



Neural Networks

Lecture 4

Multi-layer perceptrons

Igor Farkaš

2021

Introduction

- Generalization of a simple perceptron
- MLP features:
 - contains hidden layer(s)
 - neurons have a nonlinear differentiable activation function
 - full connectivity between layers
- (supervised) error “back-propagation” learning algorithm introduced
- originated after 1985: Rumelhart & McClelland: *Parallel distributed processing* (described earlier by Werbos, 1974, 1982)
- theoretical analysis difficult
- response to earlier critique of perceptrons (Minsky & Papert, 1969)

2

Two-layer perceptron

- Inputs x , weights w , v , outputs y
- Nonlinear activation function f
- Unit activation:



$$h_k = f\left(\sum_{j=1}^{n+1} v_{kj} x_j\right)$$

$$y_i = f\left(\sum_{k=1}^{q+1} w_{ik} h_k\right)$$

- Bias input: $x_{n+1} = h_{q+1} = -1$
- Examples of activation functions:

$$f(\text{net}) = \sigma(\text{net}) = \frac{1}{1 + e^{-\text{net}}}$$

$$f(\text{net}) = \tanh(\text{net}) = \frac{e^{\text{net}} - e^{-\text{net}}}{e^{\text{net}} + e^{-\text{net}}} = \frac{2}{1 + e^{-2\text{net}}} - 1$$



3

How to use output error?

- Output layer – application of delta rule
- How to compute error at hidden layer(s)?
- Instantaneous output error: $e^{(p)} = \frac{1}{2} \sum_i (d_i^{(p)} - y_i^{(p)})^2$
- We will show that error can be back-propagated across layers backwards
- At each layer the (local) weight correction has this form:
 (weight change) = (learning rate)*(unit error)*(input)
- We will derive equations for BP algorithm.


4

Derivation of BP equations

We assume an MLP with one hidden layer (but results would hold for any number of hidden layers). We only need to find the rule for updating the weights of hidden units, because these do not have targets available:

Instantaneous error at time t is:
$$e(t) = \sum_i e_i = \frac{1}{2} \sum_i (d_i(t) - y_i(t))^2$$

We need to calculate the effect of each hidden unit weight on the error. We apply gradient descent optimization as for a simple continuous perceptron. Using the chain rule, we get:



$$\frac{\partial e}{\partial v_{kj}} = \sum_i \underbrace{\frac{de_i}{dy_i}} \cdot \underbrace{\frac{\partial y_i}{\partial h_k}} \cdot \underbrace{\frac{\partial h_k}{\partial v_{kj}}} = - \sum_i (d_i - y_i) \cdot \underbrace{f'(net_i)} \cdot \underbrace{w_{ik}} \cdot \underbrace{f'(net_k)} \cdot x_j$$

$f'(net)$ – is the derivative of the activation function w.r.t. to its argument (net)

5

Learning equations for original BP


Hidden-output weights:

$$w_{ik}(t+1) = w_{ik}(t) + \alpha \delta_i h_k \quad \text{where} \quad \delta_i = (d_i - y_i) f'(net_i)$$

Input-hidden weights:

$$v_{kj}(t+1) = v_{kj}(t) + \alpha \delta_k x_j \quad \text{where} \quad \delta_k = (\sum_i w_{ik} \delta_i) f'(net_k)$$

- BP provides an “approximation” to the trajectory in weight space computed by the method of steepest descent: the smaller α , the smoother the trajectory
- BP addresses credit-assignment problem (hidden units’ responsibility for error)



6


Summary of back-propagation algorithm

Given: training data: input-target $\{x^{(p)}, d^{(p)}\}$ patterns

Initialization: randomize weights, set learning parameters

Training:

1. choose input $x^{(p)}$, compute outputs $y^{(p)}$ (**forward pass**),
2. evaluate chosen error function $e(t)$, $E \leftarrow E + e(t)$
3. compute δ_i , δ_k (**backward pass**)
4. adjust weights Δw_{ik} and Δv_{kj}
5. if all patterns used, then goto 6, else go to 1
6. if `stopping_criterion` is met, then end
else permute inputs and go to 1



7

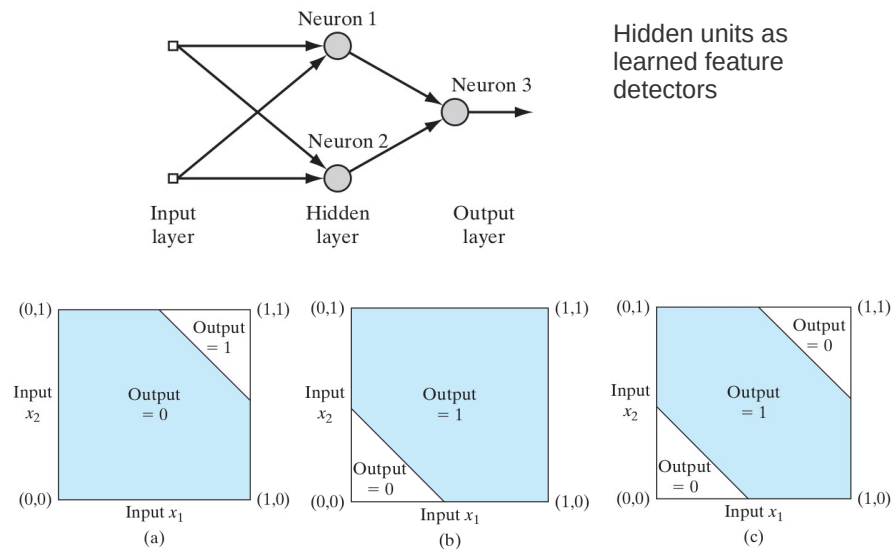
Stopping criteria

- In general, the BP algorithm cannot be shown to converge,
- there are no well-defined criteria for stopping its operation.
- However, there are some reasonable criteria, each with its own practical merit, that may be used:
- “BP algorithm is considered to have converged when
 - ... the Euclidean norm of the gradient vector reaches a sufficiently small gradient threshold.” (Kramer and Sangiovanni-Vincentelli, 1989)
 - ... the absolute rate of change in the average squared error per epoch is sufficiently small.”

(Haykin, 2009)

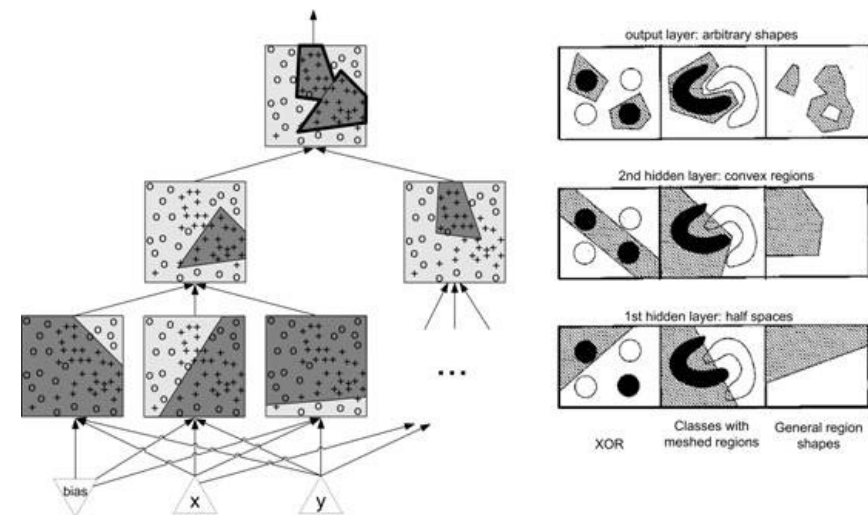
8

XOR problem solved with MLP



9

Decision regions in MLP



10


Heuristics to improve BP

1. Stochastic versus batch update (faster)
2. Shuffle the patterns before each epoch
3. Maximize information content (use 'difficult' and different inputs)
4. Consider "emphasizing scheme" with care
5. Consider bipolar activation function (e.g. \tanh), or no saturation (ReLU)
6. Use appropriate target values (ϵ -tolerance)
7. Input normalization
8. Parameter initialization: smaller weights and unit thresholds
9. Learn from hints: a priori information about function to be learned
10. Proper learning rate: the same for all units, or larger for lower layers.

(Haykin, 2009)

11

Normalization of inputs



12

Output representation and decision rule

- for binary (0/1) targets, use logistic function
- for categorical targets, use 1-of-M coding and softmax activation
- for continuous-valued targets with a bounded range, logistic and tanh functions can be used (with proper scaling)
- if target values > 0 , but have no known upper bound, you can use an exponential output activation function (beware of overflow)
- for continuous-valued targets with no known bounds, use the identity or linear activation function

M-class classification problem:

if softmax: outputs estimate a posteriori class probabilities $P(C_i | \mathbf{x})$

13

MLP as a universal approximator

Theorem: Let's have $A_{\text{train}} = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(p)}, \dots, \mathbf{x}^{(N)}\}$, $\mathbf{x}^{(p)} \in \mathbb{R}^n$. For $\epsilon > 0$ and arbitrary continuous function $F: \mathbb{R}^n \rightarrow (0,1)$ defined on discrete set A_{train} there exists such a function G :

$$G(\mathbf{x}^{(p)}) = f\left(\sum_{k=1}^{q+1} w_k f\left(\sum_{j=1}^{n+1} v_{kj} x_j^{(p)}\right)\right)$$

where parameters $w_k, v_{kj} \in \mathbb{R}$ and $f(z) = \mathbb{R} \rightarrow (0,1)$ is a continuous and monotone-increasing function satisfying $f(-\infty) = 0$ and $f(\infty) = 1$, such that:

$$\sum_p |F(\mathbf{x}^{(p)}) - G(\mathbf{x}^{(p)})| < \epsilon.$$

We say that G approximates F on A_{train} with accuracy ϵ .

G can be interpreted as a **2-layer feedforward NN with one output neuron**.

- it is an existence theorem
- **curse of dimensionality** – sparsity problem, how to get a dense sample for large n and complex F

Hecht-Nielsen (1987), Hornik, Stinchcombe & White (1989)

14

Generalization



Data set:

$$A = A_{\text{estim}} \cup A_{\text{val}} \cup A_{\text{test}}$$

- Validation set is used for model selection.
- Generalization (= testing set performance) is important in using ANNs.

Generalization depends on:

- size of A_{estim} and its representativeness
- architecture of NN
- complexity of the problem



15

Cross-validation

- The goal of cross-validation is to test the model's ability to predict new data that was not used in estimating it.
- in statistics (Stone, 1974): data partitioning = train + test
- **training set** (estimation subset + validation subset):
 - train the model using estimation subset
 - validate the model using (a smaller) validation subset
 - find the best model (model selection) using early stopping
- **test set** – used for assessing the performance of the chosen model
- Typical choice: valid. set = 10-20% of training set
- There exist variations of CV approaches

16

k-fold cross-validation

- useful for smaller data sets, to get a statistically more accurate model
- Split $A_{\text{train}} = A_{\text{val}}^1 \cup A_{\text{val}}^2 \cup \dots \cup A_{\text{val}}^k$ (A_{val}^i and A_{val}^j are disjunct)
- train each model M k -times, with $A_{\text{estim}}^i = A_{\text{train}} \setminus A_{\text{val}}^i$, $i=1, 2, \dots, k$.
Use **early stopping** in training.


- For each M compute the cross-validation coefficient:

$$CV(M) = 1/k \sum_{i=1}^k E_{\text{val}}^i(M)$$

- Choose M with smallest CV; computationally more demanding

17

Bias–variance tradeoff




$$d = f(x) \quad E[(d - \tilde{f}(x))^2] = \text{Bias}[\tilde{f}(x)]^2 + \text{Var}[\tilde{f}(x)] + \sigma^2$$

$$\text{Bias}[\tilde{f}(x)] = E[\tilde{f}(x) - f(x)] \quad \text{Var}[\tilde{f}(x)] = E[\tilde{f}(x)^2] - E[\tilde{f}(x)]^2$$

18

Quantifying the model complexity

- Based on statistical learning theory: The Vapnik–Chervonenkis (VC) dimension measures the capacity of a binary classifier (model M).
- VC dimension is defined as the **maximum** number of points that M can classify correctly, for arbitrary assignments of labels to those points.
- Discrepancy b/w TrainErr and TestError is bounded from above by a quantity that grows with growing VCdim and shrinks with Train dataset size increase (Goodfellow et al, 2015).
- e.g. for a general linear classifier, VCdim = 3.



19

Summary

- BP = standard supervised learning algorithm for MLPs in classification and/or regression tasks
- computationally efficient, complexity $O(W)$
- finds locally optimal solutions
- MLP as universal approximator
- Model selection: min. validation error (\Rightarrow best generalization)
- Bias–variance tradeoff
- VC dimension
- No free lunch theorems

20