

Fundamentals

$$z = \underbrace{\mathbf{w}[0] \times \mathbf{d}[0] + \mathbf{w}[1] \times \mathbf{d}[1] + \cdots + \mathbf{w}[m] \times \mathbf{d}[m]}_{\text{weighted sum}} \quad (1)$$

$$= \sum_{j=0}^m \mathbf{w}[j] \times \mathbf{d}[j]$$

$$= \underbrace{\mathbf{w} \cdot \mathbf{d}}_{\text{dot product}} = \underbrace{\mathbf{w}^T \mathbf{d}}_{\text{matrix product}} = [w_0, w_1, \dots, w_m] \begin{bmatrix} d_0 \\ d_1 \\ \vdots \\ d_m \end{bmatrix} \quad (2)$$

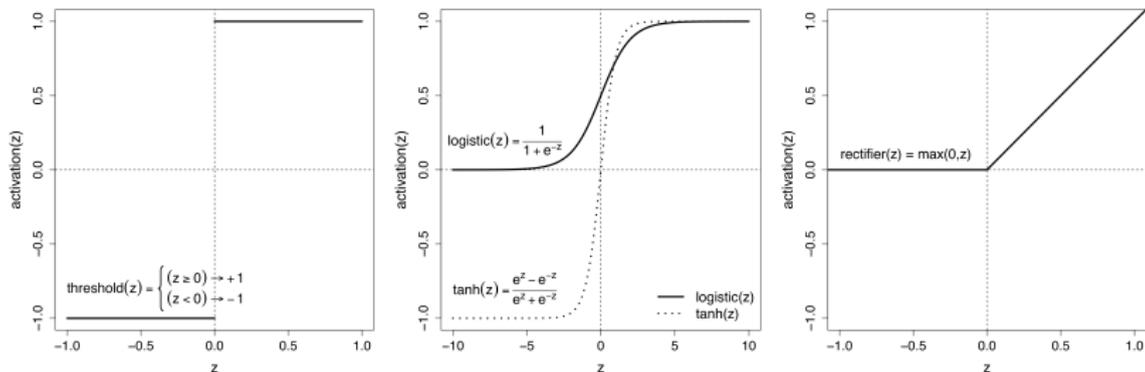


Figure 2: Plots for activation functions that have been popular in the history of neural networks.



$$\begin{aligned} M_{\mathbf{w}}(\mathbf{d}) &= \varphi (\mathbf{w} [0] \times \mathbf{d} [0] + \mathbf{w} [1] \times \mathbf{d} [1] + \dots + \mathbf{w} [m] \times \mathbf{d} [m]) \\ &= \varphi \left(\sum_{i=0}^m w_i \times d_i \right) = \varphi \left(\underbrace{\mathbf{w} \cdot \mathbf{d}}_{\text{dot product}} \right) \\ &= \varphi \left(\underbrace{\mathbf{w}^T \mathbf{d}}_{\text{matrix product}} \right) = \varphi \left([w_0, w_1, \dots, w_m] \begin{bmatrix} d_0 \\ d_1 \\ \vdots \\ d_m \end{bmatrix} \right) \end{aligned} \tag{5}$$

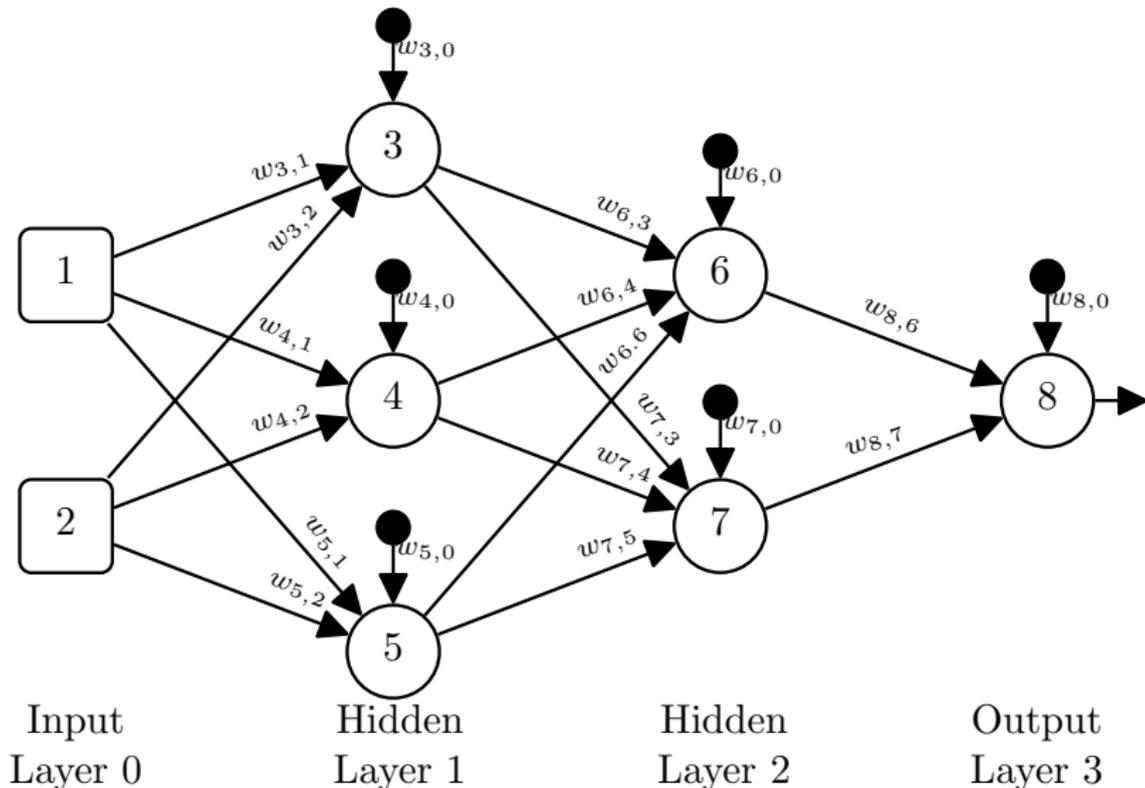
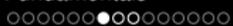


Figure 4: A schematic of a feedforward artificial neural network.



$$\mathbf{z}^{(2)} = \mathbf{W}^{(2)} \mathbf{a}^{(1)} \tag{6}$$



Neural Networks as Matrix Operations

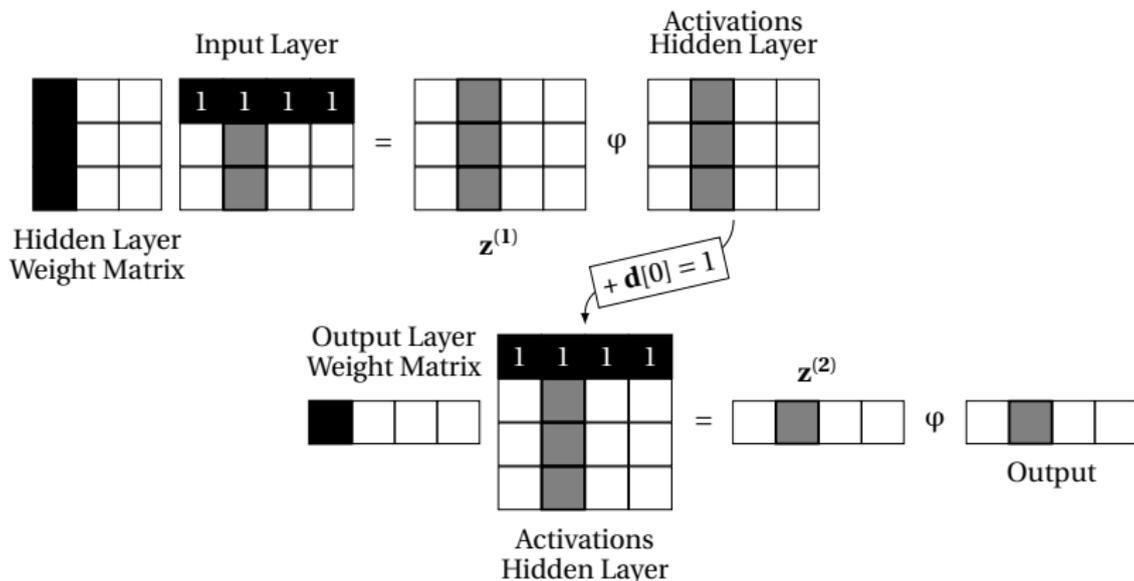


Figure 6: An illustration of how a batch of examples can be processed in parallel using matrix operations.

Why Is Network Depth Important?

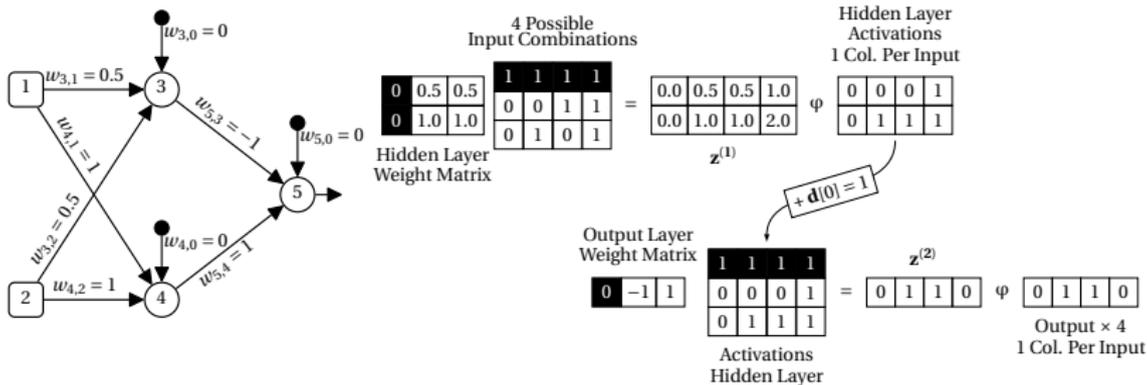


Figure 9: (left) The XOR function implemented as a two-layer neural network. (right) The network processing the four possible input combinations, one combination plus bias input per column: $[bias, FALSE, FALSE] \rightarrow [1, 0, 0]$; $[bias, FALSE, TRUE] \rightarrow [1, 0, 1]$; $[bias, TRUE, FALSE] \rightarrow [1, 1, 0]$; $[bias, TRUE, TRUE] \rightarrow [1, 1, 1]$.



Standard Approach: Backpropagation and Gradient Descent

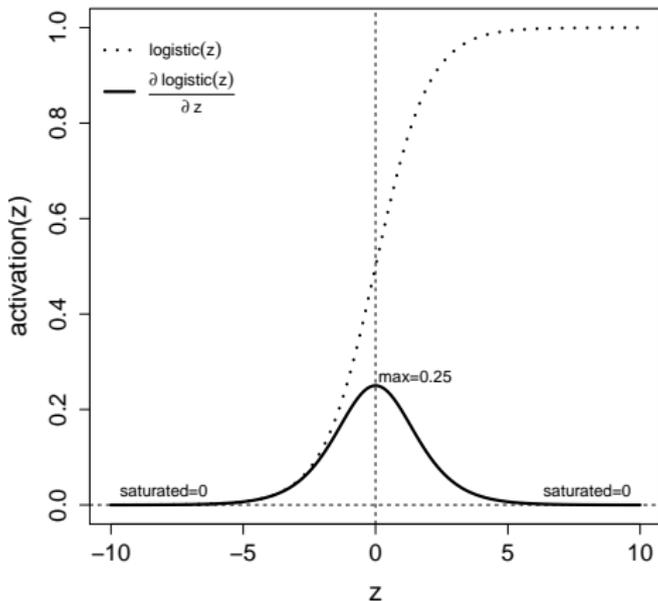


Figure 13: Plots of the logistic function and its derivative. This figure is Figure 4.6 of (Kelleher, 2019) and is used here with permission.

$$L_2(\mathbb{M}_{\mathbf{w}}, \mathcal{D}) = \frac{1}{2} \sum_{i=1}^n (t_i - \mathbb{M}_{\mathbf{w}}(\mathbf{d}_i))^2 \quad (17)$$

$$\frac{\partial}{\partial \mathbf{w}[j]} L_2(\mathbb{M}_{\mathbf{w}}, \mathbf{d}) = (t - \mathbb{M}_{\mathbf{w}}(\mathbf{d})) \times -\mathbf{d}[j] \quad (18)$$

$$\frac{\partial \mathcal{E}}{\partial a_k} = \frac{\partial L_2(\mathbb{M}_{\mathbf{w}}, \mathbf{d})}{\partial \mathbb{M}_{\mathbf{w}}(\mathbf{d})} = t - \mathbb{M}_{\mathbf{w}}(\mathbf{d}) = t_k - a_k \quad (19)$$

$$\frac{\partial \mathcal{E}}{\partial a_k} = -(t_k - a_k) \quad (20)$$

$$\begin{aligned}
 \delta_k &= \frac{\partial a_k}{\partial z_k} \times \frac{\partial \mathcal{E}}{\partial a_k} \\
 &= \frac{\partial a_k}{\partial z_k} \times -(t_k - a_k) \\
 &= \underbrace{\frac{d}{dz} \text{logistic}(z)}_{\text{Assuming a logistic activation function}} \times -(t_k - a_k) \\
 &= \underbrace{(\text{logistic}(z) \times (1 - \text{logistic}(z)))}_{\text{Assuming a logistic activation function}} \times -(t_k - a_k) \quad (21)
 \end{aligned}$$

Require: set of training instances \mathcal{D}

Require: a learning rate α that controls how quickly the algorithm converges

Require: a batch size B specifying the number of examples in each batch

Require: a convergence criterion

1: Shuffle \mathcal{D} and create the mini-batches: $[(\mathbf{X}^{(1)}, \mathbf{Y}^{(1)}), \dots, (\mathbf{X}^k, \mathbf{Y}^k)]$

2: Initialize the weight matrices for each layer: $\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(L)}$

3: **repeat** ▷ Each repeat loop is one epoch

4: **for** $t=1$ to number of mini-batches **do** ▷ Each for loop is one iteration

5: $\mathbf{A}^{(0)} \leftarrow \mathbf{X}^{(t)}$

6: **for** $l=1$ to L **do**

7: $\mathbf{v} \leftarrow [1_0, \dots, 1_m]$ ▷ Create \mathbf{v} the vector of bias terms

8: $\mathbf{A}^{(l-1)} \leftarrow [\mathbf{v}; \mathbf{A}^{(l-1)}]$ ▷ Insert \mathbf{v} into the activation matrix

9: $\mathbf{Z}^{(l)} \leftarrow \mathbf{W}^l \mathbf{A}^{(l-1)}$

10: $\mathbf{A}^{(l)} \leftarrow \varphi(\mathbf{Z}^{(l)})$ ▷ Elementwise application of φ to $\mathbf{Z}^{(l)}$

11: **end for**

12: **for** each weight $w_{i,k}$ in the network **do**

13: $\Delta w_{i,k} = 0$

14: **end for**

15: **for** each example in the mini-batch **do** ▷ Backpropagate the δ s

16: **for** each neuron i in the output layer **do**

17: $\delta_i = \frac{\partial \mathcal{E}}{\partial a_i} \times \frac{\partial a_i}{\partial z_i}$ ▷ See Equation (21)^[28]

18: **end for**

19: **for** $l = L-1$ to 1 **do**

20: **for** each neuron i in the layer l **do**

21: $\delta_i = \frac{\partial \mathcal{E}}{\partial a_i} \times \frac{\partial a_i}{\partial z_i}$ ▷ See Equation (23)^[29]

22: **end for**

23: **end for**

24: **for** each weight $w_{i,k}$ in the network **do**

25: $\Delta w_{i,k} = \Delta w_{i,k} + (\delta_i \times a_k)$ ▷ Equation (29)^[33]

26: **end for**

27: **end for**

28: **for** each weight $w_{i,k}$ in the network **do**

29: $w_{i,k} \leftarrow w_{i,k} - \alpha \times \Delta w_{i,k}$ ▷ Equation (30)^[33]

30: **end for**

31: **end for**

32: shuffle($[(\mathbf{X}^{(1)}, \mathbf{Y}^{(1)}), \dots, (\mathbf{X}^k, \mathbf{Y}^k)]$)

33: **until** convergence occurs



Table 1: Hourly samples of ambient factors and full load electrical power output of a combined cycle power plant.

ID	AMBIENT TEMPERATURE °C	RELATIVE HUMIDITY %	ELECTRICAL OUTPUT MW
1	03.21	86.34	491.35
2	31.41	68.50	430.37
3	19.31	30.59	463.00
4	20.64	99.97	447.14



Table 3: The *range-normalized* hourly samples of ambient factors and full load electrical power output of a combined cycle power plant, rounded to two decimal places.

ID	AMBIENT TEMPERATURE °C	RELATIVE HUMIDITY %	ELECTRICAL OUTPUT MW
1	0.04	0.81	0.94
2	0.84	0.58	0.13
3	0.50	0.07	0.57
4	0.53	1.00	0.36

A Worked Example: Using Backpropagation to Train a Feedforward Network for a Regression Task

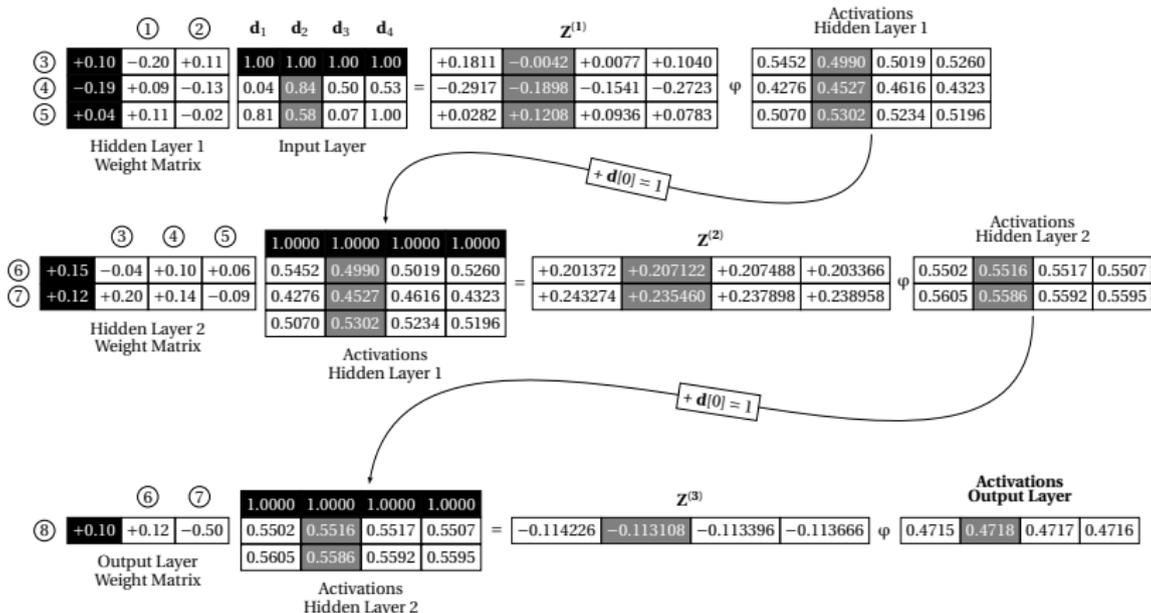


Figure 14: The forward pass of the examples listed in Table 3^[38] through the network in Figure 4^[11].

Table 4: The per example error after the forward pass illustrated in Figure 14^[39], the per example $\partial\mathcal{E}/\partial a_8$, and the **sum of squared errors** for the model over the dataset of four examples.

	d_1	d_2	d_3	d_4
Target	0.9400	0.1300	0.5700	0.3600
Prediction	0.4715	0.4718	0.4717	0.4716
Error	0.4685	-0.3418	0.0983	-0.1116
$\partial\mathcal{E}/\partial a_8$: Error $\times -1$	-0.4685	0.3418	-0.0983	0.1116
Error ²	0.21949225	0.11682724	0.00966289	0.01245456
SSE:				0.17921847

A Worked Example: Using Backpropagation to Train a Feedforward Network for a Regression Task

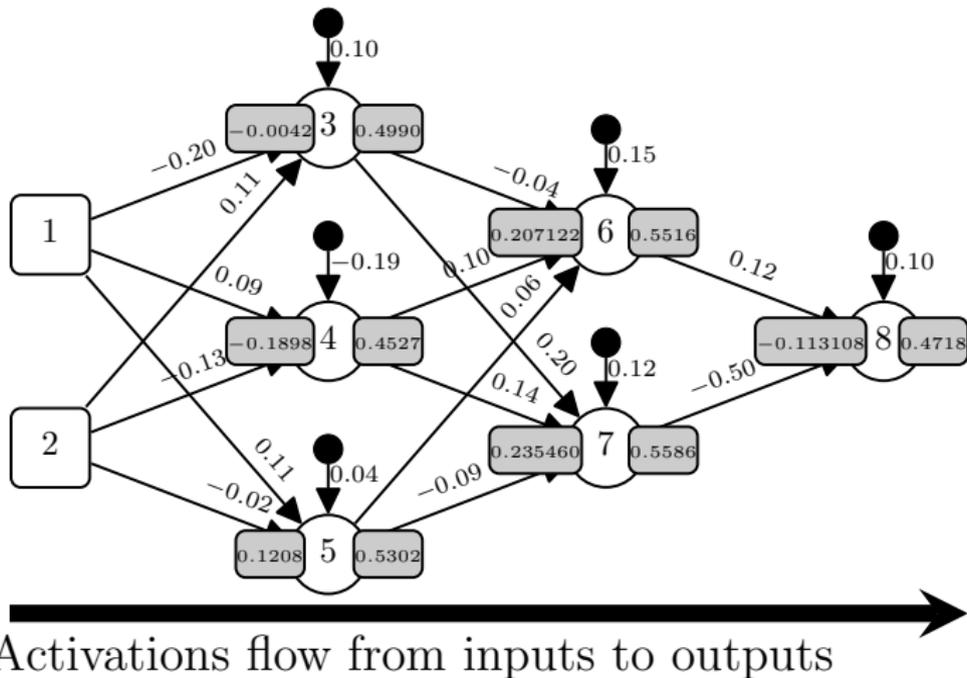


Figure 15: An illustration of the forward propagation of d_2 through the network showing the weights on each connection, and the weighted sum z and activation a value for each neuron in the network.



Table 5: The $\partial a/\partial z$ for each neuron for Example 2 rounded to four decimal places.

NEURON	z	$\partial a/\partial z$
3	-0.004200	0.2500
4	-0.189800	0.2478
5	0.120800	0.2491
6	0.207122	0.2473
7	0.235460	0.2466
8	-0.113108	0.2492



A Worked Example: Using Backpropagation to Train a Feedforward Network for a Regression Task

$$\begin{aligned}\delta_6 &= \frac{\partial \mathcal{E}}{\partial a_6} \times \frac{\partial a_6}{\partial z_6} \\ &= \left(\sum \delta_i \times w_{i,6} \right) \times \frac{\partial a_6}{\partial z_6} \\ &= (\delta_8 \times w_{8,6}) \times \frac{\partial a_6}{\partial z_6} \\ &= (0.0852 \times 0.12) \times 0.2473 \\ &= 0.0025\end{aligned}\tag{33}$$



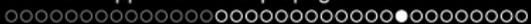
A Worked Example: Using Backpropagation to Train a Feedforward Network for a Regression Task

$$\begin{aligned}\delta_7 &= \frac{\partial \mathcal{E}}{\partial a_7} \times \frac{\partial a_7}{\partial z_7} \\ &= \left(\sum \delta_i \times w_{i,7} \right) \times \frac{\partial a_7}{\partial z_7} \\ &= (\delta_8 \times w_{8,7}) \times \frac{\partial a_6}{\partial z_6} \\ &= (0.0852 \times -0.50) \times 0.2466 \\ &= -0.0105\end{aligned}\tag{34}$$



A Worked Example: Using Backpropagation to Train a Feedforward Network for a Regression Task

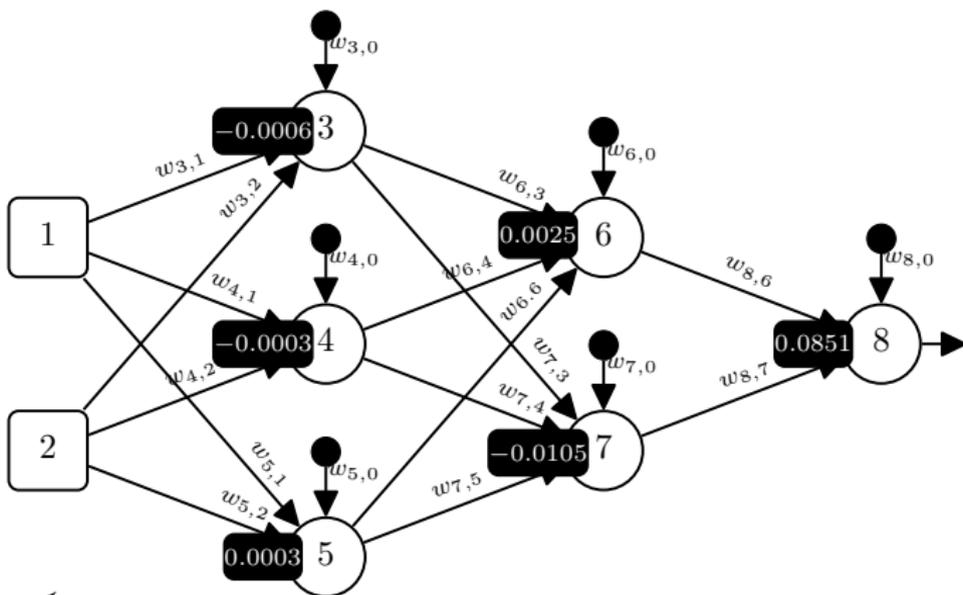
$$\begin{aligned}\delta_4 &= \frac{\partial \mathcal{E}}{\partial a_4} \times \frac{\partial a_4}{\partial z_4} \\ &= \left(\sum \delta_i \times w_{i,4} \right) \times \frac{\partial a_4}{\partial z_4} \\ &= ((\delta_6 \times w_{6,4}) + (\delta_7 \times w_{7,4})) \times \frac{\partial a_4}{\partial z_4} \\ &= ((0.0025 \times 0.10) + (-0.0105 \times 0.14)) \times 0.2478 \\ &= -0.0003 \quad (36)\end{aligned}$$



A Worked Example: Using Backpropagation to Train a Feedforward Network for a Regression Task

$$\begin{aligned}\delta_5 &= \frac{\partial \mathcal{E}}{\partial a_5} \times \frac{\partial a_5}{\partial z_5} \\ &= \left(\sum \delta_i \times w_{i,5} \right) \times \frac{\partial a_5}{\partial z_5} \\ &= ((\delta_6 \times w_{6,5}) + (\delta_7 \times w_{7,5})) \times \frac{\partial a_5}{\partial z_5} \\ &= ((0.0025 \times 0.06) + (-0.0105 \times -0.09)) \times 0.2491 \\ &= 0.0003 \quad (37)\end{aligned}$$

A Worked Example: Using Backpropagation to Train a Feedforward Network for a Regression Task



Error gradients (δ s) flow from outputs to inputs

Figure 16: The δ s for each of the neurons in the network for Ex. 2.

Table 6: The $\partial\mathcal{E}/\partial w_{i,k}$ calculations for d_2 for every weight in the network. The neuron index 0 denotes the bias input for each neuron.

NEURON _{<i>i</i>}	NEURON _{<i>k</i>}	$w_{i,k}$	δ_i	a_k	$\partial\mathcal{E}/\partial w_{i,k}$
8	0	$w_{8,0}$	0.0852	1	$0.0852 \times 1 = 0.0852$
8	6	$w_{8,6}$	0.0852	0.5516	$0.0852 \times 0.5516 = 0.04699632$
8	7	$w_{8,7}$	0.0852	0.5586	$0.0852 \times 0.5586 = 0.04759272$
7	0	$w_{7,0}$	-0.0105	1	$-0.0105 \times 1 = -0.0105$
7	3	$w_{7,3}$	-0.0105	0.4990	$-0.0105 \times 0.4527 = -0.0052395$
7	4	$w_{7,4}$	-0.0105	0.4527	$-0.0105 \times 0.4527 = -0.00475335$
7	5	$w_{7,5}$	-0.0105	0.5302	$-0.0105 \times 0.5302 = -0.0055671$
6	0	$w_{6,0}$	0.0025	1	$0.0025 \times 1 = 0.0025$
6	3	$w_{6,3}$	0.0025	0.4990	$0.0025 \times 0.4527 = 0.0012475$
6	4	$w_{6,4}$	0.0025	0.4527	$0.0025 \times 0.4527 = 0.00113175$
6	5	$w_{6,5}$	0.0025	0.5302	$0.0025 \times 0.5302 = 0.0013255$
5	0	$w_{5,0}$	0.0003	1	$0.0003 \times 1 = 0.0003$
5	1	$w_{5,1}$	0.0003	0.84	$0.0003 \times 0.84 = 0.000252$
5	2	$w_{5,2}$	0.0003	0.58	$0.0003 \times 0.58 = 0.000174$
4	0	$w_{4,0}$	-0.0003	1	$-0.0003 \times 1 = -0.0003$
4	1	$w_{4,1}$	-0.0003	0.84	$-0.0003 \times 0.84 = -0.000252$
4	2	$w_{4,2}$	-0.0003	0.58	$-0.0003 \times 0.58 = -0.000174$
3	0	$w_{3,0}$	-0.0006	1	$-0.0006 \times 1 = -0.0006$
3	1	$w_{3,1}$	-0.0006	0.84	$-0.0006 \times 0.84 = -0.000504$
3	2	$w_{3,2}$	-0.0006	0.58	$-0.0006 \times 0.58 = -0.000348$

A Worked Example: Using Backpropagation to Train a Feedforward Network for a Regression Task

$$\begin{aligned}
 w_{7,5} &= w_{7,5} - \alpha \times \delta_7 \times a_5 \\
 &= w_{7,5} - \alpha \times \frac{\partial \mathcal{E}}{\partial w_{i,k}} \\
 &= -0.09 - 0.2 \times -0.0055671 \\
 &= -0.08888658 \qquad (38)
 \end{aligned}$$

Table 7: The calculation of $\Delta w_{7,5}$ across our four examples.

MINI-BATCH EXAMPLE	$\frac{\partial \mathcal{E}}{\partial w_{7,5}}$
\mathbf{d}_1	0.00730080
\mathbf{d}_2	-0.00556710
\mathbf{d}_3	0.00157020
\mathbf{d}_4	-0.00176664
$\Delta w_{7,5} =$	0.00153726

A Worked Example: Using Backpropagation to Train a Feedforward Network for a Regression Task

$$\begin{aligned}w_{7,5} &= w_{7,5} - \alpha \times \Delta w_{i,k} \\ &= -0.09 - 0.2 \times 0.00153726 \\ &= -0.0903074520\end{aligned}\tag{39}$$

Table 8: The per example error after each weight has been updated once, the per example $\partial\mathcal{E}/\partial a_8$, and the **sum of squared errors** for the model.

	d_1	d_2	d_3	d_4
Target	0.9400	0.1300	0.5700	0.3600
Prediction	0.4738	0.4741	0.4740	0.4739
Error	0.4662	-0.3441	0.0960	-0.1139
$\partial\mathcal{E}/\partial a_8$: Error $\times -1$	-0.4662	0.3441	-0.0960	0.1139
Error ²	0.21734244	0.11840481	0.009216	0.01297321
SSE:				0.17896823

Table 9: The per example prediction, error, and the sum of squared errors after training has converged to an $SSE < 0.0001$.

	d_1	d_2	d_2	d_2
Target	0.9400	0.1300	0.5700	0.3600
Prediction	0.9266	0.1342	0.5700	0.3608
Error	0.0134	-0.0042	0.0000	-0.0008
Error ²	0.00017956	0.00001764	0.00000000	0.00000064
SSE:				0.00009892

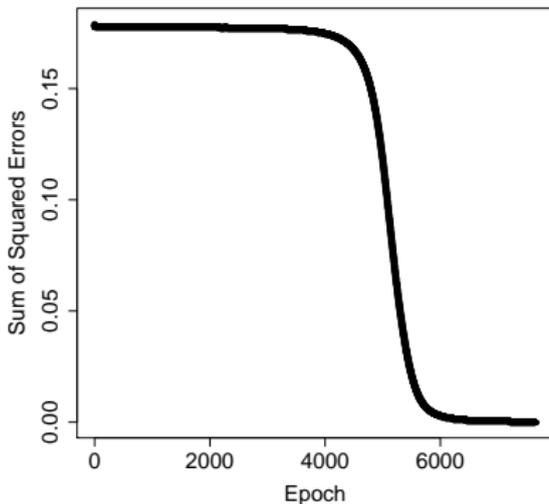
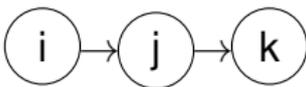


Figure 17: A plot showing how the sum of squared errors of the network changed during training.



$$\begin{aligned}
 \delta_i &= \frac{\partial \mathcal{E}}{\partial a_i} \times \frac{\partial a_i}{\partial z_i} \\
 &= w_{j,i} \times \overbrace{\frac{\partial \mathcal{E}}{\partial a_i}}^{\delta_j} \times \frac{\partial a_i}{\partial z_i} \\
 &= w_{j,i} \times w_{k,j} \times \overbrace{\frac{\partial \mathcal{E}}{\partial a_j}}^{\delta_k} \times \frac{\partial a_j}{\partial z_j} \times \frac{\partial a_i}{\partial z_i} \\
 &= w_{j,i} \times w_{k,j} \times \overbrace{\frac{\partial \mathcal{E}}{\partial a_k} \times \frac{\partial a_k}{\partial z_k}}^{\delta_k} \times \frac{\partial a_j}{\partial z_j} \times \frac{\partial a_i}{\partial z_i}
 \end{aligned}$$

(41)

Vanishing Gradients and ReLUs

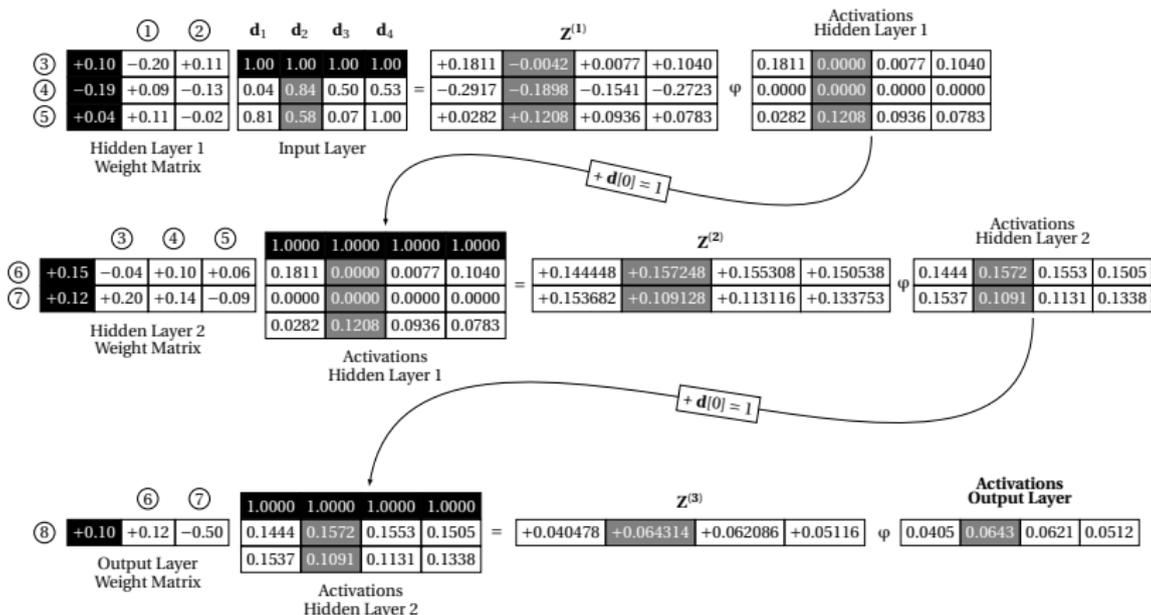


Figure 18: The forward pass of the examples listed in Table 3^[38] through the network in Figure 4^[11] when all the neurons are ReLUs.

$$\begin{aligned}
\delta_k &= \frac{\partial \mathcal{E}}{\partial a_k} \times \frac{\partial a_k}{\partial z_i} \\
\delta_8 &= -0.0657 \times 1.0 \\
&= -0.0657 \\
\delta_7 &= (\delta_8 \times w_{8,7}) \times \frac{\partial a_7}{\partial z_7} \\
&= (-0.0657 \times -0.50) \times 1 \\
&= 0.0329 \\
\delta_6 &= (\delta_8 \times w_{8,6}) \times \frac{\partial a_6}{\partial z_6} \\
&= (-0.0657 \times 0.12) \times 1 \\
&= -0.0079 \\
\delta_5 &= ((\delta_6 \times w_{6,5}) + (\delta_7 \times w_{7,5})) \times \frac{\partial a_5}{\partial z_5} \\
&= ((-0.0079 \times 0.06) + (0.0329 \times -0.09)) \times 1 \\
&= -0.0034 \\
\delta_4 &= ((\delta_6 \times w_{6,4}) + (\delta_7 \times w_{7,4})) \times \frac{\partial a_4}{\partial z_4} \\
&= ((-0.0079 \times 0.10) + (0.0329 \times 0.14)) \times 0 \\
&= 0 \\
\delta_3 &= ((\delta_6 \times w_{6,3}) + (\delta_7 \times w_{7,3})) \times \frac{\partial a_3}{\partial z_3} \\
&= ((-0.0079 \times -0.04) + (0.0329 \times 0.20)) \times 0 \\
&= 0
\end{aligned}$$

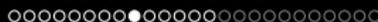


Table 12: The ReLU network's per example prediction, error, and the sum of squared errors after training has converged to an $SSE < 0.0001$.

	d_1	d_2	d_3	d_4
Target	0.9400	0.1300	0.5700	0.3600
Prediction	0.9487	0.1328	0.5772	0.3679
Error	-0.0087	-0.0028	-0.0072	-0.0079
Error ²	0.00007569	0.00000784	0.00005184	0.00006241
SSE:				0.00009889



$$\mathit{rectifier}_{\mathit{parametric}}(z_i) = \begin{cases} z_i & \text{if } z_i > 0 \\ \lambda_i \times z_i & \text{otherwise} \end{cases} \quad (47)$$

$$\frac{d}{dz} \mathit{rectifier}_{\mathit{parametric}}(z_i) = \begin{cases} 1 & \text{if } z_i > 0 \\ \lambda_i & \text{otherwise} \end{cases} \quad (48)$$

Weight Initialization and Unstable Gradients

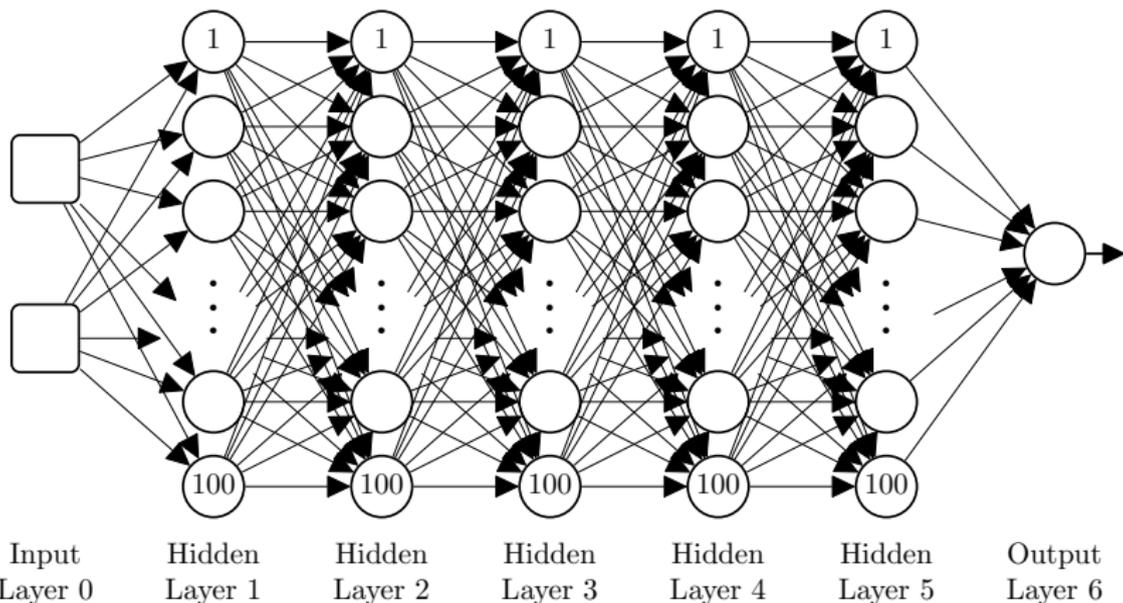


Figure 22: The architecture of the neural network used in the weight initialization experiments. Note that the neurons in this network use a linear activation function: $a_i = z_i$.



Weight Initialization and Unstable Gradients

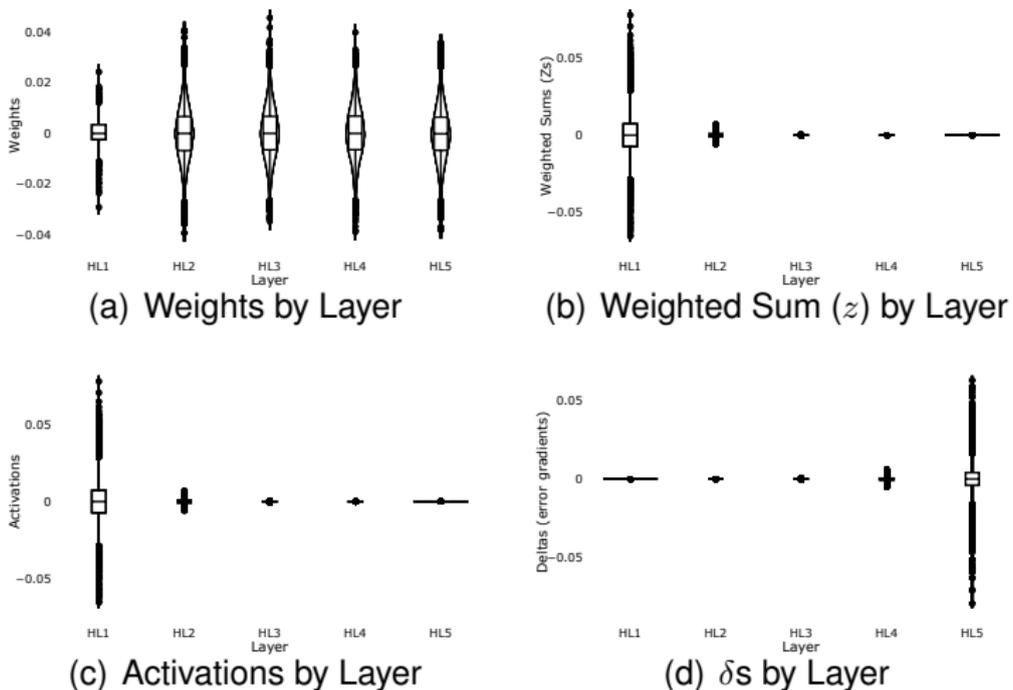


Figure 23: The internal dynamics of the network in Figure 22^[73] during the first training iteration when the weights were initialized using a normal distribution with $\mu=0.0$, $\sigma=0.01$.



$$z = (w_1 \times d_1) + (w_2 \times d_2) + \dots + (w_{n_{in}} \times d_{n_{in}}) \quad (53)$$



$$\text{var} \left(\sum_{i=1}^n X_i \right) = \sum_{i=1}^n \text{var} (X_i) \quad (54)$$



$$\text{var}(z) = \sum_{i=1}^{n_{in}} \text{var}(w_i \times d_i) = n_{in} \text{var}(\mathbf{W}) \text{var}(\mathbf{d}) \quad (58)$$



$$\begin{aligned} \text{var}(Z^{(HL2)}) &= n_{in}^{(HL2)} \times \text{var}(\mathbf{W}^{(HL2)}) \times \text{var}(\mathbf{d}^{(HL2)}) & (60) \\ &= 100 \times 0.0001 \times 0.0002 \\ &= 0.000002 \end{aligned}$$



$$\text{var}(\mathbf{W}^{(k)}) = \frac{2}{n_{in}^{(k)} + n_{out}^{(k)}} \quad (61)$$

$$\text{var}(\mathbf{W}^{(k)}) = \frac{1}{n_{in}^{(k)}} \quad (62)$$

Weight Initialization and Unstable Gradients

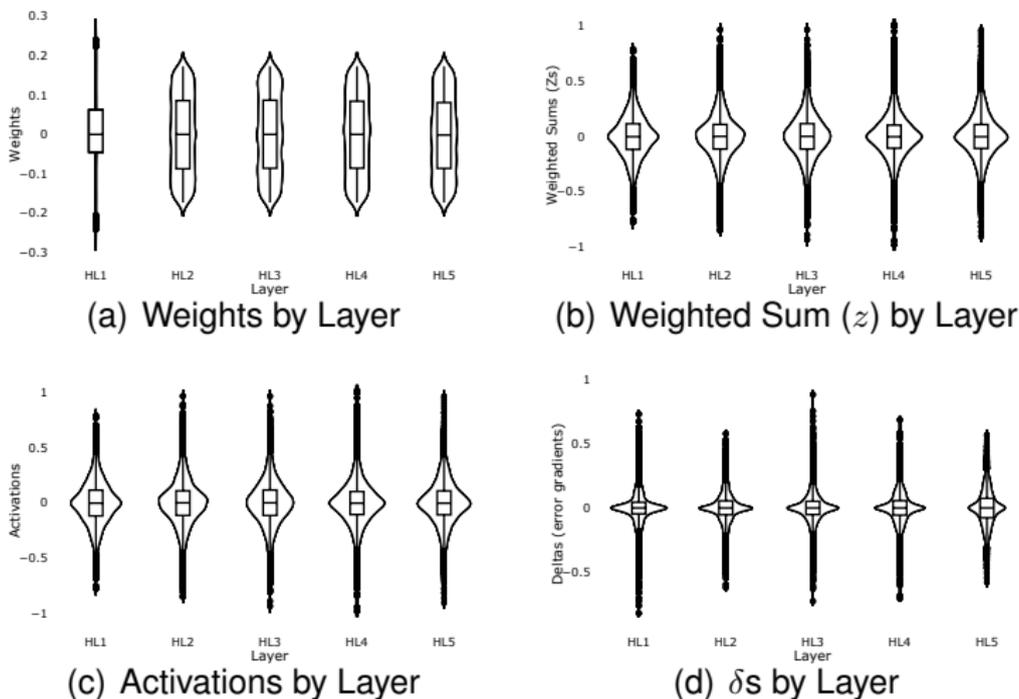
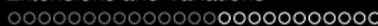


Figure 25: The internal dynamics of the network in Figure 22^[73] during the first training iteration when the weights were initialized using Xavier initialization.



$$\text{var}(\mathbf{W}^{(k)}) = \frac{2}{n_{in}^{(k)}} \quad (63)$$



$$\begin{aligned}W^{(1)} &\sim \mathcal{N}\left(0, \sqrt{\frac{1}{100}}\right) \\W^{(2)} &\sim \mathcal{N}\left(0, \sqrt{\frac{2}{80}}\right) \\W^{(3)} &\sim \mathcal{N}\left(0, \sqrt{\frac{2}{50}}\right)\end{aligned}\tag{64}$$



Weight Initialization and Unstable Gradients

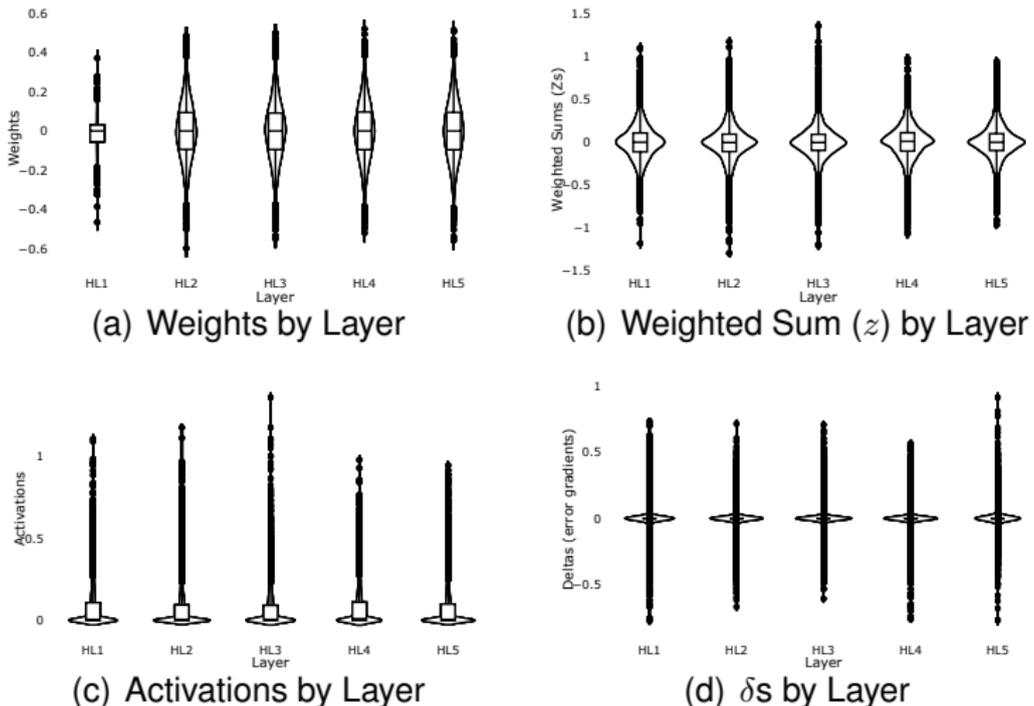


Figure 26: The internal dynamics of the network in Figure 22^[73], using ReLUs, during the first training iteration when the weights were initialized using He initialization.



- 1 represent the target feature using **one-hot encoding**;
- 2 change the output layer of the network to be a **softmax layer**; and
- 3 change the error (or loss) function we use for training to be the **cross-entropy** function.

Table 13: The *range-normalized* hourly samples of ambient factors and full load electrical power output of a combined cycle power plant, rounded to two decimal places, and with the (binned) target feature represented using one-hot encoding.

ID	AMBIENT TEMPERATURE	RELATIVE HUMIDITY	Electrical Output		
	°C	%	<i>low</i>	<i>medium</i>	<i>high</i>
1	0.04	0.81	0	0	1
2	0.84	0.58	1	0	0
3	0.50	0.07	0	1	0
4	0.53	1.00	0	1	0



$$\varphi_{sm}(z_i) = \frac{e^{z_i}}{\sum_{j=1}^m e^{z_j}} \quad (65)$$

Table 14: The calculation of the softmax activation function φ_{sm} over a vector of three logits \mathbf{l} .

	\mathbf{l}_0	\mathbf{l}_1	\mathbf{l}_2
\mathbf{l}	1.5	-0.9	0.6
$e^{\mathbf{l}_i}$	4.48168907	0.40656966	1.8221188
$\sum_i e^{\mathbf{l}_i}$			6.71037753
$\varphi_{sm}(\mathbf{l}_i)$	0.667874356	0.060588195	0.27153745

Handling Categorical Target Features: Softmax Output Layers and Cross-Entropy Loss Functions

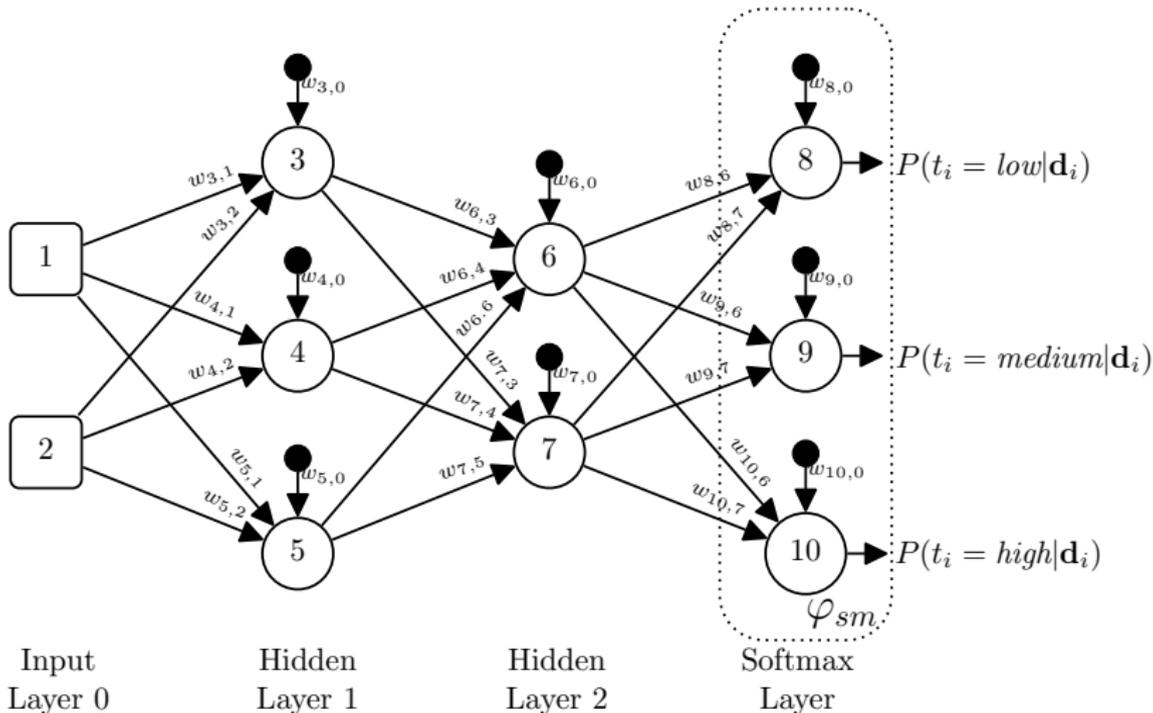


Figure 27: A schematic of a feedforward artificial neural network with a three-neuron softmax output layer.



$$L_{CE}(\mathbf{t}, \hat{\mathbf{P}}) = - \sum_j \mathbf{t}_j \ln(\hat{\mathbf{P}}_j) \quad (66)$$

$$L_{CE}(\mathbf{t}, \hat{\mathbf{P}}) = - \ln(\hat{\mathbf{P}}_*) \quad (67)$$

$$\begin{aligned}L_{CE}(\mathbf{t}, \hat{\mathbf{P}}) &= - \sum_j \mathbf{t}_j \ln(\hat{\mathbf{P}}_j) \\&= - \left(\left(\mathbf{t}_0 \ln(\hat{\mathbf{P}}_0) \right) + \left(\mathbf{t}_1 \ln(\hat{\mathbf{P}}_1) \right) + \left(\mathbf{t}_2 \ln(\hat{\mathbf{P}}_2) \right) \right) \\&= - \left(\left(0 \ln(\hat{\mathbf{P}}_0) \right) + \left(1 \ln(\hat{\mathbf{P}}_1) \right) + \left(0 \ln(\hat{\mathbf{P}}_2) \right) \right) \\&= -1 \ln(\hat{\mathbf{P}}_1) \tag{68}\end{aligned}$$

$$\frac{d \ln x}{dx} = \frac{1}{x} \quad (73)$$

$$\frac{\partial -\ln(\hat{\mathbf{P}}_{\star})}{\partial(\hat{\mathbf{P}}_{\star})} = -\frac{1}{\hat{\mathbf{P}}_{\star}} \quad (74)$$

$$\frac{\partial(\hat{\mathbf{P}}_{\star})}{\partial \mathbf{l}_k} = \begin{cases} \hat{\mathbf{P}}_{\star} (1 - \hat{\mathbf{P}}_k) & \text{if } k = \star \\ -\hat{\mathbf{P}}_{\star} \hat{\mathbf{P}}_k & \text{otherwise} \end{cases} \quad (75)$$

$$\delta_k = \frac{\partial -\ln(\hat{\mathbf{P}}_\star)}{\partial(\hat{\mathbf{P}}_\star)} \times \frac{\partial(\hat{\mathbf{P}}_\star)}{\partial \mathbf{l}_k} \quad (76)$$

$$= -\frac{1}{\hat{\mathbf{P}}_\star} \times \frac{\partial(\hat{\mathbf{P}}_\star)}{\partial \mathbf{l}_k} \quad (77)$$

$$= -\frac{1}{\hat{\mathbf{P}}_\star} \times \begin{cases} \hat{\mathbf{P}}_\star (1 - \hat{\mathbf{P}}_k) & \text{if } k = \star \\ -\hat{\mathbf{P}}_\star \hat{\mathbf{P}}_k & \text{otherwise} \end{cases} \quad (78)$$

$$= \begin{cases} -(1 - \hat{\mathbf{P}}_k) & \text{if } k = \star \\ \hat{\mathbf{P}}_k & \text{otherwise} \end{cases} \quad (79)$$

Handling Categorical Target Features: Softmax Output Layers and Cross-Entropy Loss Functions

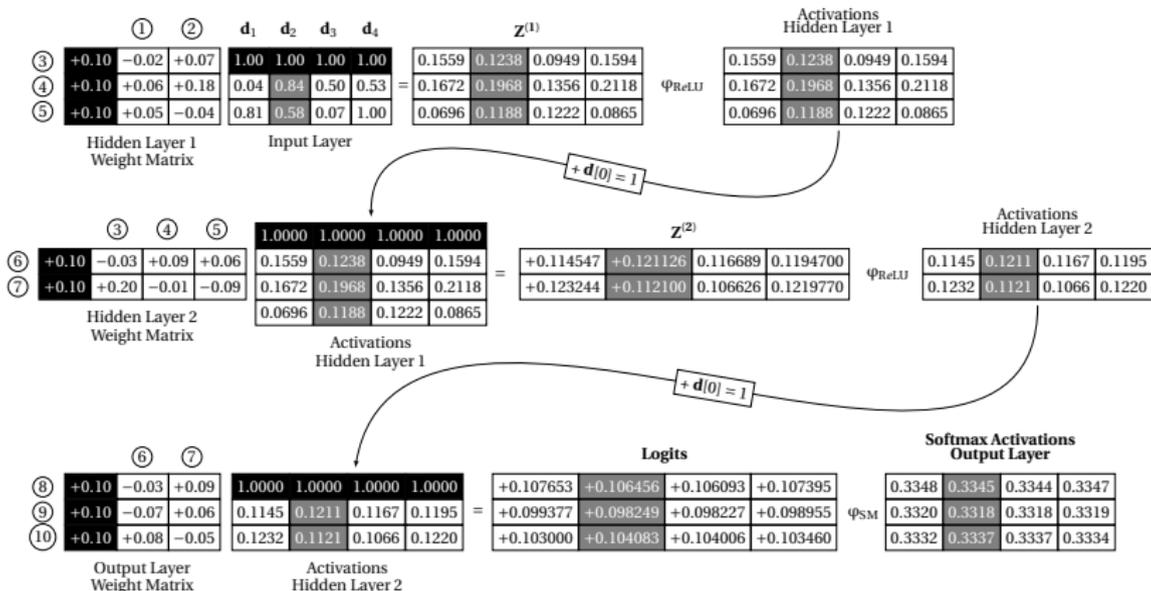


Figure 28: The forward pass of the mini-batch of examples listed in Table 13^[88] through the network in Figure 27^[91].

Table 15: The calculation of the softmax activations for each of the neurons in the output layer for each example in the mini-batch, and the calculation of the δ for each neuron in the output layer for each example in the mini-batch.

	d_1	d_2	d_3	d_4
Per Neuron Per Example logits				
Neuron 8	0.107653	0.106456	0.106093	0.107395
Neuron 9	0.099377	0.098249	0.098227	0.098955
Neuron 10	0.103000	0.104083	0.1040060	0.103460
Per Neuron Per Example e^{l_i}				
Neuron 8	1.113661238	1.112328983	1.111925281	1.11337395
Neuron 9	1.104482611	1.103237457	1.103213186	1.104016618
Neuron 10	1.108491409	1.109692556	1.109607113	1.109001432
$\sum_i e^{l_i}$	3.326635258	3.325258996	3.324745579	3.326392
Per Neuron Per Example Softmax Activations				
Neuron 8	0.3348	0.3345	0.3344	0.3347
Neuron 9	0.3320	0.3318	0.3318	0.3319
Neuron 10	0.3332	0.3337	0.3337	0.3334
Per Neuron Target One-Hot Encodings				
Neuron 8	0	1	0	0
Neuron 9	0	0	1	1
Neuron 10	1	0	0	0
Per Neuron Per Example δ s				
Neuron 8	0.3348	-0.6655	0.3344	0.3347
Neuron 9	0.3320	0.3318	-0.6682	-0.6681
Neuron 10	-0.6668	0.3337	0.3337	0.3334



$$\begin{aligned}\Delta w_{9,6} &= \sum_{j=1}^4 \delta_{9,j} \times a_{6,j} \\ &= (0.3320 \times 0.1145) + (0.3318 \times 0.1211) \\ &\quad + (-0.6682 \times 0.1167) + (-0.6681 \times 0.1195) \\ &= 0.038014 + 0.04018098 + -0.07797894 + -0.07983795 \\ &= -0.07962191 \qquad (82)\end{aligned}$$



$$\begin{aligned}w_{9,6} &= w_{9,6} - \alpha \times \Delta w_{9,6} \\ &= -0.07 - 0.01 \times -0.07962191 \\ &= -0.07 - (-0.000796219) \\ &= -0.069203781\end{aligned}\tag{83}$$

Algorithm 1 The early stopping algorithm

Require: p the patience parameter

Require: \mathcal{D}_v a validation set

```
1: bestValidationError =  $\infty$ 
2: tmpValidationError = 0
3:  $\theta$  = initial model parameters
4:  $\theta^{best} = \theta$ 
5: patienceCount = 0
6: while patienceCount <  $p$  do
7:    $\theta$  = new model parameters after most recent weight update
8:   tmpValidationError = calculateValidationError( $\theta$ ,  $\mathcal{D}_v$ )
9:   if bestValidationError  $\geq$  tmpValidationError then
10:     bestValidationError = tmpValidationError
11:      $\theta^{best} = \theta$ 
12:   patienceCount = 0
13:   else
14:     patienceCount = patienceCount + 1
15:   end if
16: end while
17: return Best Model Parameters  $\theta^{best}$ 
```

Algorithm 2 Extensions to Backpropagation to Use Inverted Dropout

Require: ρ probability that a neuron in a layer will not be dropped

▷ Forward Pass

1: **for** each input or hidden layer l **do**

2: $\mathbf{DropMask}^{(l)} = (m_1, \dots, m_{\text{size}(l)}) \sim \text{Bernoulli}(\rho)$

3: $\mathbf{a}^{(l)'} = \mathbf{a}^{(l)} \odot \mathbf{DropMax}^{(l)}$

4: $\mathbf{a}^{(l)''} = \frac{1}{\rho} \mathbf{a}^{(l)'}$

5: **end for**

▷ Backward Pass

6: **for** each layer l in backward pass **do**

7: $\delta^{(l)} = \delta^{(l)} \odot \mathbf{DropMax}^{(l)}$

8: **end for**

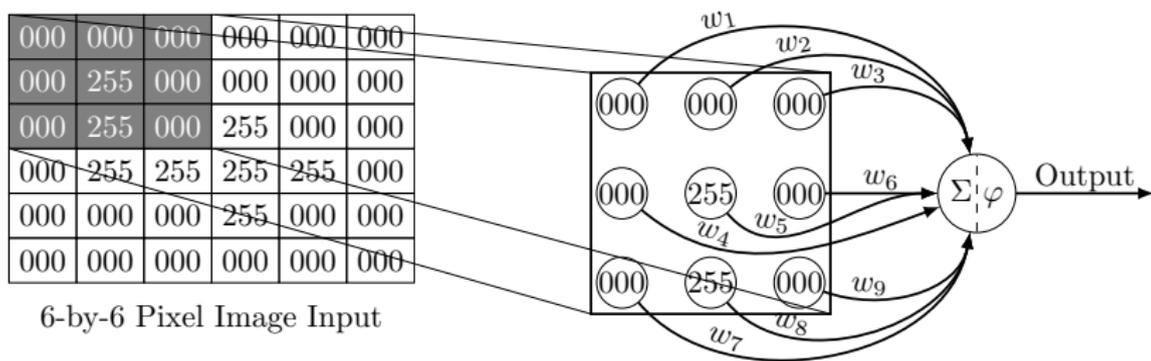


Figure 31: A 6-by-6 matrix representation of a grayscale image of a 4, and a neuron with a receptive field that covers the top-left corner of the image. This figure was inspired by Figure 2 of (Kelleher and Dobnik, 2017).

$$\begin{bmatrix} 0 & 0 & 0 \\ 1 & 1 & 1 \\ 0 & 0 & 0 \end{bmatrix} \quad (85)$$

$$\begin{aligned} a_i &= \text{rectifier}((w_1 \times 000) + (w_2 \times 000) + (w_3 \times 000) \\ &\quad + (w_4 \times 000) + (w_5 \times 255) + (w_6 \times 000) \\ &\quad + (w_7 \times 000) + (w_8 \times 255) + (w_9 \times 000)) \\ &= \text{rectifier}((0 \times 000) + (0 \times 000) + (0 \times 000) \\ &\quad + (1 \times 000) + (1 \times 255) + (1 \times 000) \\ &\quad + (0 \times 000) + (0 \times 255) + (0 \times 000)) \\ &= 255 \end{aligned} \quad (86)$$

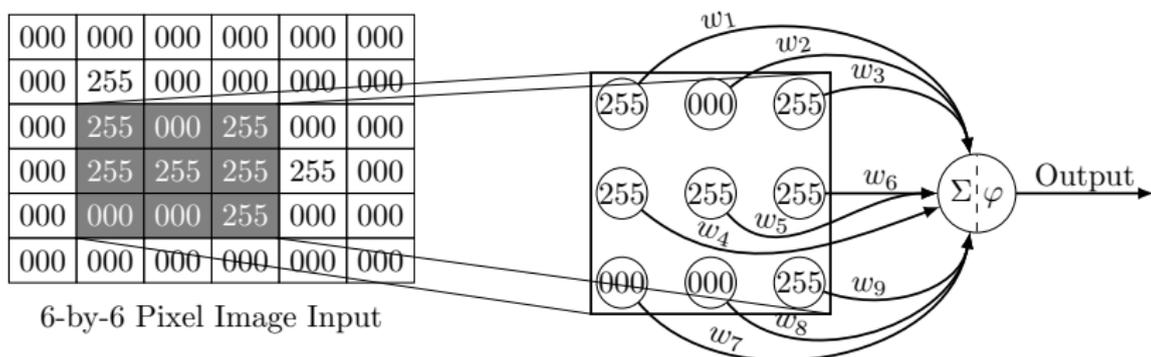
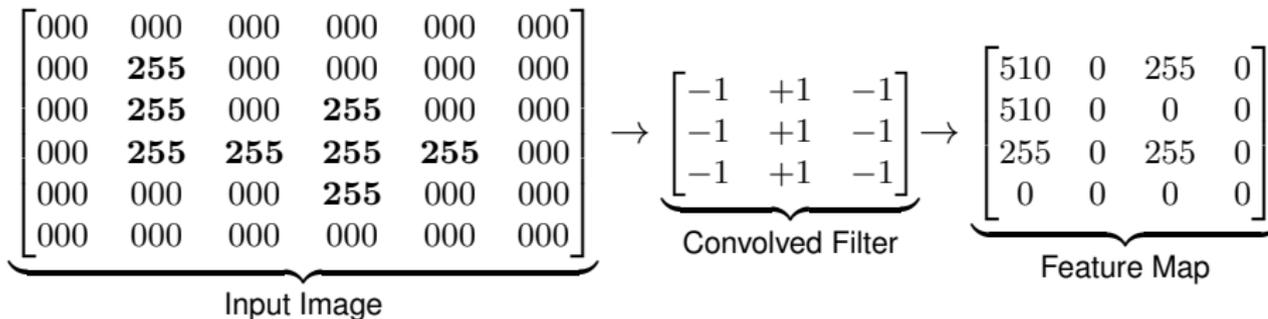
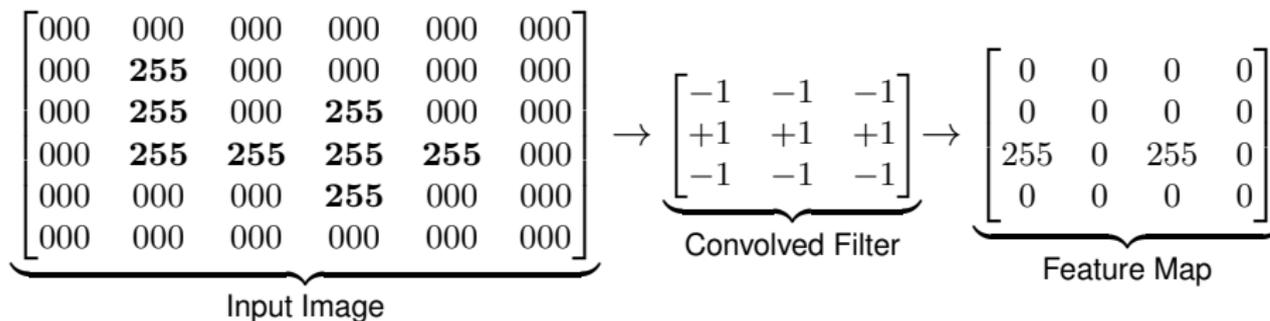


Figure 32: A 6-by-6 matrix representation of a grayscale image of a 4, and a neuron with a different receptive field from the neuron in Figure 31^[107]. This figure was inspired by Figure 2 of (Kelleher and Dobnik, 2017).



(90)



(91)

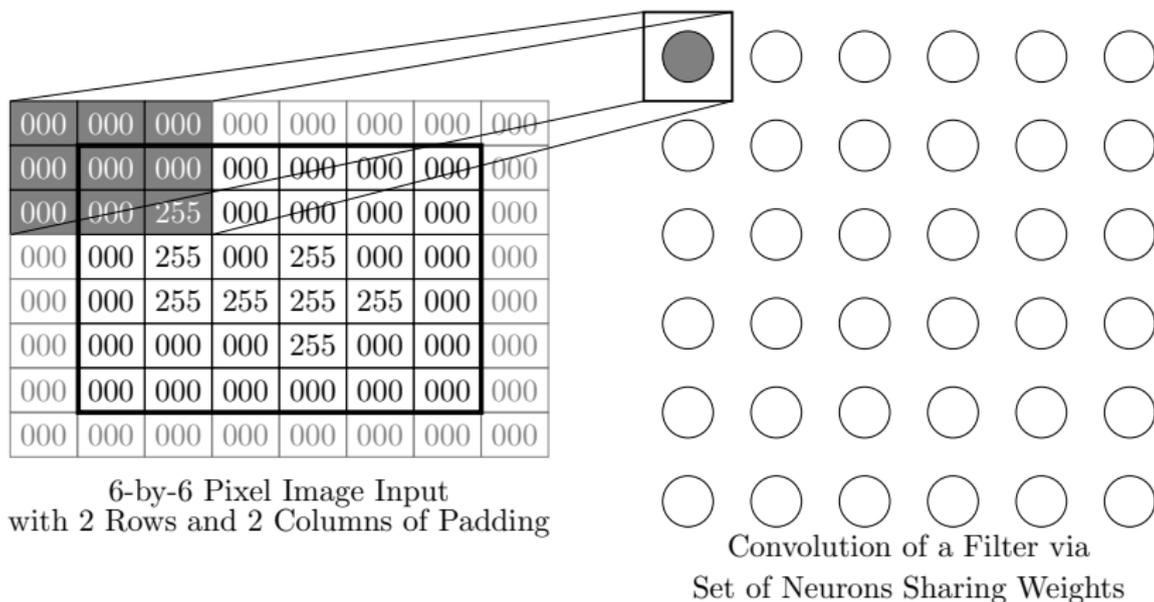
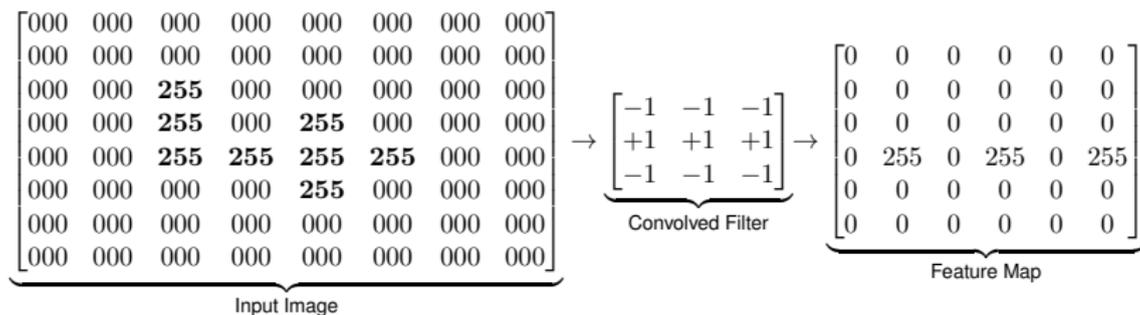


Figure 34: A grayscale image of a 4 after padding has been applied to the original 6-by-6 matrix representation, and the local receptive field of a neuron that includes both valid and padded pixels.

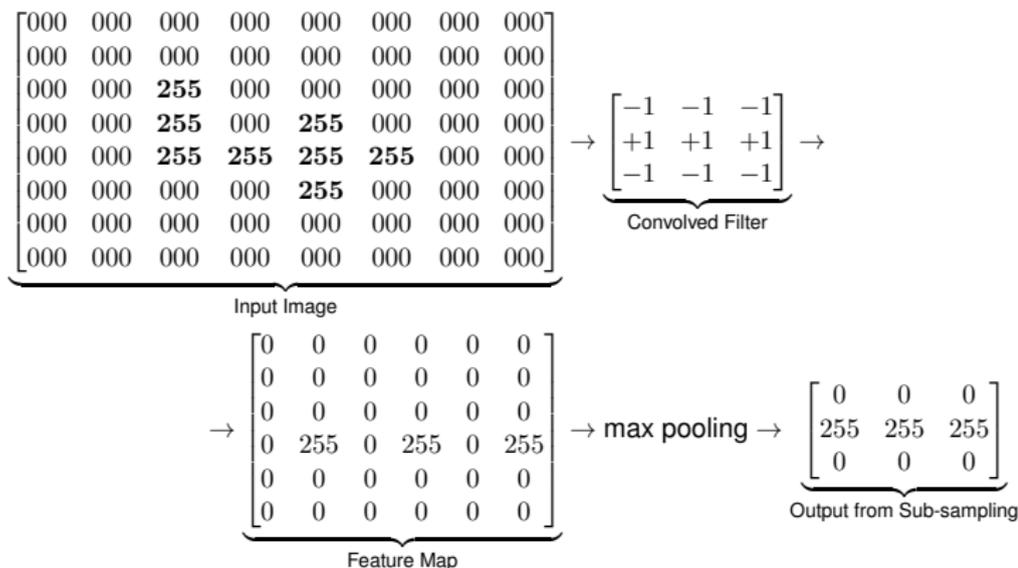
$$\underbrace{\begin{bmatrix} 000 & 000 & 000 & 000 & 000 & 000 & 000 & 000 \\ 000 & 000 & 000 & 000 & 000 & 000 & 000 & 000 \\ 000 & 000 & \mathbf{255} & 000 & 000 & 000 & 000 & 000 \\ 000 & 000 & \mathbf{255} & 000 & \mathbf{255} & 000 & 000 & 000 \\ 000 & 000 & \mathbf{255} & \mathbf{255} & \mathbf{255} & \mathbf{255} & 000 & 000 \\ 000 & 000 & 000 & 000 & \mathbf{255} & 000 & 000 & 000 \\ 000 & 000 & 000 & 000 & 000 & 000 & 000 & 000 \\ 000 & 000 & 000 & 000 & 000 & 000 & 000 & 000 \end{bmatrix}}_{\text{Input Image}} \rightarrow \underbrace{\begin{bmatrix} -1 & +1 & -1 \\ -1 & +1 & -1 \\ -1 & +1 & -1 \end{bmatrix}}_{\text{Convolved Filter}} \rightarrow \underbrace{\begin{bmatrix} 0 & 255 & 0 & 0 & 0 & 0 \\ 0 & 510 & 0 & 255 & 0 & 0 \\ 0 & 510 & 0 & 0 & 0 & 0 \\ 0 & 255 & 0 & 255 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 255 & 0 & 0 \end{bmatrix}}_{\text{Feature Map}}$$

(92)

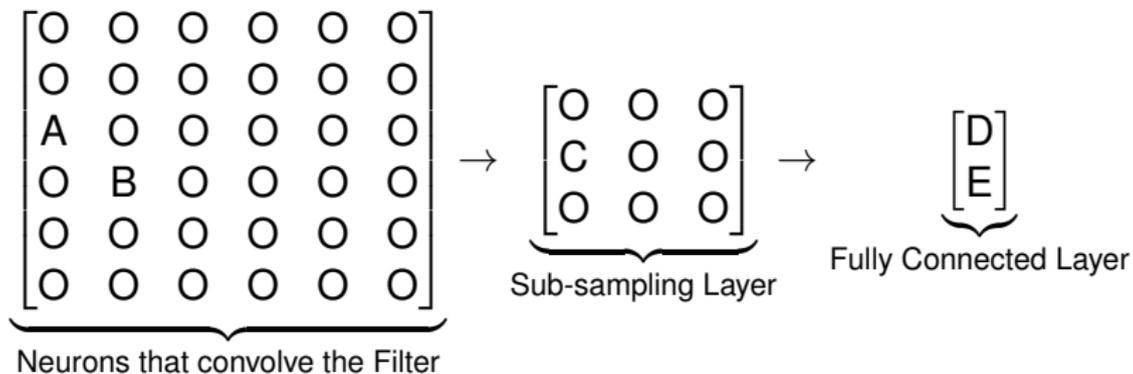


(93)

Convolutional Neural Networks



(94)



$$\left[\begin{array}{c} \left[\begin{array}{ccc} 1 & 1 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{array} \right] \\ \left[\begin{array}{ccc} 0 & 0 & 2 \\ 0 & 0 & 2 \\ 0 & 0 & 2 \end{array} \right] \\ \left[\begin{array}{ccc} 3 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 3 \end{array} \right] \end{array} \right] \quad (100)$$

Red Channel Green Channel Blue Channel

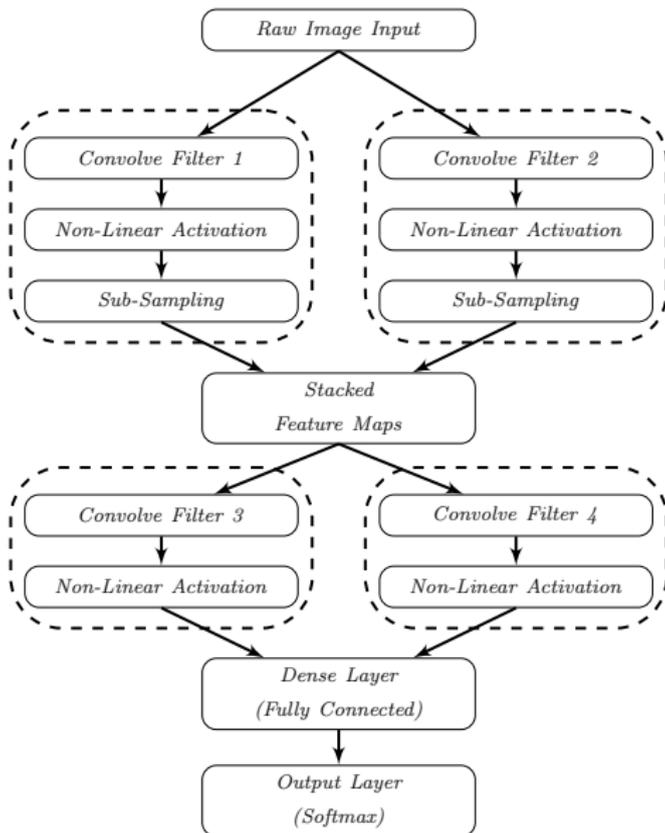


Figure 35: Schematic of the typical sequences of layers found in a convolutional neural network.

Convolutional Neural Networks

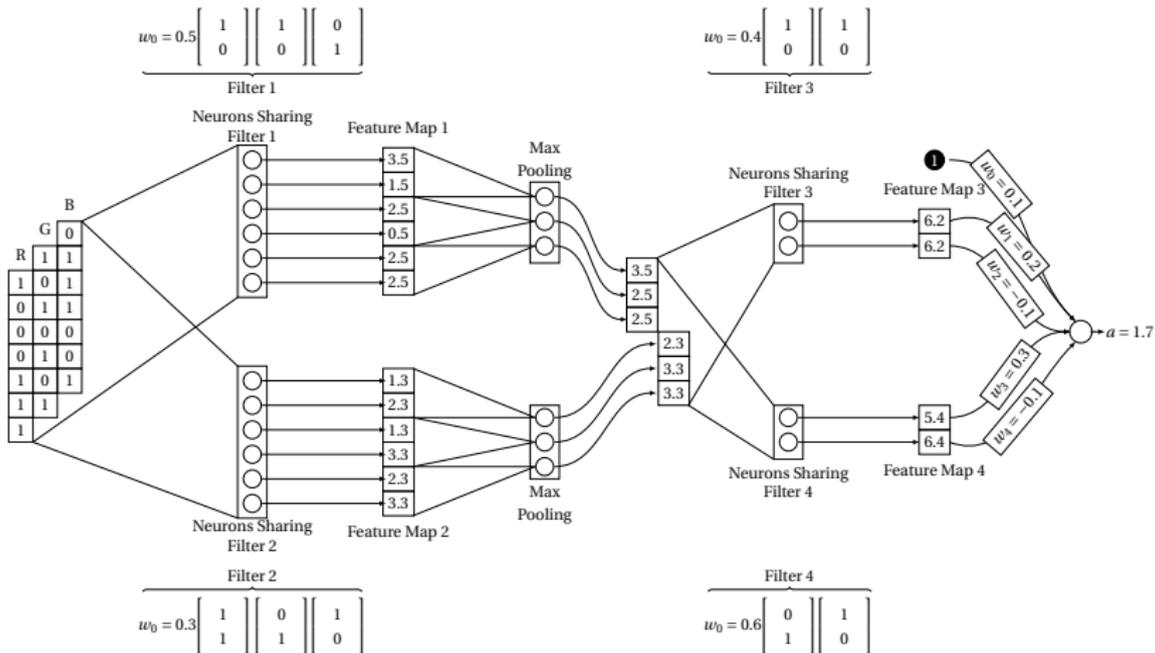


Figure 36: Worked example illustrating the dataflow through a multilayer, multifilter CNN.

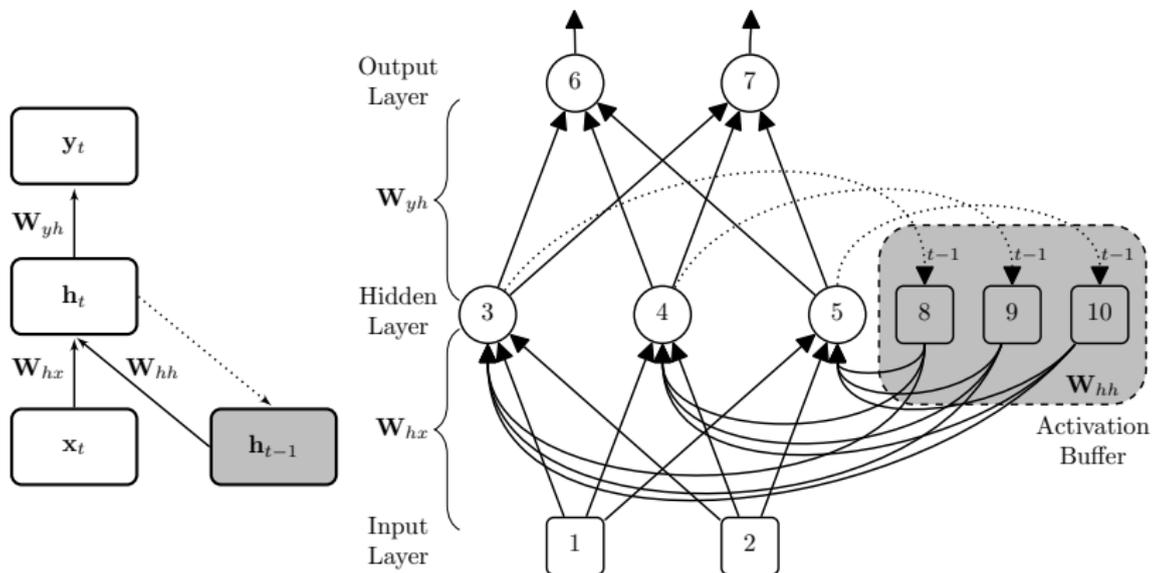


Figure 37: Schematic of the simple recurrent neural architecture.



$$\mathbf{h}_t = \varphi((\mathbf{W}_{hh} \cdot \mathbf{h}_{t-1}) + (\mathbf{W}_{hx} \cdot \mathbf{x}_t) + \mathbf{w}_0) \quad (104)$$

$$\mathbf{y}_t = \varphi(\mathbf{W}_{yh} \cdot \mathbf{h}_t) \quad (105)$$

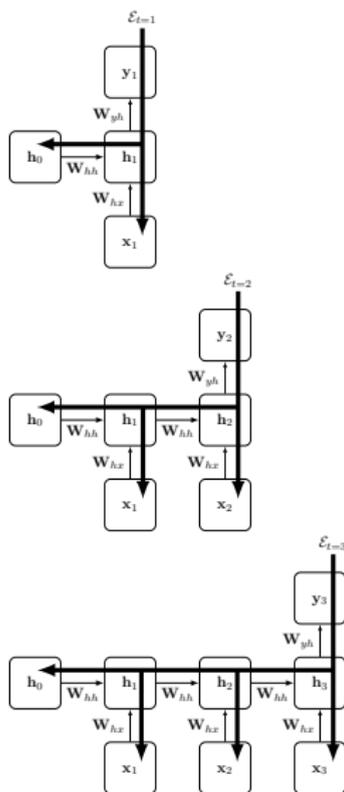


Figure 39: An illustration of the different iterations of backpropagation during backpropagation through time.

Algorithm 3 The Backpropagation Through Time Algorithm

Require: h_0 initialized hidden state

Require: \mathbf{x} a sequence of inputs

Require: \mathbf{y} a sequence of target outputs

Require: n length of the input sequence

Require: Initialized weight matrices (with associated biases)

Require: $\Delta \mathbf{w}$ a data structure to accumulate the summed weight updates for each weight across time-steps

```
1: for  $t = 1$  to  $n$  do
2:    $Inputs = [x_0, \dots, x_t]$ 
3:    $h_{tmp} = h_0$ 
4:   for  $i = 0$  to  $t$  do                                ▷ Unroll the network through  $t$  steps
5:      $h_{tmp} = ForwardPropagate(Inputs[i], h_{tmp})$ 
6:   end for
7:    $\hat{y}_t = OutputLayer(h_{tmp})$                     ▷ Generate the output for time-step  $t$ 
8:    $\mathcal{E}_t = \mathbf{y}[t] - \hat{y}_t$                       ▷ Calculate the error at time-step  $t$ 
9:    $Backpropagate(\mathcal{E}_t)$                           ▷ Backpropagate  $\mathcal{E}_t$  through  $t$  steps
10:  For each weight, sum the weight updates across the unrolled network and update  $\Delta \mathbf{w}$ 
11: end for
12: Update the network weights using  $\Delta \mathbf{w}$ 
```

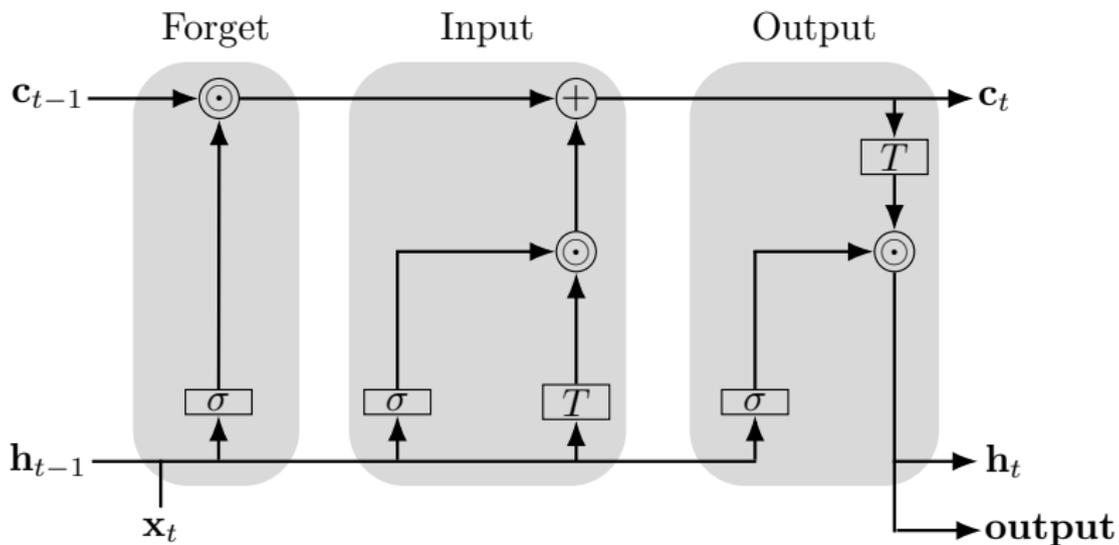


Figure 40: A schematic of the internal structure of a long short-term memory unit. This figure is based on Figure 5.4 of (Kelleher, 2019), which in turn was inspired by an image by Christopher Olah (available at: <http://colah.github.io/posts/2015-08-Understanding-LSTMs/>).



$$\mathbf{f}_t = \varphi_{sigmoid}(\mathbf{W}^{(f)} \cdot \mathbf{h}\mathbf{x}_t) \quad (106)$$

$$\mathbf{c}_t^\dagger = \mathbf{c}_{t-1} \odot \mathbf{f}_t \quad (107)$$



$$\mathbf{c}_{t-1} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \quad \mathbf{h}_{t-1} = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \quad \mathbf{x}_t = [4]$$
$$\mathbf{W}^{(f)} = \begin{bmatrix} 0 & 1 & 1 & 1 \\ 0 & -1 & -1 & -1 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

(108)

$$\underbrace{\begin{bmatrix} 0 & 1 & 1 & 1 \\ 0 & -1 & -1 & -1 \\ 0 & 0 & 0 & 0 \end{bmatrix}}_{\mathbf{W}^{(f)}} \times \underbrace{\begin{bmatrix} 1 \\ 1 \\ 1 \\ 4 \end{bmatrix}}_{\mathbf{hx}_t} = \underbrace{\begin{bmatrix} 6 \\ -6 \\ 0 \end{bmatrix}}_{\mathbf{z}_t^{(f)}} \rightarrow \varphi_{\text{sigmoid}} \rightarrow \underbrace{\begin{bmatrix} 0.997527377 \\ 0.002472623 \\ 0.500000000 \end{bmatrix}}_{\mathbf{f}_t}$$

$$\underbrace{\begin{bmatrix} 0.997527377 \\ 0.002472623 \\ 0.500000000 \end{bmatrix}}_{\mathbf{f}_t} \odot \underbrace{\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}}_{\mathbf{c}_{t-1}} = \underbrace{\begin{bmatrix} 0.997527377 \\ 0.002472623 \\ 0.500000000 \end{bmatrix}}_{\mathbf{c}_t^\dagger}$$

(109)



$$\mathbf{i}_{\uparrow t} = \varphi_{\text{sigmoid}}(\mathbf{W}^{(i\uparrow)} \cdot \mathbf{h}\mathbf{x}_t) \quad (110)$$

$$\mathbf{i}_{\downarrow t} = \varphi_{\text{tanh}}(\mathbf{W}^{(i\downarrow)} \cdot \mathbf{h}\mathbf{x}_t) \quad (111)$$

$$\mathbf{i}_t = \mathbf{i}_{\uparrow t} \odot \mathbf{i}_{\downarrow t} \quad (112)$$

$$\mathbf{c}_t = \mathbf{c}_{\uparrow} + \mathbf{i}_t \quad (113)$$

$$\mathbf{f}_t = \varphi_{\text{sigmoid}}(\mathbf{W}^{(f)} \cdot \mathbf{h}\mathbf{x}_t)$$

$$\mathbf{c}_t^\dagger = \mathbf{c}_{t-1} \odot \mathbf{f}_t$$

$$\mathbf{i}_t^\dagger = \varphi_{\text{sigmoid}}(\mathbf{W}^{(i^\dagger)} \cdot \mathbf{h}\mathbf{x}_t)$$

$$\mathbf{i}_{t+}^\dagger = \varphi_{\text{tanh}}(\mathbf{W}^{(i_{t+}^\dagger)} \cdot \mathbf{h}\mathbf{x}_t)$$

$$\mathbf{i}_t = \mathbf{i}_t^\dagger \odot \mathbf{i}_{t+}^\dagger$$

$$\mathbf{c}_t = \mathbf{c}_t^\dagger + \mathbf{i}_t$$

$$\mathbf{o}_t^\dagger = \varphi_{\text{sigmoid}}(\mathbf{W}^{(o^\dagger)} \cdot \mathbf{h}\mathbf{x}_t)$$

$$\mathbf{o}_{t+}^\dagger = \varphi_{\text{tanh}}(\mathbf{W}^{(o_{t+}^\dagger)} \cdot \mathbf{c}_t)$$

$$\mathbf{o}_t = \mathbf{o}_t^\dagger \odot \mathbf{o}_{t+}^\dagger$$

$$\mathbf{h}_{t+1} = \mathbf{o}_t$$

Sequential Models: Recurrent Neural Networks and Long Short-Term Memory Networks

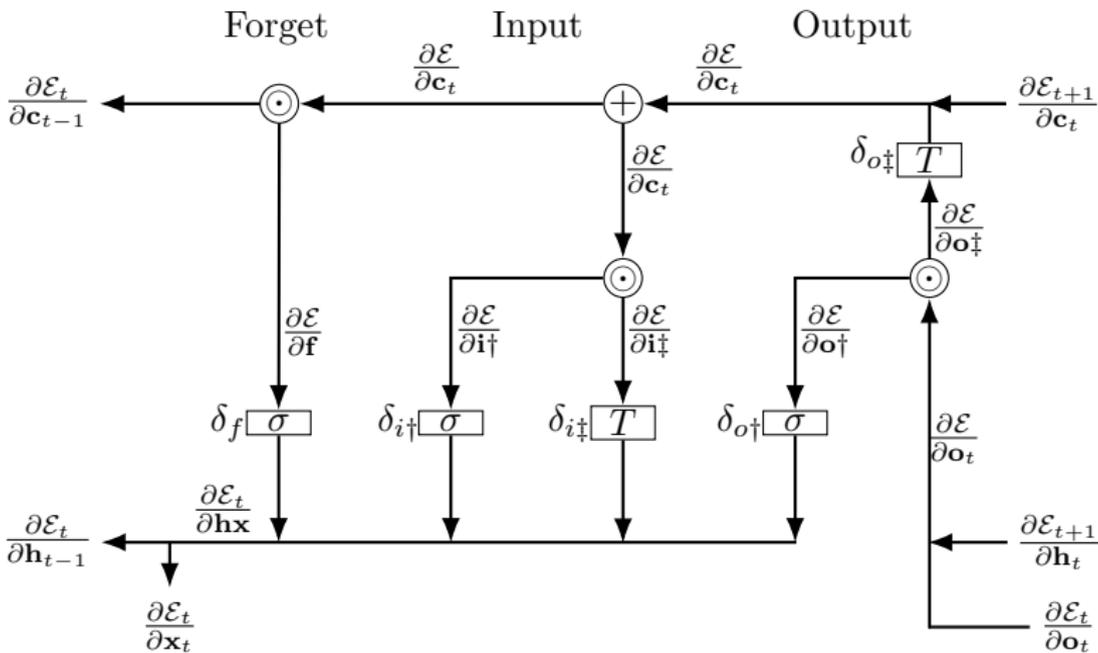
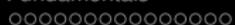


Figure 42: The flow of error gradients through a long short-term memory unit during backpropagation.



$$\frac{\partial \mathcal{E}}{\partial \mathbf{o}_t} = \frac{\partial \mathcal{E}_t}{\partial \mathbf{o}_t} + \frac{\partial \mathcal{E}_{t+1}}{\partial \mathbf{h}_t} \quad (118)$$



$$\Delta \mathbf{W}^{(f)} = \delta_f \cdot \mathbf{h}\mathbf{x}^\top \quad (131)$$

$$\frac{\partial \mathcal{E}}{\partial \mathbf{o}_t} = \underbrace{\begin{bmatrix} 0.75 \\ 0.25 \end{bmatrix}}_{\partial \mathcal{E}_{t+1} / \partial \mathbf{h}_t} + \underbrace{\begin{bmatrix} 0.15 \\ 0.60 \end{bmatrix}}_{\partial \mathcal{E}_t / \partial \mathbf{o}_t} = \begin{bmatrix} 0.9 \\ 0.85 \end{bmatrix} \quad (132)$$

$$\frac{\partial \mathcal{E}}{\partial \mathbf{o}_t^\dagger} = \underbrace{\begin{bmatrix} 0.9 \\ 0.85 \end{bmatrix}}_{\partial \mathcal{E} / \partial \mathbf{o}_t} \odot \underbrace{\begin{bmatrix} 0.575420058 \\ 0.52098661 \end{bmatrix}}_{\mathbf{o}_t^\dagger} = \begin{bmatrix} 0.517878052 \\ 0.442839512 \end{bmatrix} \quad (133)$$

$$\delta_{\mathbf{o}_t^\dagger} = \underbrace{\begin{bmatrix} 0.517878052 \\ 0.442839512 \end{bmatrix}}_{\partial \mathcal{E} / \partial \mathbf{o}_t^\dagger} \odot \underbrace{\begin{bmatrix} 0.999184559 \\ 0.997048635 \end{bmatrix}}_{1 - \tanh(\mathbf{c}_t)} = \begin{bmatrix} 0.517455753 \\ 0.441532531 \end{bmatrix} \quad (134)$$

$$\frac{\partial \mathcal{E}}{\partial \mathbf{c}_t} = \underbrace{\begin{bmatrix} 0.517455753 \\ 0.441532531 \end{bmatrix}}_{\delta_{\mathbf{o}_t^\dagger}} + \underbrace{\begin{bmatrix} 0.35 \\ 0.50 \end{bmatrix}}_{\partial \mathcal{E}_{t+1} / \partial \mathbf{c}_t} = \begin{bmatrix} 0.867455753 \\ 0.941532531 \end{bmatrix} \quad (135)$$

$$\frac{\partial \mathcal{E}}{\partial \mathbf{f}} = \underbrace{\begin{bmatrix} 0.867455753 \\ 0.941532531 \end{bmatrix}}_{\partial \mathcal{E} / \partial \mathbf{c}_t} \odot \underbrace{\begin{bmatrix} 0.3 \\ 0.6 \end{bmatrix}}_{\mathbf{c}_{t-1}} = \begin{bmatrix} 0.260236726 \\ 0.564919518 \end{bmatrix} \quad (136)$$

$$\delta_f = \underbrace{\begin{bmatrix} 0.260236726 \\ 0.564919518 \end{bmatrix}}_{\partial \mathcal{E} / \partial \mathbf{f}} \odot \underbrace{\begin{bmatrix} 0.248743973 \\ 0.249694 \end{bmatrix}}_{\mathbf{f}_t \odot (1 - \mathbf{f}_t)} = \begin{bmatrix} 0.064732317 \\ 0.141057014 \end{bmatrix} \quad (137)$$

$$\begin{aligned} \Delta \mathbf{W}^{(f)} &= \underbrace{\begin{bmatrix} 0.064732317 \\ 0.141057014 \end{bmatrix}}_{\delta_f} \cdot \underbrace{\begin{bmatrix} 1.00 & 0.10 & 0.80 & 0.90 \end{bmatrix}}_{\mathbf{h}\mathbf{x}^\top} \\ &= \begin{bmatrix} 0.064732317 & 0.006473232 & 0.051785854 & 0.058259085 \\ 0.141057014 & 0.014105701 & 0.112845611 & 0.126951313 \end{bmatrix} \end{aligned} \quad (138)$$

$$\begin{aligned}
\frac{\partial \mathcal{E}}{\partial \mathbf{c}_{t-1}} &= \mathbf{f}_t \odot \frac{\partial \mathcal{E}}{\partial \mathbf{c}_t} \\
&= \mathbf{f}_t \odot \left(\frac{\partial \mathcal{E}_{t+1}}{\partial \mathbf{c}_t} + \delta_{o_t^\dagger} \right) \\
&= \mathbf{f}_t \odot \left(\frac{\partial \mathcal{E}_{t+1}}{\partial \mathbf{c}_t} + \left((\mathbf{1} - \tanh^2(\mathbf{c}_{t+1}^\dagger)) \odot \frac{\partial \mathcal{E}}{\partial \mathbf{o}_{t+1}^\dagger} \right) \right) \\
&= \mathbf{f}_t \odot \left(\frac{\partial \mathcal{E}_{t+1}}{\partial \mathbf{c}_t} + \left((\mathbf{1} - \tanh^2(\mathbf{c}_{t+1}^\dagger)) \odot \left(\mathbf{o}_{t+1}^\dagger \odot \frac{\partial \mathcal{E}}{\partial \mathbf{o}_t} \right) \right) \right) \\
&= \mathbf{f}_t \odot \left(\frac{\partial \mathcal{E}_{t+1}}{\partial \mathbf{c}_t} + \left((\mathbf{1} - \tanh^2(\mathbf{c}_{t+1}^\dagger)) \odot \left(\mathbf{o}_{t+1}^\dagger \odot \left(\frac{\partial \mathcal{E}_{t+1}}{\partial \mathbf{h}_t} + \frac{\partial \mathcal{E}_t}{\partial \mathbf{o}_t} \right) \right) \right) \right)
\end{aligned}$$

(140)



reducing internal covariate shift. In *International conference on machine learning*, 448–456.

Kelleher, John D. 2019. *Deep learning. Essential knowledge series*. MIT Press.

Kelleher, John D, and Simon Dobnik. 2017. What is not where: the challenge of integrating spatial representations into deep learning architectures. *CLASP Papers in Computational Linguistics*.

Kingma, Diederik P, and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Marsland, Stephen. 2011. *Machine learning: an algorithmic perspective*. CRC Press.

Reagen, Brandon, Robert Adolf, Paul Whatmough, Gu-Yeon Wei, and David Brooks. 2017. *Deep learning for computer architects*. Morgan and Claypool.

