


Neural Networks

Part 1 - Introduction

CSE 4309 – Machine Learning
Vassilis Athitsos
Computer Science and Engineering Department
University of Texas at Arlington

Perceptrons




- A perceptron is a function that maps D-dimensional vectors to real numbers.
- For notational convenience, we add an extra input, called the **bias input**. The bias input is **always equal to 1**.

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \dots \\ x_D \end{bmatrix}$$

- b is called the **bias weight**. It is optimized during training.
- w_1, \dots, w_D are also weights that are optimized during training.

Perceptrons



- A perceptron computes its output z in two steps:


Step 1: $a = b + \mathbf{w}^T \mathbf{x} = b + \sum_{i=0}^D (w_i x_i)$

Step 2: $z = h(a)$

- h is called an **activation function**.

- For example, h could be the sigmoid function $\sigma(a) = \frac{1}{1+e^{-a}}$

Perceptrons




- A perceptron computes its output z in two steps:

Step 1: $a = b + \mathbf{w}^T \mathbf{x} = b + \sum_{i=0}^D (w_i x_i)$

Step 2: $z = h(a)$


- In a single formula: $z = h\left(b + \sum_{i=0}^D (w_i x_i)\right)$

Notation for Bias Weight



- There is an alternative representation that we will not use, where b is denoted as w_0 , and weight vector $\mathbf{w} = (w_0, w_1, w_2, \dots, w_D)$.
- Then, instead of writing $z = h(b + \mathbf{w}^T \mathbf{x})$ we can simply write $z = h(\mathbf{w}^T \mathbf{x})$.
- In our slides, we will denote the bias weight as b and treat it separately from the other weights. That will make life easier later.

Perceptrons and Neurons




- Perceptrons are inspired by neurons.
 - Neurons are the cells forming the nervous system, and the brain.
 - Neurons somehow sum up their inputs, and if the sum exceeds a threshold, they "fire".
- Since brains are "intelligent", computer scientists have been hoping that perceptron-based systems can be used to model intelligence.

Activation Functions

- A perceptron produces output $z = h(b + \mathbf{w}^T \mathbf{x})$.
- One choice for the activation function h : the **step function**.

$$h(a) = \begin{cases} 0, & \text{if } a < 0 \\ 1, & \text{if } a \geq 0 \end{cases}$$




- The step function is useful for providing some intuitive examples.
- It is not useful for actual real-world systems.
 - Reason: it is not differentiable, it does not allow optimization via gradient descent.

Activation Functions

- A perceptron produces output $z = h(b + \mathbf{w}^T \mathbf{x})$.
- Another choice for the activation function $h(a)$: the **sigmoidal function**.

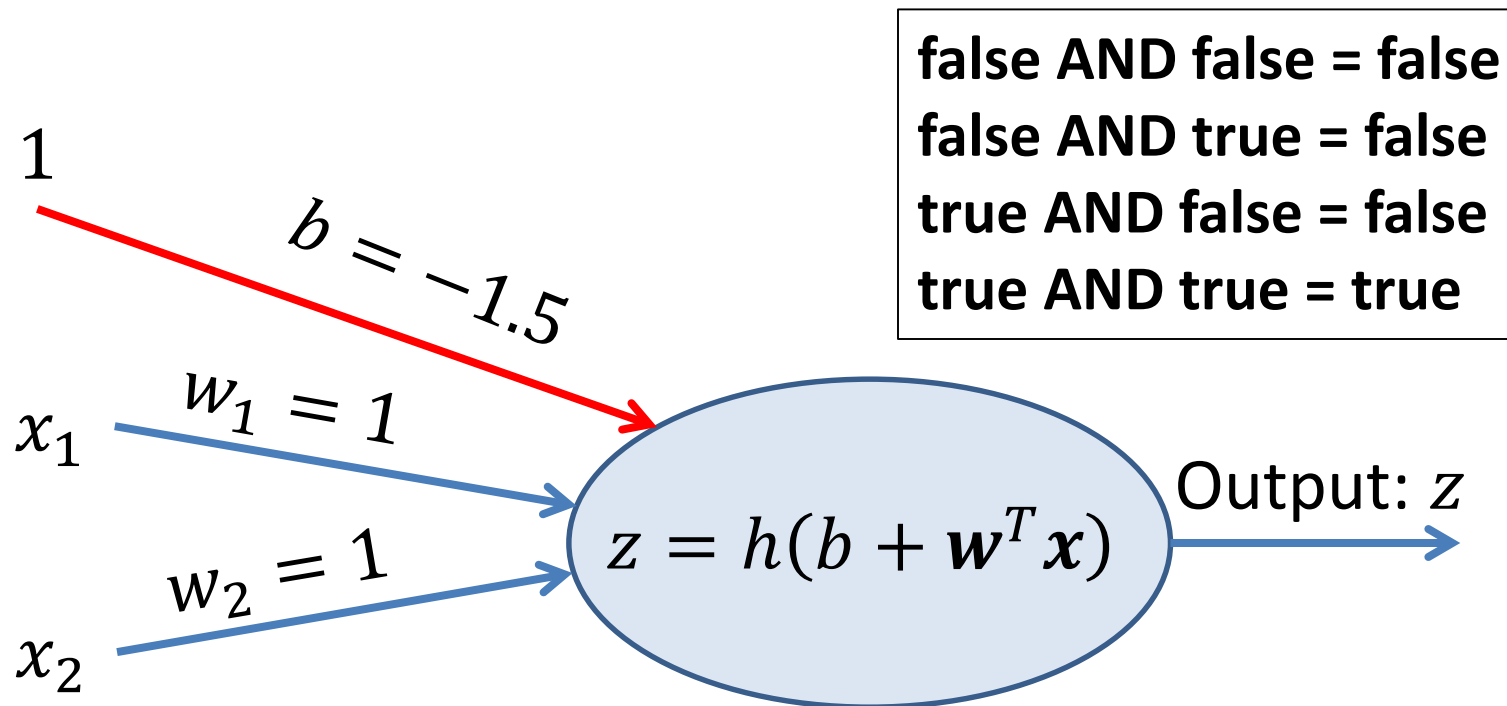
$$\sigma(a) = \frac{1}{1+e^{-a}}$$

- The sigmoidal is often used in real-world systems.
- It is a differentiable function, it allows use of gradient descent.



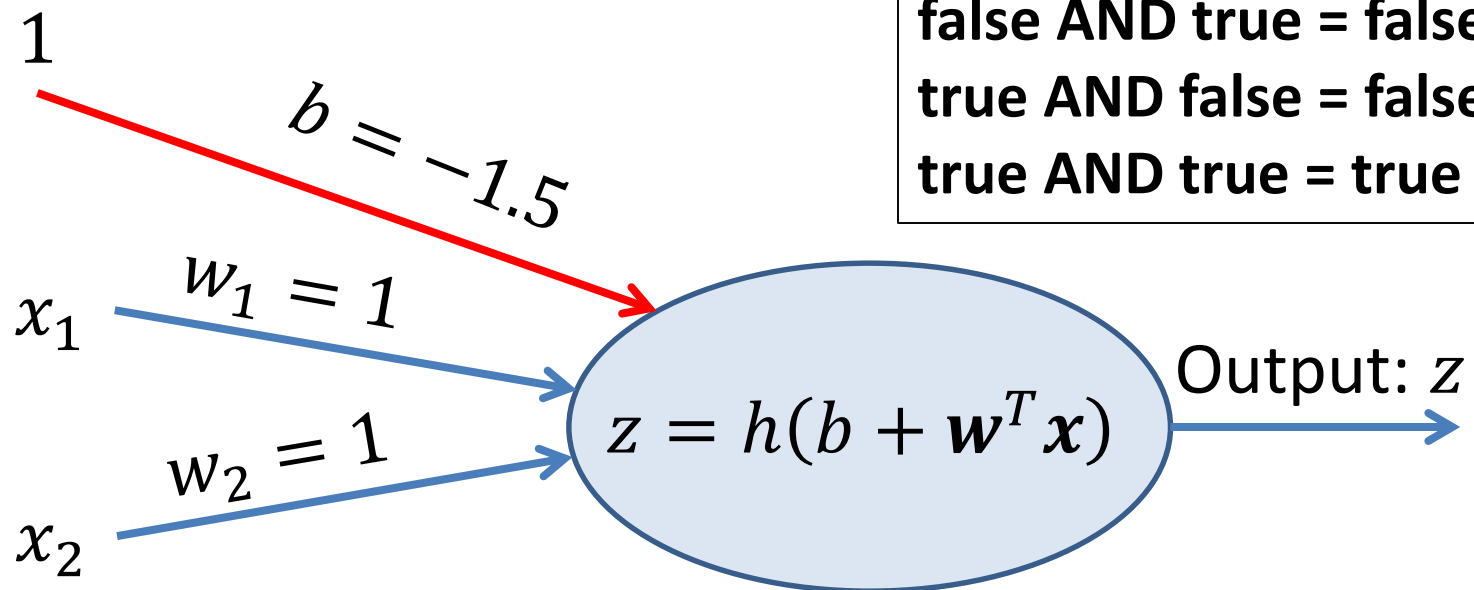
Example: The AND Perceptron

- Suppose we use the **step function** for activation.
- Suppose boolean value **false** is represented as number 0.
- Suppose boolean value **true** is represented as number 1.
- Then, the perceptron below computes the boolean AND function:



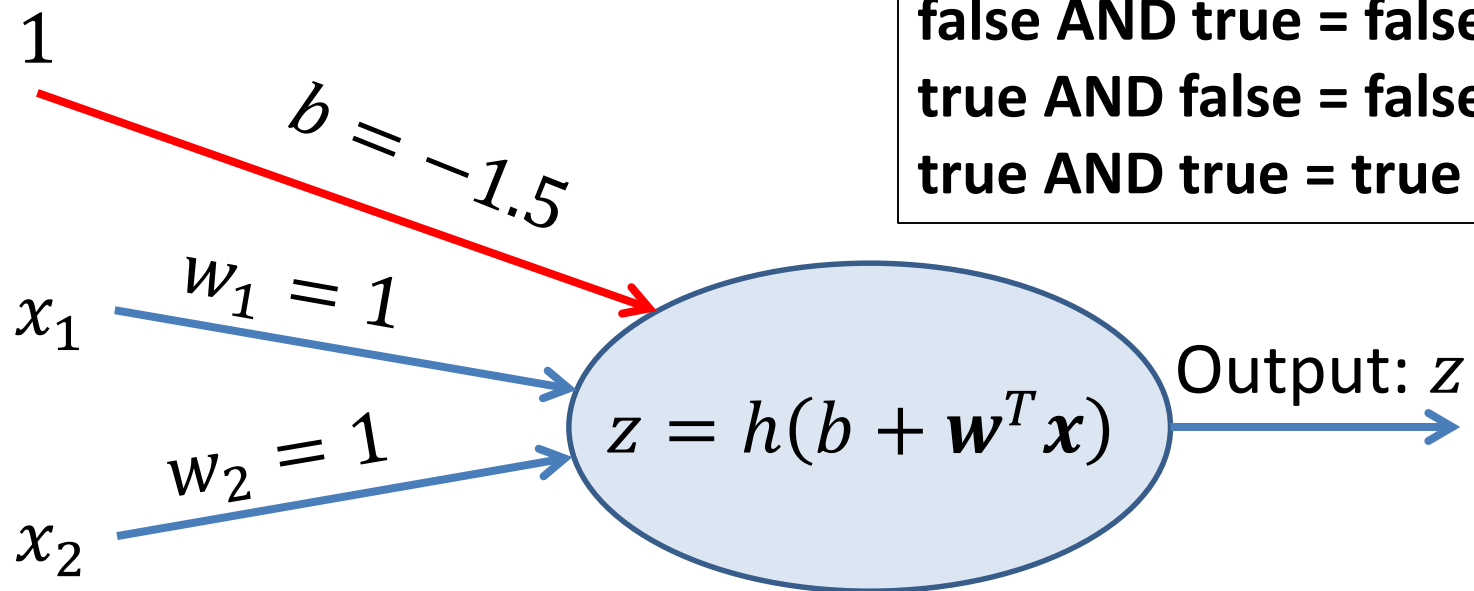
Example: The AND Perceptron

- Verification: If $x_1 = 0$ and $x_2 = 0$:
 - $b + \mathbf{w}^T \mathbf{x} = -1.5 + 1 * 0 + 1 * 0 = -1.5$.
 - $h(b + \mathbf{w}^T \mathbf{x}) = h(-1.5) = 0$.
- Corresponds to case **false AND false = false**.




Example: The AND Perceptron

- Verification: If $x_1 = 0$ and $x_2 = 1$:
 - $b + \mathbf{w}^T \mathbf{x} = -1.5 + 1 * 0 + 1 * 1 = -0.5$.
 - $h(b + \mathbf{w}^T \mathbf{x}) = h(-0.5) = 0$.
- Corresponds to case **false AND true = false**.




Example: The AND Perceptron

- Verification: If $x_1 = 1$ and $x_2 = 0$:
 - $b + \mathbf{w}^T \mathbf{x} = -1.5 + 1 * 1 + 1 * 0 = -0.5$.
 - $h(b + \mathbf{w}^T \mathbf{x}) = h(-0.5) = 0$.
- Corresponds to case **true AND false = false**.



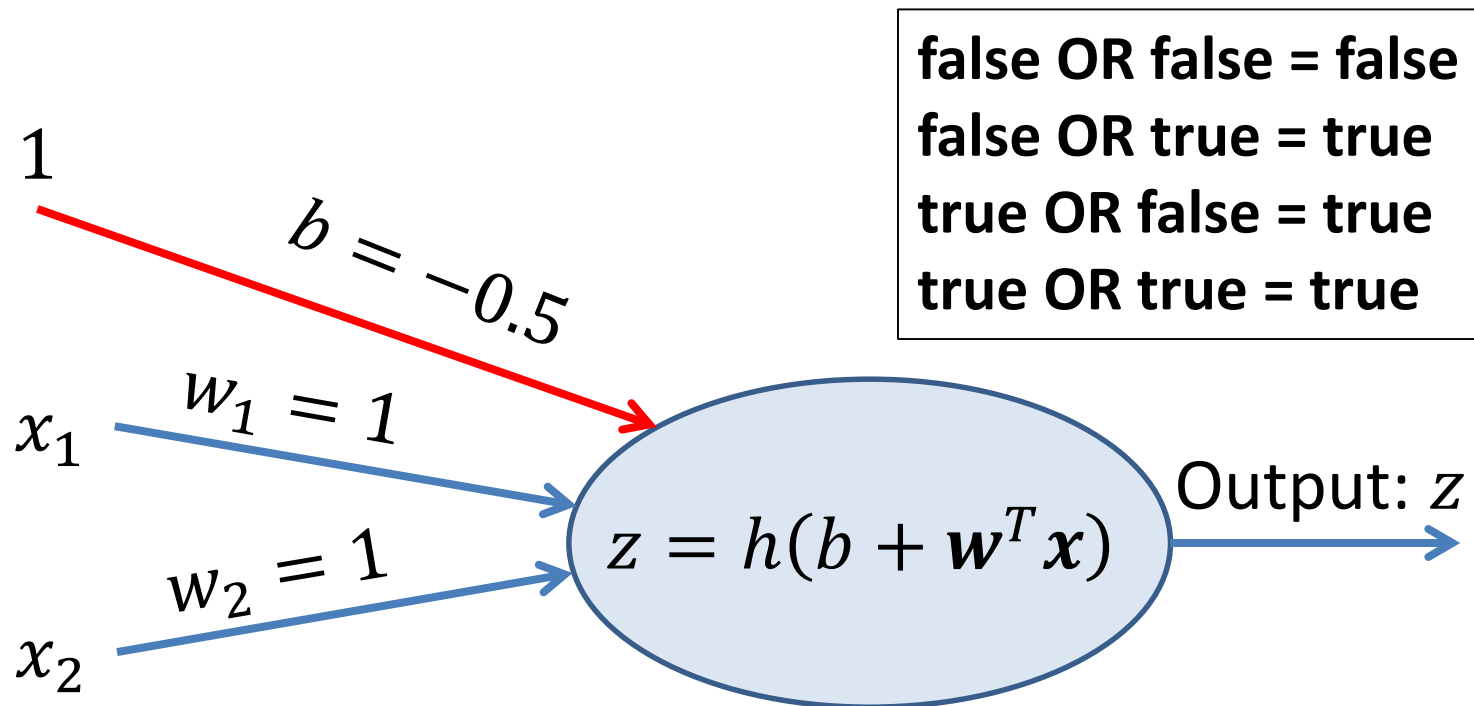
Example: The AND Perceptron

- Verification: If $x_1 = 1$ and $x_2 = 1$:
 - $b + \mathbf{w}^T \mathbf{x} = -1.5 + 1 * 1 + 1 * 1 = 0.5$.
 - $h(b + \mathbf{w}^T \mathbf{x}) = h(0.5) = 1$.
- Corresponds to case **true AND true = true**.



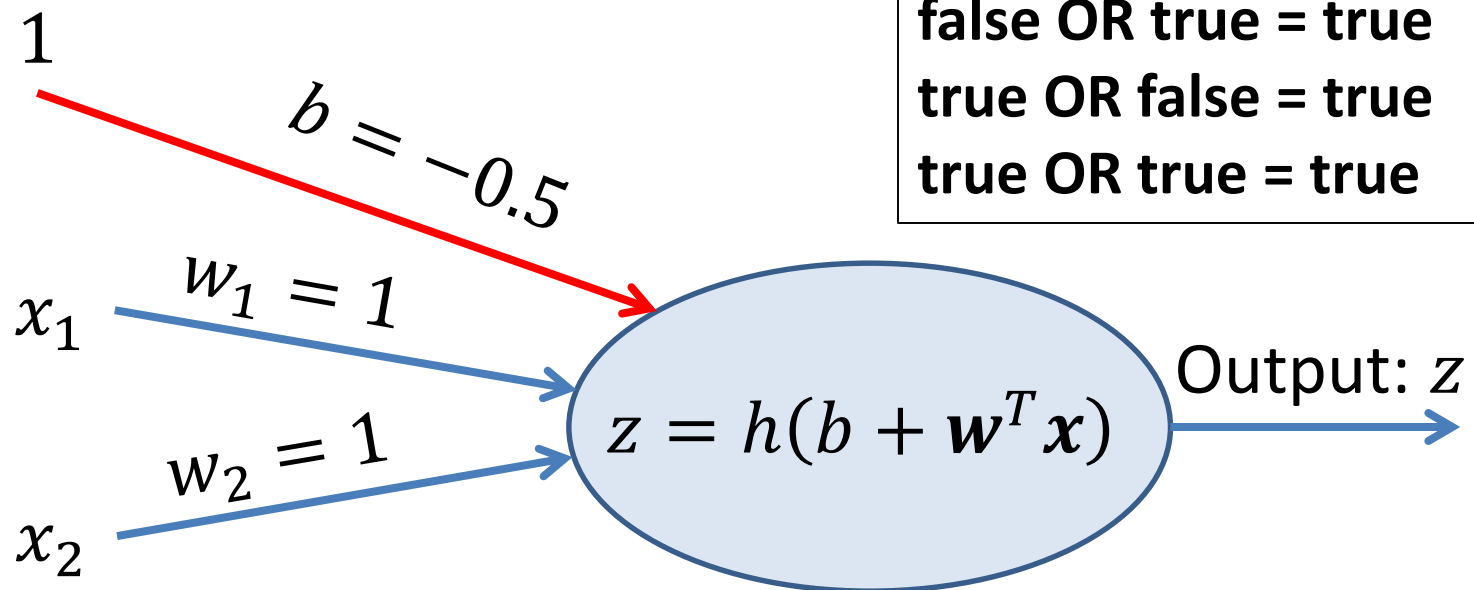
Example: The OR Perceptron

- Suppose we use the **step function** for activation.
- Suppose boolean value **false** is represented as number 0.
- Suppose boolean value **true** is represented as number 1.
- Then, the perceptron below computes the boolean OR function:




Example: The OR Perceptron

- Verification: If $x_1 = 0$ and $x_2 = 0$:
 - $b + \mathbf{w}^T \mathbf{x} = -0.5 + 1 * 0 + 1 * 0 = -0.5$.
 - $h(b + \mathbf{w}^T \mathbf{x}) = h(-0.5) = 0$.
- Corresponds to case **false OR false = false**.




Example: The OR Perceptron

- Verification: If $x_1 = 0$ and $x_2 = 1$:
 - $b + \mathbf{w}^T \mathbf{x} = -0.5 + 1 * 0 + 1 * 1 = 0.5$.
 - $h(b + \mathbf{w}^T \mathbf{x}) = h(0.5) = 1$.
- Corresponds to case **false OR true = true**.




Example: The OR Perceptron

- Verification: If $x_1 = 1$ and $x_2 = 0$:
 - $b + \mathbf{w}^T \mathbf{x} = -0.5 + 1 * 1 + 1 * 0 = 0.5$.
 - $h(b + \mathbf{w}^T \mathbf{x}) = h(0.5) = 1$.
- Corresponds to case **true OR false = true**.



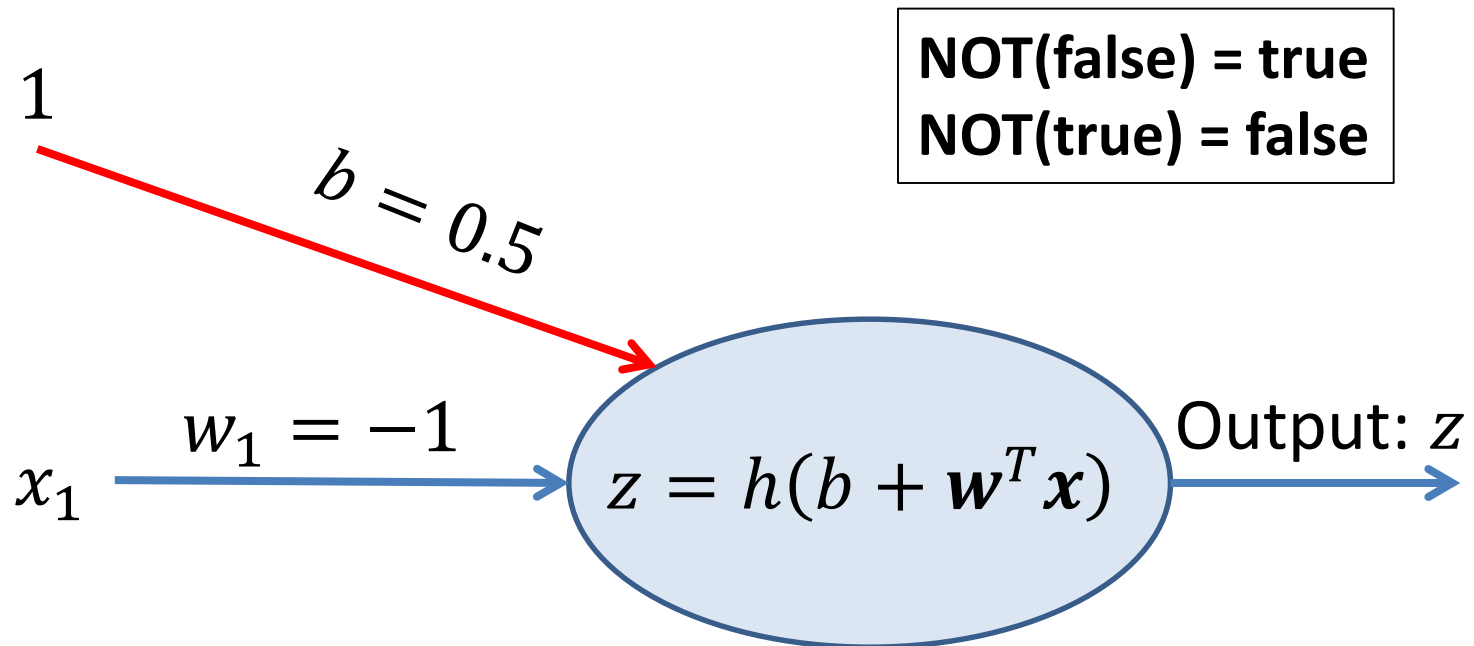
Example: The OR Perceptron

- Verification: If $x_1 = 1$ and $x_2 = 1$:
 - $b + \mathbf{w}^T \mathbf{x} = -0.5 + 1 * 1 + 1 * 1 = 1.5$.
 - $h(b + \mathbf{w}^T \mathbf{x}) = h(1.5) = 1$.
- Corresponds to case **true OR true = true**.




Example: The NOT Perceptron

- Suppose we use the **step function** for activation.
- Suppose boolean value **false** is represented as number 0.
- Suppose boolean value **true** is represented as number 1.
- Then, the perceptron below computes the boolean NOT function:




Example: The NOT Perceptron

- Verification: If $x_1 = 0$:
 - $b + \mathbf{w}^T \mathbf{x} = 0.5 - 1 * 0 = 0.5$.
 - $h(b + \mathbf{w}^T \mathbf{x}) = h(0.5) = 1$.
- Corresponds to case **NOT(false) = true**.




Example: The NOT Perceptron

- Verification: If $x_1 = 1$:
 - $b + \mathbf{w}^T \mathbf{x} = 0.5 - 1 * 1 = -0.5$.
 - $h(b + \mathbf{w}^T \mathbf{x}) = h(-0.5) = 0$.
- Corresponds to case **NOT(true) = false**.



The XOR Function

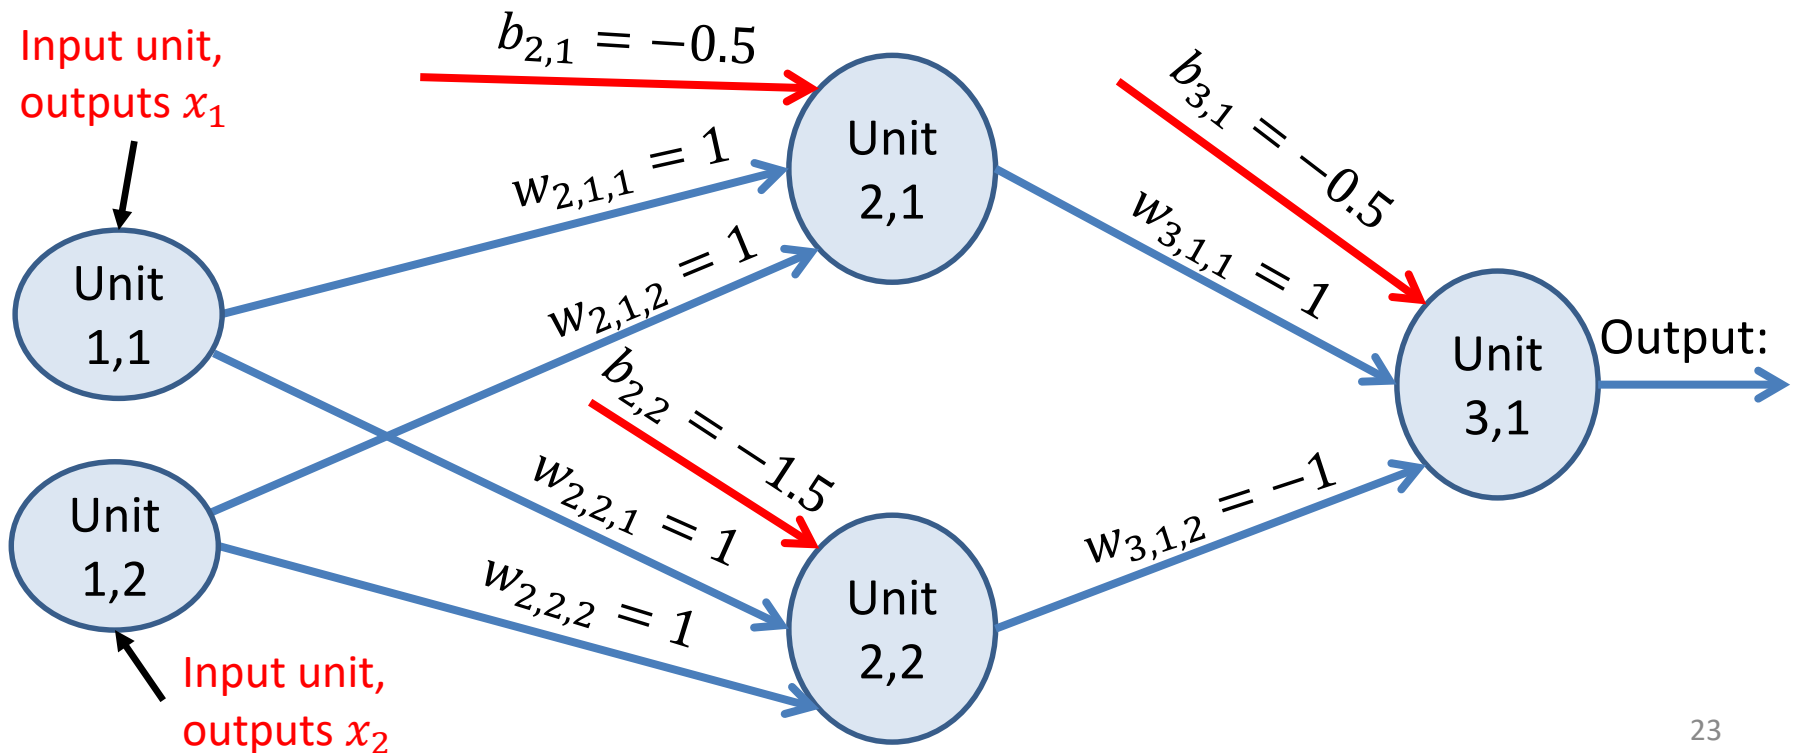
false XOR false = false
false XOR true = true
true XOR false = true
true XOR true = false



- As before, we represent **false** with 0 and **true** with 1.
- The figure shows the four input points of the XOR function.
 - red corresponds to output value **true**.
 - green corresponds to output value **false**.
- The two classes (true and false) are not linearly separable.
- Therefore, no perceptron can compute the XOR function.

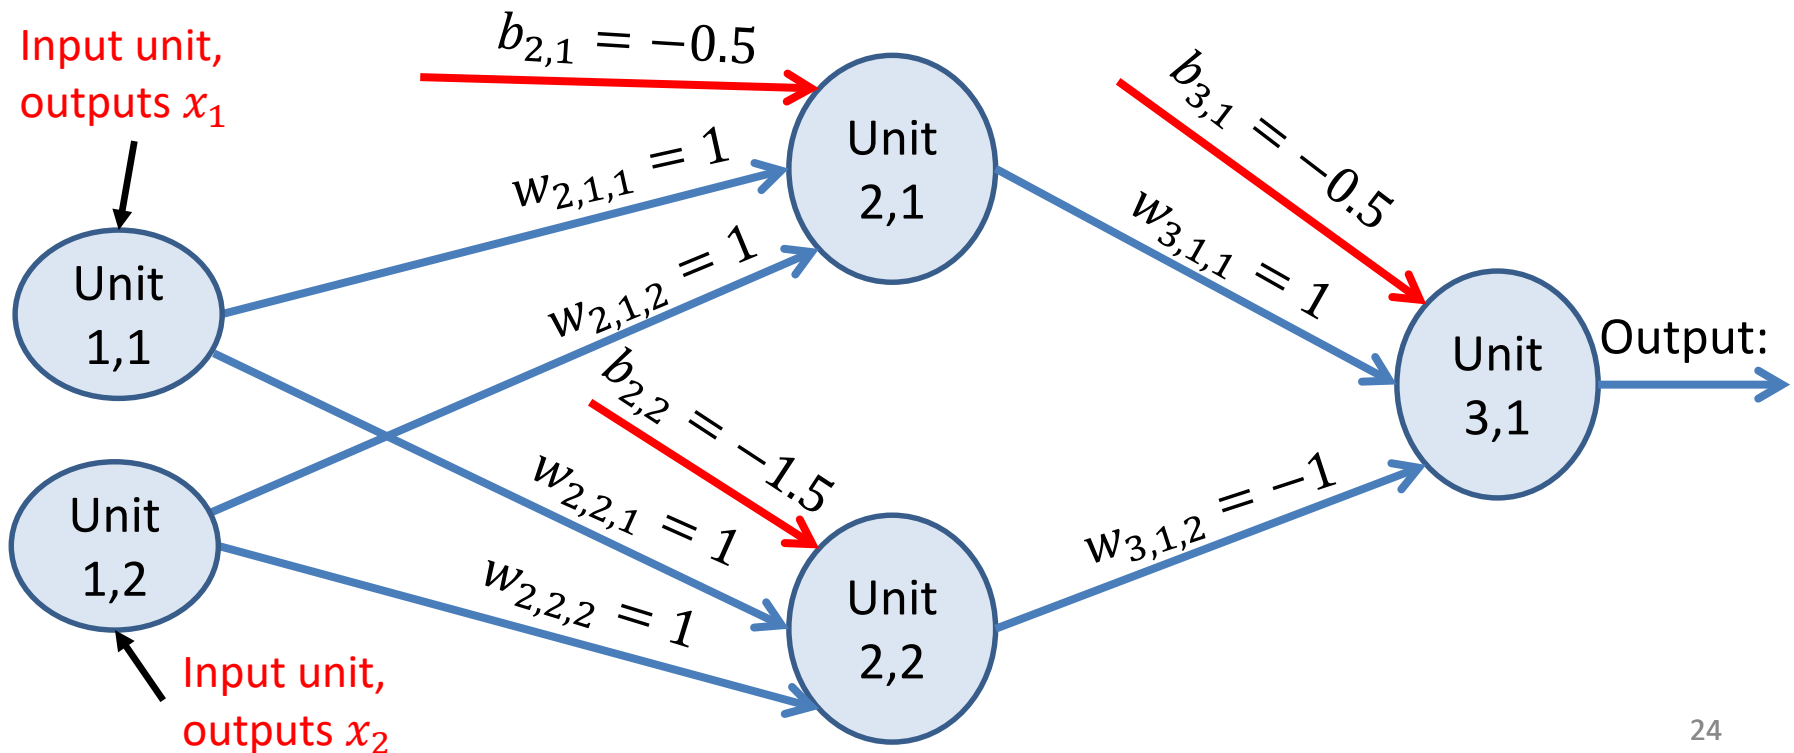
Our First Neural Network: XOR

- A neural network is built using perceptrons as building blocks.
- The inputs to some perceptrons are outputs of other perceptrons.
- Here is an example neural network computing the XOR function.



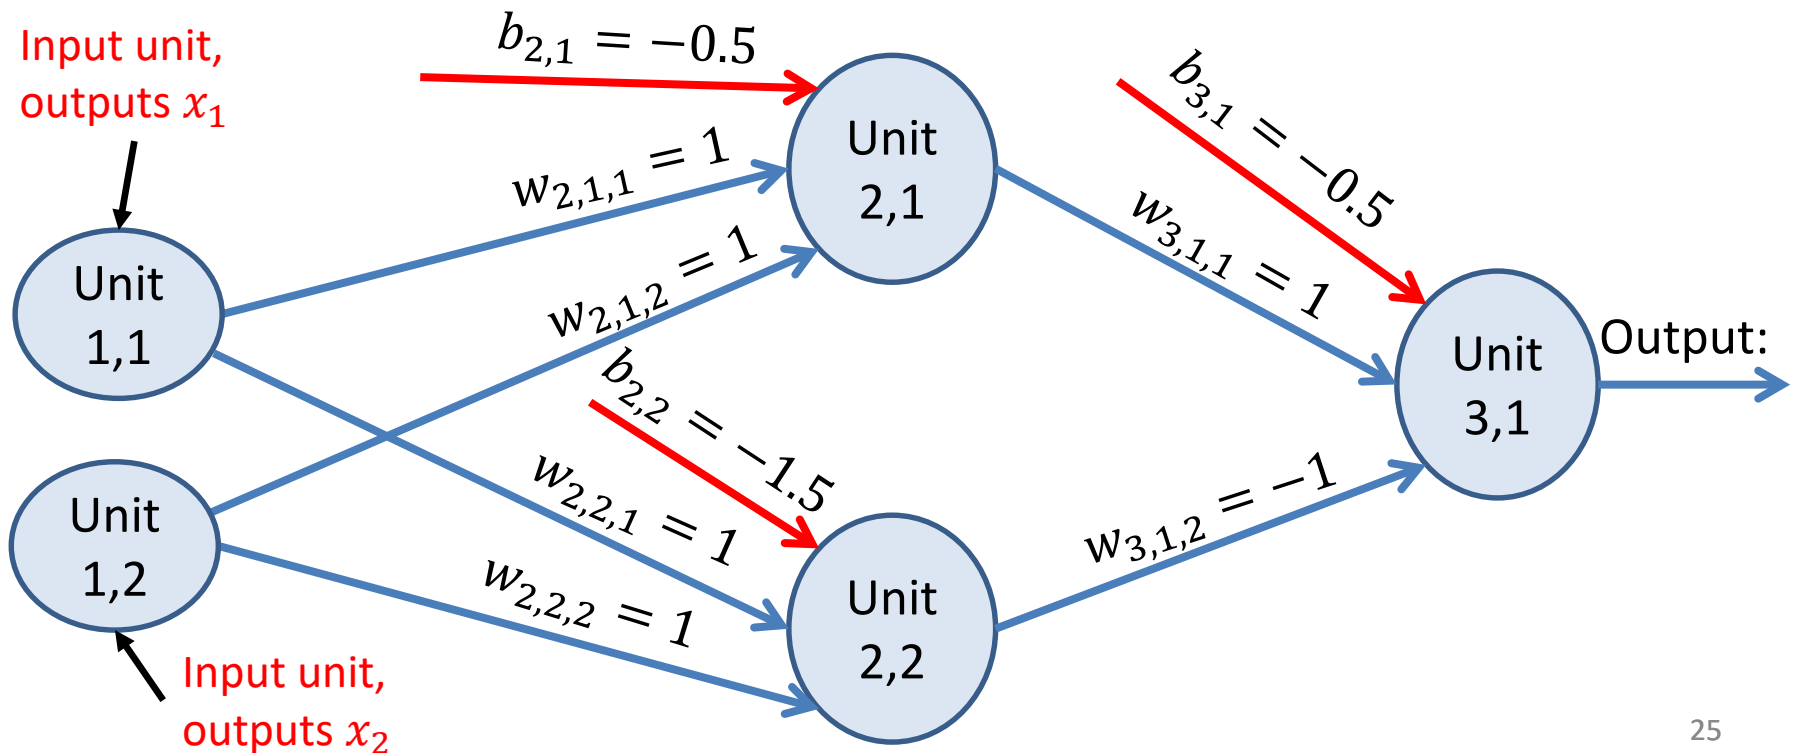
Our First Neural Network: XOR

- Terminology: inputs and perceptrons are all called “units”.
- Units are grouped in layers: layer 1 (input), layer 2, layer 3 (output).
- The input layer just represents the inputs to the network.
 - There are two inputs: x_1 and x_2 .



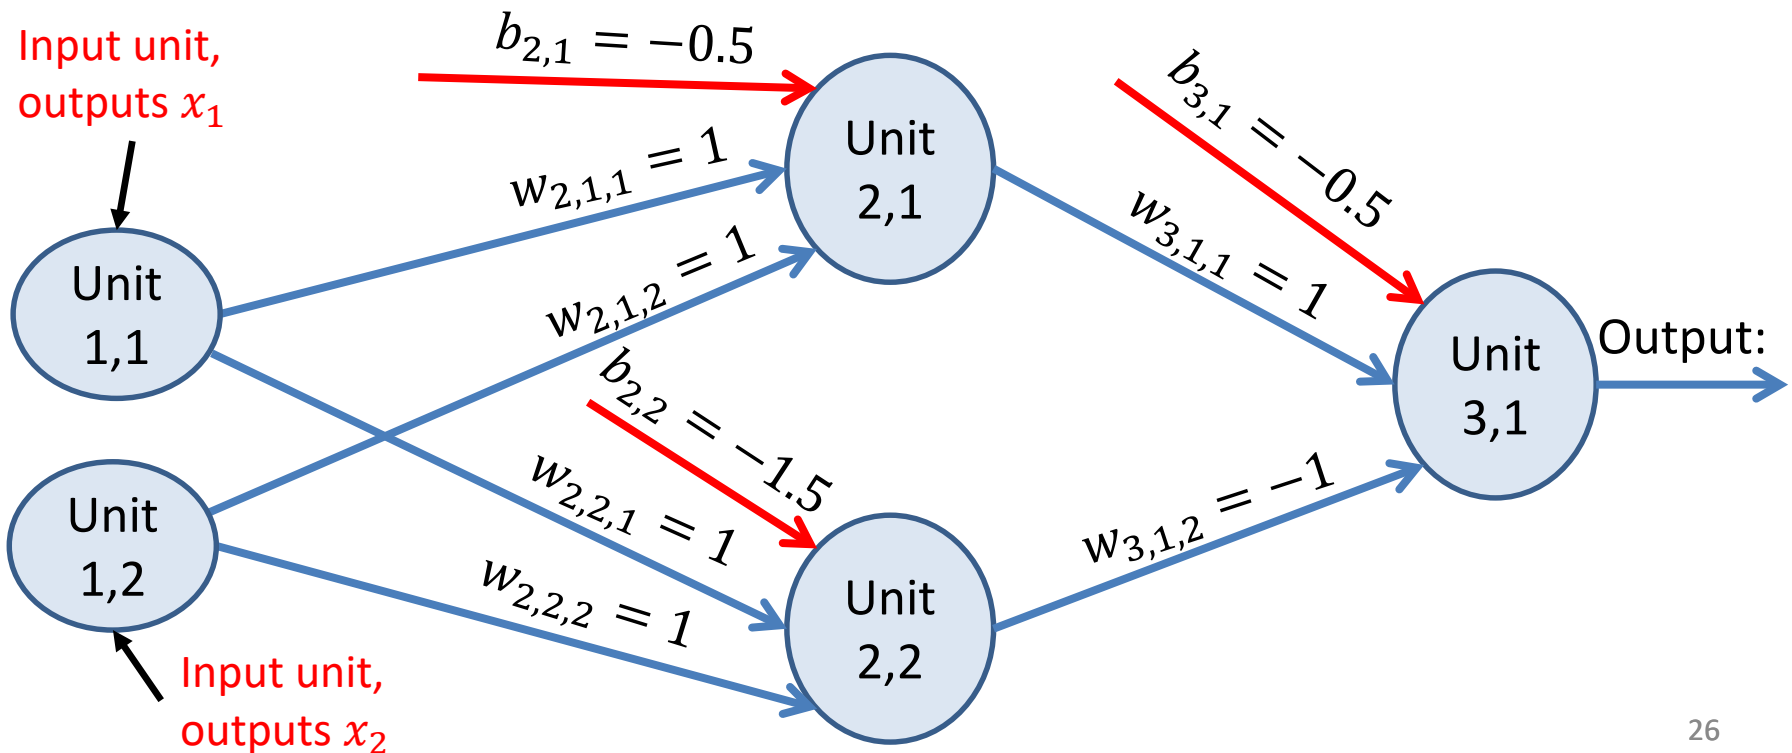
Our First Neural Network: XOR

- Such networks are called **layered** networks, more details later.
- Each unit is indexed by two numbers (layer index, unit index).
- Each bias weight b is indexed by the same two numbers as its unit.
- Each weight w is indexed by three numbers (layer, unit, weight).



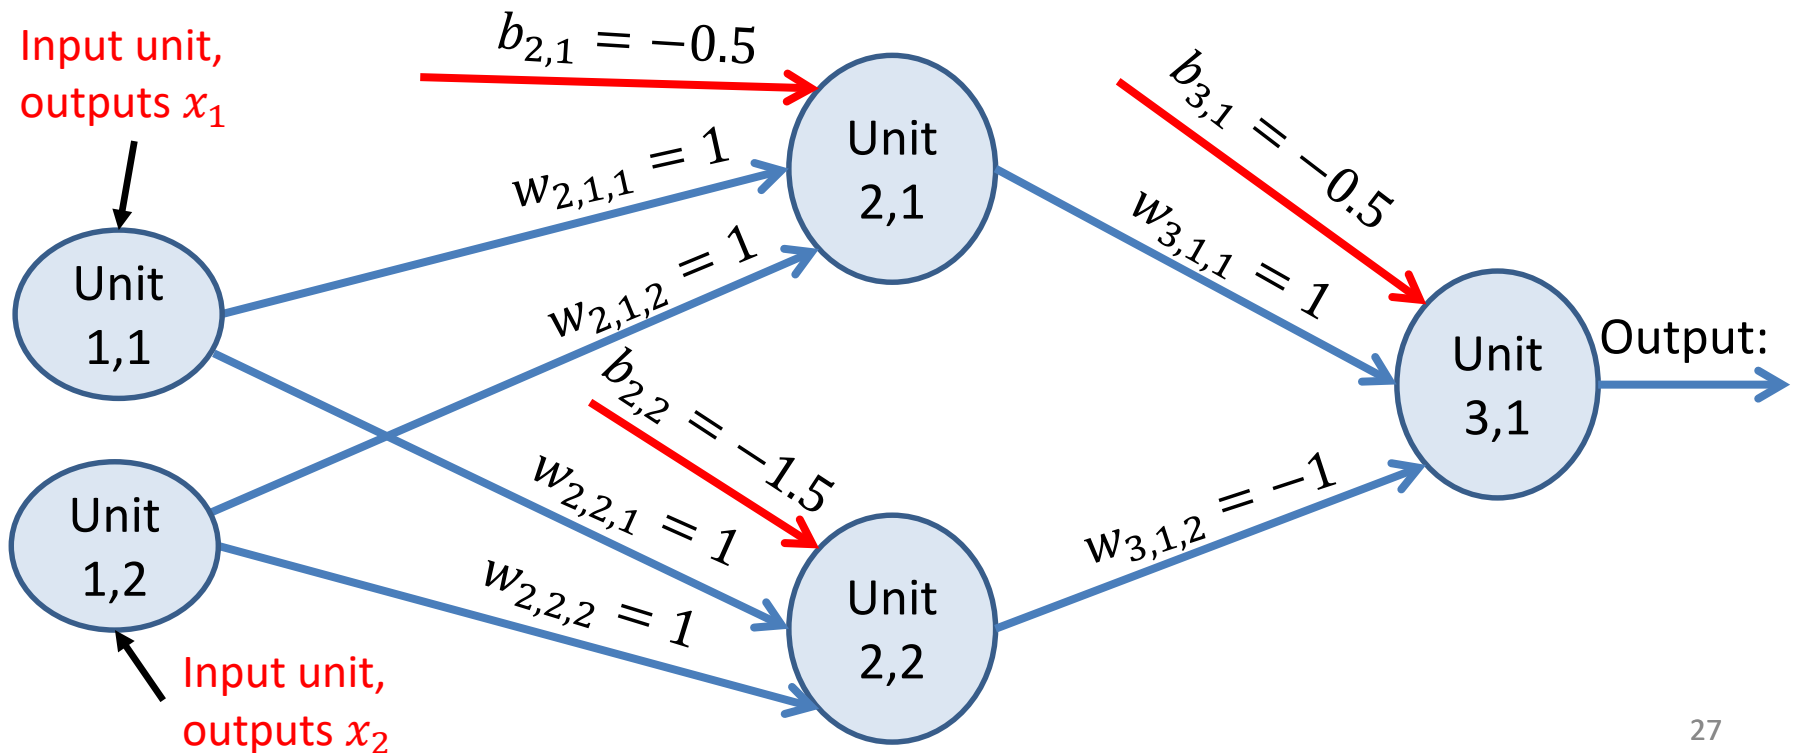
Our First Neural Network: XOR

- Note: every weight is associated with two units: it connects the output of a unit with an input of another unit.
 - Which of the two units do we use to index the weight?



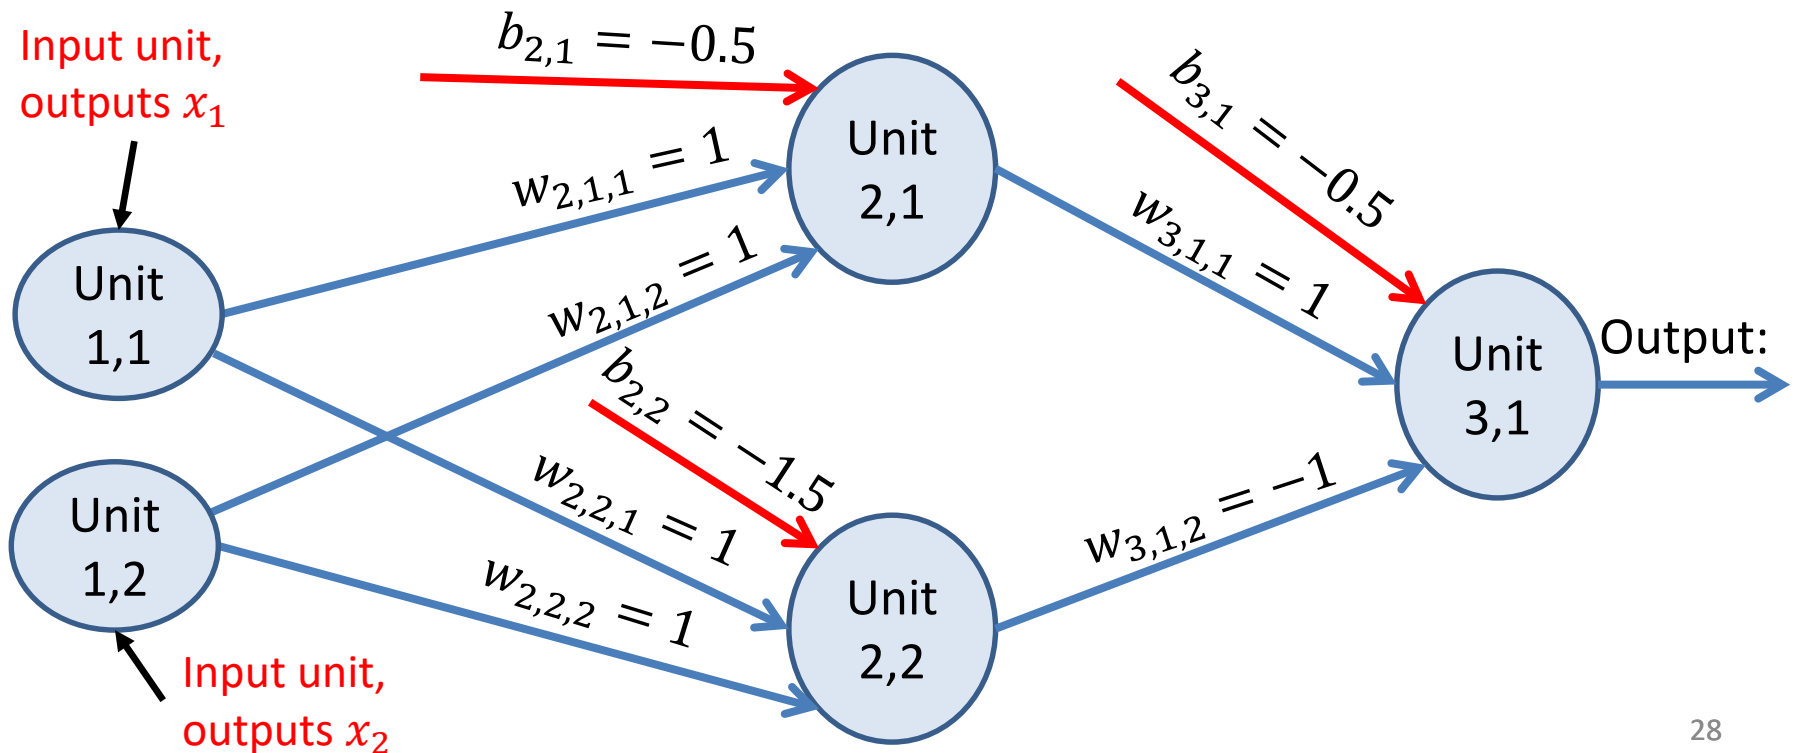
Our First Neural Network: XOR

- To index a weight w , we use the layer number and unit number of the unit for which w is an **incoming weight**.
- Weights incoming to unit l, i are indexed as l, i, j , where j ranges from 1 to the number of incoming weights for unit l, i .



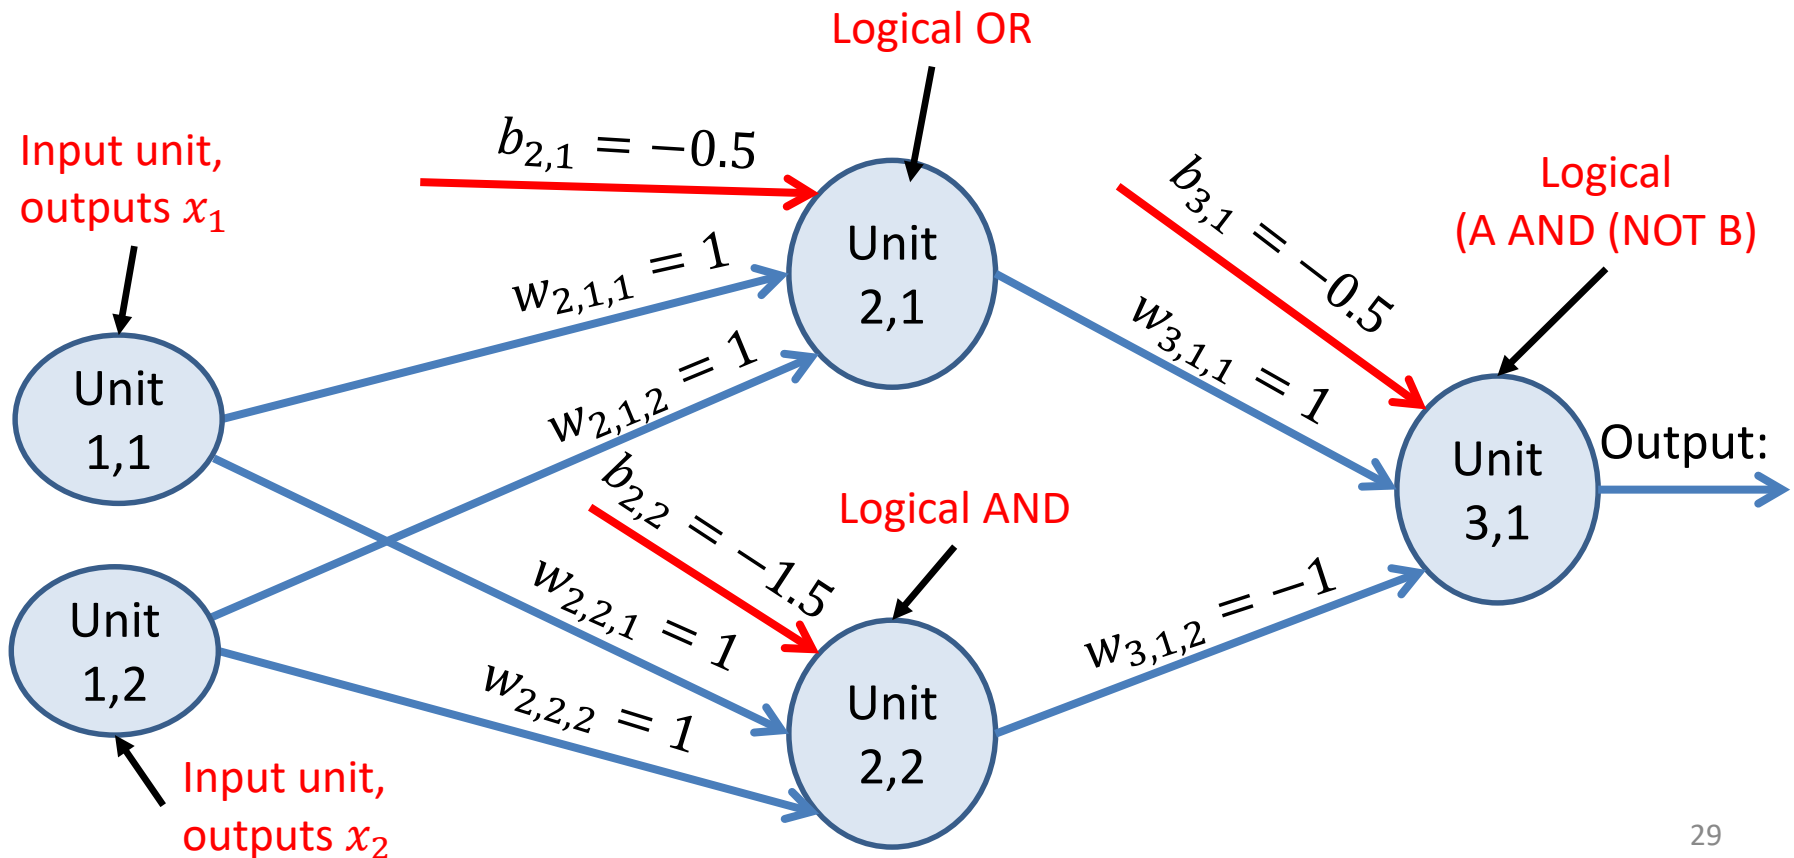
Our First Neural Network: XOR

- Weights incoming to unit l, i are indexed as l, i, j , where j ranges from 1 to the number of incoming weights for unit l, i .
- Since the input layer (which is layer 1) has no incoming weights, there are no weights indexed as $w_{1,i,j}$.



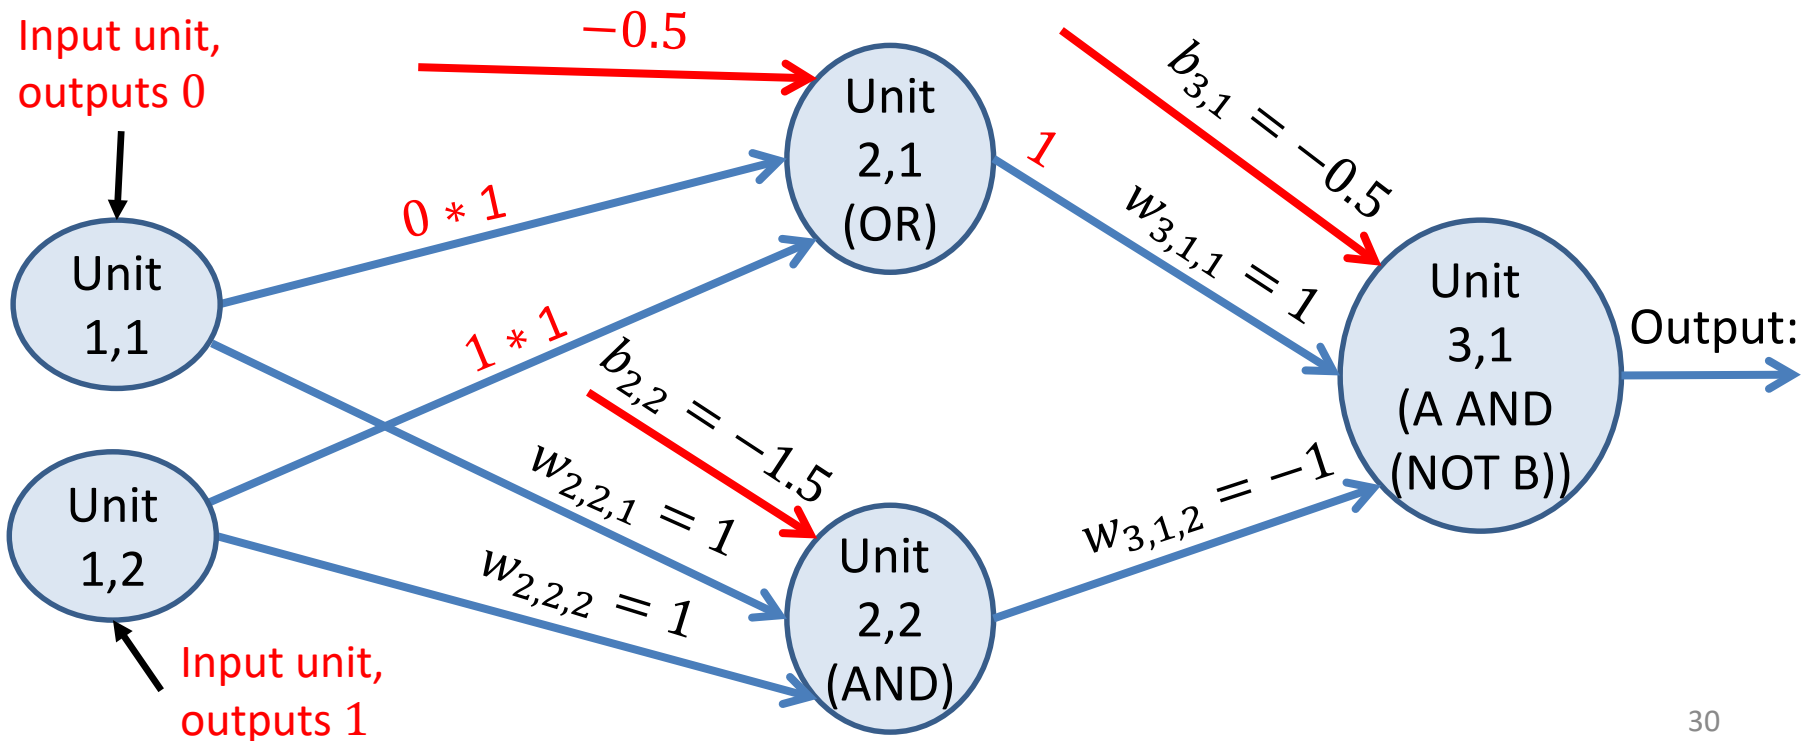
Our First Neural Network: XOR

- The XOR network shows how individual perceptrons can be combined to perform more complicated functions.



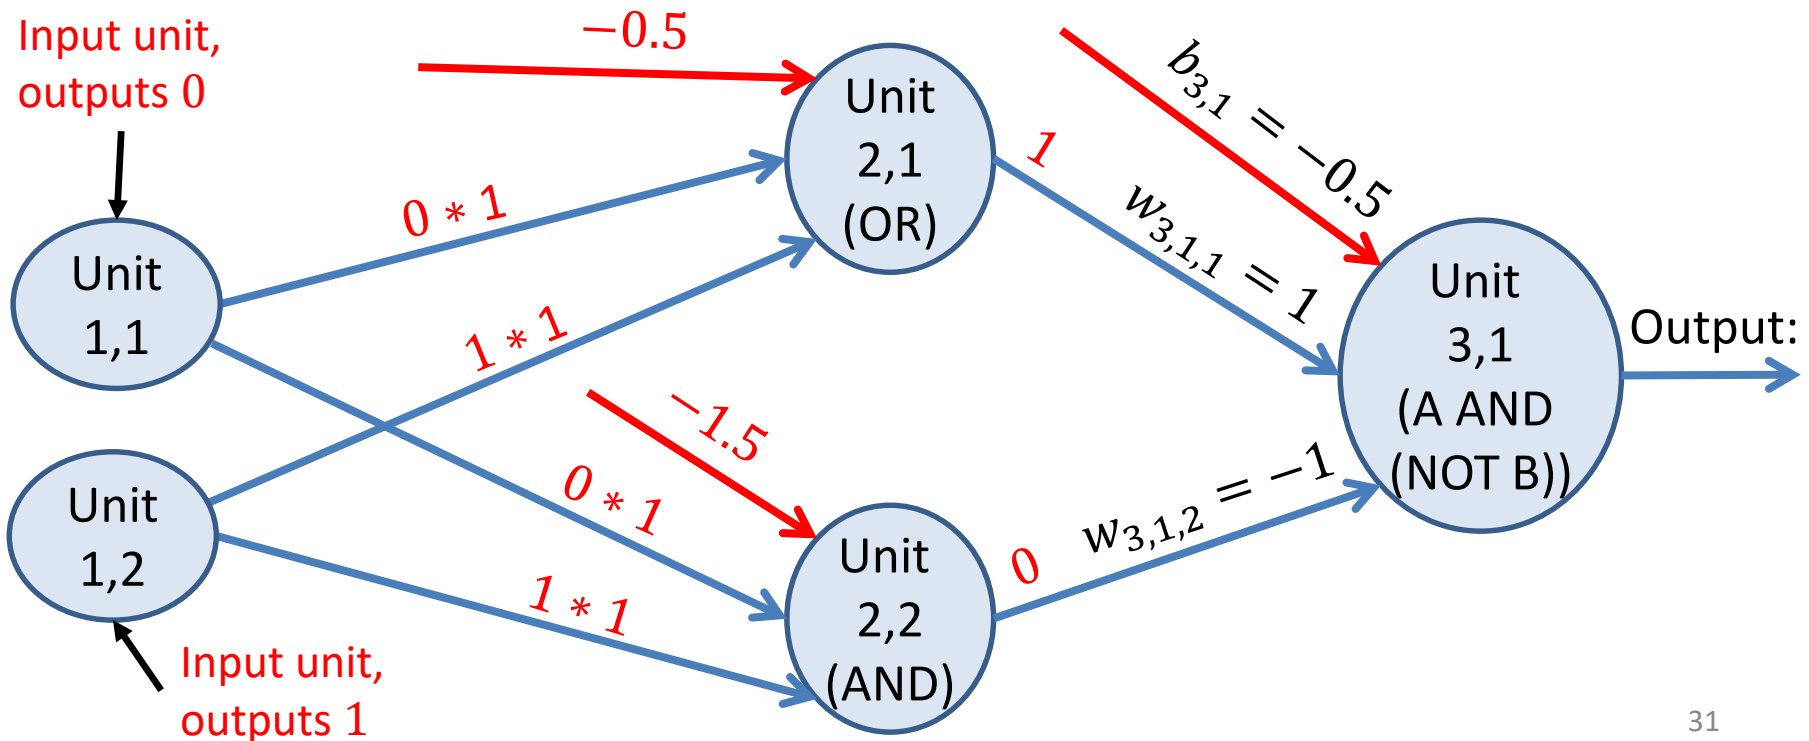
Computing the Output: An Example

- Suppose that $x_1 = 0, x_2 = 1$ (corresponding to **false XOR true**).
- For Unit 2,1, which performs a logical OR:
 - The output is $h(-0.5 + 0 * 1 + 1 * 1) = h(0.5)$.
 - Assuming that h is the step function, $h(0.5) = 1$, so Unit 2,1 outputs 1.



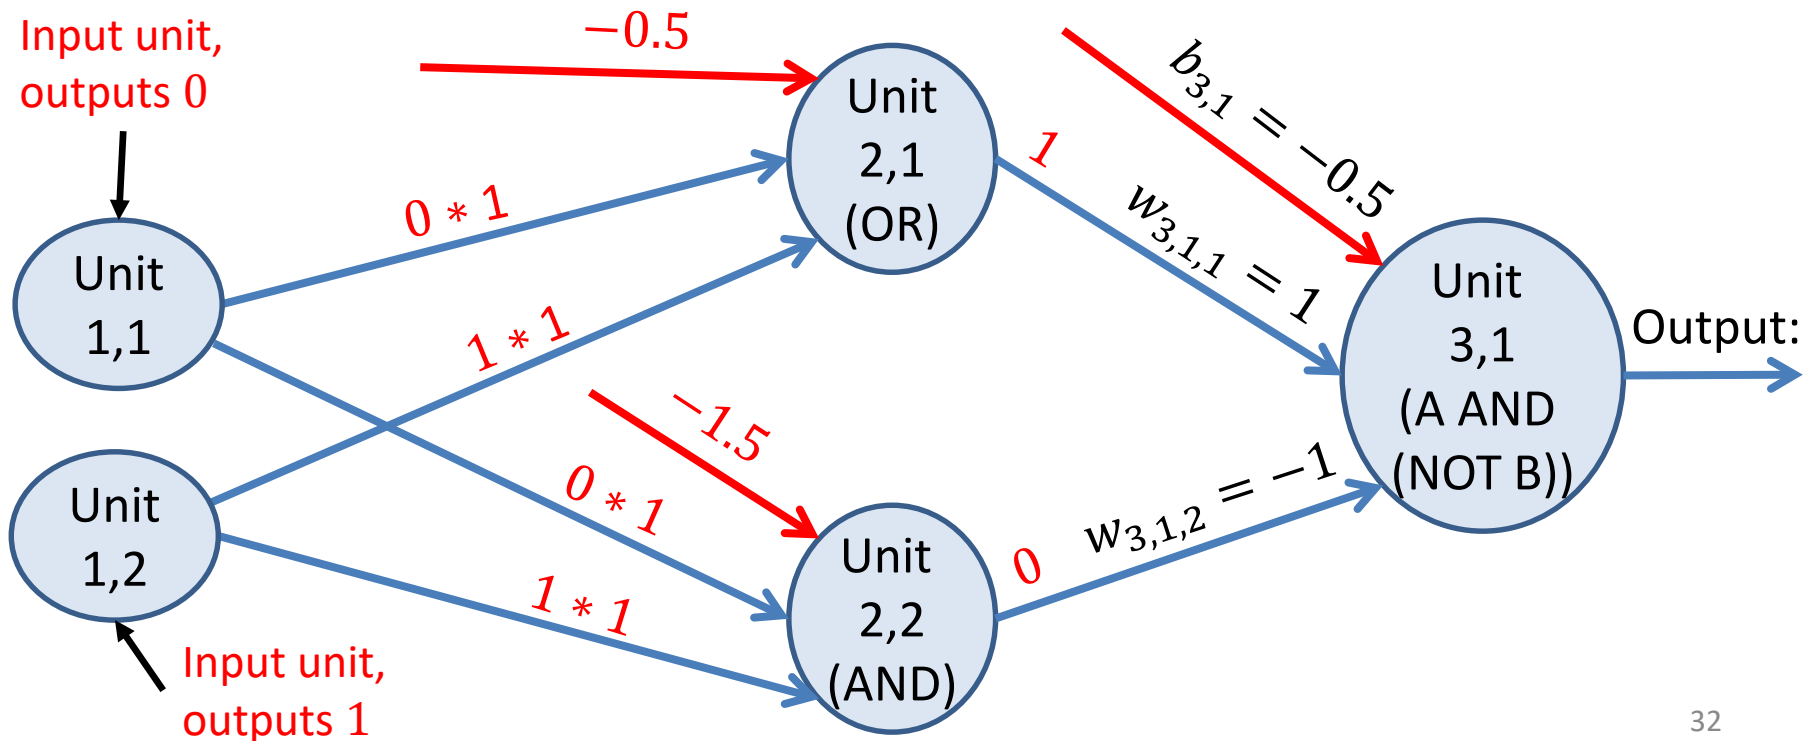
Computing the Output: An Example

- Suppose that $x_1 = 0, x_2 = 1$ (corresponding to **false XOR true**).
- For Unit 2,2, which performs a logical AND:
 - The output is $h(-1.5 + 0 * 1 + 1 * 1) = h(-0.5)$.
 - Since h is the step function, $h(-0.5) = 0$, so Unit 2,2 outputs 0.



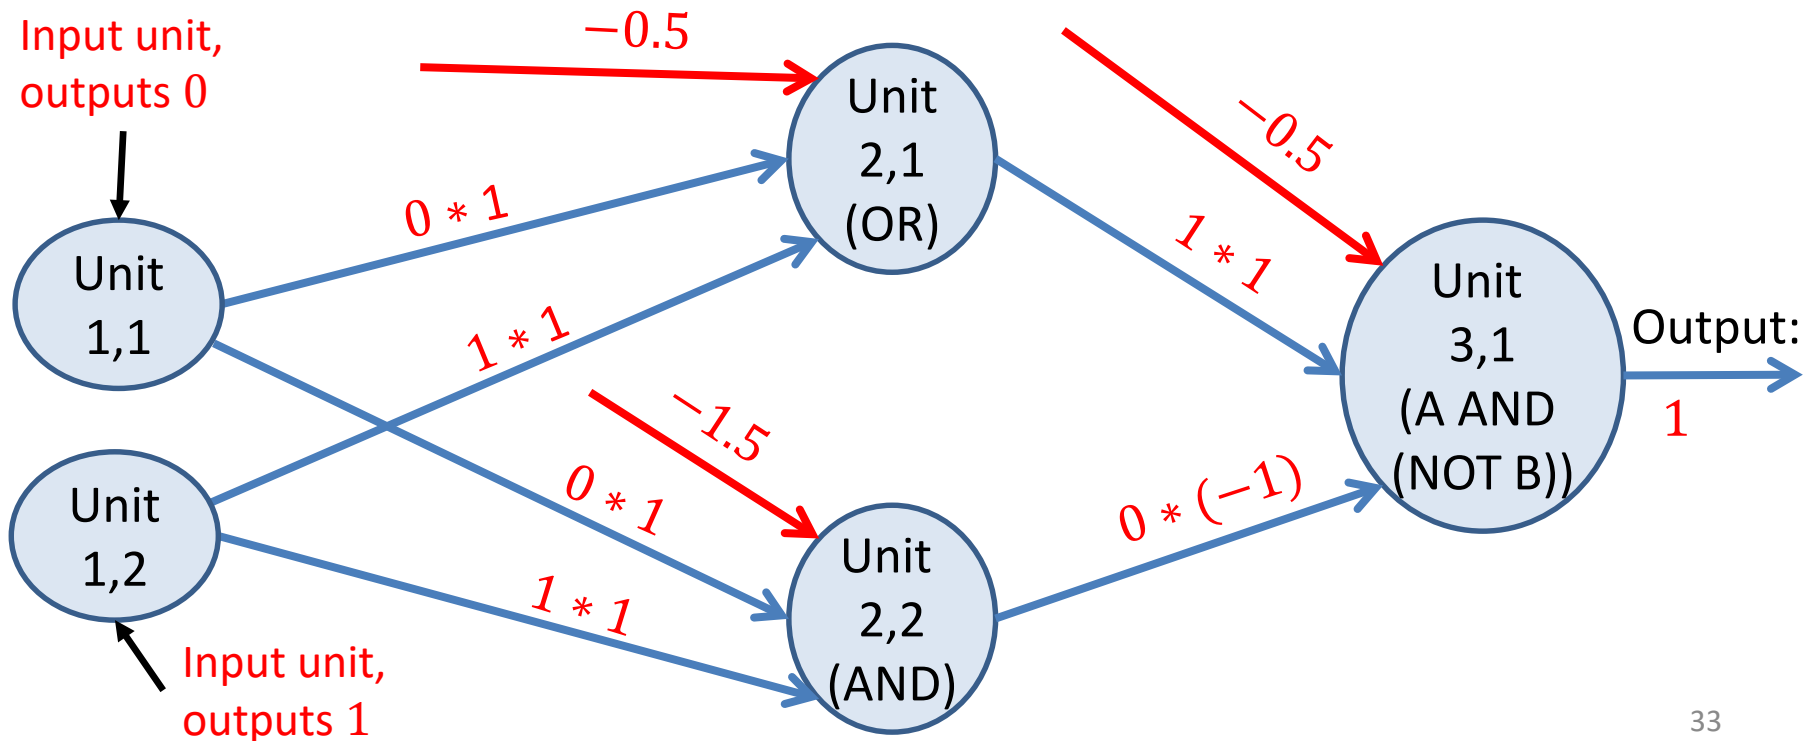
Computing the Output: An Example

- Suppose that $x_1 = 0, x_2 = 1$ (corresponding to **false XOR true**).
- Unit 3,1 is the **output unit**, computing the A AND (NOT B) function:
 - One input is the output of the OR unit, which is 1.
 - The other input is the output of the AND unit, which is 0.



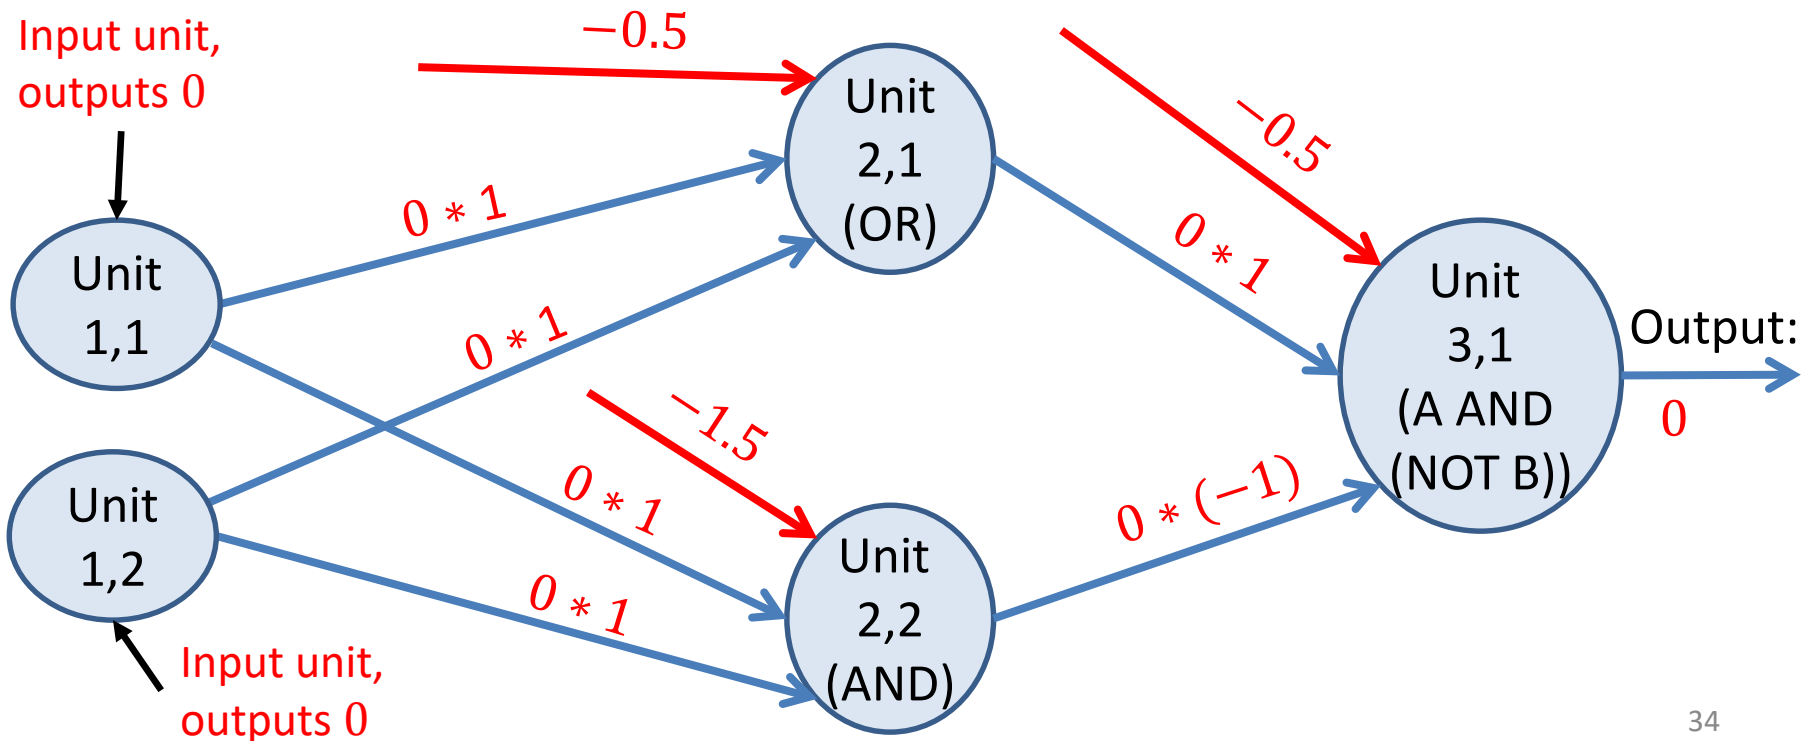
Computing the Output: An Example

- Suppose that $x_1 = 0, x_2 = 1$ (corresponding to **false XOR true**).
- For the output unit (computing the A AND (NOT B) function):
 - The output is $h(-0.5 + 1 * 1 + 0 * (-1)) = h(0.5)$.
 - Since h is the step function, $h(0.5) = 0$, so Unit 3,1 outputs 1.



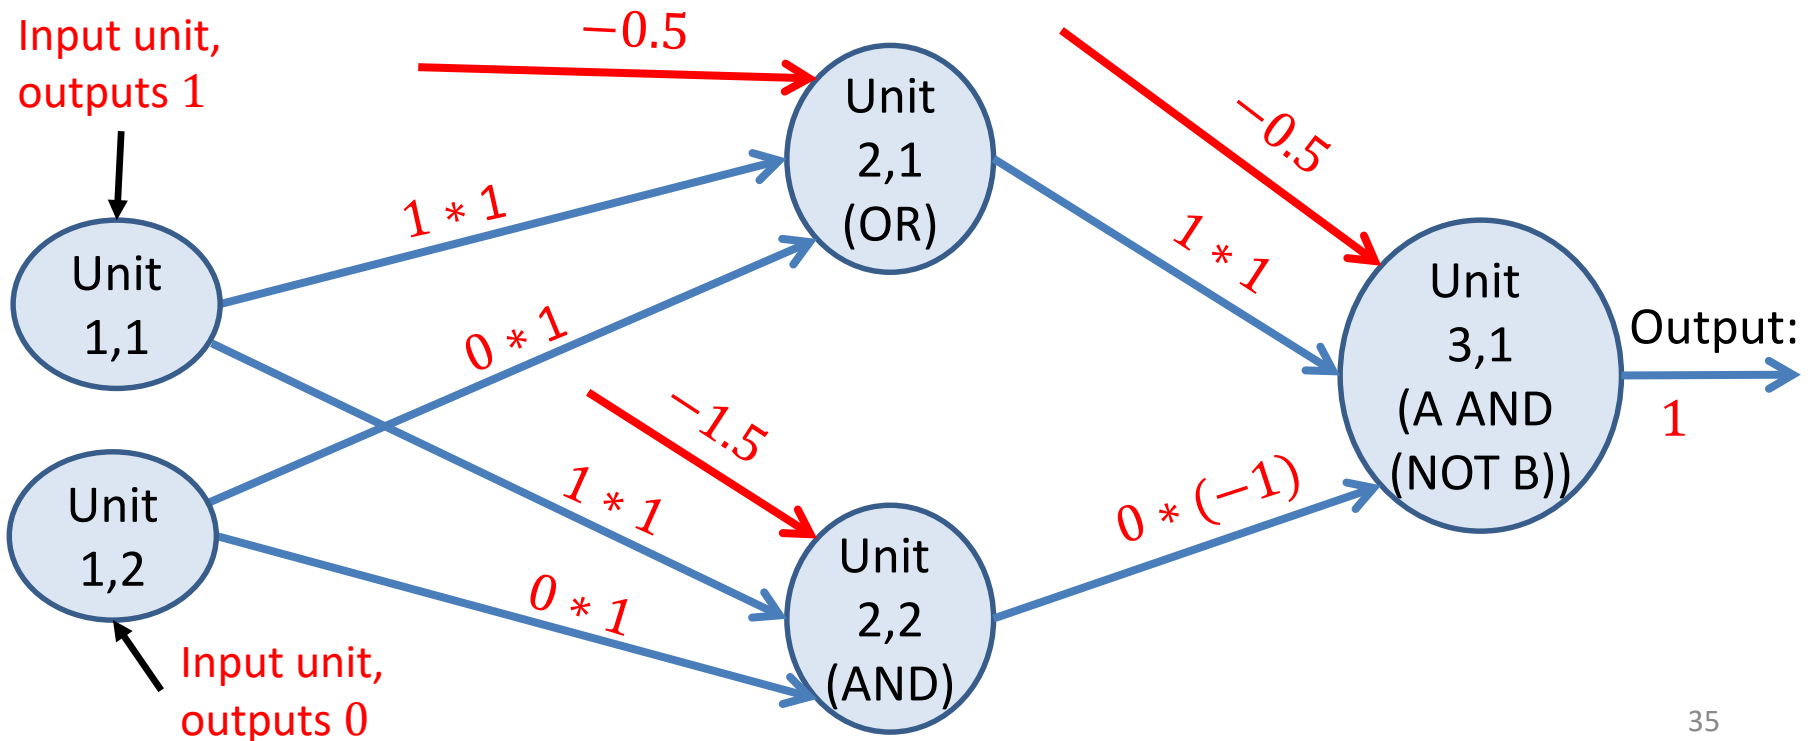
Verifying the XOR Network

- We can follow the same process to compute the output of this network for the other three cases.
 - Here we consider the case where $x_1 = 0, x_2 = 0$ (corresponding to **false XOR false**).
 - The output is 0, as it should be.



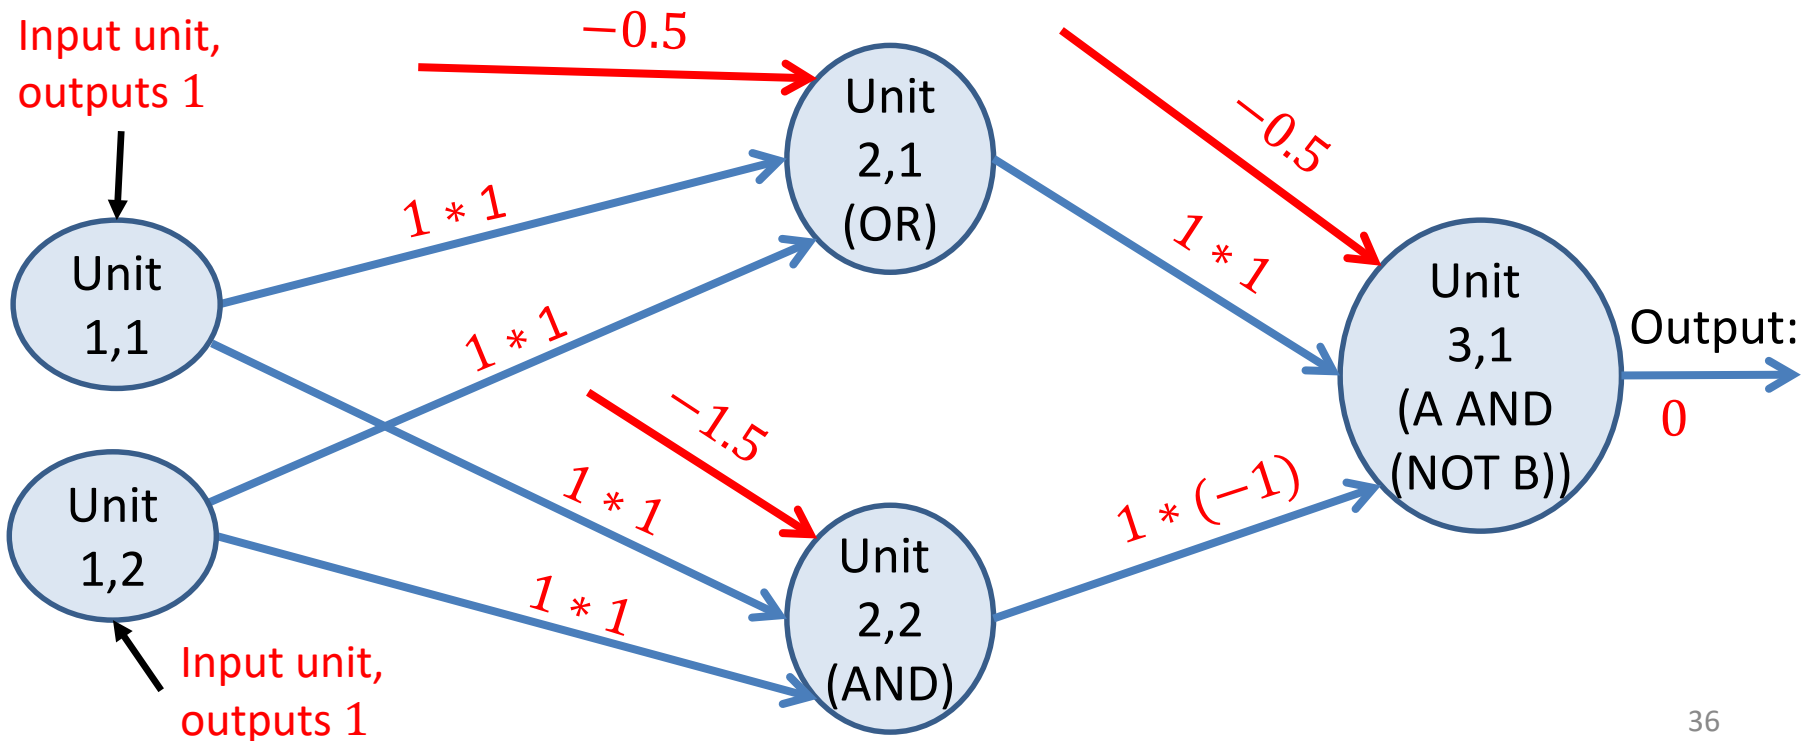
Verifying the XOR Network

- We can follow the same process to compute the output of this network for the other three cases.
 - Here we consider the case where $x_1 = 1, x_2 = 0$ (corresponding to **true XOR false**).
 - The output is 1, as it should be.



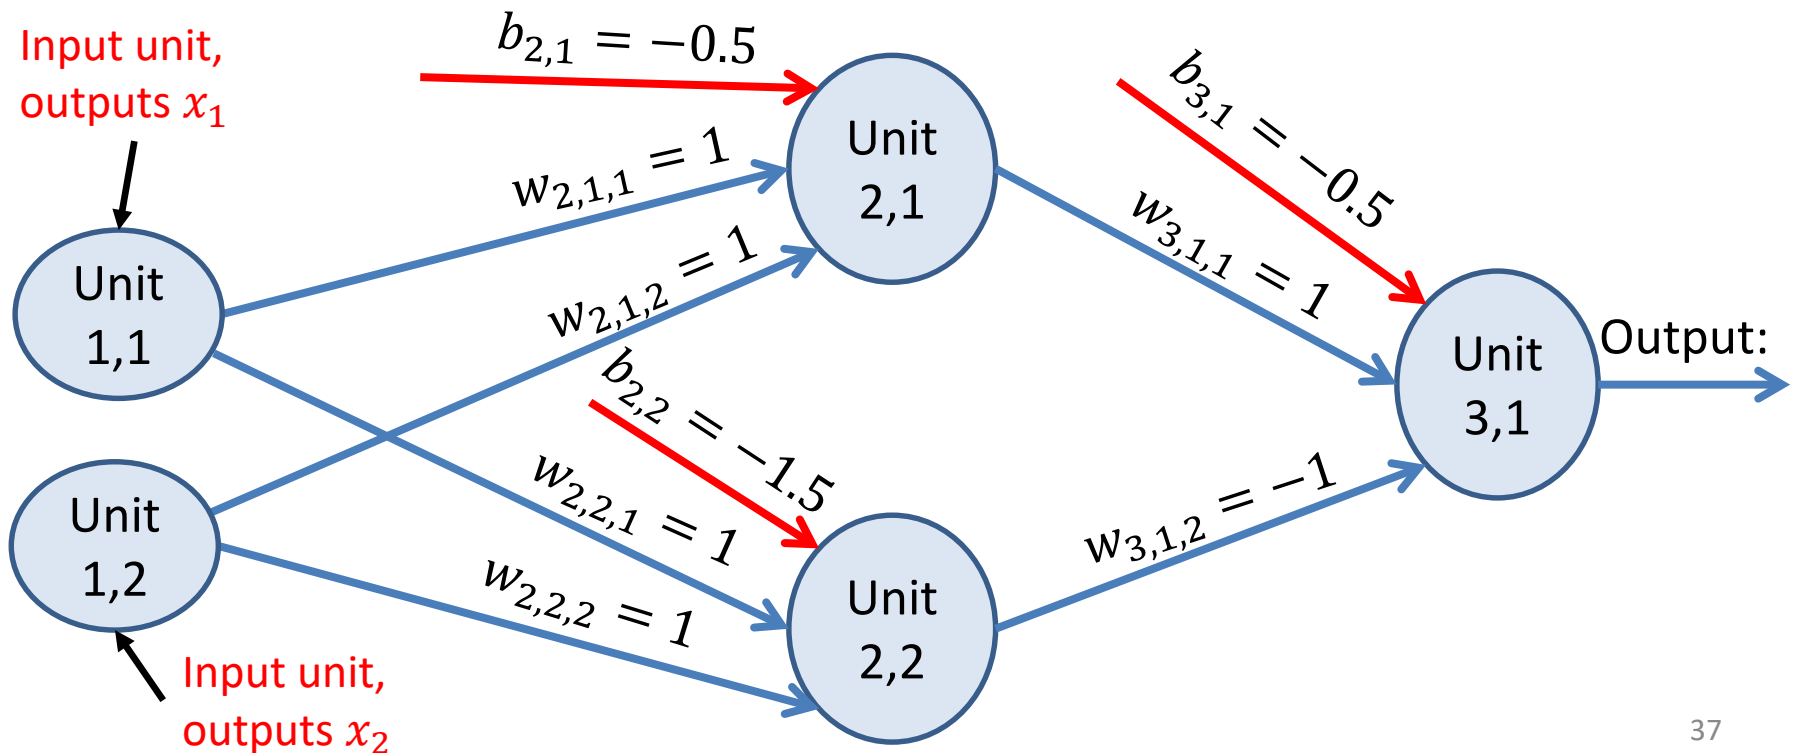
Verifying the XOR Network

- We can follow the same process to compute the output of this network for the other three cases.
 - Here we consider the case where $x_1 = 1, x_2 = 1$ (corresponding to **true XOR true**).
 - The output is 0, as it should be.



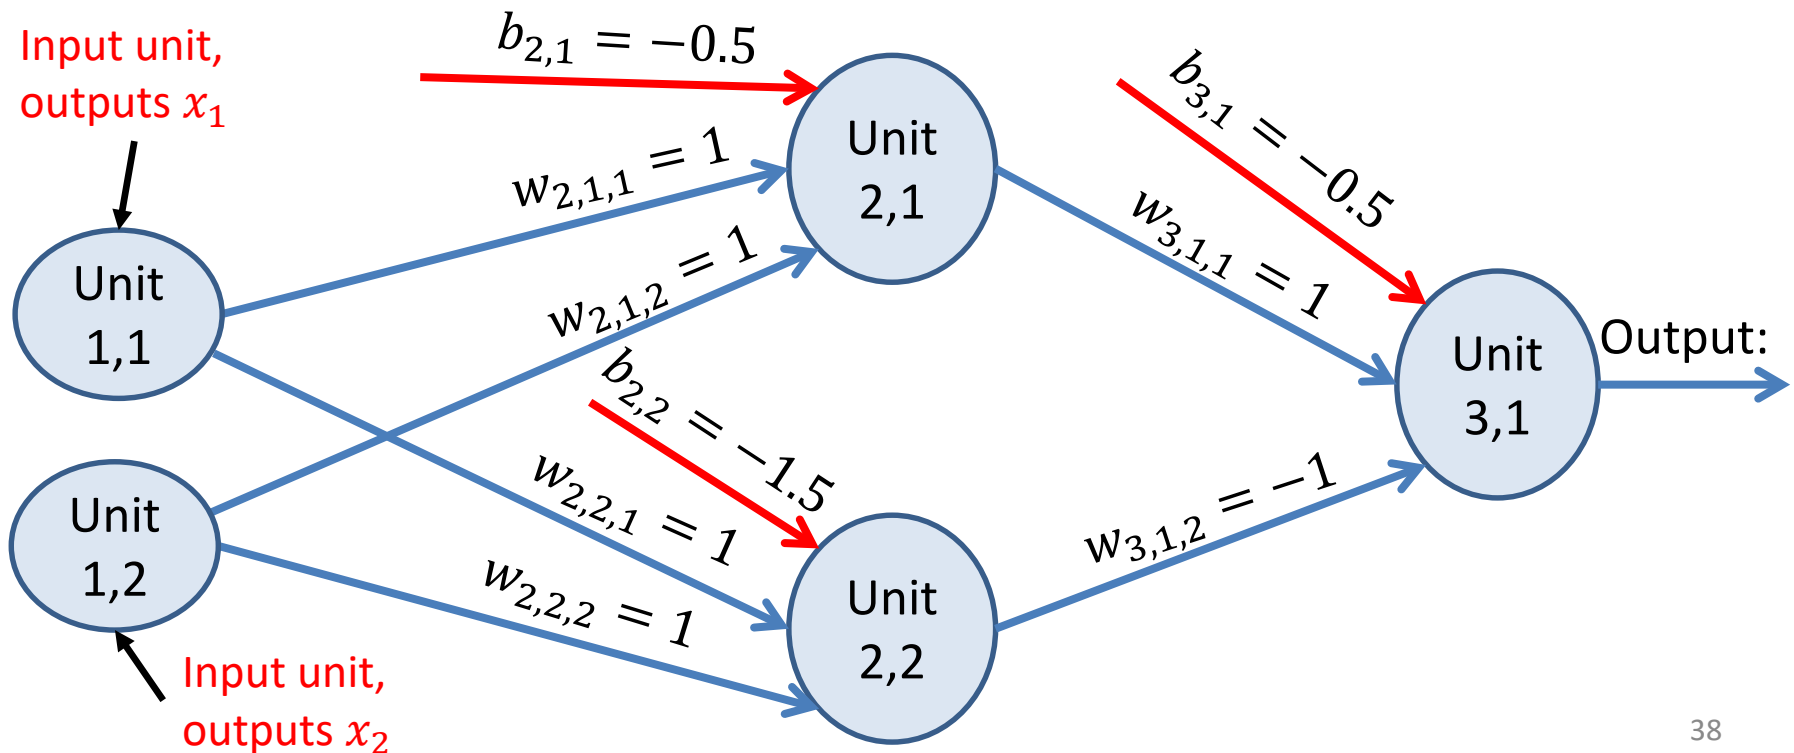
Neural Networks

- Our XOR neural network consists of five units:
 - Two input units, that just represent the two inputs to the network.
 - Three perceptrons.



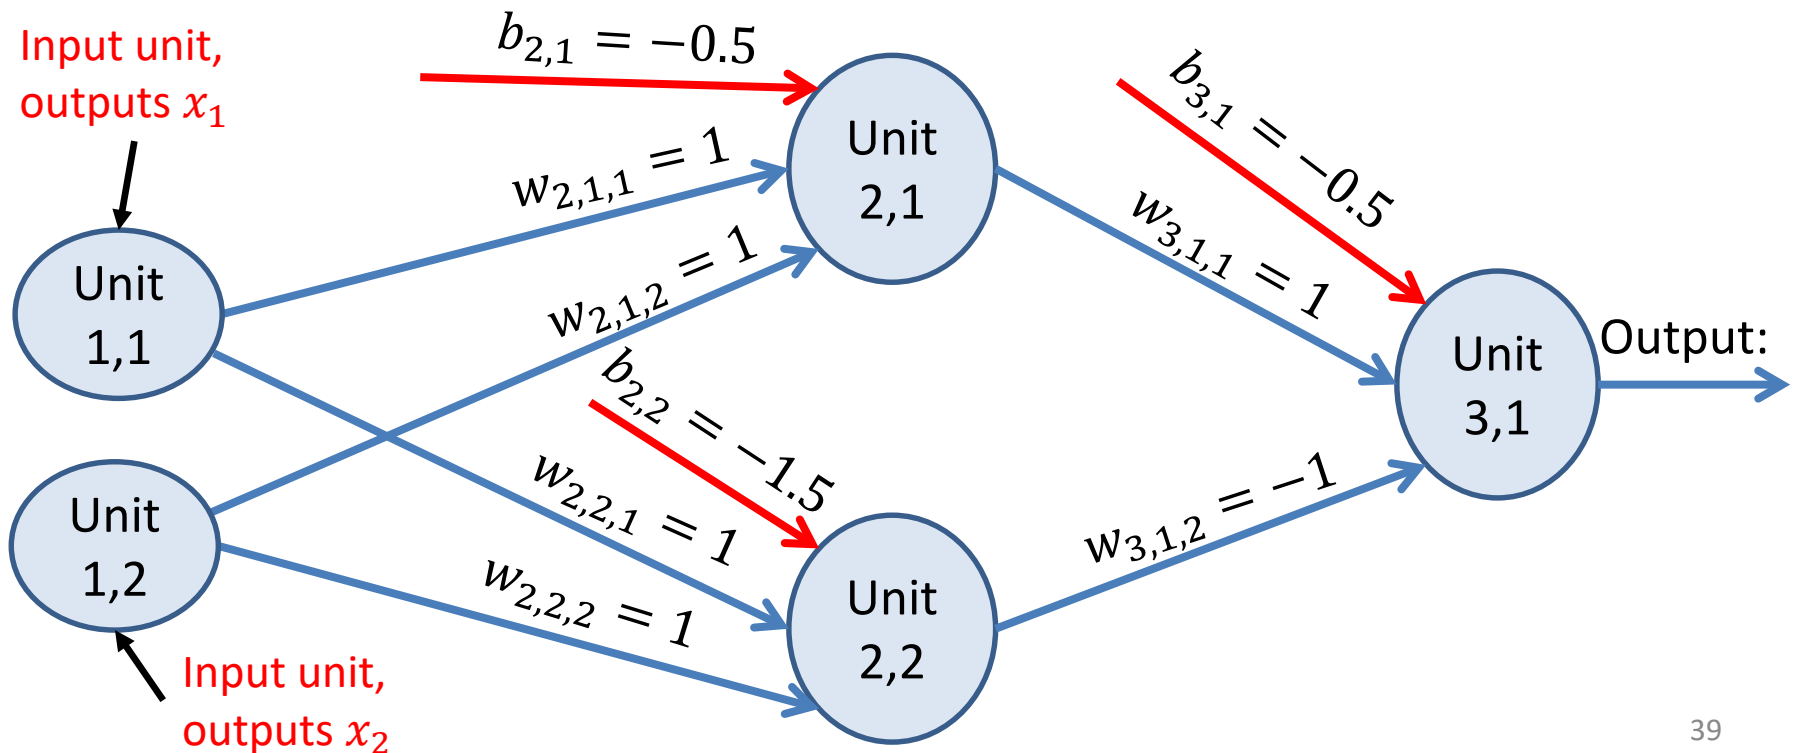
Neural Network Layers

- Oftentimes, as in the XOR example, neural networks are organized into layers.
- The **input layer** is the initial layer of input units (units 1,1 and 1,2 in our example).
- The output layer is at the end (unit 3,1 in our example).
- Zero, one or more hidden layers can be between the input and output layers.



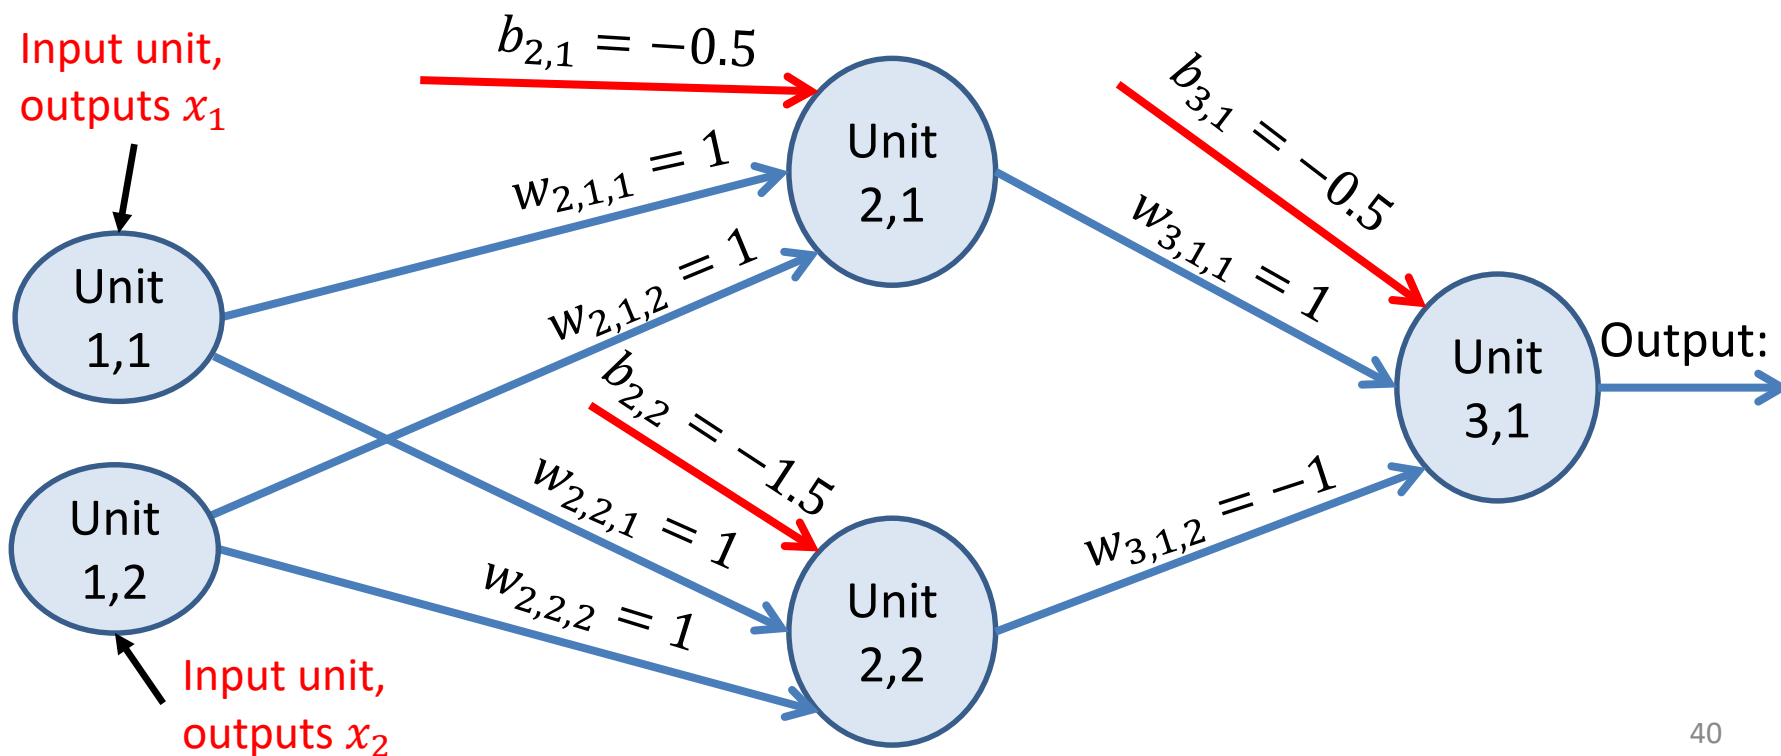
Neural Network Layers

- There is only one hidden layer in our example, containing units 2,1 and 2,2.
- Each hidden layer's inputs are outputs from the previous layer.
- Each hidden layer's outputs are inputs to the next layer.
- The first hidden layer's inputs come from the input layer.
- The last hidden layer's outputs are inputs to the output layer.



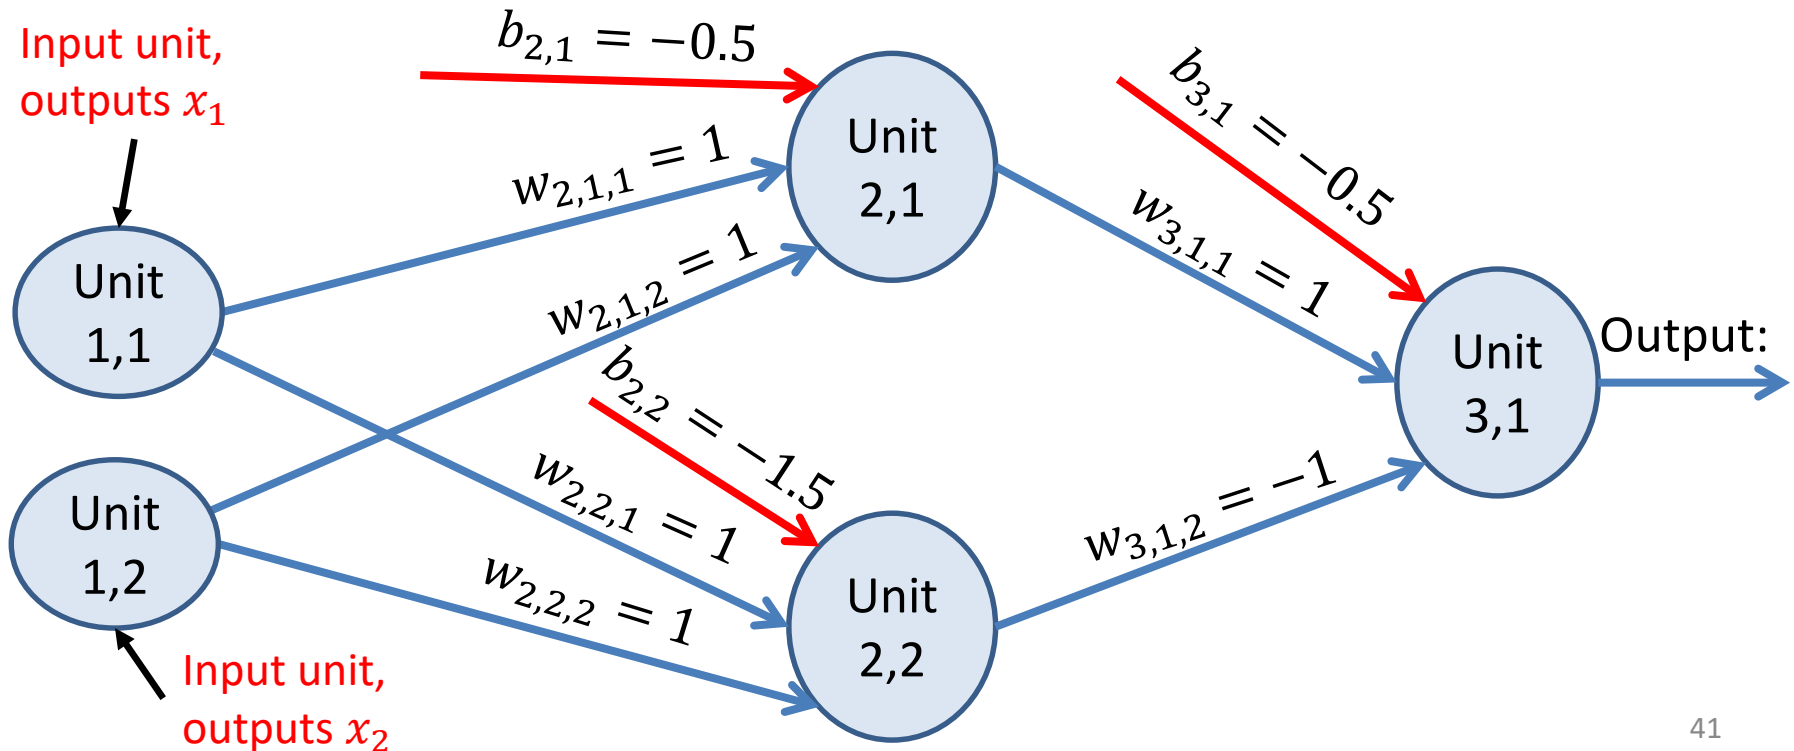
Feedforward Networks

- Feedforward networks are networks where there are no directed loops.
- If there are no loops, the output of a unit cannot (directly or indirectly) influence its input.
- While there are varieties of neural networks that are not feedforward or layered, our main focus will be layered feedforward networks.



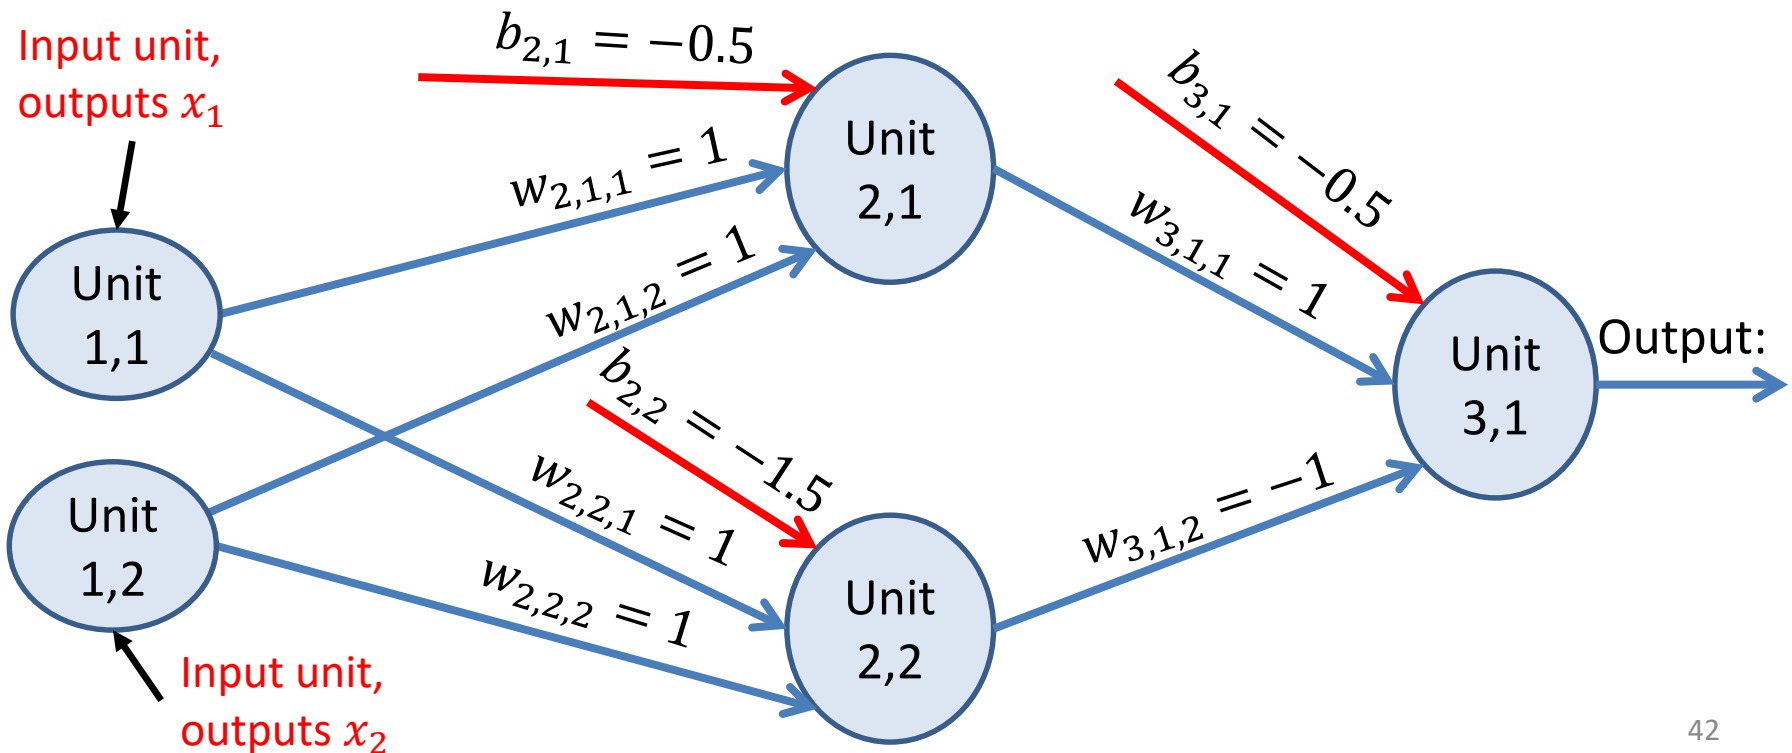
Computing the Output

- Notation: L is the number of layers. Layer 1 is the input layer, layer L is the output layer.
- The outputs of the units of layer 1 are simply the inputs to the network.
- For (layer $l = 2; l \leq L; l = l + 1$):
 - Compute the outputs of layer l , given the outputs of layer $l - 1$.



Computing the Output

- To compute the outputs of layer l (where $l > 1$), we simply need to compute the output of each perceptron belonging to layer l .
 - For each such perceptron, its inputs are coming from outputs of units at layer $l - 1$, which we have already computed.
 - Remember, we compute layer outputs in increasing order of l .



What Neural Networks Can Compute

- An individual perceptron is a linear classifier.
 - The weights of the perceptron define a linear boundary between two classes.
- Layered feedforward neural networks with one hidden layer can compute any continuous function.
- Layered feedforward neural networks with two hidden layers can compute any mathematical function.
- This has been known for decades, and is one reason scientists have been optimistic about the potential of neural networks to model intelligent systems.
- Another reason is the analogy between neural networks and biological brains, which have been a standard of intelligence we are still trying to achieve.
- There is only one catch: How do we find the right weights?

Finding the Right Weights

- The goal of training a neural network is to figure out good values for the weights of the units in the network.