Neural networks

Virgil Pavlu October 1, 2008

1 The perceptron

Lets suppose we are (as with regression regression) with $(\mathbf{x}_i, y_i); i = 1, .., m$ the data points and labels. This is a classification problem with two classes $y \in \{-1, 1\}$

Like with regression we are looking for a linear predictor (classifier)

$$h_{\mathbf{w}}(\mathbf{x}) = \mathbf{x}\mathbf{w} = \sum_{d=0}^{D} x^d w^d$$

(we added the $x^0 = 1$ component so we can get the free term w^0) such that $h_{\mathbf{w}}(\mathbf{x}) \geq 0$ when $y = 1$ and $h_{\mathbf{w}}(\mathbf{x}) \leq 0$ when $y = -1$.

On $y = -1$ data points: given that all $\mathbf{x}$ and y are numerical, we will make the following transformation: when $y = -1$, we will reverse the sign of the input; that is replace $\mathbf{x}$ with -$\mathbf{x}$ and $y = -y$. Then the condition $h_{\mathbf{w}}(\mathbf{x}) \leq 0$ becomes $h_{\mathbf{w}}(\mathbf{x}) \geq 0$ for all data points.

The perceptron objective function is a combination of the number of miss-classification points and how bad the miss-classification is

$$J(\mathbf{w}) = \sum_{\mathbf{x} \in M} -h_{\mathbf{w}}(\mathbf{x}) = \sum_{\mathbf{x} \in M} -\mathbf{x}\mathbf{w}$$

where M is the set of miss-classified data points. Note that each term of the sum is positive, since miss-classified implies $\mathbf{w}\mathbf{x} < 0$. Using gradient descent, we first differentiate J

XW

$$\nabla_{\mathbf{w}} J(\mathbf{w}) = \sum_{\mathbf{x} \in M} -\mathbf{x}^T$$

$$\frac{\partial J}{\partial w_i} = \sum_{M} -x_i$$

then we write down the gradient descent update rule

$$\mathbf{w} := \mathbf{w} + \lambda \sum_{\mathbf{x} \in M} \mathbf{x}^T$$

(λ is the learning rate). The batch version looks like

1. init $\mathbf{w}$
2. LOOP
3. get M = set of missclassified data points
4. $\mathbf{w} = \mathbf{w} + \lambda \sum_{\mathbf{x} \in M} \mathbf{x}^T$
5. UNTIL $|\lambda \sum_{\mathbf{x} \in M} \mathbf{x}| < \epsilon$

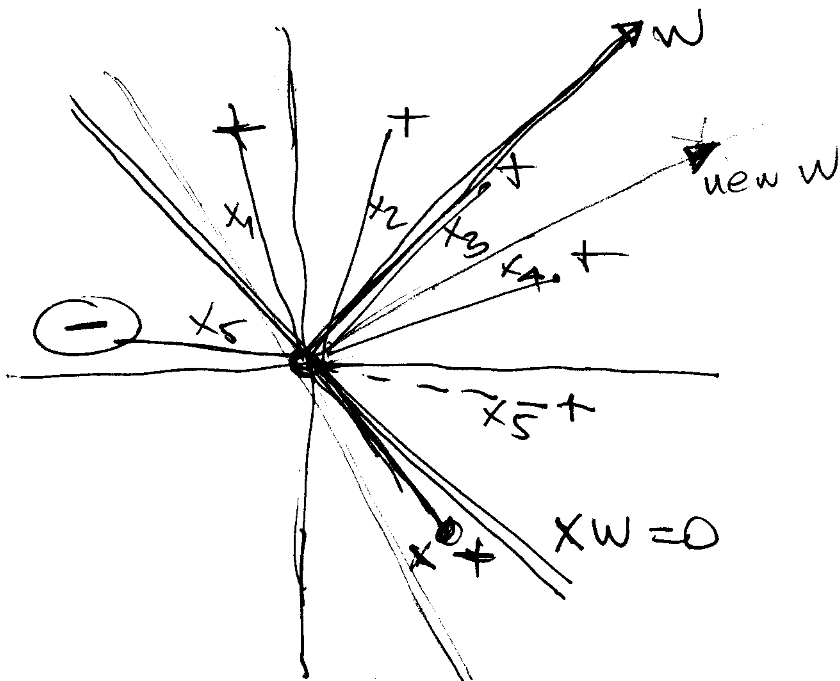

x on the right side $xw > 0$

x on the wrong side $xw < 0$

$wx < 0 \rightarrow$ move x toward x

$w \Rightarrow w + x$

Assume the instances are linearly separable. Then we can modify the algorithm

1. init $\mathbf{w}$
2. LOOP
3. get M = set of missclassified data points
4. for each $\mathbf{x} \in M$ do $\mathbf{w} = \mathbf{w} + \lambda \mathbf{x}^T$
5. UNTIL M is empty

Proof of perceptron convergence Assuming data is linearly separable , or there is a solution $\bar{\mathbf{w}}$ such that $\mathbf{x}\bar{\mathbf{w}} > 0$ for all $\mathbf{x}$.
Lets call $\mathbf{w}_k$ the $\mathbf{w}$ obtained at the k-th iteration (update). Fix an $\alpha > 0$. Then

$$\mathbf{w}_{k+1} - \alpha\bar{\mathbf{w}} = (\mathbf{w}_k - \alpha\bar{\mathbf{w}}) + \mathbf{x}_k^T$$

where $\mathbf{x}_k$ is the datapoint that updated $\mathbf{w}$ at iteration k. Then

$$||\mathbf{w}_{k+1} - \alpha\bar{\mathbf{w}}||^2 = ||\mathbf{w}_k - \alpha\bar{\mathbf{w}}||^2 + 2\mathbf{x}_k(\mathbf{w}_k - \alpha\bar{\mathbf{w}}) + ||\mathbf{x}_k||^2 \leq ||\mathbf{w}_k - \alpha\bar{\mathbf{w}}||^2 - 2\mathbf{x}_k\alpha\bar{\mathbf{w}} + ||\mathbf{x}_k||^2$$

Since $\mathbf{x}_k\bar{\mathbf{w}} > 0$ all we need is an α sufficiently large to show that this update process cannot go on forever. When it stops, all datapoints must be classified correctly.

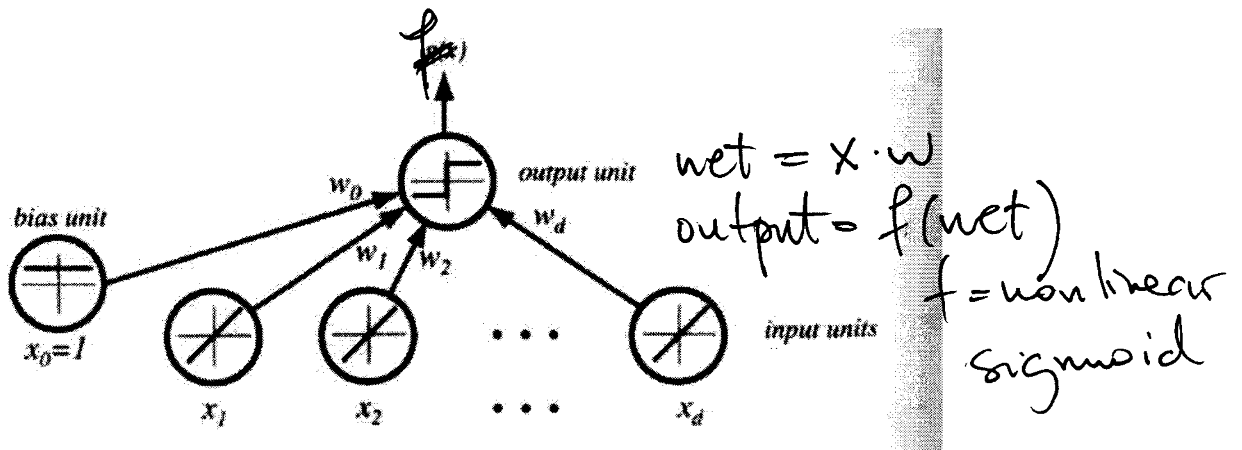

Figure 1: bias unit

2 Multilayer perceptrons

Also called *feed-forward networks.*

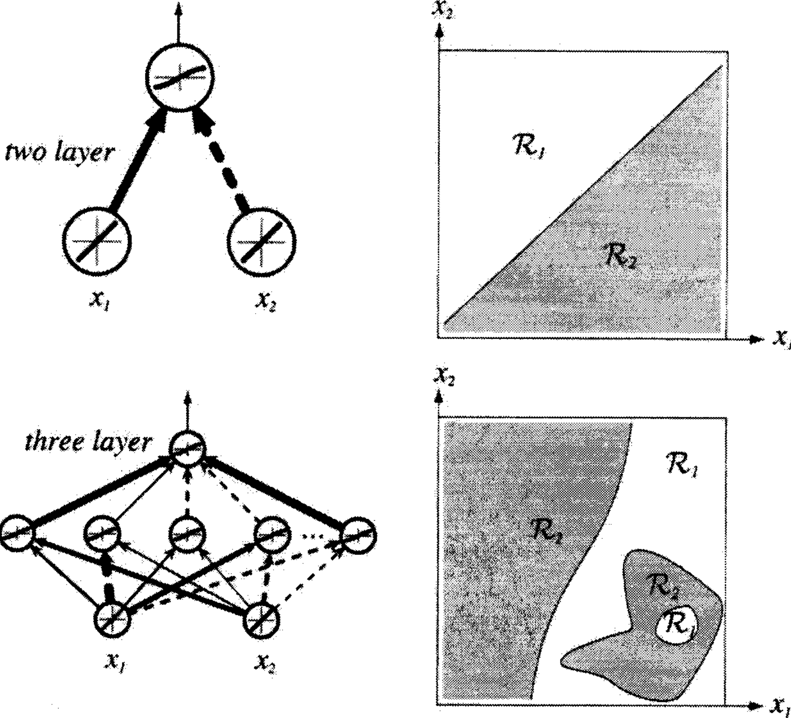

Figure 2: multilayer perceptron

2.1 More than linear functions, example: XOR

- activation functions, XOR example

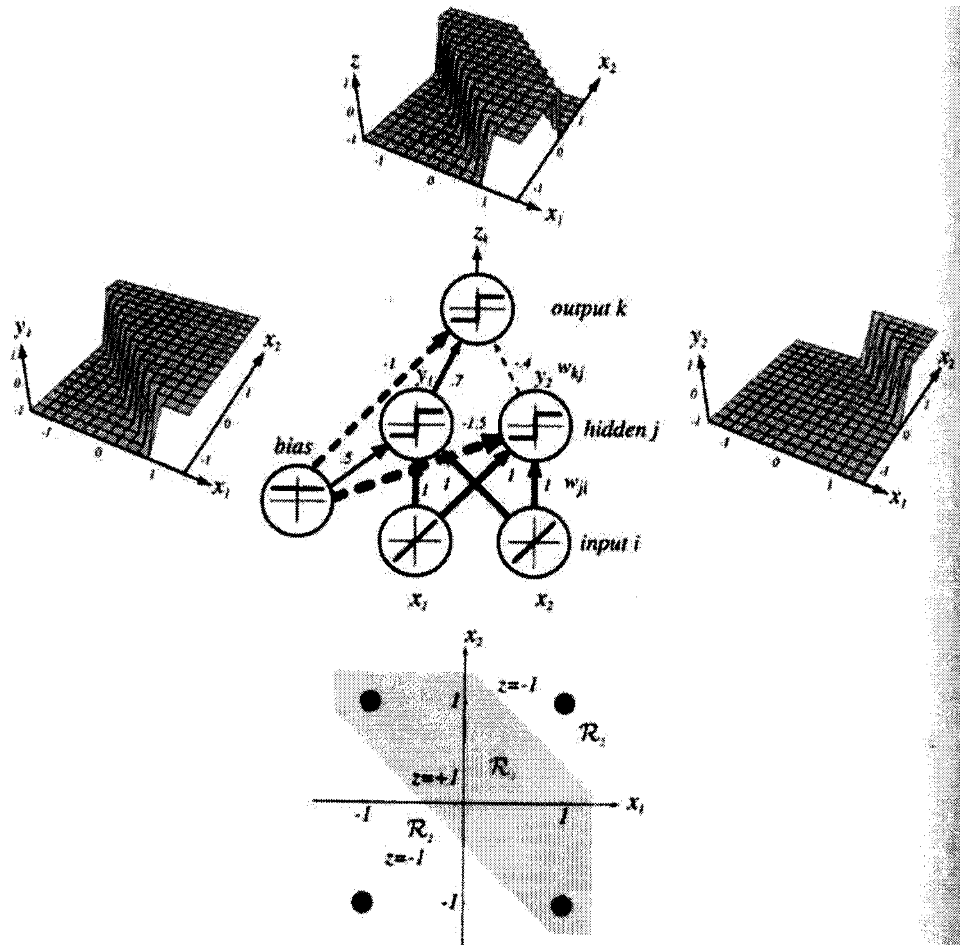

Figure 3: XOR NNet

$Y_1 = X_1 \text{ OR } X_2$ $\quad$ $Y_2 = X_1 \text{ AND } X_2$

$Z = Y_1 \text{ AND NOT } Y_2$

2.2 Construction and structure of NNets

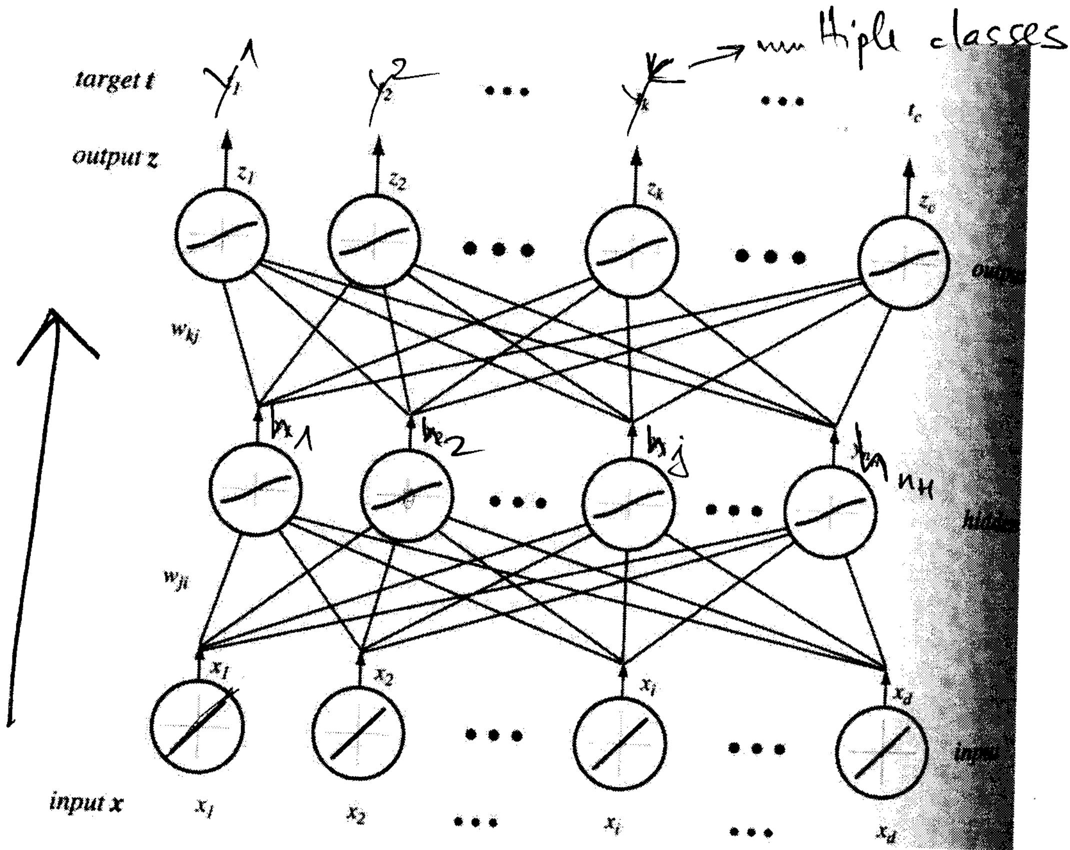

Figure 4: NNet fully connected

The oputpit units can be written as

$$g_k(\mathbf{x}) = z_k = f\left(\sum_j w_{kj} f\left(\sum_i w_{ij}x^i + w_{j0}\right) + w_{k0}\right) = F\left(F(\mathbf{x}\mathbf{w}_j)\mathbf{w}_k\right)$$

2.3 Kolmogorov theorem, expressive power of NNet

Any function g can be written

$$g(\mathbf{x}) = \sum_j \Xi_j \left(\sum_d \Psi_{dj}(x^d)\right)$$

but there is no practical way to use this theorem in practice. Usually Ξ and Ψ are very complex and not smooth.

def net = linear combination = $\sum_i w_{ij} x^i + w_{j0}$ (third layer)

$= \sum_j w_{kj} f(\,) + w_{k0}$ (second layer)

def output = sigmoid(net) = $f(\text{net})$

3 Training, Error backpropagation

- error for one ~~data~~point x

$$J(w) = \frac{1}{2}\sum_{K}(y_t^k - z^k)^2 = \frac{1}{2}\|y - z\|^2$$

$$\nabla w_{pq} = -\lambda \frac{\partial J}{\partial w_{pq}}$$

- propagation to last set of weights (close to output)

$$\frac{\partial J}{\partial w_{kj}} = \frac{\partial J}{\partial net_k} \cdot \frac{\partial net_k}{\partial w_{kj}} = -\delta_k \cdot \frac{\partial net_k}{\partial w_{kj}}$$

(δ_k: sensitivity of unit k)

$$\delta_k = -\frac{\partial J}{\partial net_k} = -\frac{\partial J}{\partial z_k}\frac{\partial z_k}{\partial net_k} = (y^k - z^k) f'(net_k)$$

$$\frac{\partial net_k}{\partial w_{kj}} = h^j$$

$$w_{kj} = w_{kj} + \lambda (y^k - z^k) f'(net_k) h_j$$

$$\lambda \delta_k h_j$$

- propagation to first set of weights (close to input)

$$\frac{\partial J}{\partial w_{ji}} = \frac{\partial J}{\partial h_j} \cdot \frac{\partial h_j}{\partial net_j} \cdot \frac{\partial net_j}{\partial w_{ji}}$$

$$\frac{\partial J}{\partial h_j} = \frac{\partial \left[\frac{1}{2}\sum_k (y^k - z^k)^2\right]}{\partial h^j} = -\sum_k (y^k - z^k) \frac{\partial z^k}{\partial net_k} \cdot \frac{\partial net_k}{\partial h_j} =$$

$$= \sum_k (y^k - z^k) f'(net_k) w_{kj}$$

$$\frac{\partial h_j}{\partial net_j} = f'(net_j) \qquad \frac{\partial net_j}{\partial w_{ji}} = x^i$$

- s[illegible]

$$w_{ji} : w_{ji} - \lambda \left[\sum_k (y^k - z^k) f'(net_k) w_{kj}\right] \cdot f(net_j) \cdot x_i$$

δ_j : sensitivity at unit j

$$\delta_j = f'(net_j) \sum_k w_{kj} \delta_k$$

δ_k: sensitivity at k

STOCHASTIC BACKP

~~Select~~ x_t (randomly chosen)

$$w_{ji} = w_{ji} + \lambda \delta_j x_i$$

$$w_{kj} = w_{kj} + \lambda \delta_k h_j$$

until $|\nabla_w J| < \varepsilon$

BATCH

$$J = \sum_{t=1}^{m} J(w, x_t)$$

for each iteration

for each x_t

$$\Delta w_{ji} = \Delta w_{ji} + \lambda \delta_j x_i$$

$$\Delta w_{kj} = \Delta w_{kj} + \lambda \delta_k h_j$$

$$w_{ji} = w_{ji} + \Delta w_{ji}$$

$$w_{kj} = w_{kj} + \Delta w_{kj}$$

until $\|\nabla J(w)\| < \varepsilon$

4 How to improve backpropagation

4.1 Activation function

- continuous, differentiable (smoothness)
  - nonlinear
  - saturation
  - monotonicity

4.2 Scale input
mass (grams) vs length (meters) STANDARDIZE (OK)

4.3 Target values
for classification use 1, -1 etc. sync with f

4.4 Noise

4.5 Number of hidden units
= complexity, overfit, test error, occam razor, REGULARIZATION

4.6 Weights initialization
→ choose weights randomly from a single dist (avoid learn one class first)

4.7 Learning rates

5 Network size and structure. Regularization

6 Jacobian and Hessian

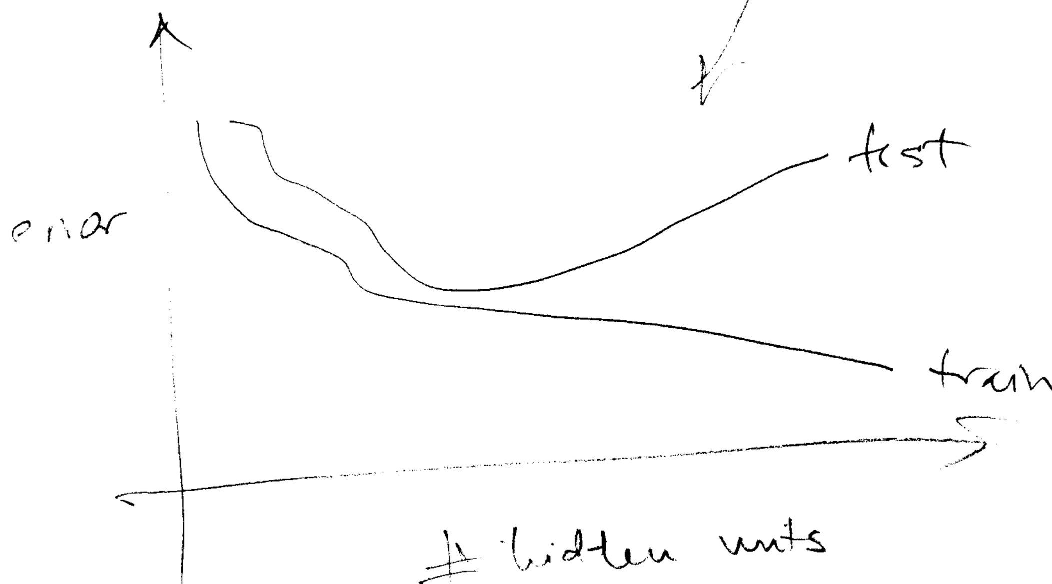

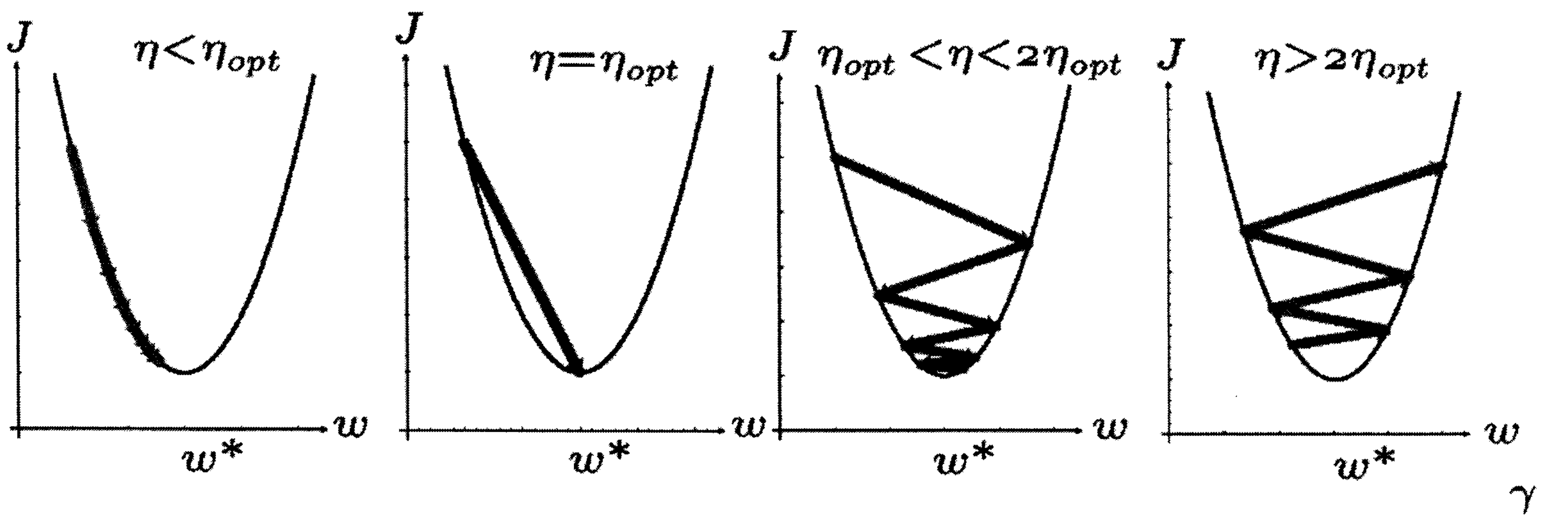

Figure 6.18: Gradient descent in a one-dimensional quadratic criterion with different learning rates. If $\eta < \eta_{opt}$, convergence is assured, but training can be needlessly slow. If $\eta = \eta_{opt}$, a single learning step suffices to find the error minimum. If $\eta_{opt} < \eta < 2\eta_{opt}$, the system will oscillate but nevertheless converge, but training is needlessly slow. If $\eta > 2\eta_{opt}$, the system diverges.