Neural Networks — April 20, 2022

We want to use a *Neural Network* (NN) to learn, e.g., the XOR function. Apparently, in this case we have two input nodes (`inputs` $= 2$) and one output node.

Formulas

First we fix an activation function $g\colon x \mapsto g(x)$ and its derivative g'. For instance the sigmoid function σ with $\sigma(x) = 1/(1+e^{-\beta x})$ and $\sigma'(x) = \beta\, \sigma(x)(1-\sigma(x))$ (usually $\beta = 1$), or the Rectified Linear Unit ReLU with $\text{ReLU}(x) = x$ if $x \geq 0$; and 0 otherwise and $\text{ReLU}'(x) = 1$ if $x > 0$; and 0 otherwise.

For weights W_j ($j = 0, 1, 2, \ldots, \texttt{hiddens}$) from hidden layer to output layer (with one output node) the update rule is:

$$W_j \longleftarrow W_j + \alpha \cdot a_j \cdot \Delta \quad \text{with} \quad \Delta = \texttt{error} \cdot g'(\underline{\text{in}})$$

Here α is the learning rate, a_j is the activation of the jth hidden node, and $\underline{\text{in}} = \sum_{\ell=0}^{\texttt{hiddens}} W_\ell\, a_\ell$ is the input for the single output node (in general there can be more than one); `error` is defined as the target value t minus the net output $g(\underline{\text{in}})$. Always keep the hidden bias node 0 at $a_0 = -1$.

And for weights $W_{k,j}$ ($k = 0, 1, , \ldots, \texttt{inputs}$; $j = 1, 2, \ldots, \texttt{hiddens}$) from input layer to hidden layer the update rule is:

$$W_{k,j} \longleftarrow W_{k,j} + \alpha \cdot x_k \cdot \Delta_j \quad \text{with} \quad \Delta_j = g'(\underline{\text{in}}_j) \cdot W_j \cdot \Delta$$

Here x_k is the kth input, and $\underline{\text{in}}_j = \sum_{\ell=0}^{\texttt{inputs}} W_{\ell,j}\, x_\ell$ is the input for the jth hidden node, and $a_j = g(\underline{\text{in}}_j)$. Always keep the input bias node 0 at $x_0 = -1$.

Finally, the *Backpropagation algorithm* reads like this:

repeat
  for each $(*)$ $e = (x_1, x_2, \ldots, x_{\texttt{inputs}}, t)$ **in** training set **do**
    compute $\underline{\text{in}}_j$'s, a_j's, $\underline{\text{in}}$ and $g(\underline{\text{in}})$
    compute `error`, Δ and Δ_j's
    update W_j's and $W_{k,j}$'s
until network "converged"

$(*)$ in random order
In the figure we have `inputs` $= 3$ en `hiddens` $= 2$.

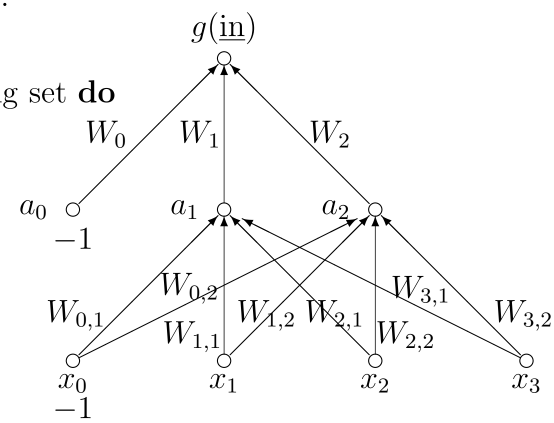

Implementation

On the website www.liacs.leidenuniv.nl/~kosterswa/AI/ a simple skeleton program called `nnskelet.cc` is available. The variables are: `input[k]` $\leftrightarrow x_k$, `inhidden[j]` $\leftrightarrow \underline{\text{in}}_j$, `acthidden[j]` $\leftrightarrow a_j$, `inoutput` $\leftrightarrow \underline{\text{in}}$, `netoutput` $\leftrightarrow g(\underline{\text{in}})$, `target` $\leftrightarrow t$, `delta` $\leftrightarrow \Delta$, `deltahidden[j]` $\leftrightarrow \Delta_j$, `inputtohidden[k][j]` $\leftrightarrow W_{k,j}$, `hiddentooutput[j]` $\leftrightarrow W_j$, `BETA` $\leftrightarrow \beta$ and `ALPHA` $\leftrightarrow \alpha$. Furthermore, `inputs` $<$ `MAX` and `hiddens` $<$ `MAX`.