

Artificial Neural Networks (Gerstner). Solutions for week 7

Error function and Optimization

Exercise 1. Weight space symmetries

Suppose you have found a minimum for some set of weights that achieves zero loss in a network with m hidden layers of n neurons each.

- Show that there is always at least $(n!)^m$ equivalent solutions.
- Assume that a network with one hidden layer of n neurons and 1 output neuron finds a zero-loss solution. Assume that the minimum loss attained by the network with one hidden layer of $n - 1$ neurons is bigger than zero. In other words, no network with $n - 1$ neurons can find a solution.

Now imagine the original network with n neurons, constructed from $n - 1$ hidden neurons as follows: the input vector of the new neuron copies one of the $n - 1$ other input vectors, and its output weight is zero. Note that this neuron addition does not change the predictions of the network with $n - 1$ neurons on any of the data points.

Counting all possible permutations, calculate the number of the equivalent weight configurations corresponding to the network of $n - 1$ neurons by duplicating one of the input vectors.

- Now consider the duplicated neurons. Adjust the output weights of these two neurons such that their summation is constant. Note that on this line of points, the predictions of the network with one hidden layer of n neurons do not change.

Counting all possible permutations, calculate the number of equivalent lines corresponding to the network of $n - 1$ neurons.

- Show that all such configurations of the network with n neurons that correspond to the minimum of a network with $n - 1$ neurons are critical points (if the gradient descent reaches this point, it does not move further). These critical points are called symmetry-induced saddles.

Solution:

- Given a solution, and assuming the same activation functions across all neurons in a layer, we can swap any two neurons simply by swapping their input and output weights. There are therefore $(n \text{ permute } n)$ or $n!$ ways to arrange which neuron has which weights in one layer.

In addition, we can choose combinations across different layers independently. For instance, in a 2-layer network with 3 neurons in each layer, we have 6 arrangements of layer 1 and 6 arrangements of layer 2, and $6 \cdot 6 = 36$ unique combinations of the two layers together. In general, this gives us $(n!)^m$ equivalent solutions.

- Choose one of the $n - 1$ neurons for duplication in $\binom{n-1}{1} = n - 1$ ways. Permutations of n neurons yields a factor of $n!$, overall we find $(n - 1)n!$ such points.
- Choose one of the $n - 1$ neurons for duplication in $\binom{n-1}{1} = n - 1$ ways. In this case, the outgoing weights of the duplicated input vectors are indistinguishable. Permutations of n neurons yield a factor of $n!$, and factoring out the swapping among the duplicated input vectors yields a division by $2!$. Overall, we find $(n - 1)n!/2$ lines. Each one of the lines connects one pair of points in (b).
- We need to check that all partial derivatives in the network with n neurons are zero. Since the network predictions do not change, the partial derivatives of the old neurons are 0. For the new neuron, the partial derivatives of the input weight are zero since the 0 outgoing weight is backpropagated. The partial derivative of the outgoing weight is also zero since the corresponding input weight is a duplicate.

Exercise 2. Unitwise learning rates

Consider minimizing the *narrow valley* function $E(w_1, w_2) = |w_1| + 75|w_2|$ by gradient descent.

- Sketch the equipotential lines of E , i.e. the points in the $w_1 - w_2$ -plane, where $E(w_1, w_2) = c$ for different values of c .
- Start at the point $\mathbf{w}^{(0)} = (10, 10)$ and make a gradient descent step, i.e. $\mathbf{w}^{(1)} = \mathbf{w}^{(0)} - \eta(\partial E/\partial w_1, \partial E/\partial w_2)$ with $\eta = 0.1$.
Hint: Use the numeric definition of $\partial|x|/\partial x = \text{sgn}(x)$ if $x \neq 0$ and 0 otherwise.
- Continue gradient descent, i.e. compute $\mathbf{w}^{(2)}, \mathbf{w}^{(3)}$ and $\mathbf{w}^{(4)}$ and draw the points $\mathbf{w}^{(0)}, \dots, \mathbf{w}^{(4)}$ in your sketch with the equipotential lines. What do you observe? Can you choose a better value for η such that gradient descent converges faster?
- Repeat now the gradient descent procedure with different learning rates for the different dimensions, i.e. $\mathbf{w}^{(1)} = \mathbf{w}^{(0)} - (\eta_1 \partial E/\partial w_1, \eta_2 \partial E/\partial w_2)$ with $\eta_1 = 1$ and $\eta_2 = 1/75$. What do you observe? Can you choose better values for η_1 and η_2 such that gradient descent converges faster?
- An alternative to individual learning rates is to use momentum, i.e. $\Delta \mathbf{w}^{(t+1)} = -\eta(\partial E/\partial w_1, \partial E/\partial w_2) + \alpha \Delta \mathbf{w}^{(t)}$ with $\alpha \in [0, 1)$ and $\mathbf{w}^{(t+1)} = \mathbf{w}^{(t)} + \Delta \mathbf{w}^{(t+1)}$.
Repeat the gradient descent procedure for 3 steps with $\eta = 0.2$ and $\alpha = 0.5$. What do you observe?
- Assume $\partial E/\partial w_1 = 1$ in all time steps while $\partial E/\partial w_2 = \pm 75$ switches the sign in every time step. Compute $\lim_{t \rightarrow \infty} \Delta \mathbf{w}^{(t)}$ as a function of η and α . Hint: $\sum_{s=0}^t \alpha^s = \frac{1-\alpha^{t+1}}{1-\alpha}$.
- What do you conclude from this exercise in view of training neural networks by gradient descent with or without momentum?

Solution:

- In Figure 1 is shown the solution of $|w_2| = \frac{c-|w_1|}{75}$ for three different value of c .




Figure 1: Equipotential lines of the narrow valley function. $c = \{750, 300, 150, 75\}$ from outer to inner lines. The blue cross corresponds to the first iterations of gradient descent.

-

$$\frac{\partial E}{\partial w_1} = \text{sign}(w_1) \quad (1)$$

$$\frac{\partial E}{\partial w_2} = 75 \text{sign}(w_2) \quad (2)$$

$$\mathbf{w}^{(1)} = (10, 10) - 0.1(1, 75) = (9.9, 2.5) \quad (3)$$

c.

$$\mathbf{w}^{(2)} = (9.9, 2.5) - 0.1(1, 75) = (9.8, -5) \quad (4)$$

$$\mathbf{w}^{(3)} = (9.8, -5) - 0.1(1, -75) = (9.7, 2.5) \quad (5)$$

$$\mathbf{w}^{(4)} = (9.7, 2.5) - 0.1(1, 75) = (9.6, -5) \quad (6)$$

The points are sketched on Figure 1. Because of the big difference in amplitudes of the partial derivatives, finding the correct learning rate is difficult. A smaller learning rate can prevent the oscillations in w_2 but considerably slows down changes in w_1 . For faster convergence, the oscillation along the w_2 dimension should be avoided. For this initialization, it can be achieved with a learning rate of $\eta = 10/75$, which makes w_2 to reach the minimum in 1 update step, followed by w_1 , 75 steps later.

d.

$$\mathbf{w}^{(1)} = (10, 10) - (1, \frac{1}{75}) \circ (1, 75) = (9, 9) \quad (7)$$

$$\mathbf{w}^{(2)} = (9, 9) - (1, \frac{1}{75}) \circ (1, 75) = (8, 8) \quad (8)$$

$$\mathbf{w}^{(3)} = (8, 8) - (1, \frac{1}{75}) \circ (1, 75) = (7, 7) \quad (9)$$

$$\mathbf{w}^{(4)} = (7, 7) - (1, \frac{1}{75}) \circ (1, 75) = (6, 6) \quad (10)$$

Now the learning rates are scaled with respect to the partial derivatives amplitudes. Therefore, the update steps in both dimensions are equal. $\eta_1 = 10$ and $\eta_2 = 10/75$ allows to reach the minimum after 1 update.

e.

$$\mathbf{w}^{(1)} = (10, 10) - 0.2(1, 75) + 0.5(0, 0) = (9.8, -5) \quad (11)$$

$$\mathbf{w}^{(2)} = (9.8, -5) - 0.2(1, -75) + 0.5(-0.2, -15) = (9.5, 2.5) \quad (12)$$

$$\mathbf{w}^{(3)} = (9.5, 2.5) - 0.2(1, 75) + 0.5(-0.3, 7.5) = (9.15, -8.75) \quad (13)$$

The updates in the first dimension become larger and larger, while the magnitude of the oscillation in the second dimension decreases.

f. With $\eta \partial E / \partial w_1 = \eta$, we find $\Delta \mathbf{w}_1^{(t)} = -\eta \sum_{s=0}^t \alpha^s = -\eta \frac{1-\alpha^{t+1}}{1-\alpha}$. With $\eta \partial E / \partial w_2 = \pm 75\eta$ we find $\Delta w_2^{(t)} = (-1)^t 75\eta \sum_{s=0}^t (-\alpha)^s = (-1)^t 75\eta \frac{1-(-\alpha)^{t+1}}{1+\alpha}$. Therefore

$$\lim_{t \rightarrow \infty} \Delta \mathbf{w}^{(2t)} = (-\eta / (1 - \alpha), 75\eta / (1 + \alpha)) \quad (14)$$

$$\lim_{t \rightarrow \infty} \Delta \mathbf{w}^{(2t+1)} = (-\eta / (1 - \alpha), -75\eta / (1 + \alpha)) \quad (15)$$

g. Wisely chosen unitwise learning rates can significantly speed up learning. Momentum implements an effective unitwise learning rates. It can speed up learning and dampen oscillations. But generally it does not find the optimal unitwise learning rates.

Exercise 3. Averaging of Stochastic gradients (for ADAM)

We consider stochastic gradient descent in a network with three weights, (w_1, w_2, w_3) .

Evaluating the gradient for 100 input patterns (one pattern at a time), we observe the following time series

for w_1 : observed gradients are 1.1; 0.9, 1.1; 0.9; 1.1; 0.9; ...

for w_2 : observed gradients are 0.1; 0.1; 0.1; 0.1; 0.1; ...

for w_3 : observed gradients are 1.1; -0.9; 1.1; -0.9; 1.1; -0.9; ...

- Calculate the mean gradient (first moment m_1) $\langle g_k \rangle$ for w_k , $k \in [1, 2, 3]$.
- Calculate the mean of the squared gradient (second moment m_2) $\langle g_k^2 \rangle$ for w_k , $k \in [1, 2, 3]$.
- Divide the result of (a) by that of (b) so as to calculate $\langle g_k \rangle / \langle g_k^2 \rangle$ as well as $\langle g_k \rangle / \sqrt{\langle g_k^2 \rangle}$ for w_k , $k \in [1, 2, 3]$.
- The signal-to-noise ratio (SNR) of the gradient g_k is defined as

$$\text{SNR} = \frac{m_1}{\sqrt{\sigma^2}} = \frac{m_1}{\sqrt{m_2 - m_1^2}},$$

where we used the definition of the variance $\sigma^2 = \langle (g_k - m_1)^2 \rangle = m_2 - m_1^2$ of g_k . In ADAM the weight update is proportional to $\Delta w_k \propto m_1 / \sqrt{m_2}$. With that, show that

- the update in ADAM is proportional to $\Delta w_k \propto \frac{1}{\sqrt{1+1/\text{SNR}^2}} = \frac{\text{SNR}}{\sqrt{\text{SNR}^2+1}}$.
- even though the update in ADAM is not proportional to the SNR, it is proportional to the SNR for small SNR $\rightarrow 0$ and saturates for big SNR $\rightarrow \infty$.

Solution:

- The gradients correspond to three sequences with periods of 2 (for w_1), 1 (for w_2) and 4 (for w_3). Over 100 samples, they will repeat 50, 100, and 25 times respectively. This gives us

$$\begin{aligned} \langle g_1 \rangle &= \frac{1}{100} \sum_{i=1}^{50} (1.1 + 0.9) = 1, \\ \langle g_2 \rangle &= \frac{1}{100} \sum_{i=1}^{100} 0.1 = 0.1, \\ \langle g_3 \rangle &= \frac{1}{100} \sum_{i=1}^{50} (1.1 - 0.9) = 0.1. \end{aligned}$$

- By the same approach,

$$\begin{aligned} \langle g_1^2 \rangle &= \frac{1}{100} \sum_{i=1}^{50} (1.1^2 + 0.9^2) = 1.01, \\ \langle g_2^2 \rangle &= \frac{1}{100} \sum_{i=1}^{100} 0.1^2 = 0.01, \\ \langle g_3^2 \rangle &= \frac{1}{100} \sum_{i=1}^{50} (1.1^2 + (-0.9)^2) = 1.01. \end{aligned}$$

- From the previous two results,

$$\begin{aligned} \langle g_1 \rangle / \langle g_1^2 \rangle &= 0.99, \quad \langle g_1 \rangle / \sqrt{\langle g_1^2 \rangle} = 0.995, \\ \langle g_2 \rangle / \langle g_2^2 \rangle &= 10, \quad \langle g_2 \rangle / \sqrt{\langle g_2^2 \rangle} = 1, \\ \langle g_3 \rangle / \langle g_3^2 \rangle &\approx 0.1, \quad \langle g_3 \rangle / \sqrt{\langle g_3^2 \rangle} \approx 0.1. \end{aligned}$$

- By using the definition of the variance σ^2 and the SNR we see that:

$$\begin{aligned} m_2 &= m_1^2 + \sigma^2, \\ &= m_1^2 \left(1 + \frac{\sigma^2}{m_1^2} \right), \\ &= m_1^2 \left(1 + \frac{1}{\text{SNR}^2} \right), \\ \sqrt{m_2} &= m_1 \sqrt{1 + \frac{1}{\text{SNR}^2}} \end{aligned}$$

If we plug this into the ADAM weight update we obtain:

$$\Delta w_k \propto m_1 / \sqrt{m_2} = \frac{m_1}{m_1 \sqrt{1 + 1/\text{SNR}^2}} = \frac{1}{\sqrt{1 + 1/\text{SNR}^2}} = \frac{\text{SNR}}{\sqrt{\text{SNR}^2 + 1}}.$$

- For the case $\text{SNR} \rightarrow 0$ we can Taylor-expand the above expression around $\text{SNR} = 0$ and obtain $\Delta w_k(\text{SNR})|_{\text{SNR}=0} \propto \text{SNR} + \mathcal{O}(\text{SNR}^2)$. For $\text{SNR} \rightarrow \infty$ we can quickly check that $\Delta w_k \propto \frac{1}{\sqrt{1+1/\text{SNR}^2}} \rightarrow 1$ for $\text{SNR} \rightarrow \infty$.

Exercise 4. Averaging with exponential filters (for ADAM)

In this exercise we study averaging with exponential filters as used e.g. for gradient averaging in SGD with momentum or ADAM.

- a. You use an algorithm to update a variable m :

$$m(n+1) = \rho m(n) + (1-\rho)x(n) \quad (*)$$

where $\rho \in [0, 1)$ and $x(n)$ refers to an observed time series $x(1), x(2), x(3), \dots$

Show that, if all values of x are identical (that is, $x(k) = \bar{x}$ for all k), then the algo (*) converges to $m = \bar{x}$.

- b. Assume the initial condition $m(0) = 0$. Show that, for $1 - \rho \ll 1$ the algorithm outputs in time step $n + 1$ the value

$$m(n+1) = (1-\rho) \sum_{k=0}^n \exp[-(1-\rho)k] \cdot x(n-k)$$

Hint: (i) compare $m(n+1)$ with $m(n)$ and reorder terms. (ii) At the end of your calculation you may approximate $\exp(\epsilon) = 1 + \epsilon$ (which is valid for small $\epsilon \ll 1$).

- c. Your friend makes the following statement:

The algo () performs a running average of the time series $x(n)$ with an exponentially weighted window that extends roughly over $1/(1-\rho)$ samples. Therefore, if you want to include about 100 samples in the average, you should choose $\rho = 0.99$.*

Is your friend's claim correct?

Solution:

- a. We can expand the expression for $m(n+1)$ recursively into an expression in terms of all previous data points and the initial mean $m(0)$,

$$\begin{aligned}
m(n+1) &= \rho m(n) + (1-\rho)x(n) \\
&= \rho[\rho m(n-1) + (1-\rho)x(n-1)] + (1-\rho)x(n) \\
&= \rho^{n+1}m(0) + (1-\rho) \sum_{k=0}^n \rho^k x(n-k) \\
&= \rho^{n+1}m(0) + \bar{x} \cdot (1-\rho) \sum_{k=0}^n \rho^k
\end{aligned}$$

Taking $n \rightarrow \infty$ and $\rho \in [0, 1)$, the first term goes to 0. Recognizing the geometric series in the second term, we have

$$m(n+1) \approx \bar{x} \cdot (1-\rho) \frac{1}{1-\rho} = \bar{x}.$$

- b. Since $1-\rho \ll 1$, we approximate

$$\begin{aligned}
\exp(-(1-\rho)k) &= \exp(-(1-\rho))^k \\
&\approx (1-(1-\rho))^k \\
&= \rho^k
\end{aligned}$$

From the last solution, taking $m(0) = 0$, we have

$$\begin{aligned}
m(n+1) &= (1-\rho) \sum_{k=0}^n \rho^k x(n-k) \\
&\approx (1-\rho) \sum_{k=0}^n \exp[-(1-\rho)k] \cdot x(n-k)
\end{aligned}$$

- c. From above, a sample k time steps in the past has a weight of $w_k = (1-\rho)\rho^k$. From (c), we know that $\sum_{k=0}^n w_k \approx 1$ for high values of n , in which case we can also ignore $m(0)$.

We can consider the contribution of all samples more than 100 time steps in the past to the current value of the mean, for $\rho = 0.99$:

$$\begin{aligned}
\frac{\sum_{k=100}^n w_k}{\sum_{k=0}^n w_k} &\approx (1-\rho) \sum_{k=100}^n \rho^k \\
&\approx 1 - (1-\rho) \sum_{k=0}^{99} \rho^k \\
&= 1 - (1-\rho) \left(\frac{1-\rho^{100}}{1-\rho} \right) \\
&= \rho^{100} \approx 0.366
\end{aligned}$$

so samples more than 100 steps in the past will account for roughly one third of m . The value of $\rho^{1/(1-\rho)}$ converges to $\exp(-1) \approx 0.368$ for increasing values of ρ (which follows from the fact that $x \approx -\ln(1-x)$ for small values of x). So, for large enough ρ , the most recent $n = 1/(1-\rho)$ samples account for roughly two thirds of the exponential moving average (n is the *time constant*). We can think of this as roughly the number of samples being averaged over, although we should note that samples more than $1/(1-\rho)$ time steps in the past can still make a significant contribution to m (one third).

Exercise 5. Bias and variance of gradient estimators

For training data $(\mathbf{x}^1, y^1), \dots, (\mathbf{x}^P, y^P)$ and some loss $E(\mathbf{w}) = \frac{1}{P} \sum_{\mu} \ell(f_{\mathbf{w}}(\mathbf{x}^{\mu}), y^{\mu})$, the gradient is given by $\nabla E(\mathbf{w}) = \frac{1}{P} \sum_{\mu} \nabla \ell(f_{\mathbf{w}}(\mathbf{x}^{\mu}), y^{\mu})$, with e.g. $\ell(x, y) = \frac{1}{2}(x - y)^2$.

- In each step of stochastic gradient descent one sample $(\mathbf{x}^{\mu}, y^{\mu})$ of the training data is selected. Show that $\nabla \ell(f_{\mathbf{w}}(\mathbf{x}^{\mu}), y^{\mu})$ is an unbiased estimator of $\nabla E(\mathbf{w})$, if each training sample is selected with equal probability. Hint: An estimator $\hat{\theta}$ of a quantity θ is called unbiased, if its expectation $\mathbb{E}[\hat{\theta}] = \theta$.
- Instead of single sample stochastic gradient descent it is common practice to use mini-batches. Show that the mini-batch gradient estimator $\frac{1}{M} \sum_{i=1}^M \nabla \ell(f_{\mathbf{w}}(\mathbf{x}^i), y^i)$, with $1 < M < P$, has lower variance than the single sample estimator, if the samples (\mathbf{x}^i, y^i) in each mini-batch are sampled uniformly from the training data.
- How does this exercise link to Ex. 2 of week 1?

Solution:

- Let μ be a random training sample selected uniformly at random from $\{1, \dots, P\}$. Then

$$\mathbb{E}[\nabla \ell(f_{\mathbf{w}}(\mathbf{x}^{\mu}), y^{\mu})] = \frac{1}{P} \sum_{\mu} \nabla \ell(f_{\mathbf{w}}(\mathbf{x}^{\mu}), y^{\mu}) \quad (16)$$

$$= \nabla E(\mathbf{w}) \quad (17)$$

- Let $X_{\mu} = \nabla \ell(f_{\mathbf{w}}(\mathbf{x}^{\mu}), y^{\mu})$ be the gradient from a single training sample. The variance of the mini-batch gradient of size M , $\frac{1}{M} \sum_{i=1}^M X_i$ for $\{i, \dots, M\}$ selected uniformly at random from $\{1, \dots, P\}$ is

$$\mathbb{E} \left[\left(\frac{1}{M} \sum_{i=1}^M X_i - \mathbb{E}[X] \right)^2 \right] = \mathbb{E} \left[\left(\frac{1}{M} \sum_{i=1}^M (X_i - \mathbb{E}[X]) \right)^2 \right] \quad (18)$$

$$= \frac{1}{M^2} \left(\sum_{i=1}^M \mathbb{E}[(X_i - \mathbb{E}[X])^2] + \sum_{i \neq j} \mathbb{E}[(X_i - \mathbb{E}[X])(X_j - \mathbb{E}[X])] \right) \quad (19)$$

$$= \frac{1}{M} \text{Var}[X] \quad (20)$$

Exercise 6. ADAM and minibatches.

Suppose that in a project you have already spent some time on optimizing the ADAM parameters ρ_1 and ρ_2 while you ran preliminary tests with a minibatch size of 128 on your computer.

For the final run you get access to a bigger and faster computer that allows you to run minibatches of size 512.

How should you rescale ρ_1 and ρ_2 so as to expect roughly the same behavior of the two machines on the training base?

Solution:

We assume that we need to average over the gradients from the last n samples to get a good approximation to the batch gradient; this averaging happens both within a minibatch (of size s) and across minibatches. When using exponential smoothing as in ADAM, we take $\frac{n}{s} \approx 1/(1 - \rho_1)$. To get similar

behaviour, we take

$$\begin{aligned}\frac{n}{s} &\approx \frac{1}{1 - \rho_{1,s}} \\ n &\approx \frac{s}{1 - \rho_{1,s}} \\ n &\approx \frac{128}{1 - \rho_{1,128}} = \frac{512}{1 - \rho_{1,512}} \\ 1 - \rho_{1,512} &= 4(1 - \rho_{1,128}) \\ \rho_{1,512} &= 4\rho_{1,128} - 3.\end{aligned}$$

E.g. if we used $\rho_1 = 0.99$ for minibatch size 128, we can take $\rho_1 = 0.96$ on the new machine with minibatch size 512.

Exercise 7. Simple Perceptron and Bagging

We have four data points:

Two positive examples $t^1 = t^2 = 1$ at $x^1 = (1, 0)^T$ and $x^2 = (0, 1)^T$; and Two negative examples $t^3 = t^4 = 0$ at $x^3 = (0, 0)^T$ and $x^4 = (1, 1)^T$.

- Draw (with replacement) four times randomly from this data set. What is the probability that you draw each example exactly once?
- You have generated four new data sets $1 \leq k \leq 4$ by drawing with replacement from the above set. Each set contains four points. You find that in data set k point k is missing ($1 \leq k \leq 4$).
You work with the perceptron algorithm with hard gain function $g(a) = 1$ for $a > 0$ and zero otherwise.
Make a graph in the data space (input space) and sketch in the graph a solution that the perceptron algorithm finds for data set $k = 1$. Draw the hyperplane.
- Sketch in the same graph, a solution (one each) that the perceptron algorithm finds for $k = 2, 3, 4$. Label your proposed solutions with $k = 1 \dots 4$.
- Now you perform bagging. What is the value of the (real-valued) bagged output in each region of the above graph in response to an arbitrary data point x^5 . In the above graph, give the regions a different texture and write in each region a number which indicates the amplitude of the bagged response.
- Now you perform majority voting. How many of the 4 data points are correctly classified?
- Replace the four points by four Gaussian clusters of 25 data points each (Gaussians centered x^1, x^2, x^3, x^4) with standard deviation $\sigma = 0.1$ each; labels are the same for all points inside one Gaussian cluster.) Repeat the above arguments. Assume that the resampled data set k has 20 data points from cluster k , 30 data points from another cluster $k' \neq k$, and 25 from the remaining two clusters.

Sketch a plot of this new problem on a separate page and repeat the above arguments (draw the hyperplanes etc, parts b - e). Imagine you generate new data points (from the four Gaussians) for the test set. What's the probability for one of those point of not being correctly clustered after bagging with majority vote?

Solution:

- You want to pick up each data exactly once, but you don't care about the order in which you pick them. In the first pick, any data is good, in the second pick you have probability $3/4$ to

pick a data different from the first one, in the third pick you have probability $2/4$ to pick a data different from the first two and in the final pick you have probability $1/4$ to pick the last data you need. Given that each pick is done independently, the probability of picking 4 different data is $p = 1 \cdot 3/4 \cdot 2/4 \cdot 1/4 = 6/64$. This is also equivalent to $4!/4^4 = 6/64$, where $1/4^4$ is the probability of a specific (ordered) combination of 4 data and the factor $4!$ is the number of possible permutation (since we don't care about the order).

- b. Data set $k = 1$ contains all original data, but x^1 . Refer to right hand side of Fig. 2. Since we have chosen g to be an Heavyside function, any hyperplane which correctly separates correctly the points has error $E = 0$, so they are equally good (for this choice of g).

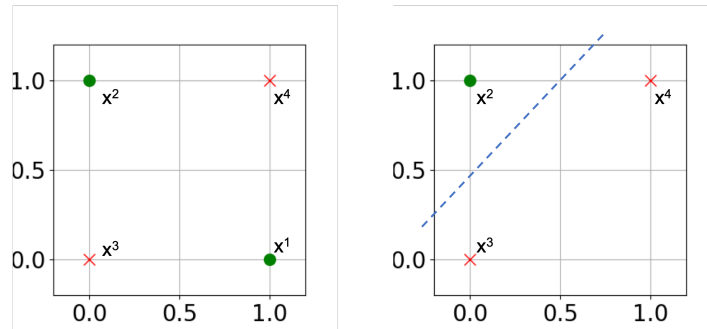


Figure 2: Left side) all 4 data points are plotted. It's the XOR problem and can't be classified with a single layer perceptron. Right side) Data space of the data set for $k = 1$. The large dotted blue line indicates one possible solution (hyperplane) of the perceptron algorithm for this data set.

- c. You can see all 4 solutions (hyperplanes) in Fig. 3.




Figure 3

- d. You can see the values of \hat{y} in the different regions of the input space in Fig. 4.
- e. Just by looking at Fig. 4, we can conclude that all 4 initial data points are correctly classified by majority voting.
- f. Fig. 5 (on the right) shows that with the suggested sampling we would still get an hyperplane similar to that of point b. In fact, by placing the hyperplane like in the figure, we have $E = 20$, corresponding to the 20 data points of type x^1 , which is the lowest error we can have in this situation. By creating analogously new data set $k = 2, 3, 4$ we get just the analogue of Fig. 3, with the same bagged output and therefore the same outcome of majority voting. Now we generate 100 new data points. How many data point will be classified correctly depends on the Euclidean distance of the hyperplanes from the mean of the cluster. For simplicity, we will consider the error probabilities associated to 1-dimensional Gaussians. In this way we can use the probabilities as in Fig. 6. We generate one data point from the Gaussian distribution $N(x^2, \sigma = 0.1)$. If the hyperplane for $k = 1$ is at a distance of 3σ from $\langle x^2 \rangle$, then the probability that a point get mis-classified is $p_{\text{error}} = 1 - 0.9973$. Now we can ask the question: what is




Figure 4




Figure 5: Dots = positive points, crosses= negative points Left) All input data are shown. 25 data points are sample from each of the 4 Gaussian distributions centered on the 4 data points of point a. All point which are sample from the same Gaussian have the same name, so x^1 are in blue, x^2 are in green, x^3 are in red and x^4 are in purple. Center) Data set $k = 1$. We sample with replacement 20 point from x^1 , 25 points from x^2 , 30 points from x^3 and 25 points from x^4 . Right) Analogue of Fig. 3.

the probability that the hyperplane $k = 1$ (drawn with the above method) lies at a distance of 3σ from center of mass of the cluster $\langle x^2 \rangle$? The cluster x^2 contains 25 data points and by construction, they are all on the same side of hyperplane $k = 1$, the hyperplane has probability $p = 0.9973^{25} = 0.93$ of being placed here.




Figure 6: thanks to Wikipedia