, vol. 8, no. 1, pp. 213–216, Feb. 2019. doi: 10.1109/LWC.2018.2867459.

[7] H. Sun, X. Chen, Q. Shi, M. Hong, X. Fu, and N. D. Sidiropoulos, "Learning to optimize: Training deep neural networks for interference management," *IEEE Trans. Signal Process.*, vol. 66, no. 20, pp. 5438–5453, Oct. 2018. doi: 10.1109/TSP.2018.2866382.

[8] N. Samuel, T. Diskin, and A. Wiesel, "Deep MIMO detection," in *Proc. IEEE 18th Int. Workshop Signal Processing Advances Wireless Communications (SPAWC)*, July 2017, pp. 1–5. doi: 10.1109/SPAWC.2017.8227772.

[9] J. R. Hershey, J. L. Roux, and F. Weninger, "Deep unfolding: Model-based inspiration of novel deep architectures," 2014. [Online]. Available: arXiv:abs/1409.2574

[10] Ö. T. Demir and E. Björnson, "Channel estimation in massive MIMO under hardware non-linearities: Bayesian methods versus deep learning," *IEEE Open J. Commun. Soc.*, vol. 1, no. 1, pp. 109–124, 2020. doi: 10.1109/OJCOMS.2019.2959913.

[11] A. D. Ellis, J. Zhao, and D. Cotter, "Approaching the non-linear Shannon limit," *J. Lightwave Technol.*, vol. 28, no. 4, pp. 423–433, Feb. 2010. doi: 10.1109/JLT.2009.2030693.

[12] J. J. Bussgang, "Crosscorrelation functions of amplitude-distorted Gaussian signals," RLE, MIT, Cambridge, MA, Tech. Rep. 216, 1952.

# Closed-Form Output Response of Discrete-Time Linear Time-Invariant Systems Using Intermediate Auxiliary Functions

Chun-Lin Liu and Soo-Chang Pei

Linear time-invariant (LTI) systems find ubiquitous applications in digital signal processing [1], image processing [2], communication [3], and array processing [4], to name a few. An LTI system can be characterized by its impulse response, which is defined as the output

response to an impulse input. For an arbitrary input signal, the output signal of an LTI system is obtained by the convolution of the input signal with the impulse response. This property is closely related to transform-based methods in signal processing, such as discrete-time Fourier transforms,  $z$ -transforms, and discrete Fourier transforms. Interested readers are referred to [1] and the references therein.

## Relevance

In the literature, the impulse response and the convolution are typically analyzed in the transform domain [1]. For instance, the convolution in the time domain is equivalent to multiplication in the  $z$ -transform domain. However, this approach requires an additional "inverse transform" step to find the time-domain signal, which could be