

Predictive Analytics using Neural Networks: Applications, Pitfalls and Beyond

Bratislava, MLMU, 24th January, 2018

Rudradeb Mitra

<https://www.linkedin.com/in/mitrar/>

Today's talk (~50-60 mins)

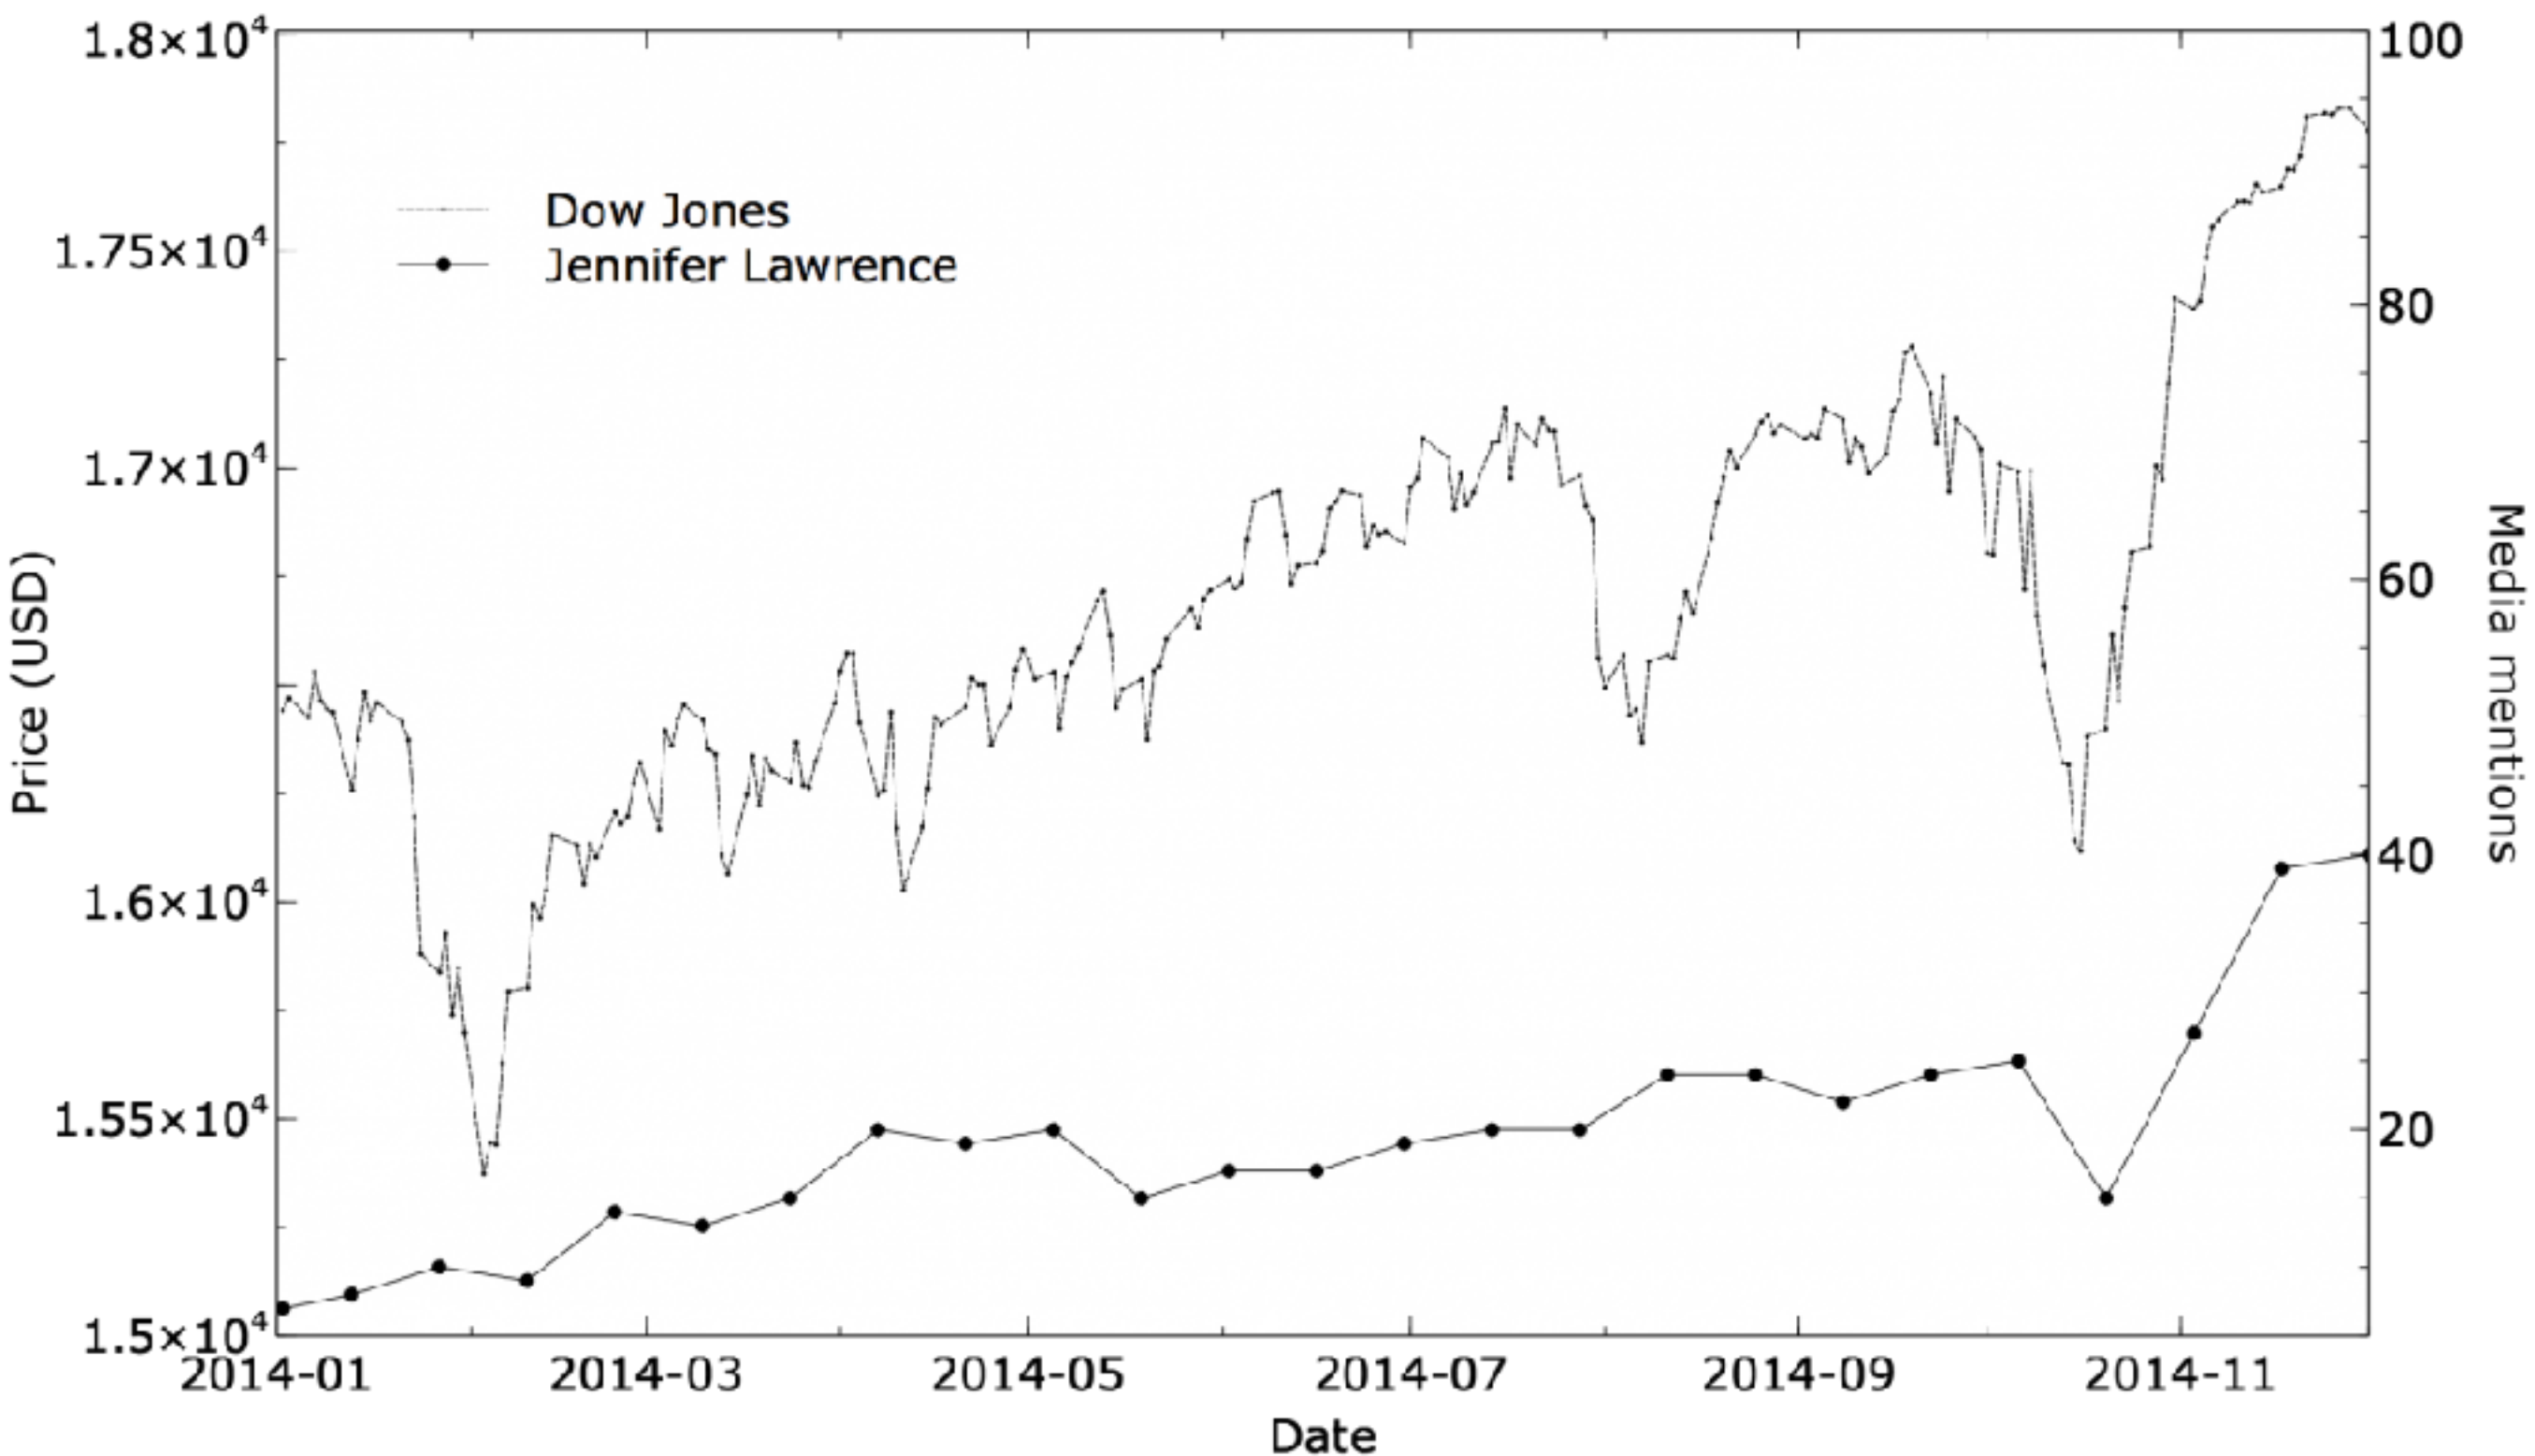
- What is Time-series data and Predictive Analytics?
- One case study - How Neural networks can be used for predicting next purchase
 - word2vec
 - Long Short-Term Memory (LSTM)
 - Neural networks using Reinforcement Learning
- Other applications and pitfalls

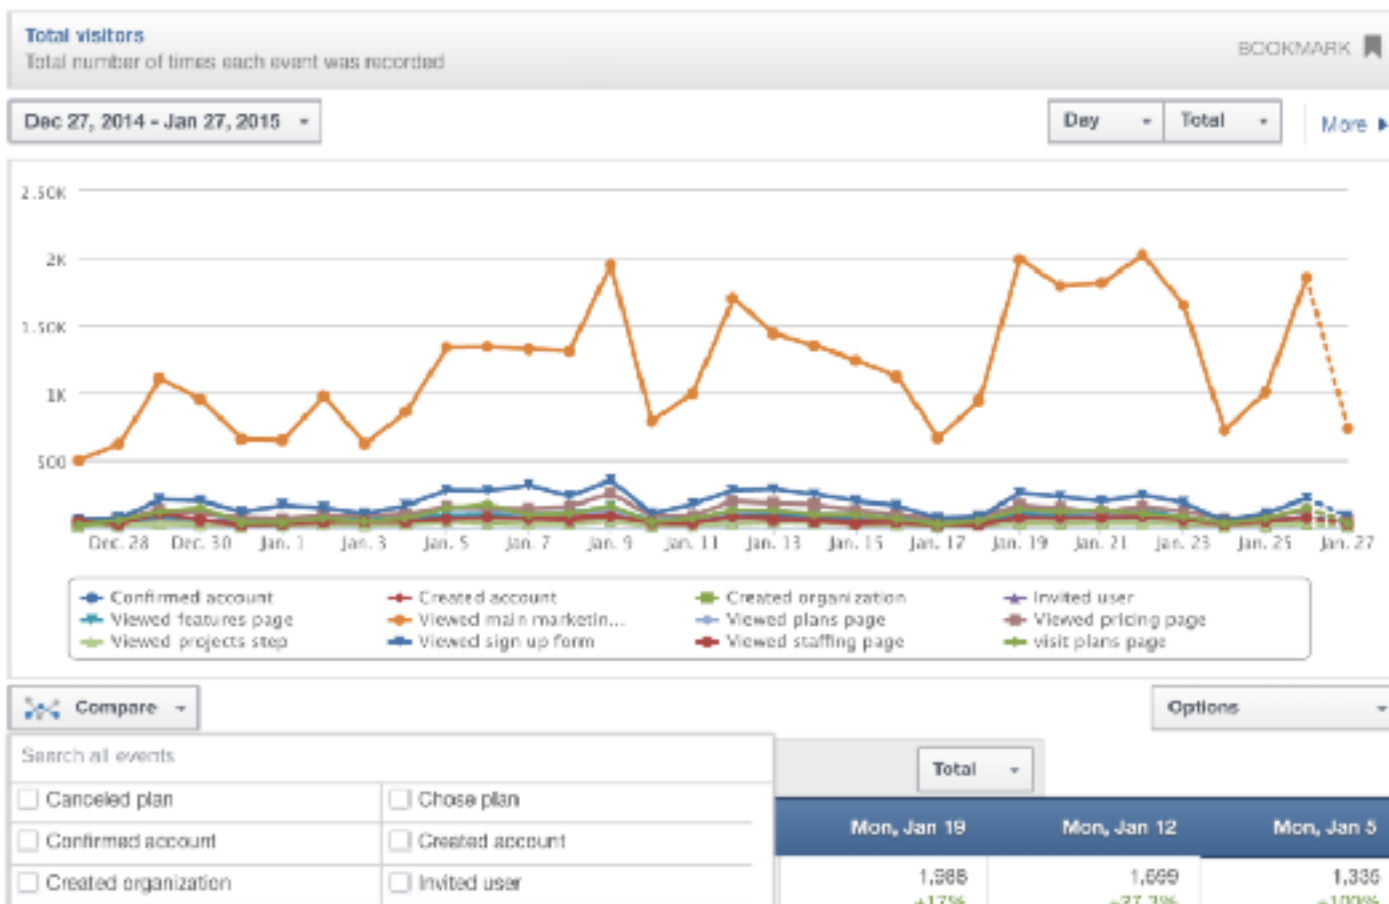
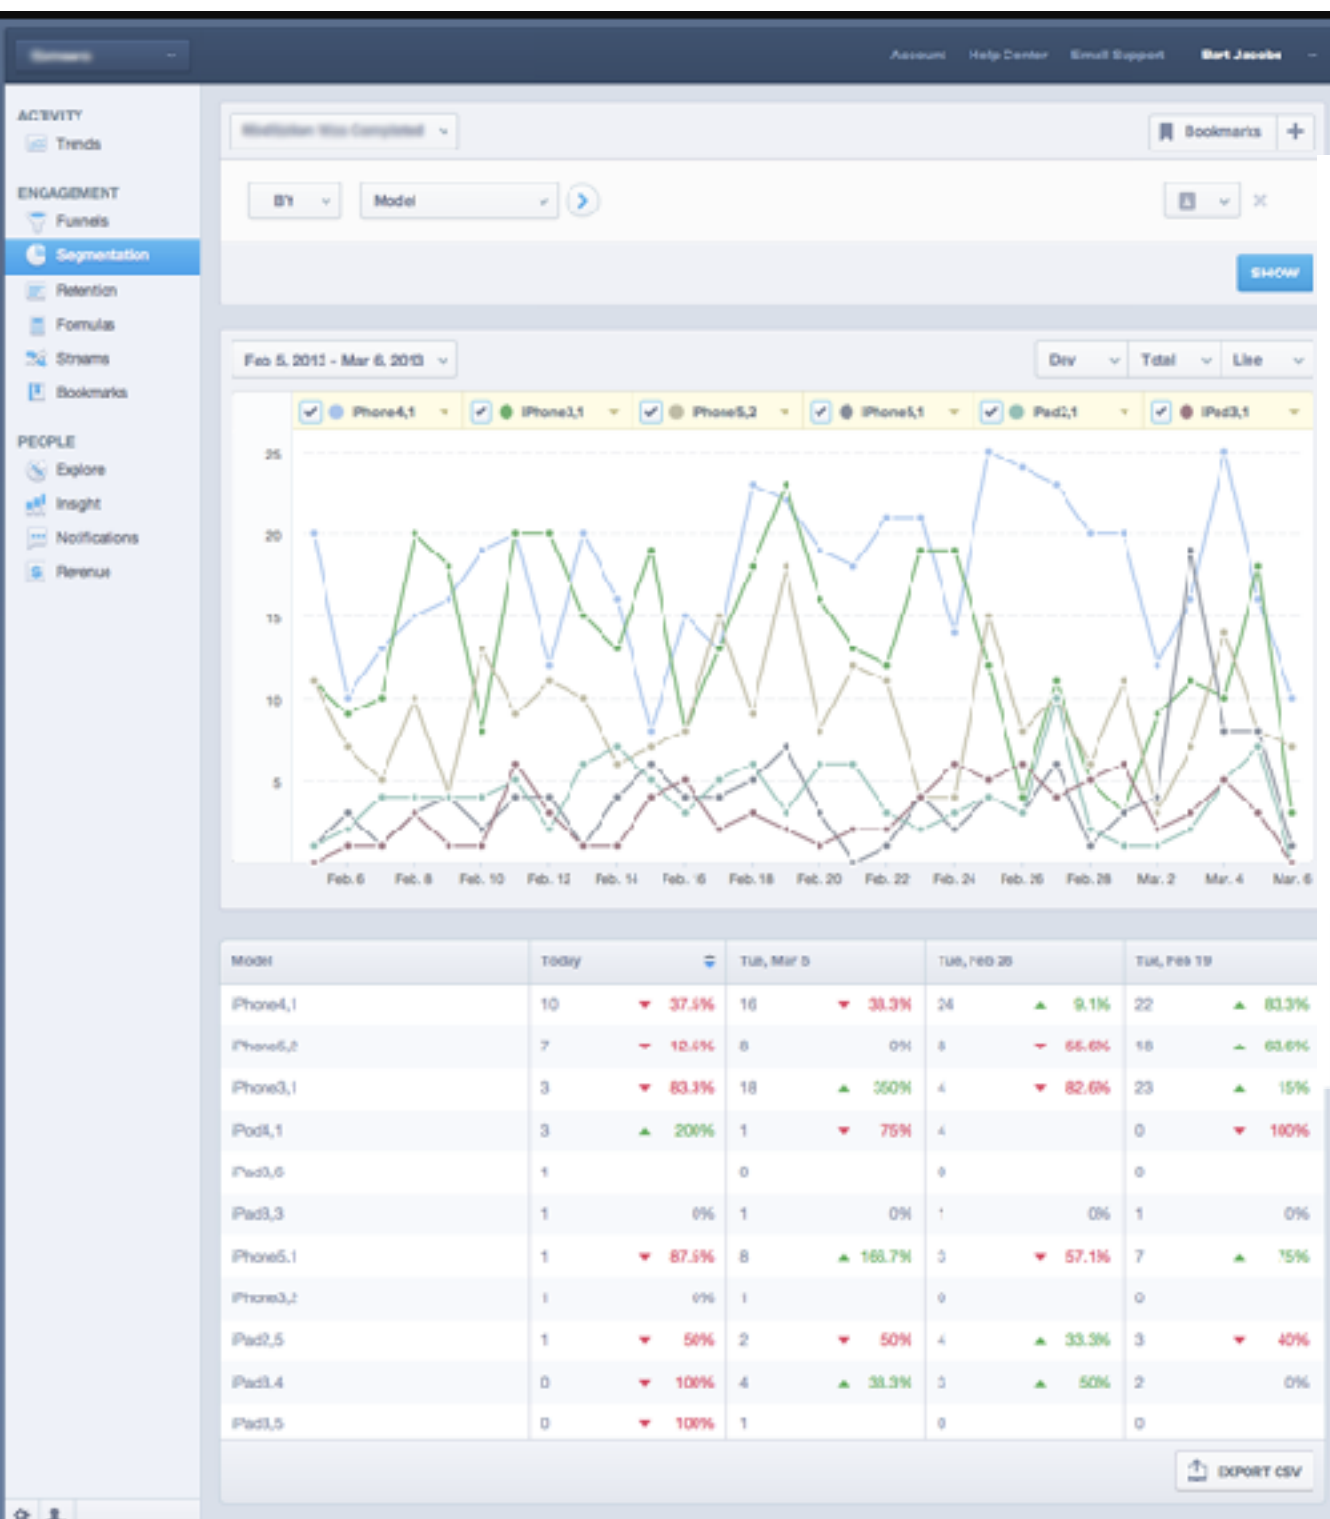
Why?

- Bridge gap between technical AI and practical AI
- You can build your own applications using Neural networks.

Time series data

Time series data





What is missing?

Predictive Analytics

Case study

Predicting next purchase



instacart

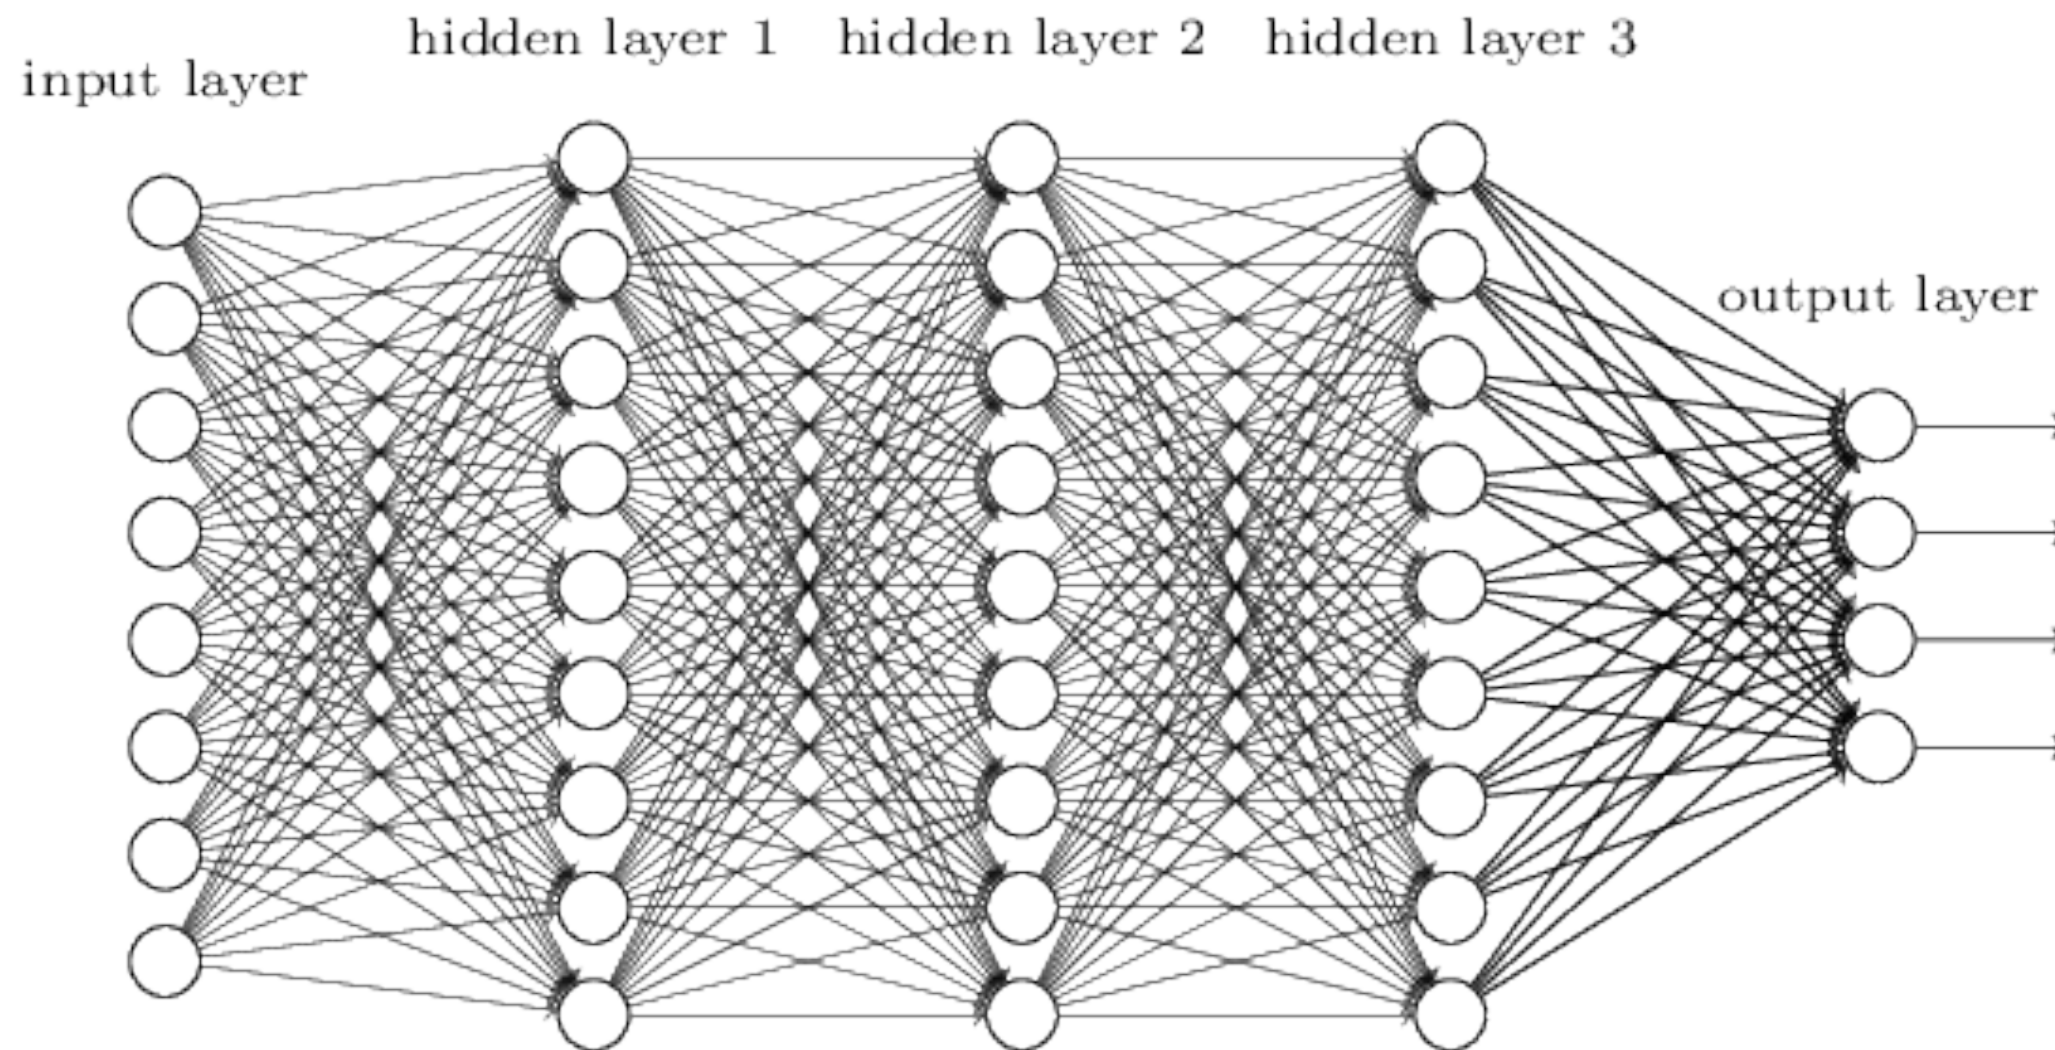
Groceries Delivered From Local Stores

Neural Networks

Neural Network

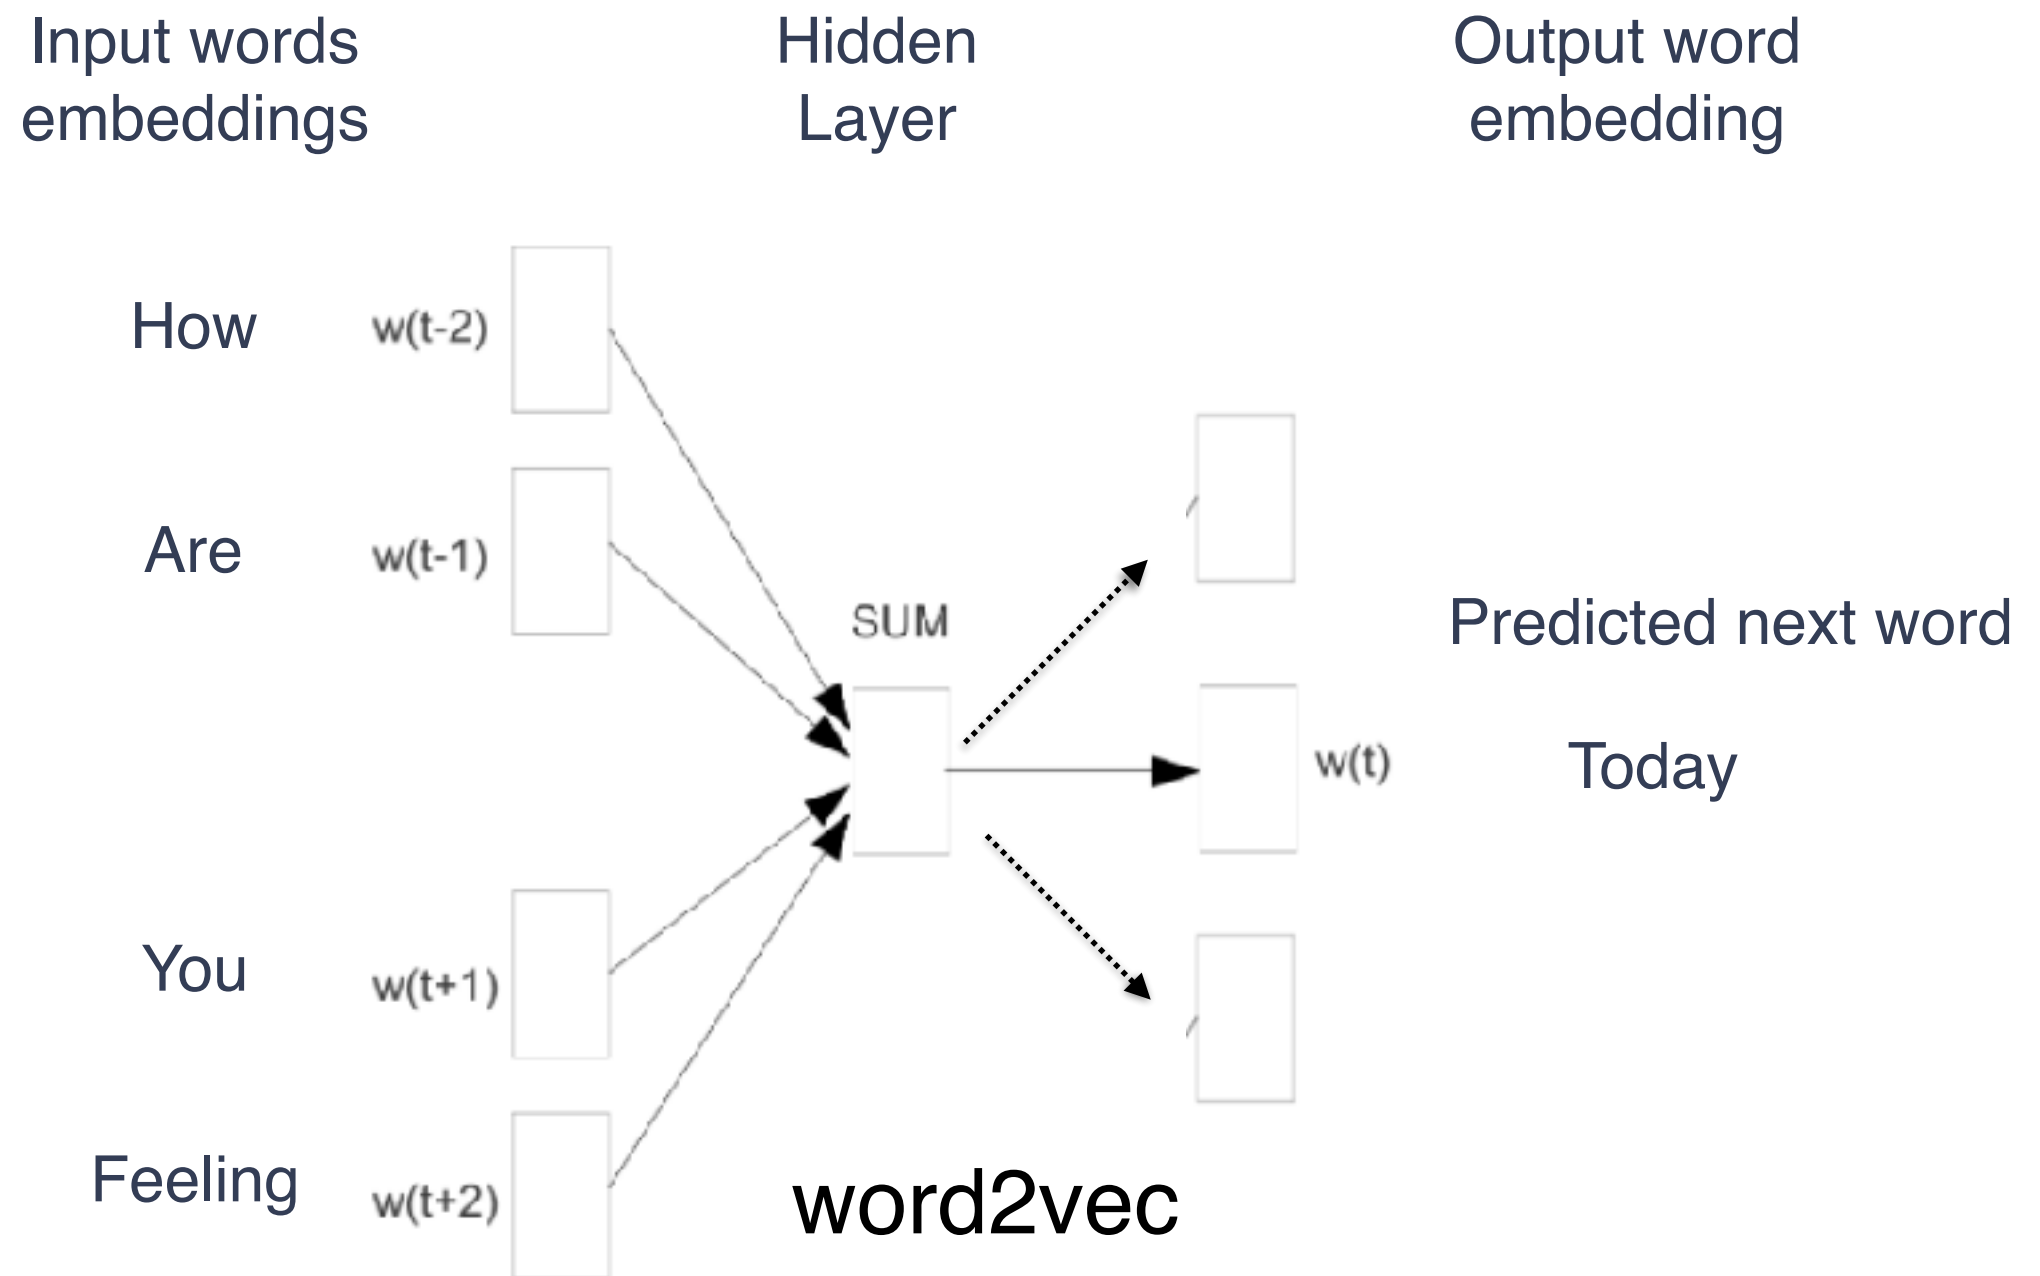
Approximation
function

Learning via Back propagation



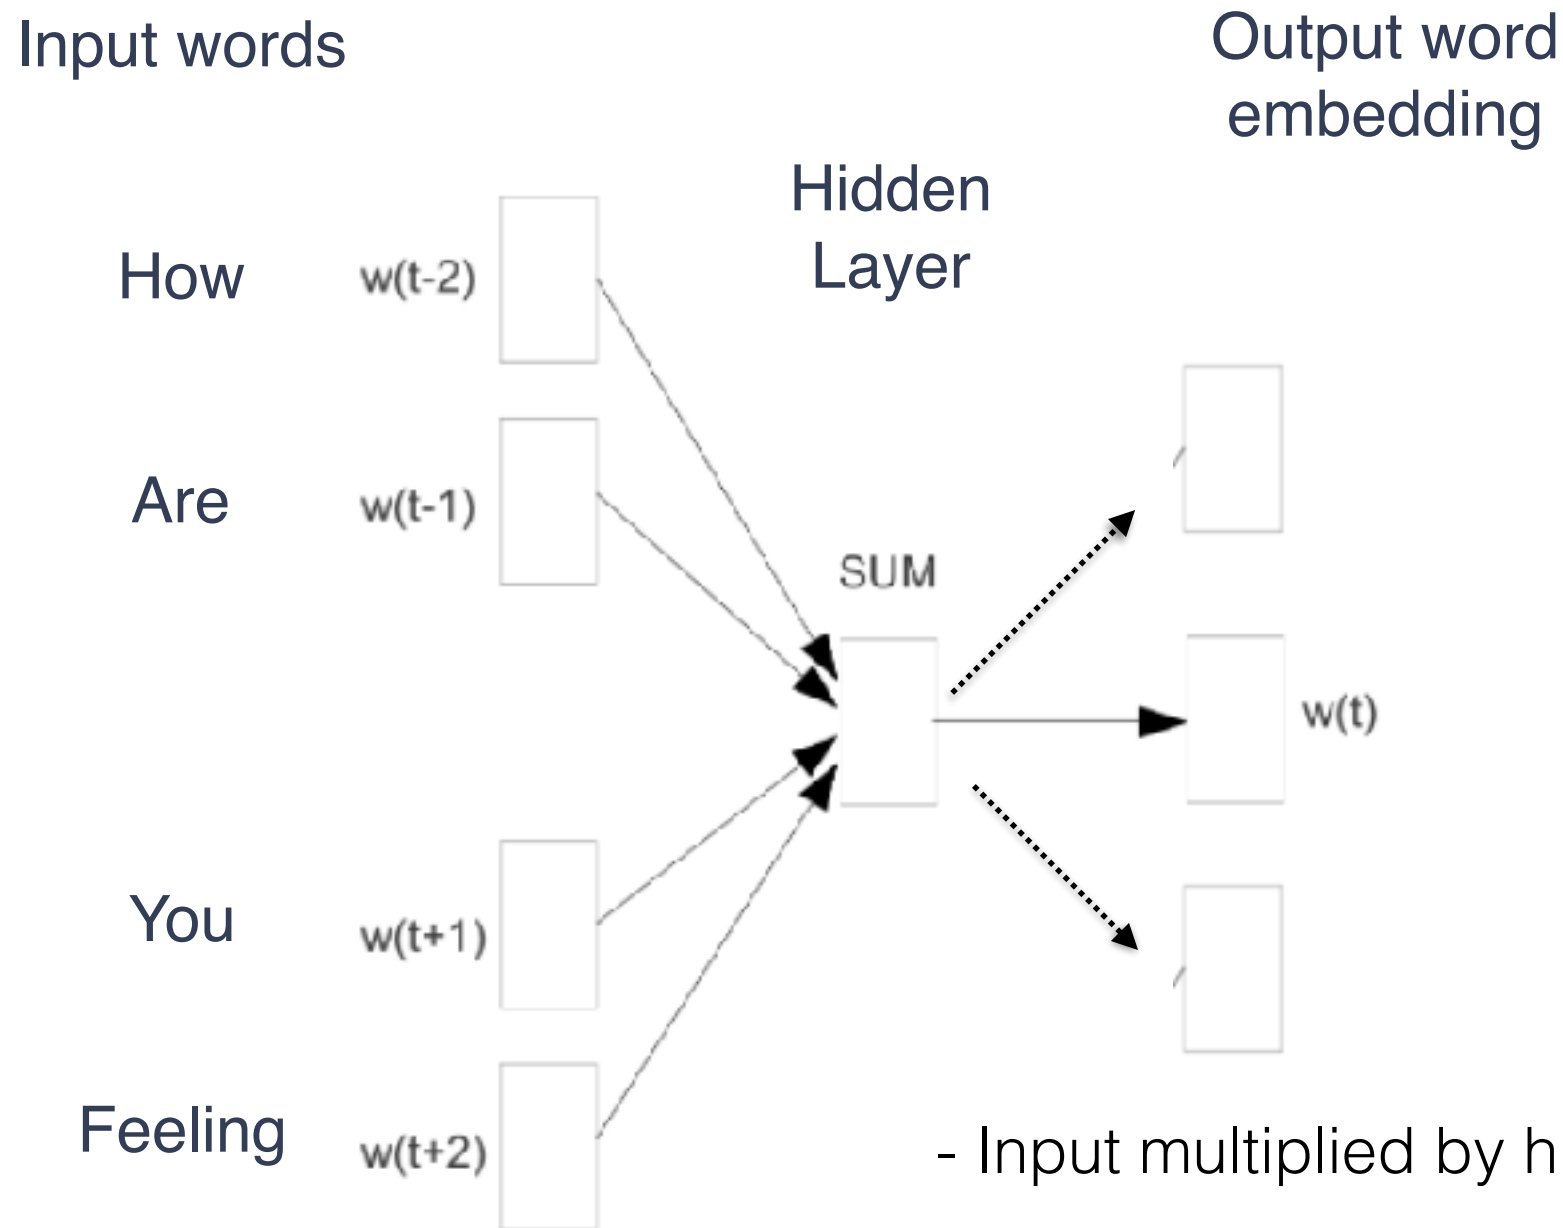
But how Neural Network can help in predicting?

word2vec- Predicting next word



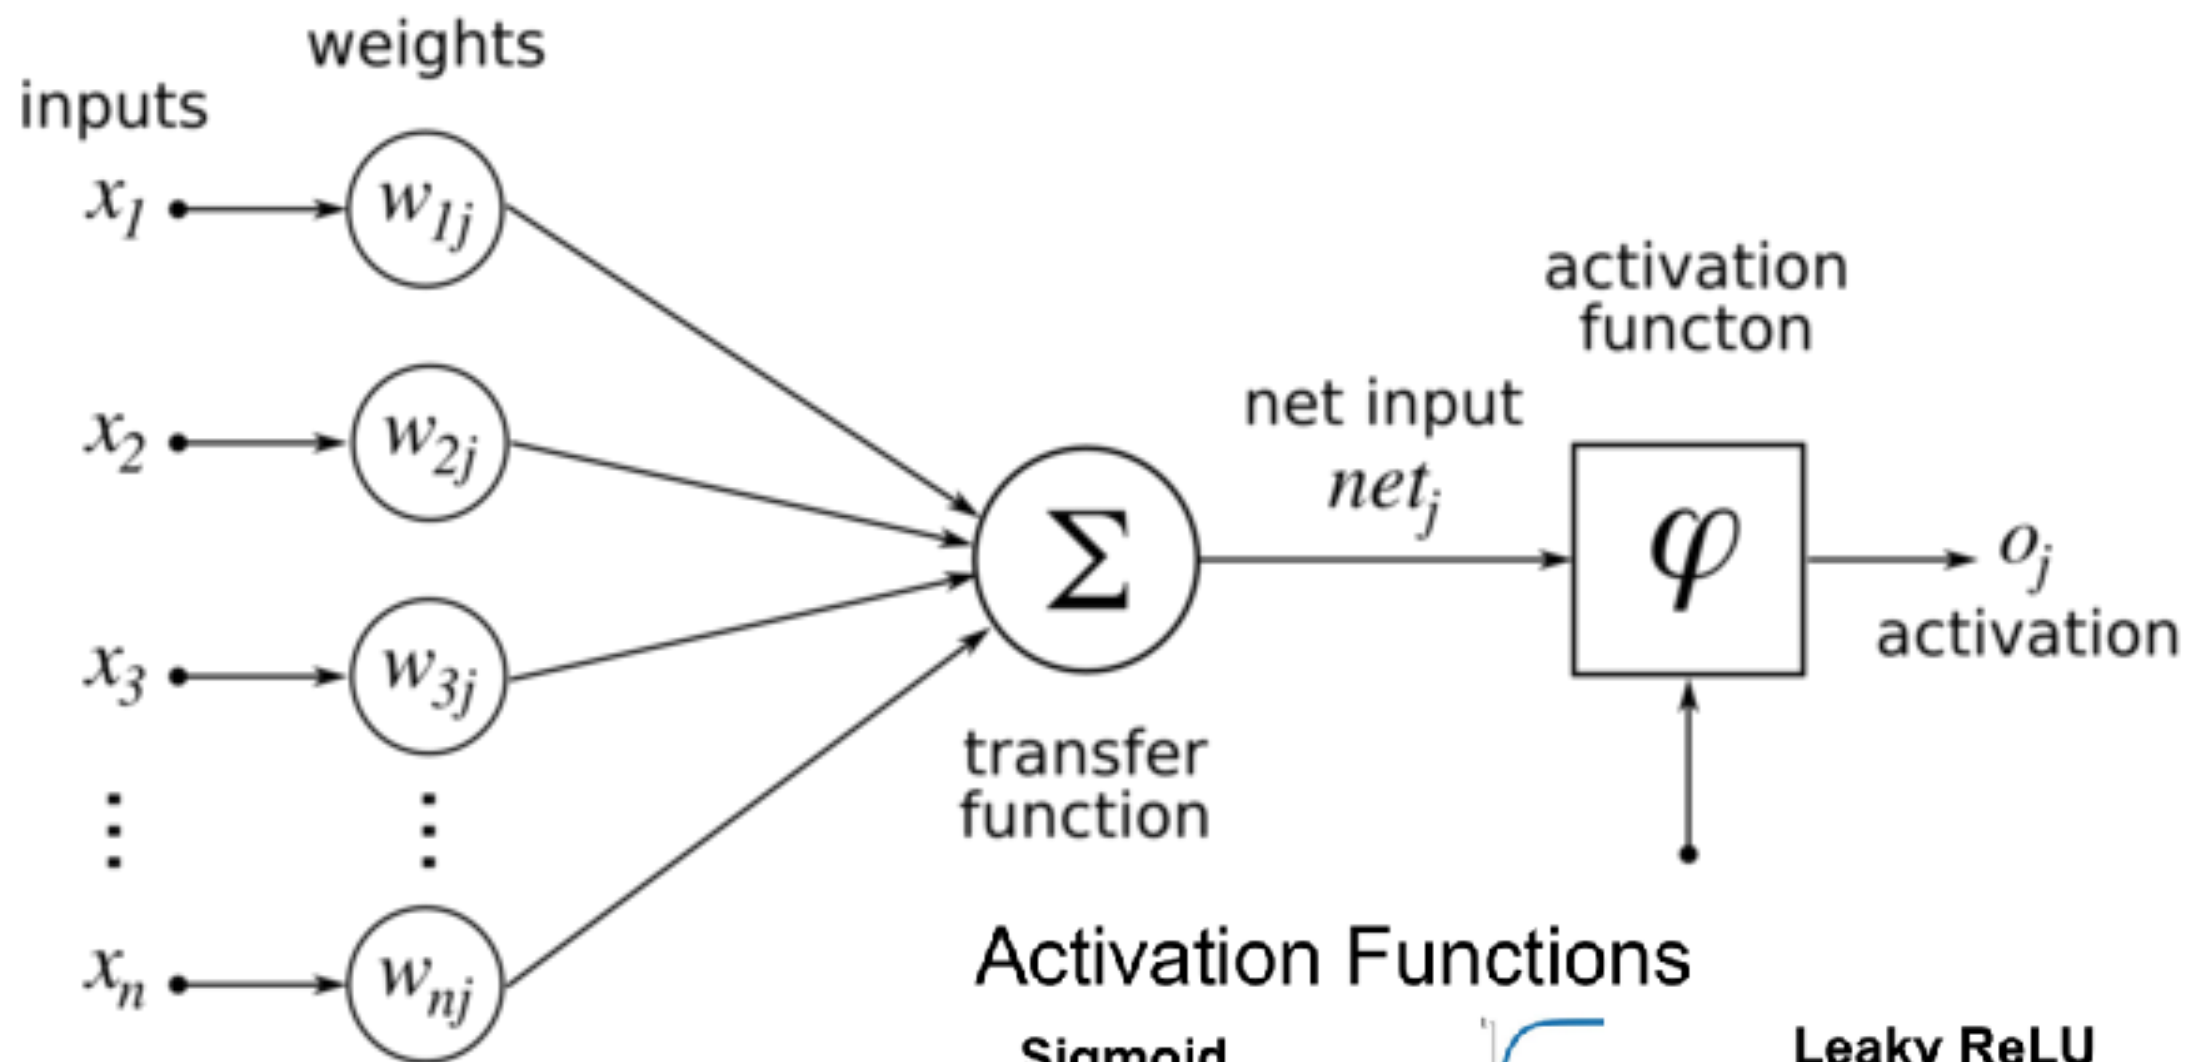
- Does not understand words or grammar

Semi supervised learning



- Input multiplied by hidden input weights
- Hidden input multiplied by hidden output weights
- Converts output to probabilities through **softmax**
- Error back propagated

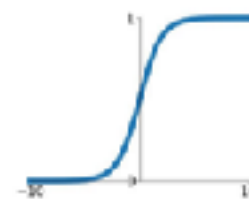
Activation function



Activation Functions

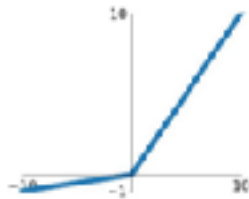
Sigmoid

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



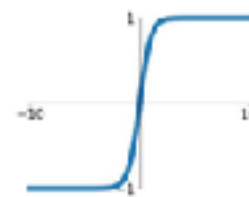
Leaky ReLU

$$\max(0.1x, x)$$



tanh

$$\tanh(x)$$

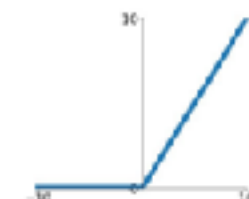


Maxout

$$\max(w_1^T x + b_1, w_2^T x + b_2)$$

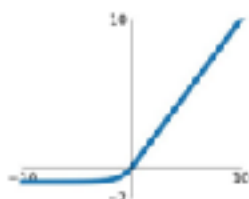
ReLU

$$\max(0, x)$$



ELU

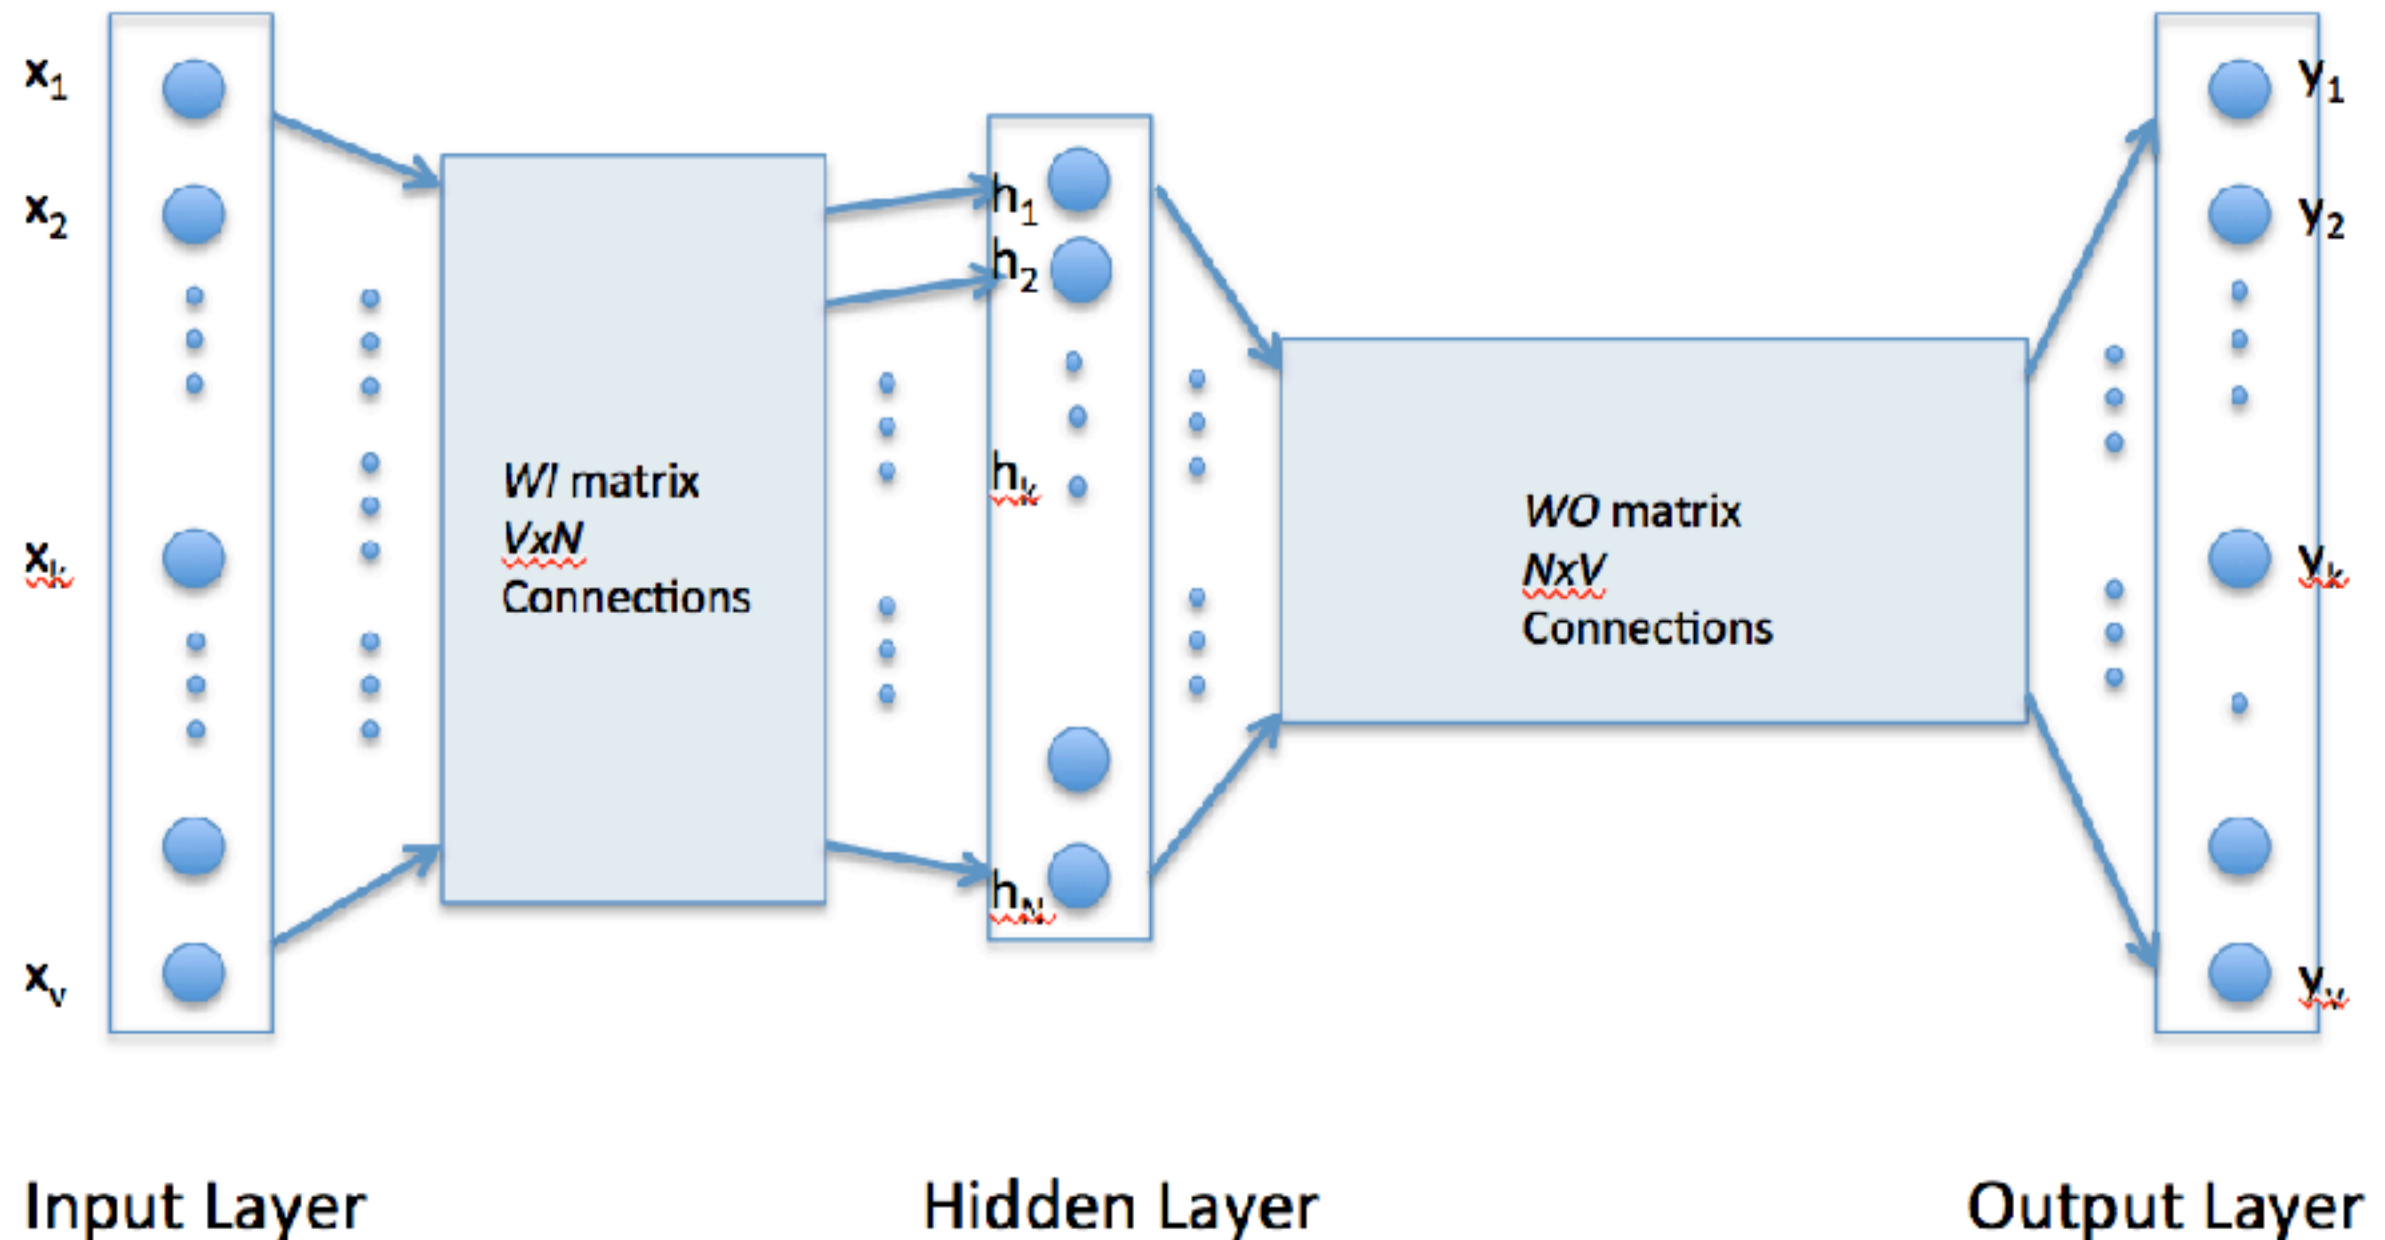
$$\begin{cases} x & x \geq 0 \\ \alpha(e^x - 1) & x < 0 \end{cases}$$



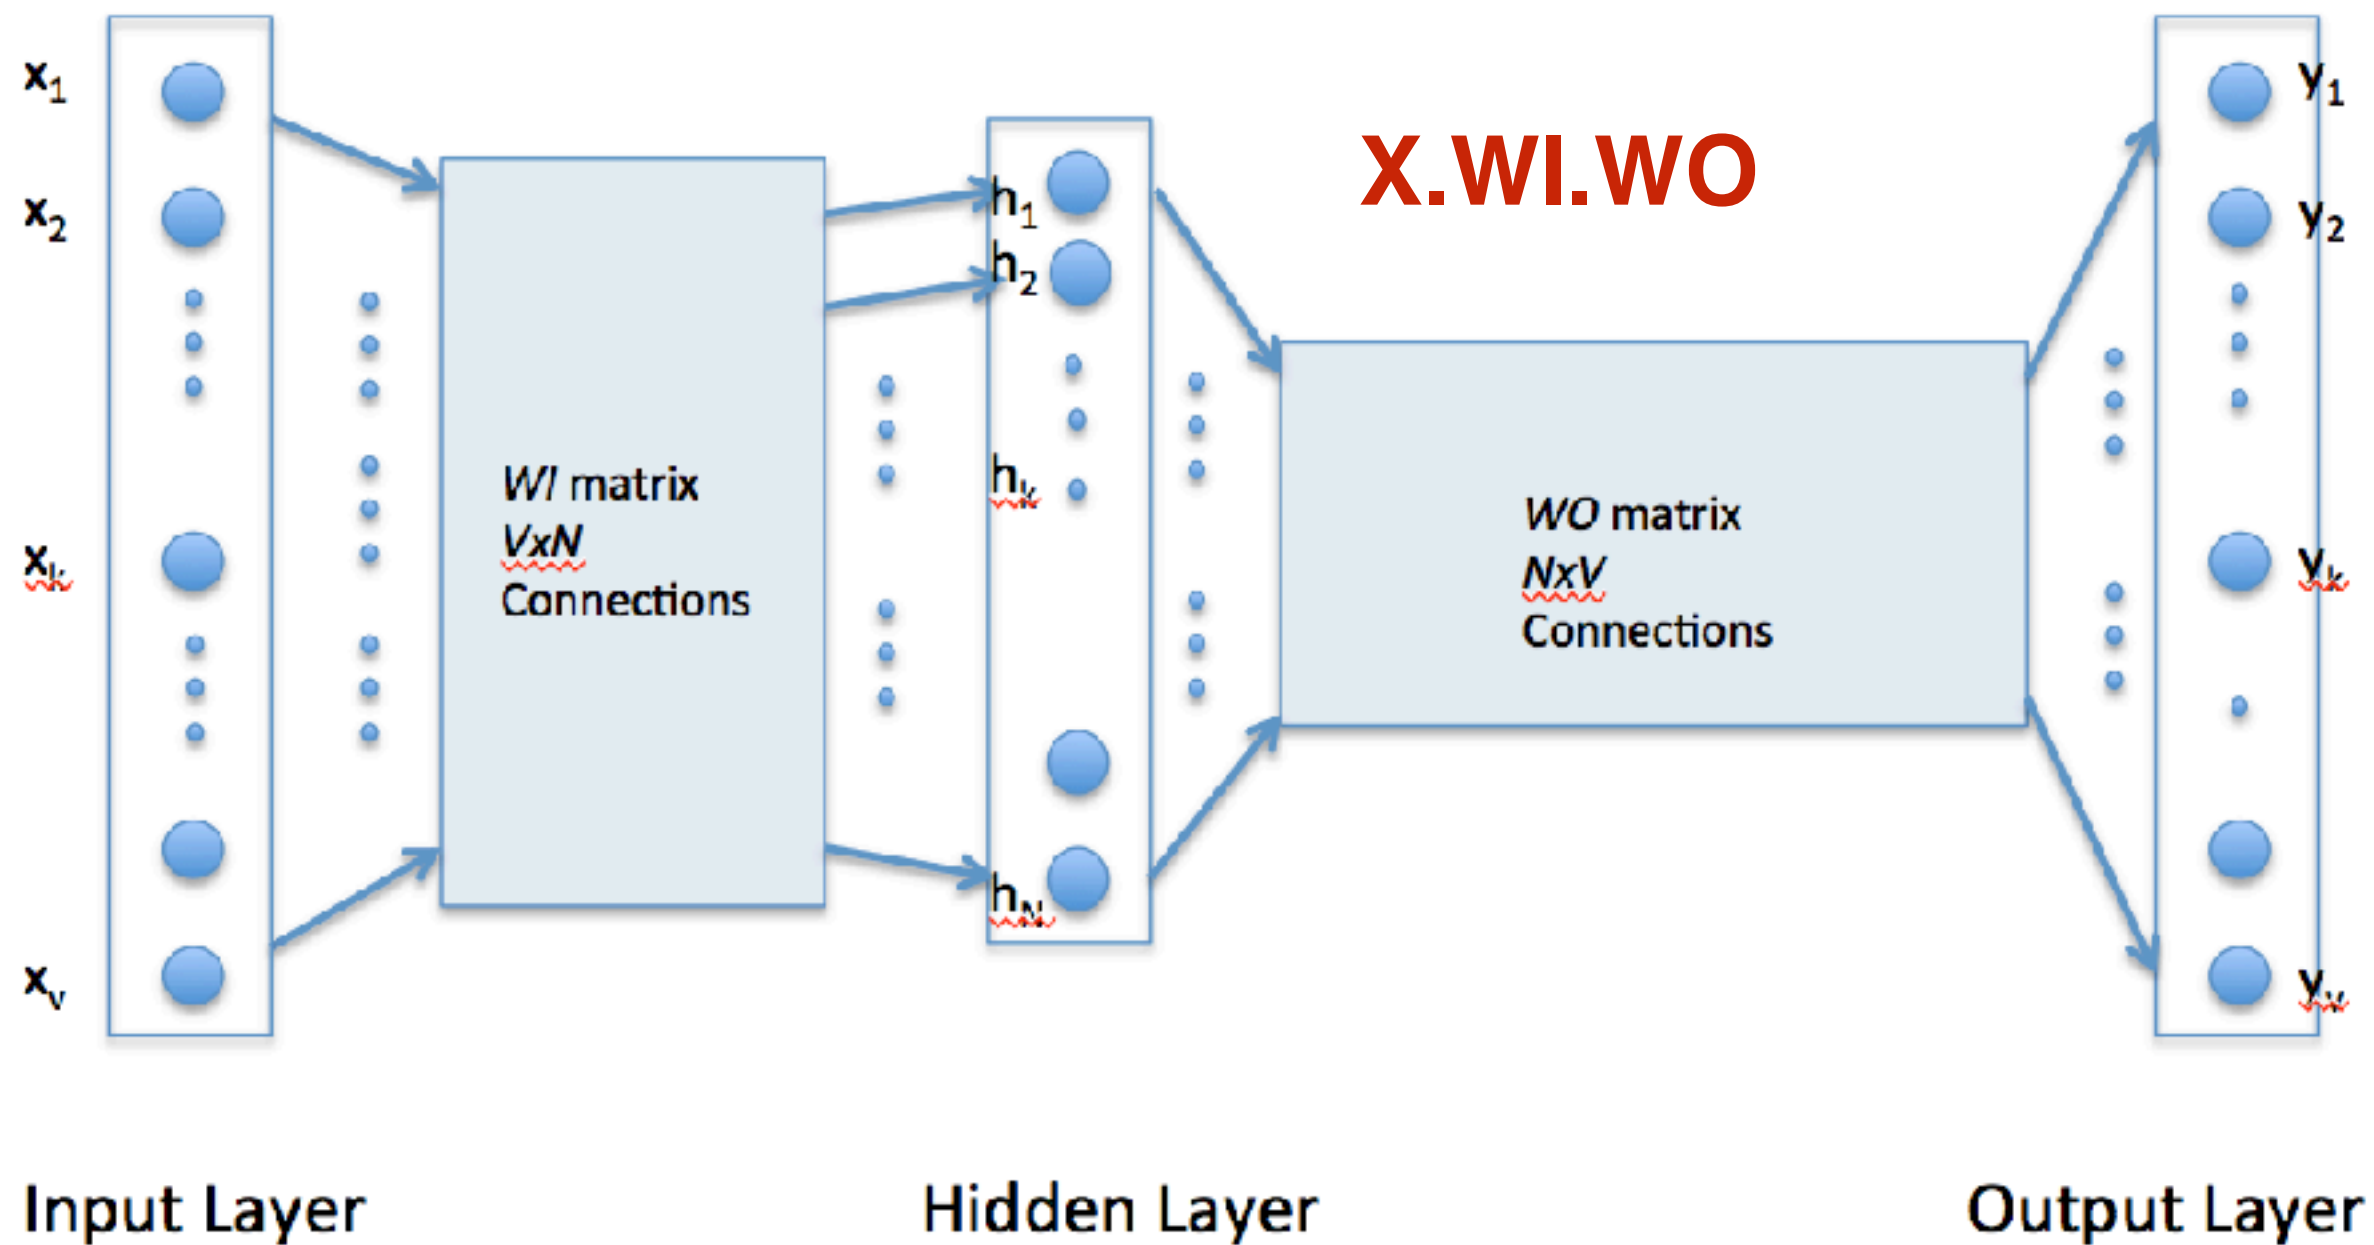
Softmax

- Multiclass regression
- Extension of sigmoid to multiclass

word2vec - Something more interesting?

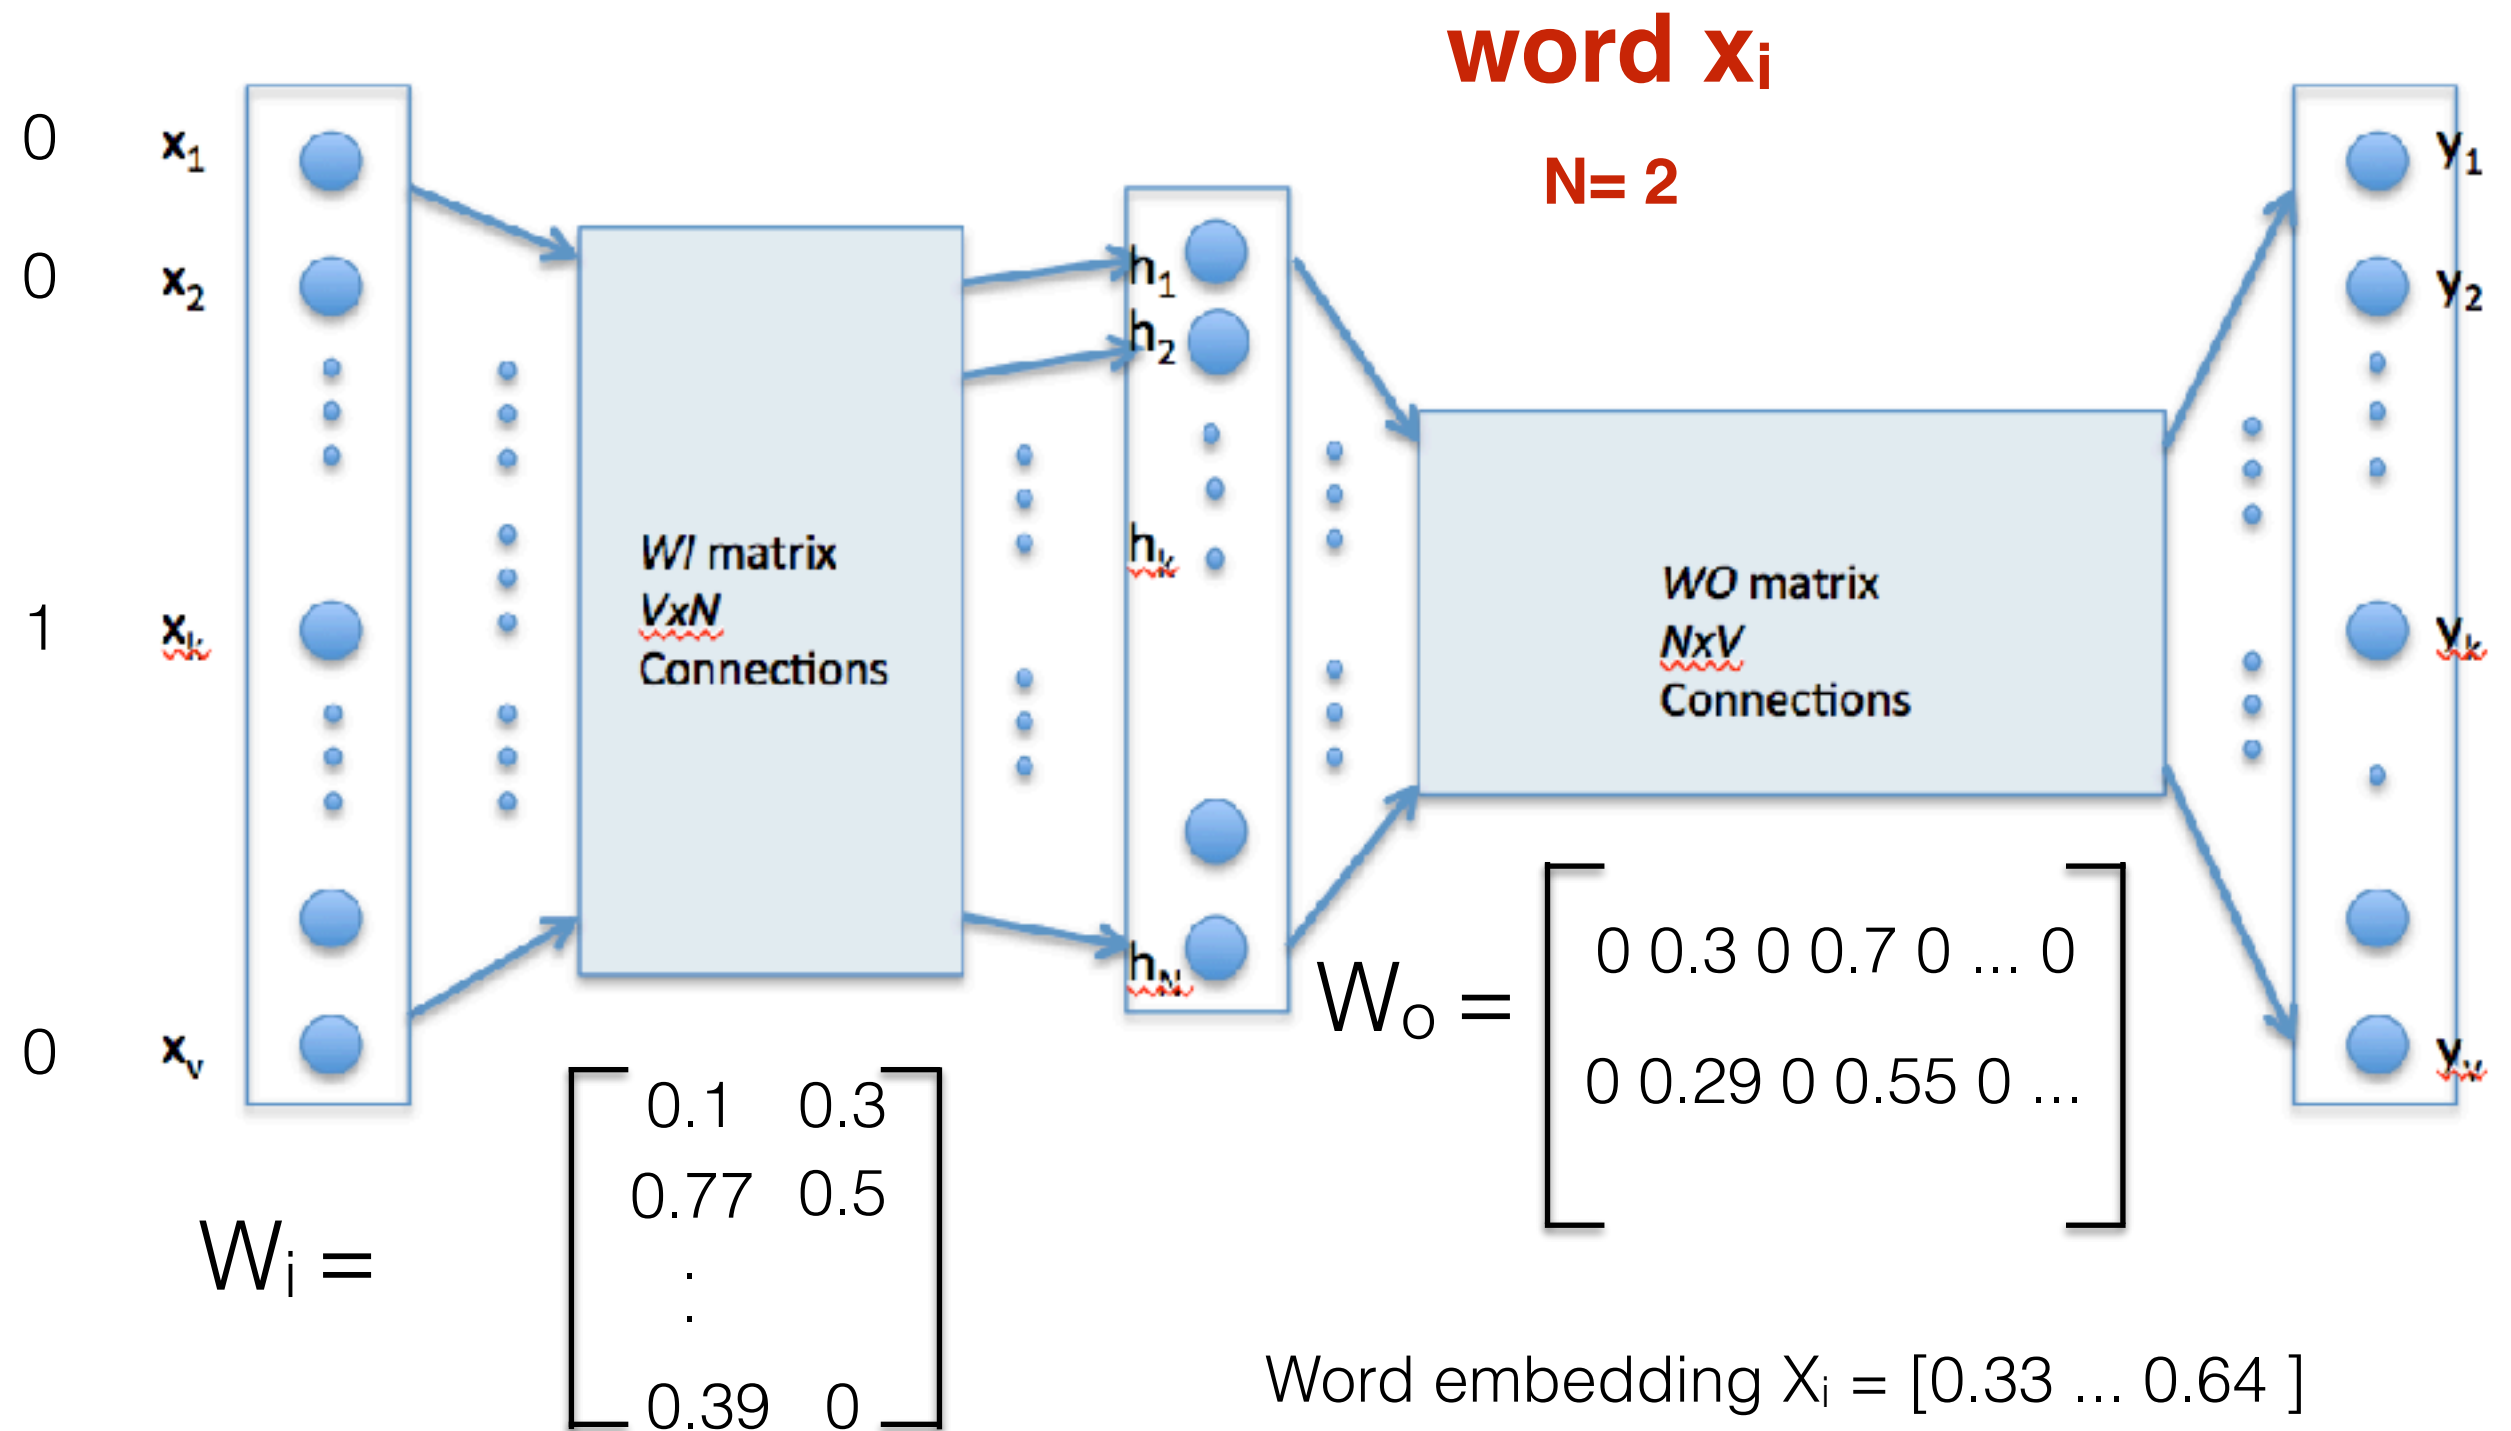


Word2vec Word embeddings



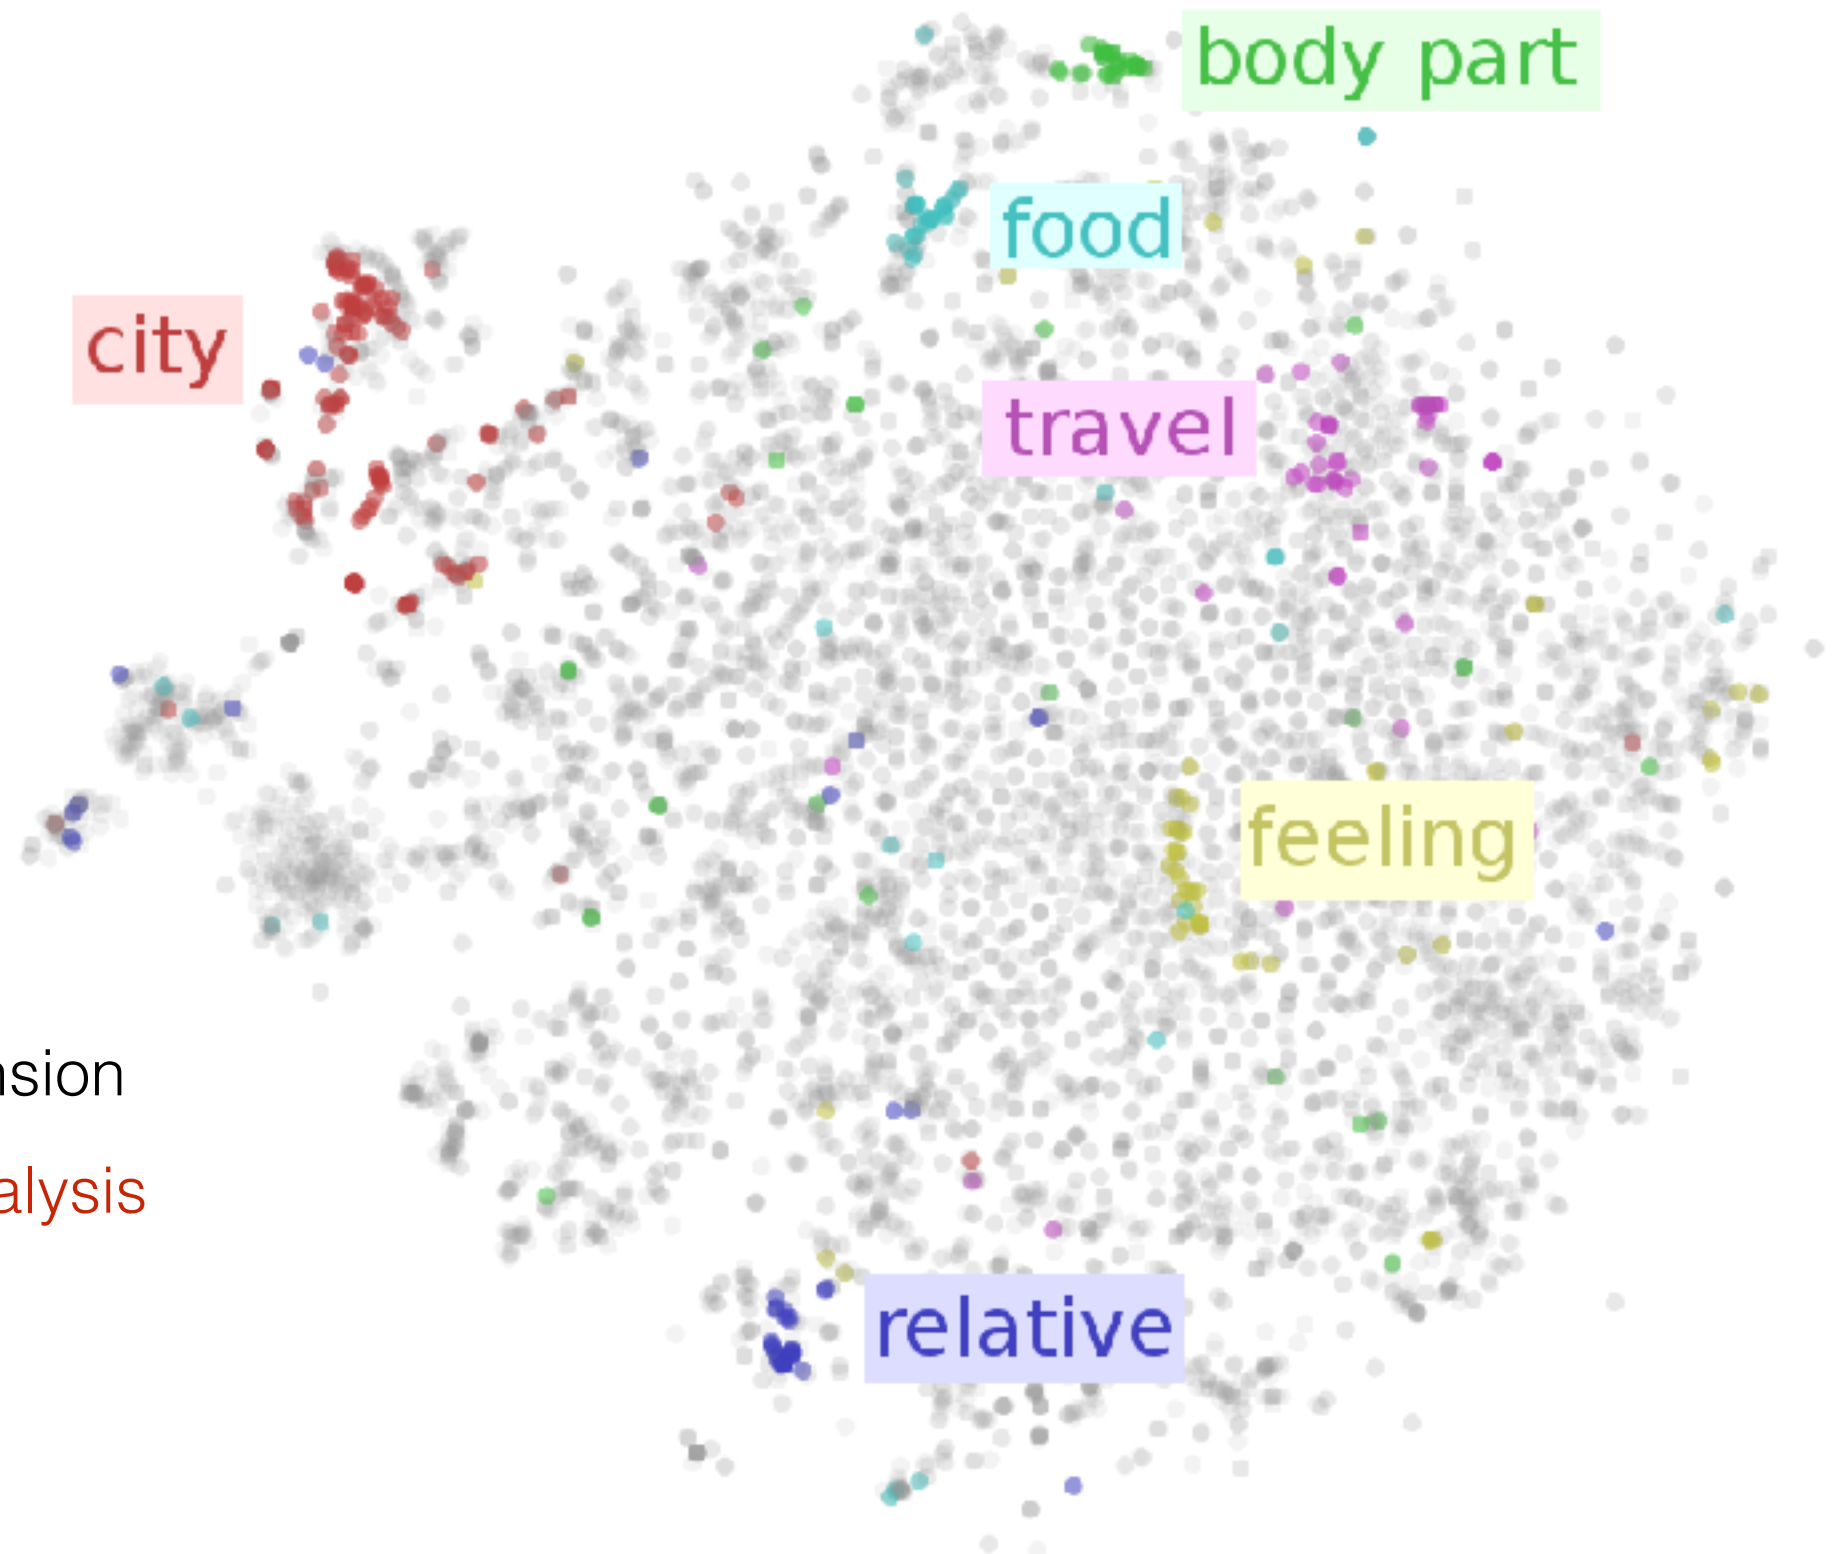
$$X.W_i = \begin{bmatrix} x_1 w_1 & x_1 w_i \\ & x_i w_i \\ & x_v w_i & x_v w_n \end{bmatrix}$$

Example



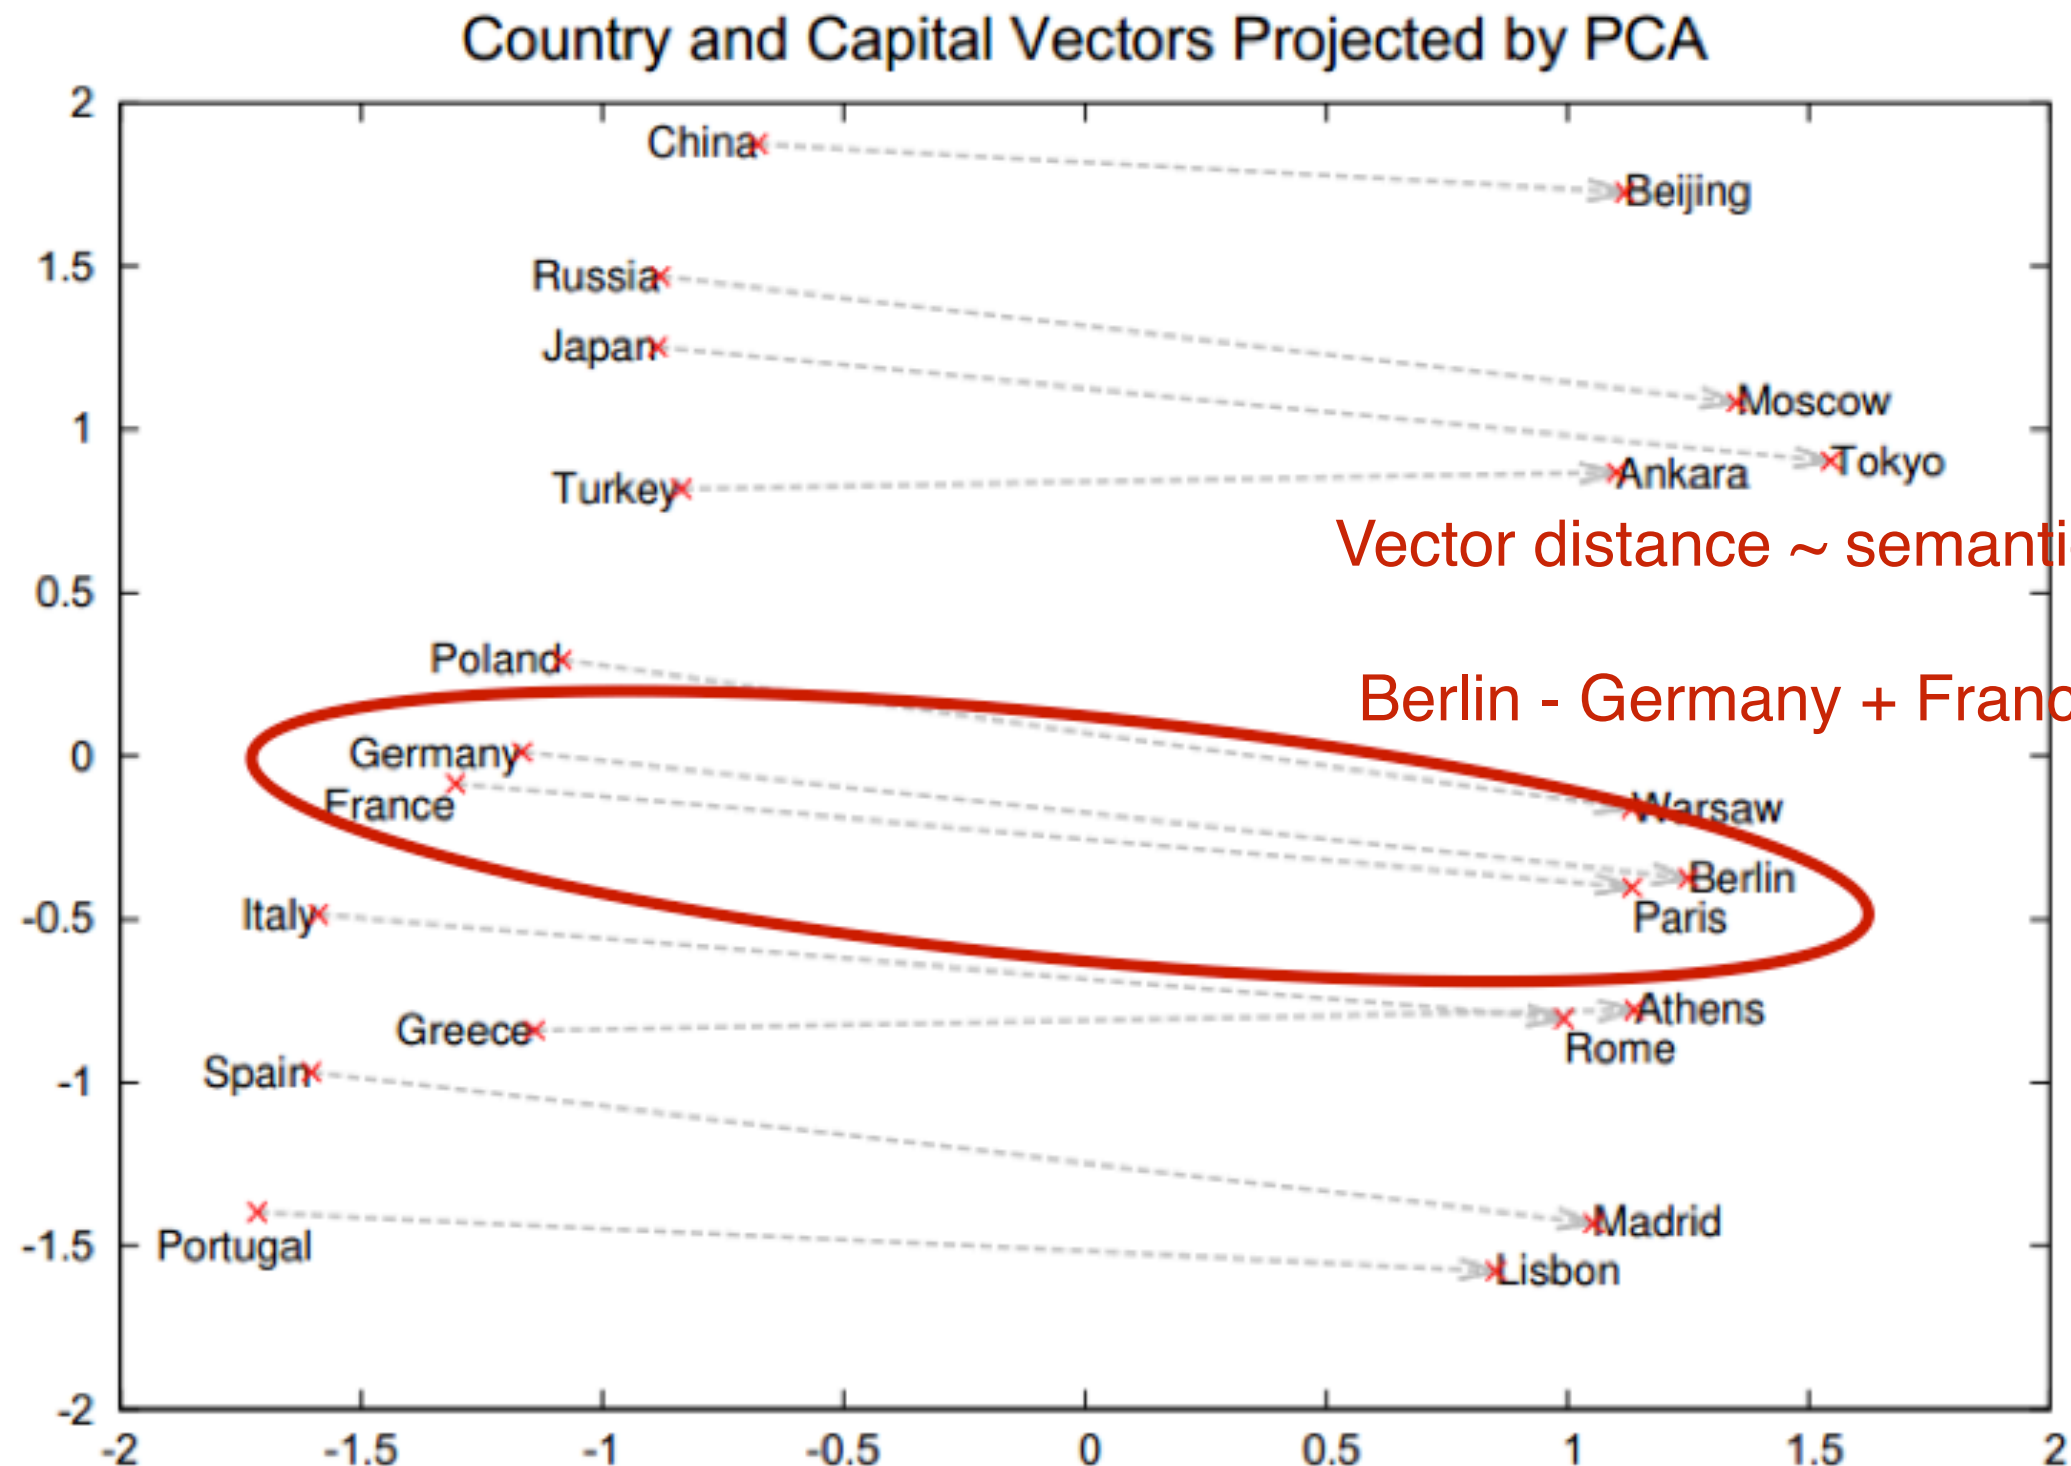
$$\text{Word embedding} = X \cdot W_i \cdot W_o$$

word2vec - 2d space



- Output matrix to 2 dimension
- Principal component Analysis

word2vec



word2vec - Predicting next order



Groceries Delivered From Local Stores

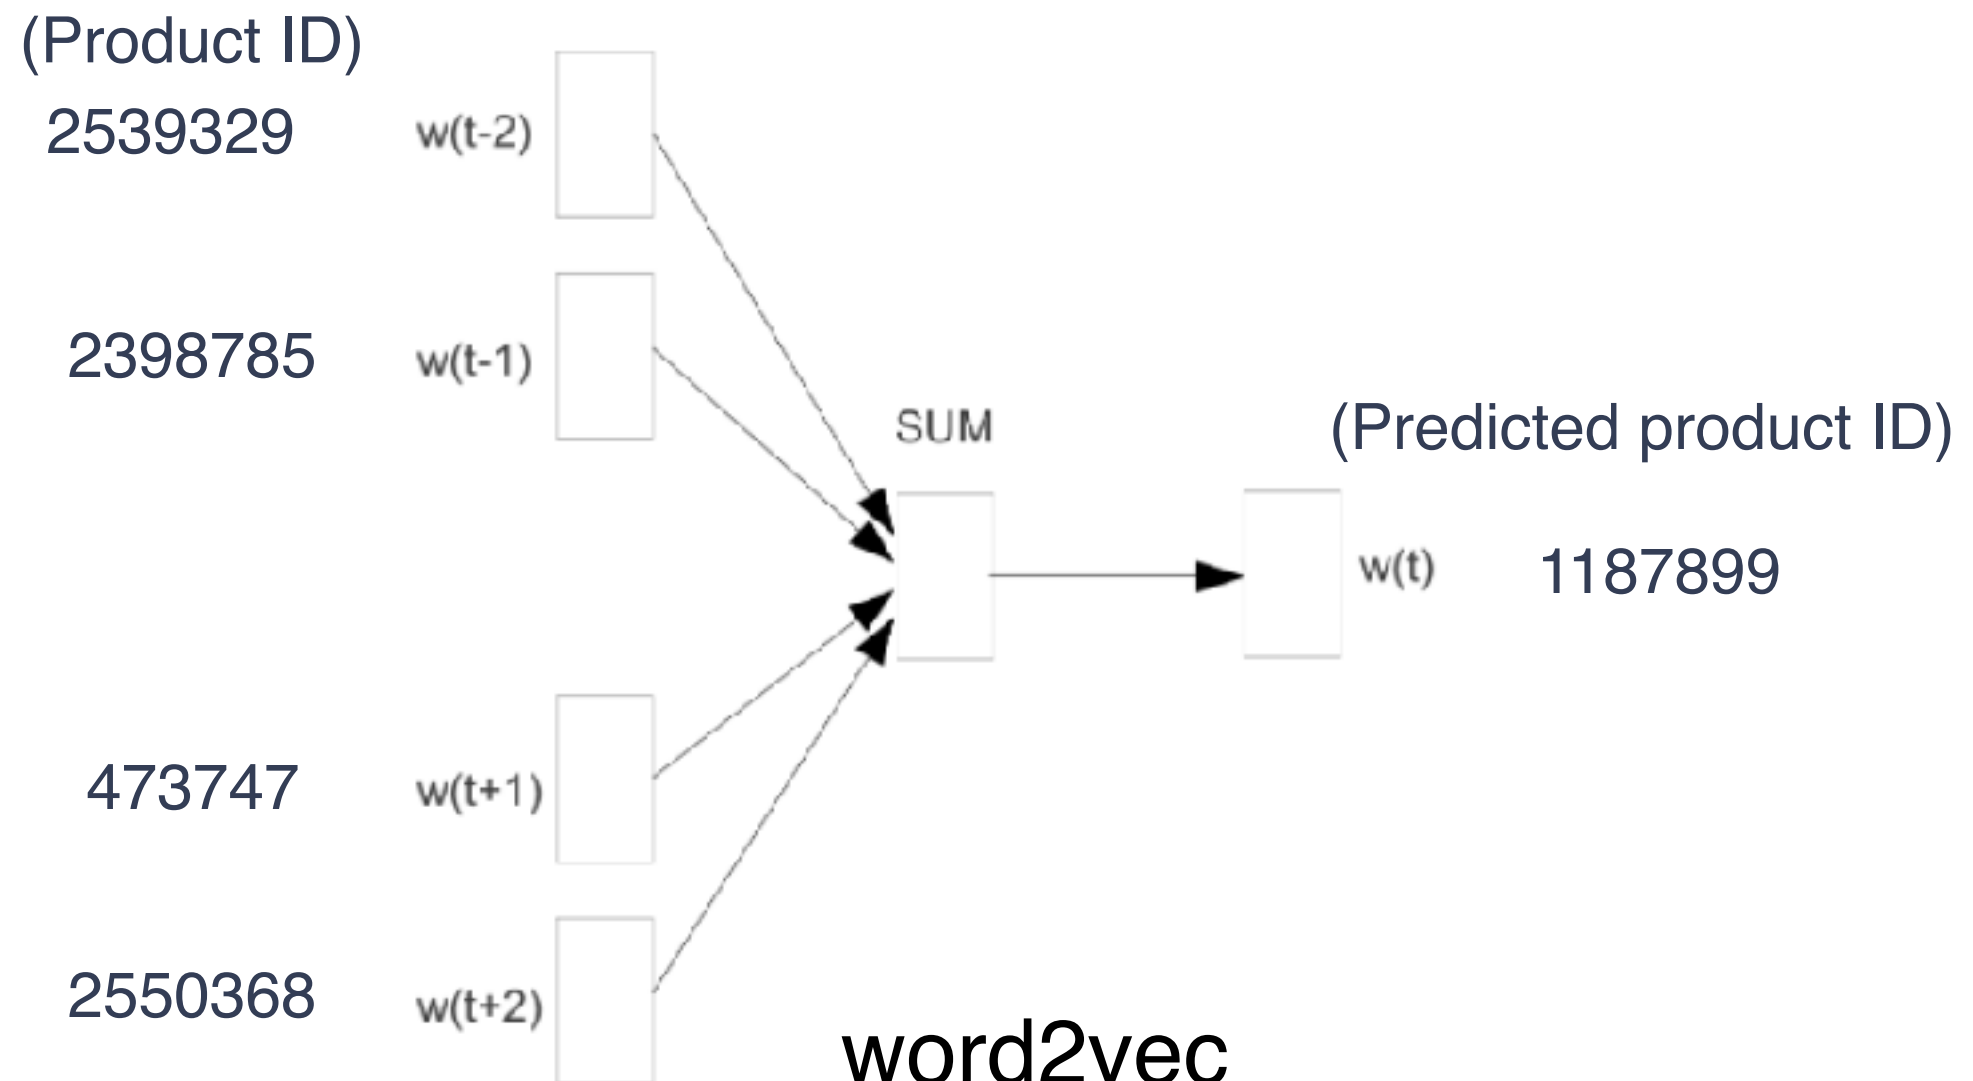
Analyzing 3M Instacart Orders

- **Prior: ~3.2m orders**
- **Train: ~131k orders**
- **Test: ~75k orders**

Instacart Data

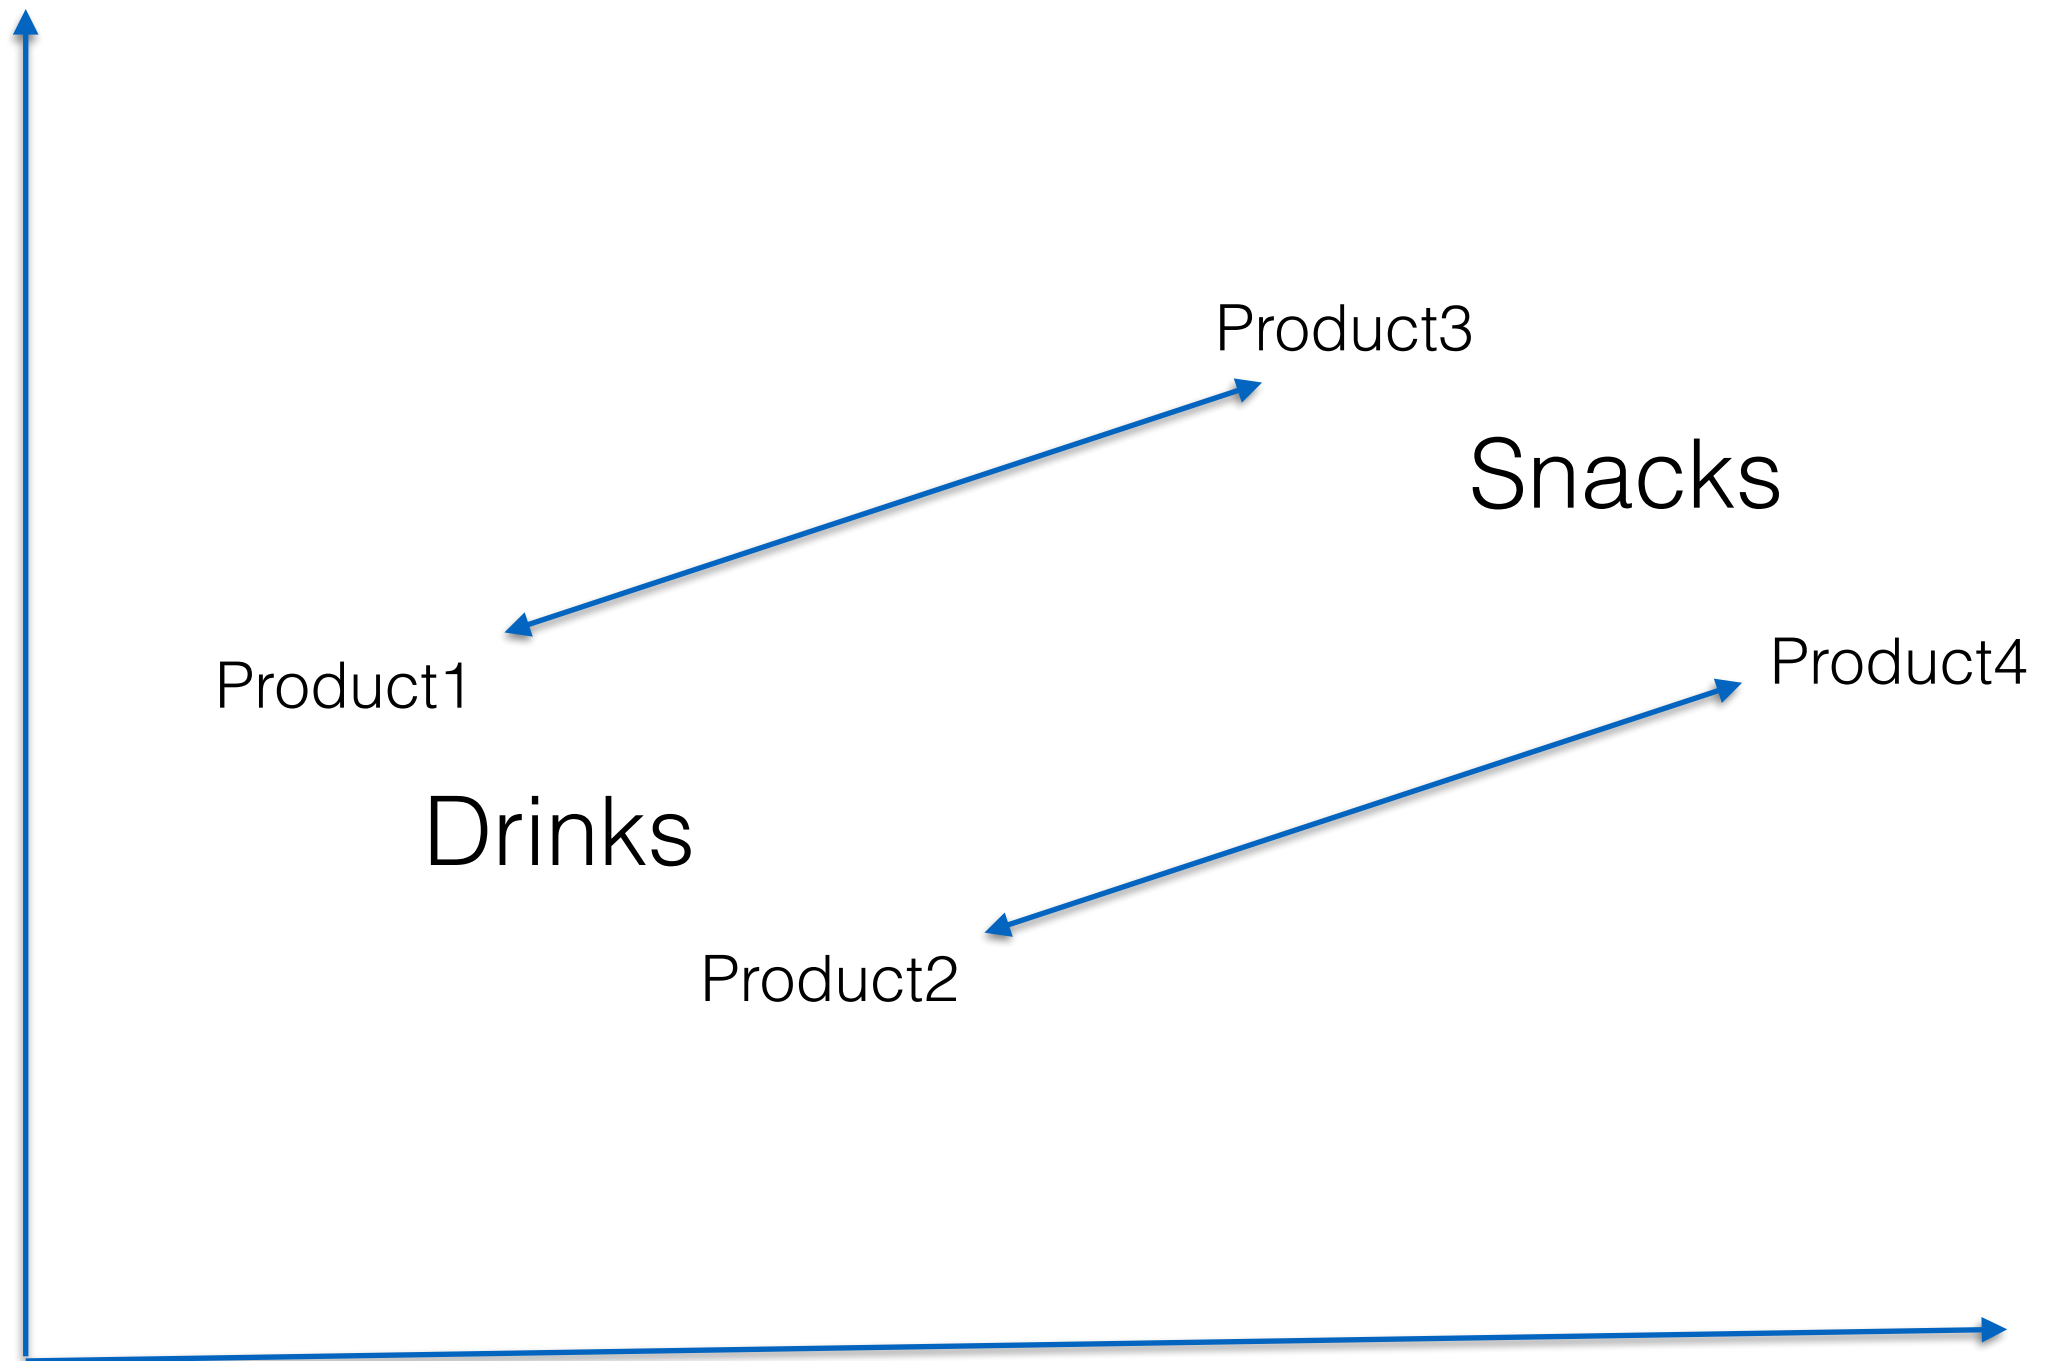
product_id	user_id	eval_set	order_number	order_dow	order_hour_of_day	days_since_prior_order
2539329	1	prior	1	2	8	7
2398795	1	prior	2	5	7	15
473747	1	prior	3	7	12	20
22544786	1	prior	4	1	7	21
4215438	1	prior	5	3	15	28
2295261	1	prior	6	2	7	19
2295261	1	prior	7	6	20	20
2550362	1	prior	8	5	14	14
1187899	1	prior	9	2	16	0
2168274	1	prior	10	2	8	30
1501582	1	train	11	1	11	10

Current approach to prediction for instacart



- Using the semantic association between orders

Vector space of products



word2vec in tensorflow

Training the Model

Training the model is then as simple as using a `feed_dict` to push data into the placeholders and calling `tf.Session.run` with this new data in a loop.

```
for inputs, labels in generate_batch(...):  
    feed_dict = {train_inputs: inputs, train_labels: labels}  
    _, cur_loss = session.run([optimizer, loss], feed_dict=feed_dict)
```

See the full example code in [tensorflow/examples/tutorials/word2vec/word2vec_basic.py](https://www.tensorflow.org/examples/tutorials/word2vec/word2vec_basic.py).

Feeding

TensorFlow's feed mechanism lets you inject data into any Tensor in a computation graph. A python computation can thus feed data directly into the graph.

Supply feed data through the `feed_dict` argument to a `run()` or `eval()` call that initiates computation.

```
with tf.Session():  
    input = tf.placeholder(tf.float32)  
    classifier = ...  
    print(classifier.eval(feed_dict={input: my_python_preprocessing_fn()}))
```

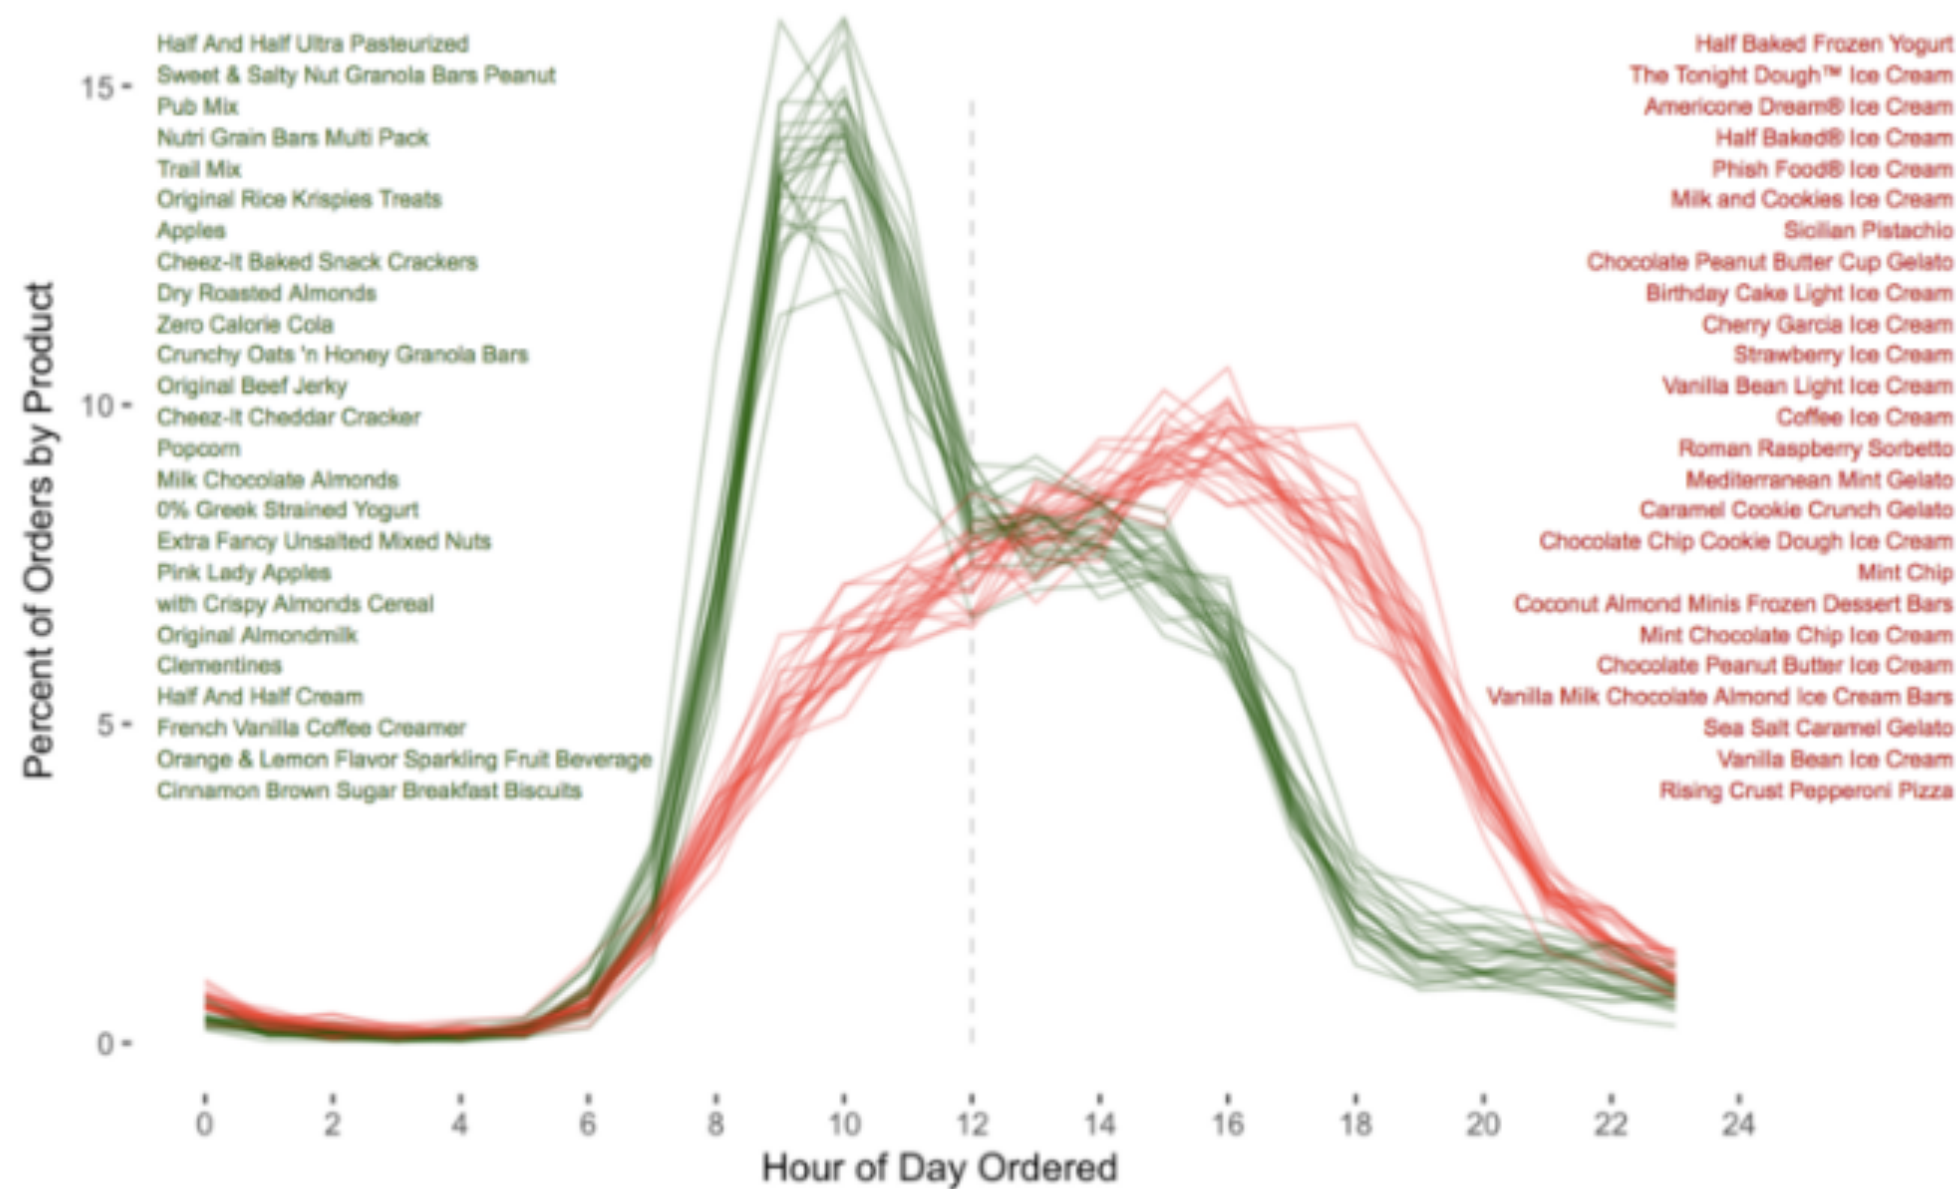
Use-case feature for using word2vec

- **Repeated patterns of human action**

Any other use case?

What is missing?

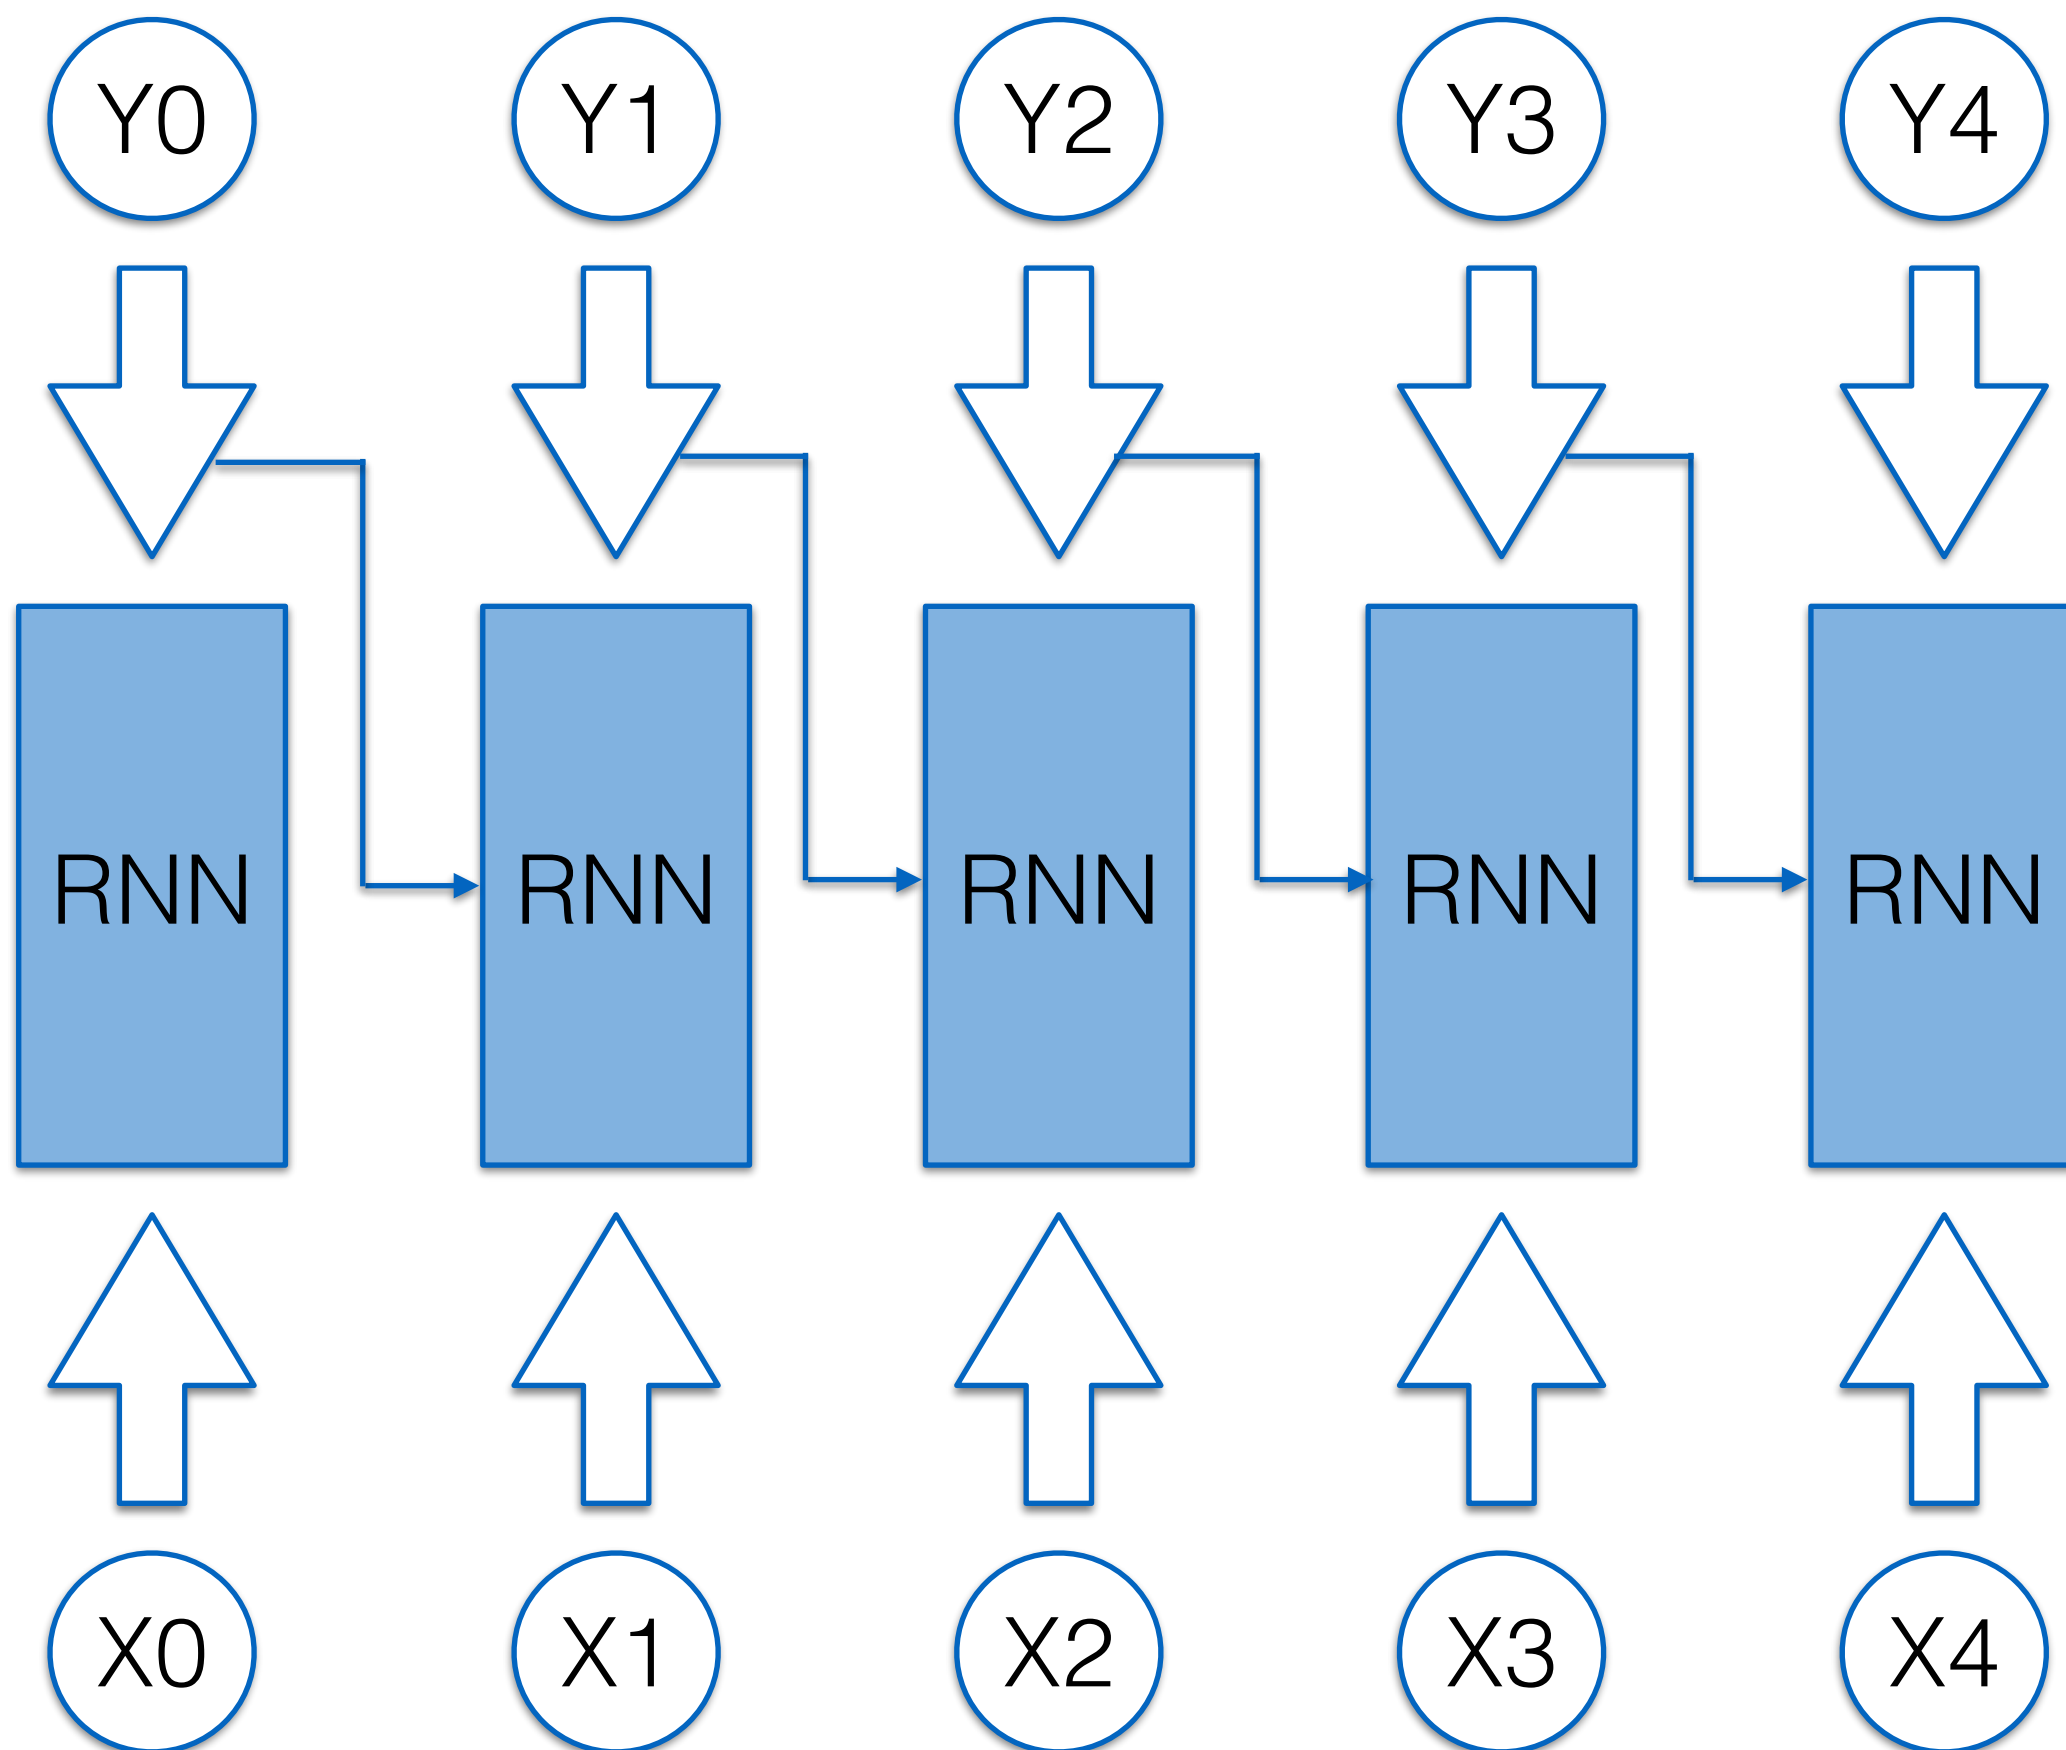
Sequence of order is important!



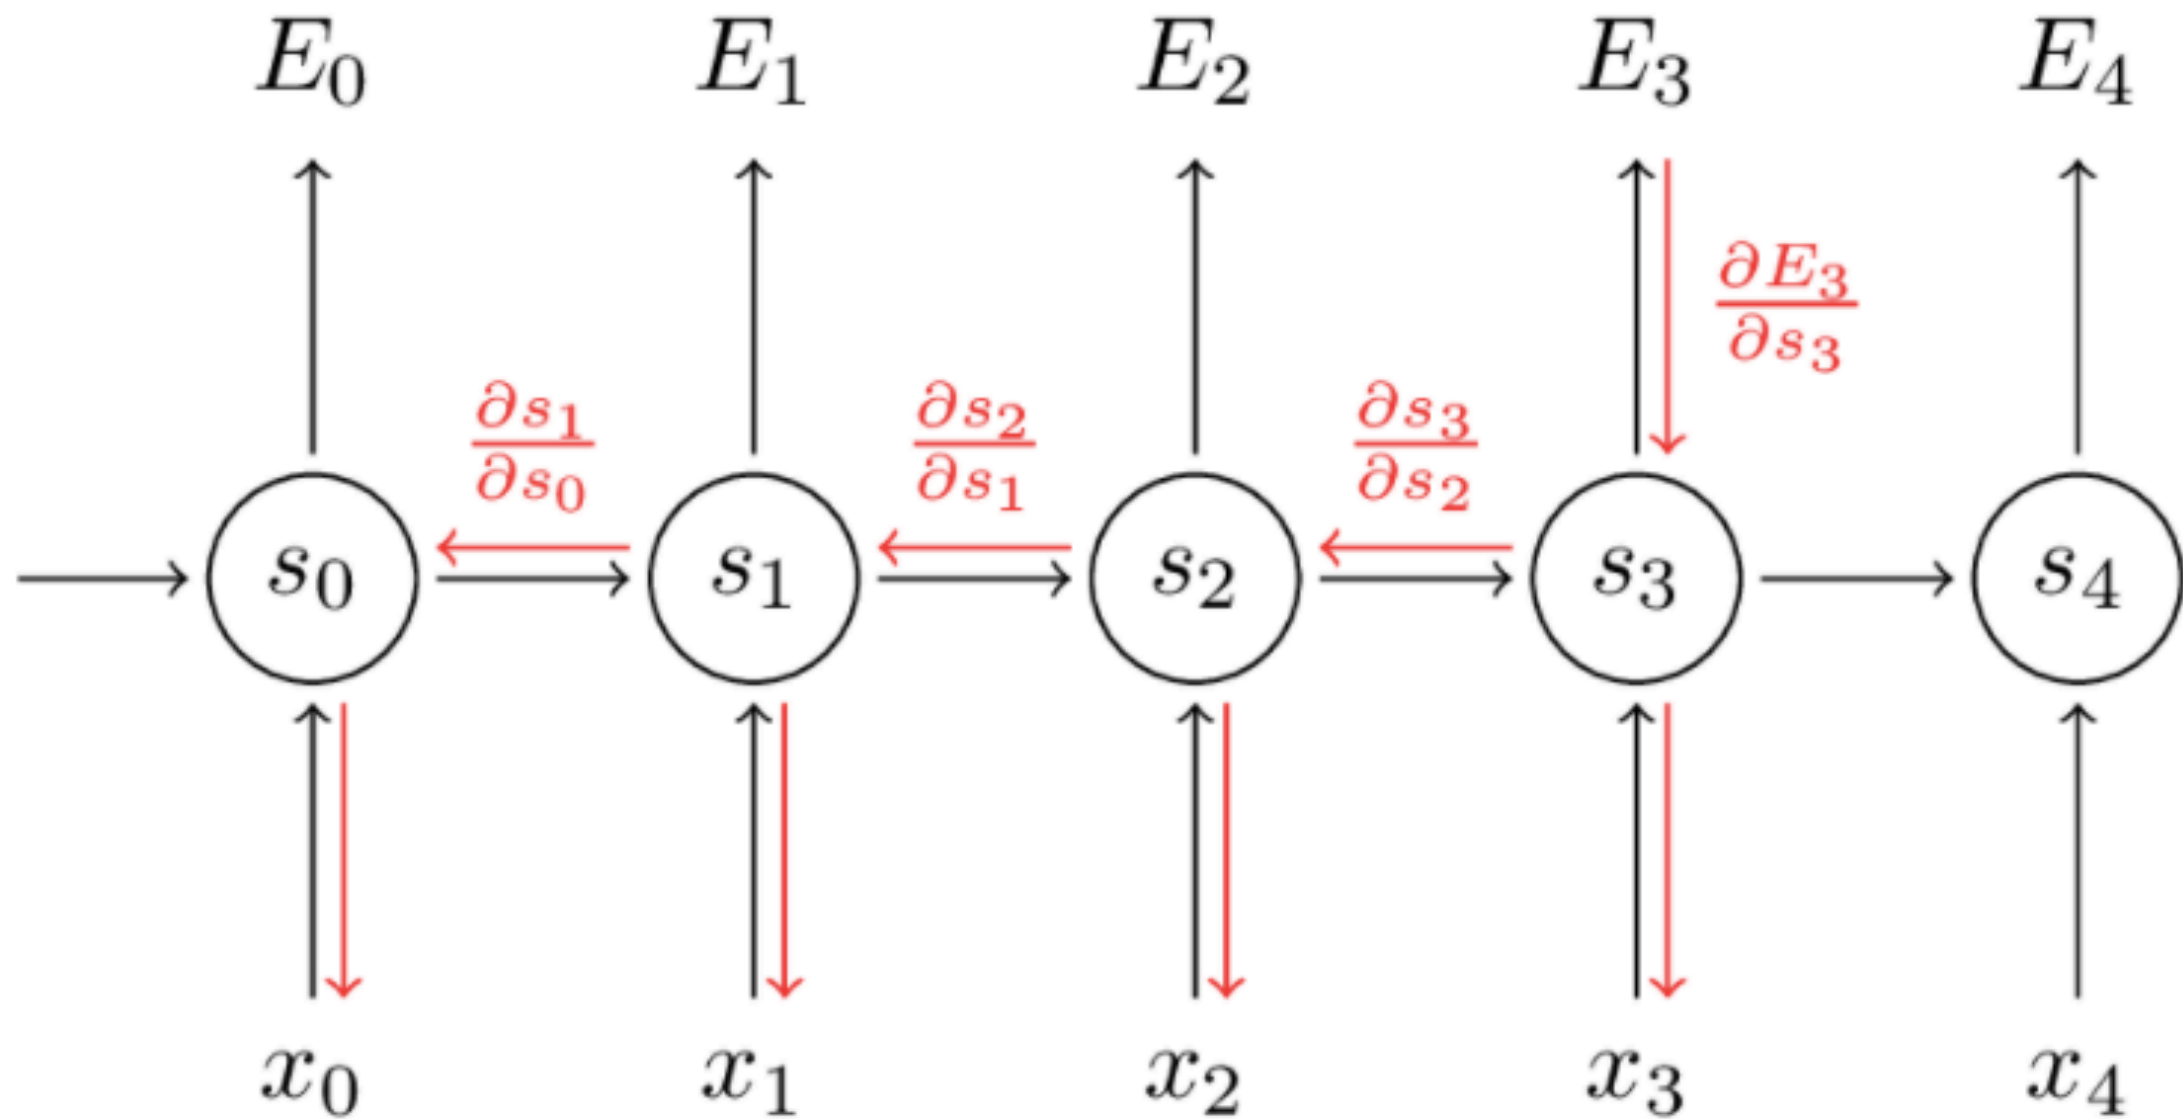
Popular products purchased earliest in the day (green) and latest in the day (red).

Instacart Data

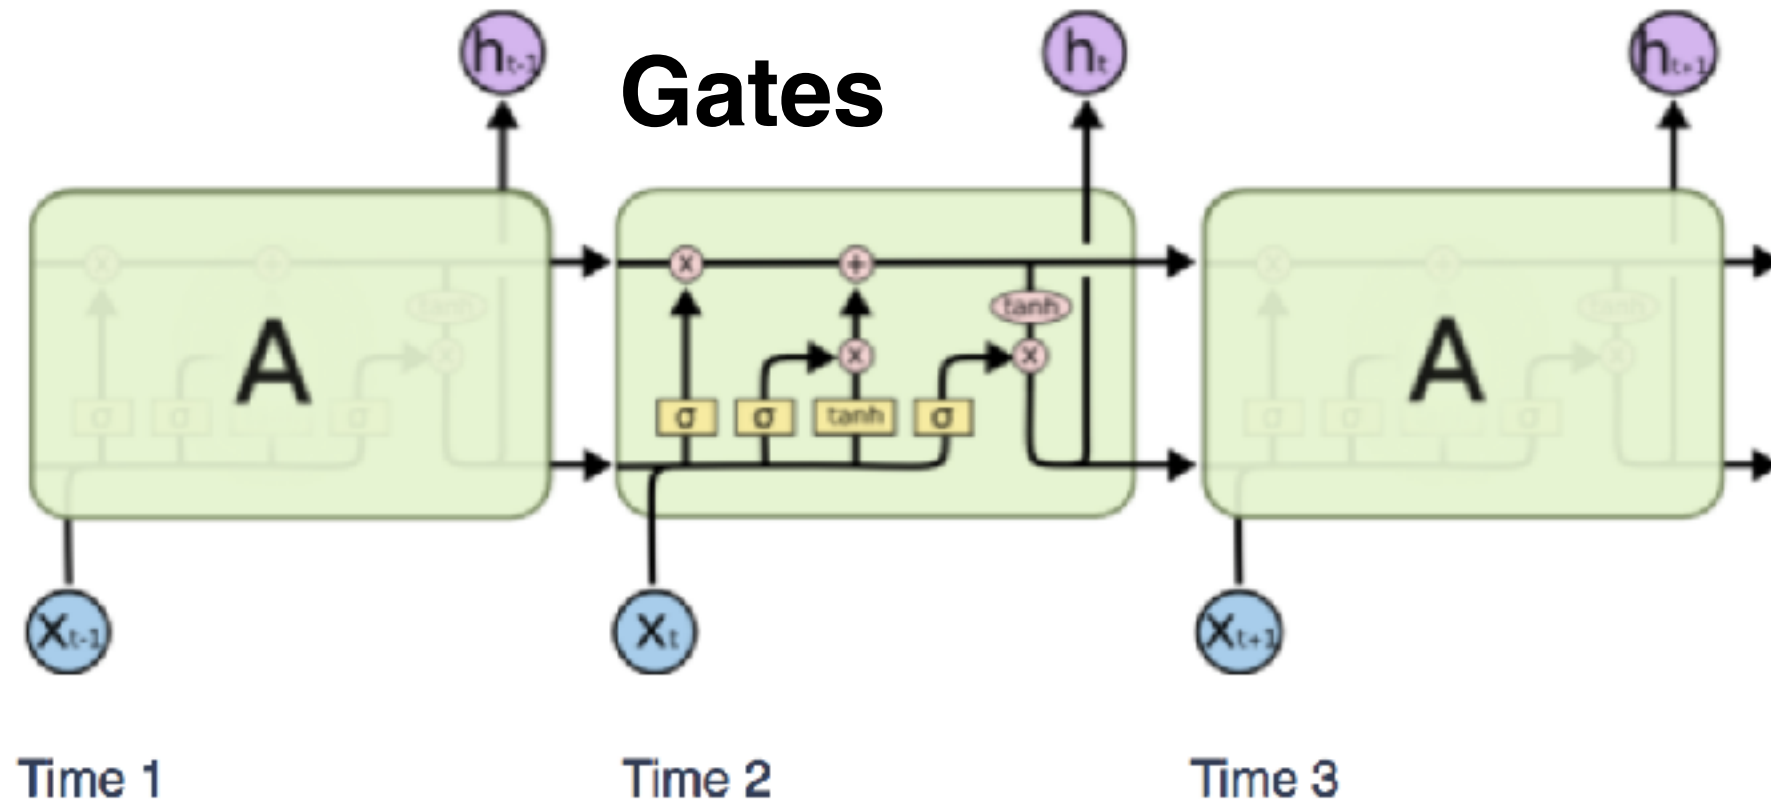
product_id	user_id	eval_set	order_number	order_dow	order_hour_of_day	days_since_prior_order
2539329	1	prior	1	2	8	
2398795	1	prior	2	5	7	15
473747	1	prior	3	7	12	20
22544786	1	prior	4	1	7	21
4215438	1	prior	5	3	15	28
2295261	1	prior	6	2	7	19
2295261	1	prior	7	6	20	20
2550362	1	prior	8	5	14	14
1187899	1	prior	9	2	16	0
2168274	1	prior	10	2	8	30
1501582	1	train	11	1	11	10



Vanishing gradient problem

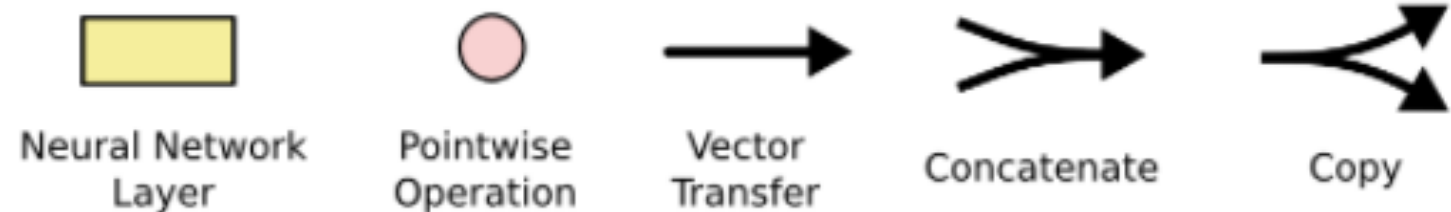


Long Short Term Memory

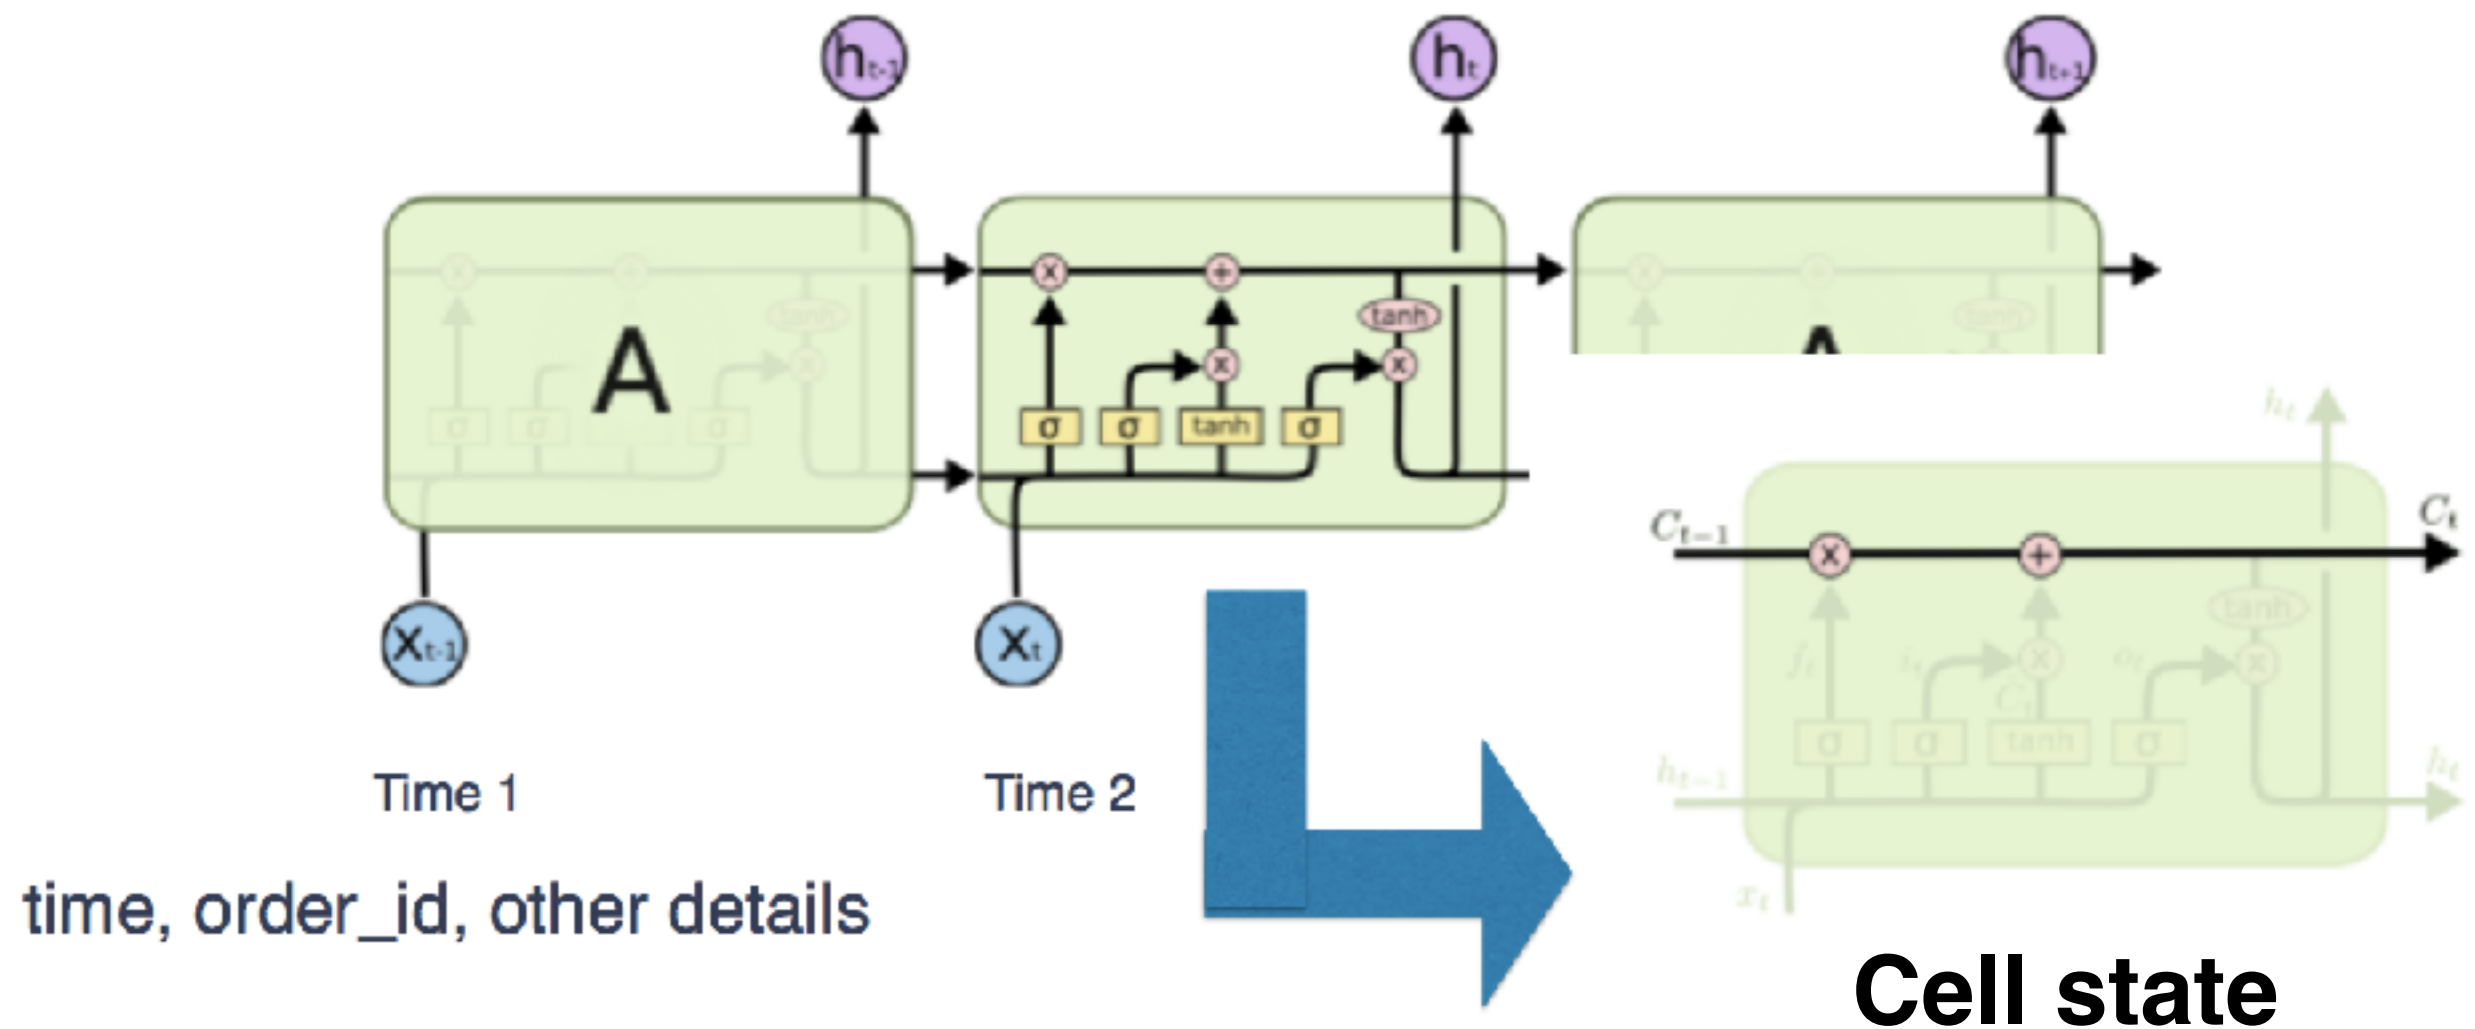


time, order_id, other details

time, order_id, other details

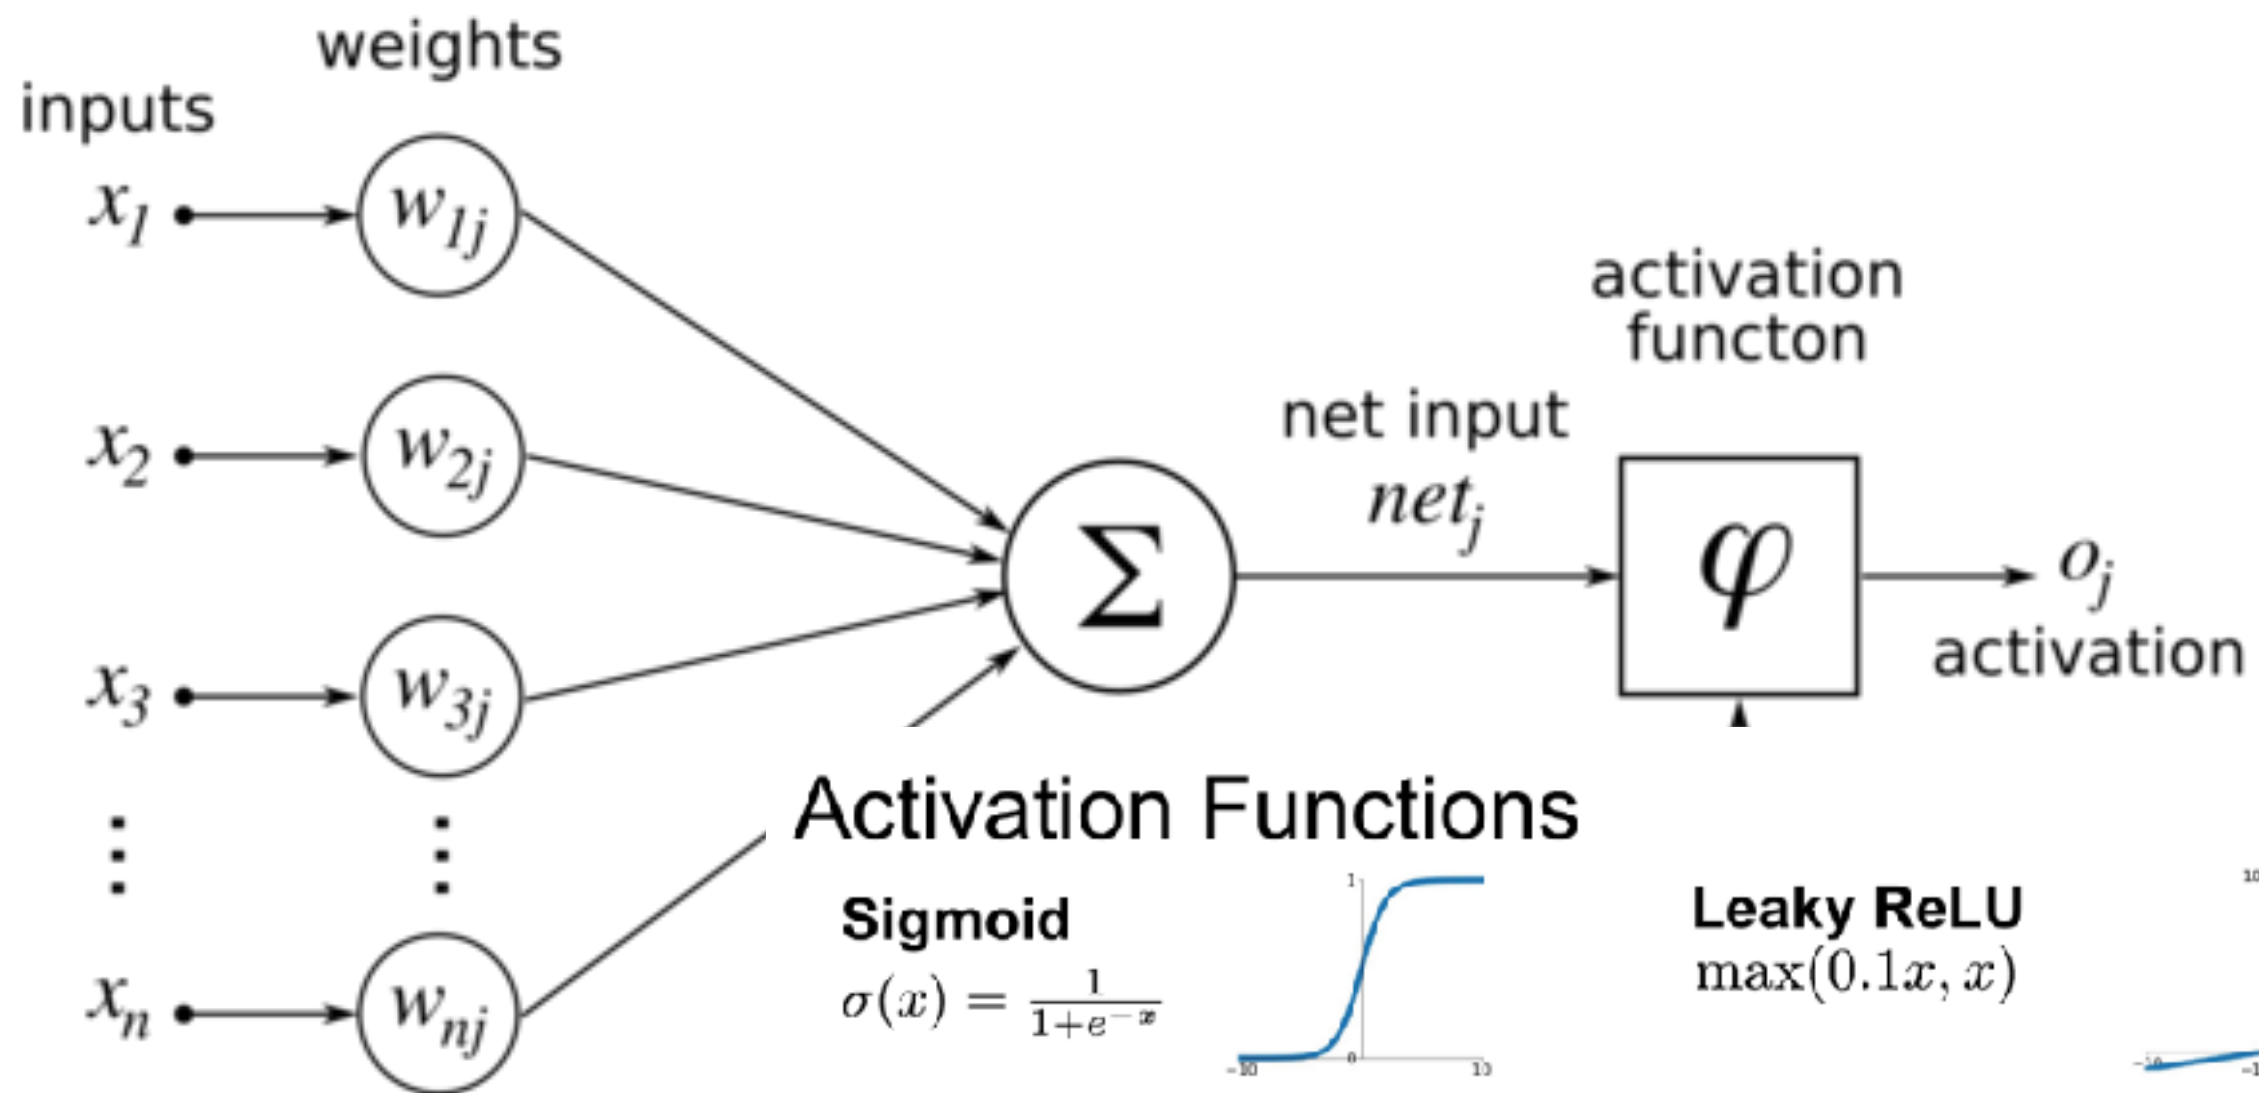


Long Short Term Memory



Forget, Update and Output

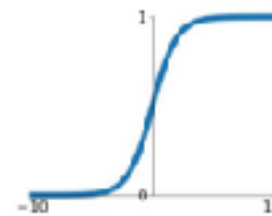
Activation function



Activation Functions

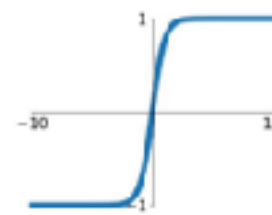
Sigmoid

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



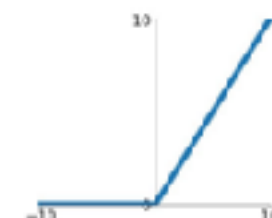
tanh

$$\tanh(x)$$



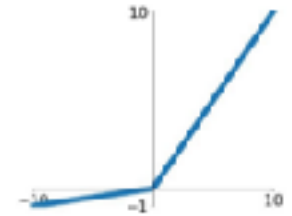
ReLU

$$\max(0, x)$$



Leaky ReLU

$$\max(0.1x, x)$$

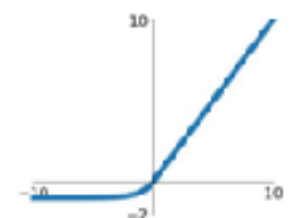


Maxout

$$\max(w_1^T x + b_1, w_2^T x + b_2)$$

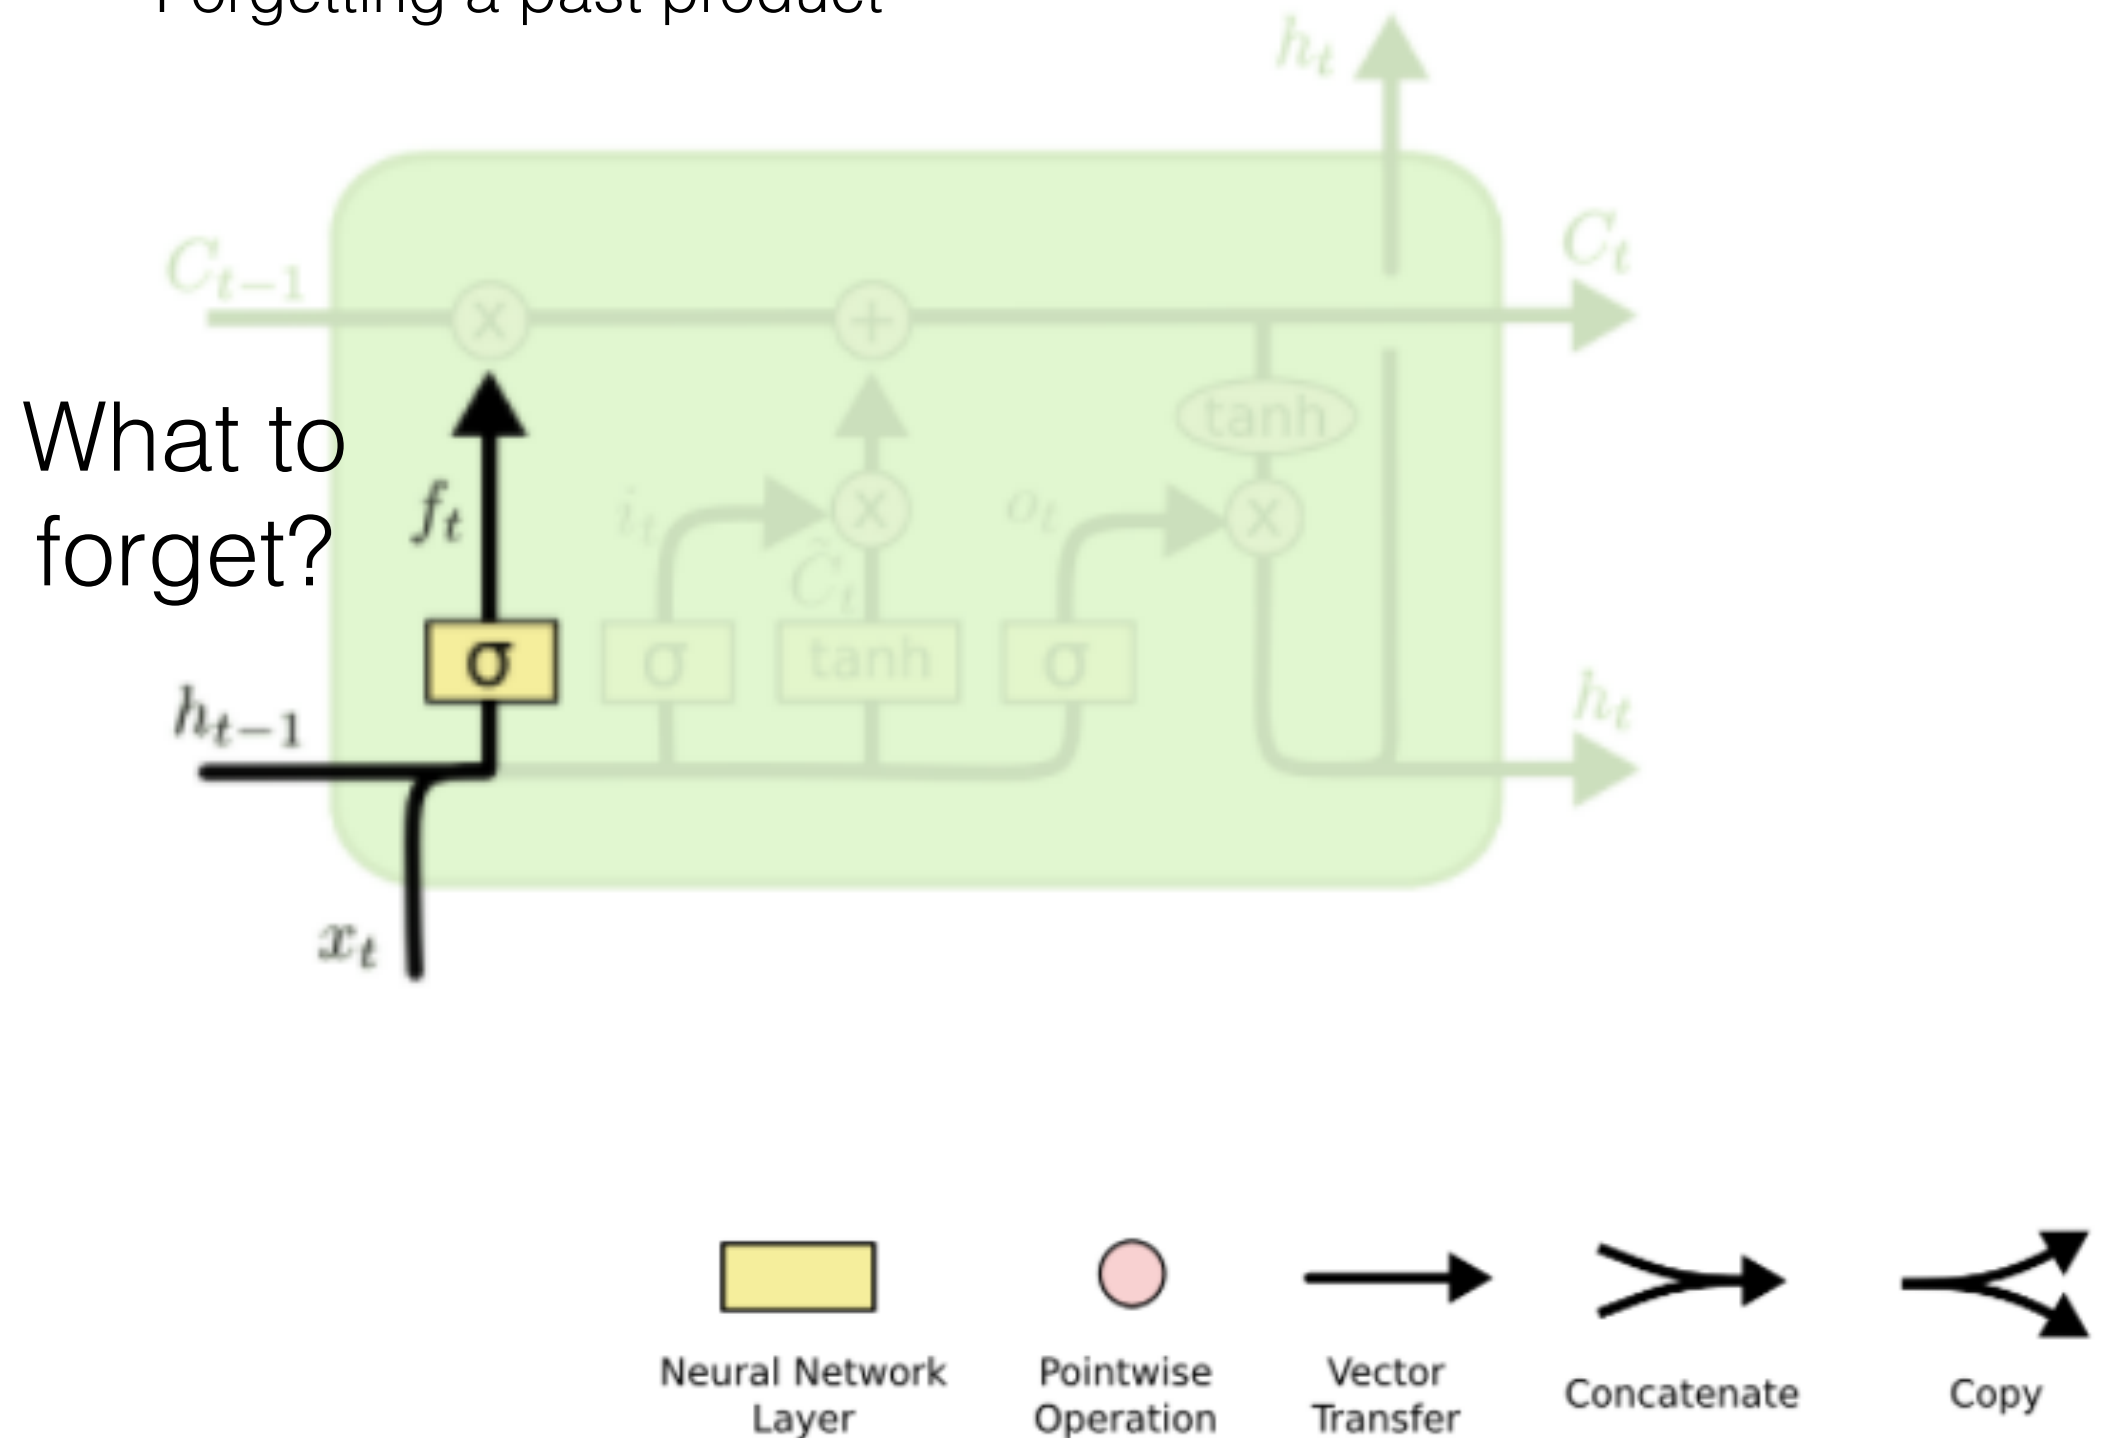
ELU

$$\begin{cases} x & x \geq 0 \\ \alpha(e^x - 1) & x < 0 \end{cases}$$

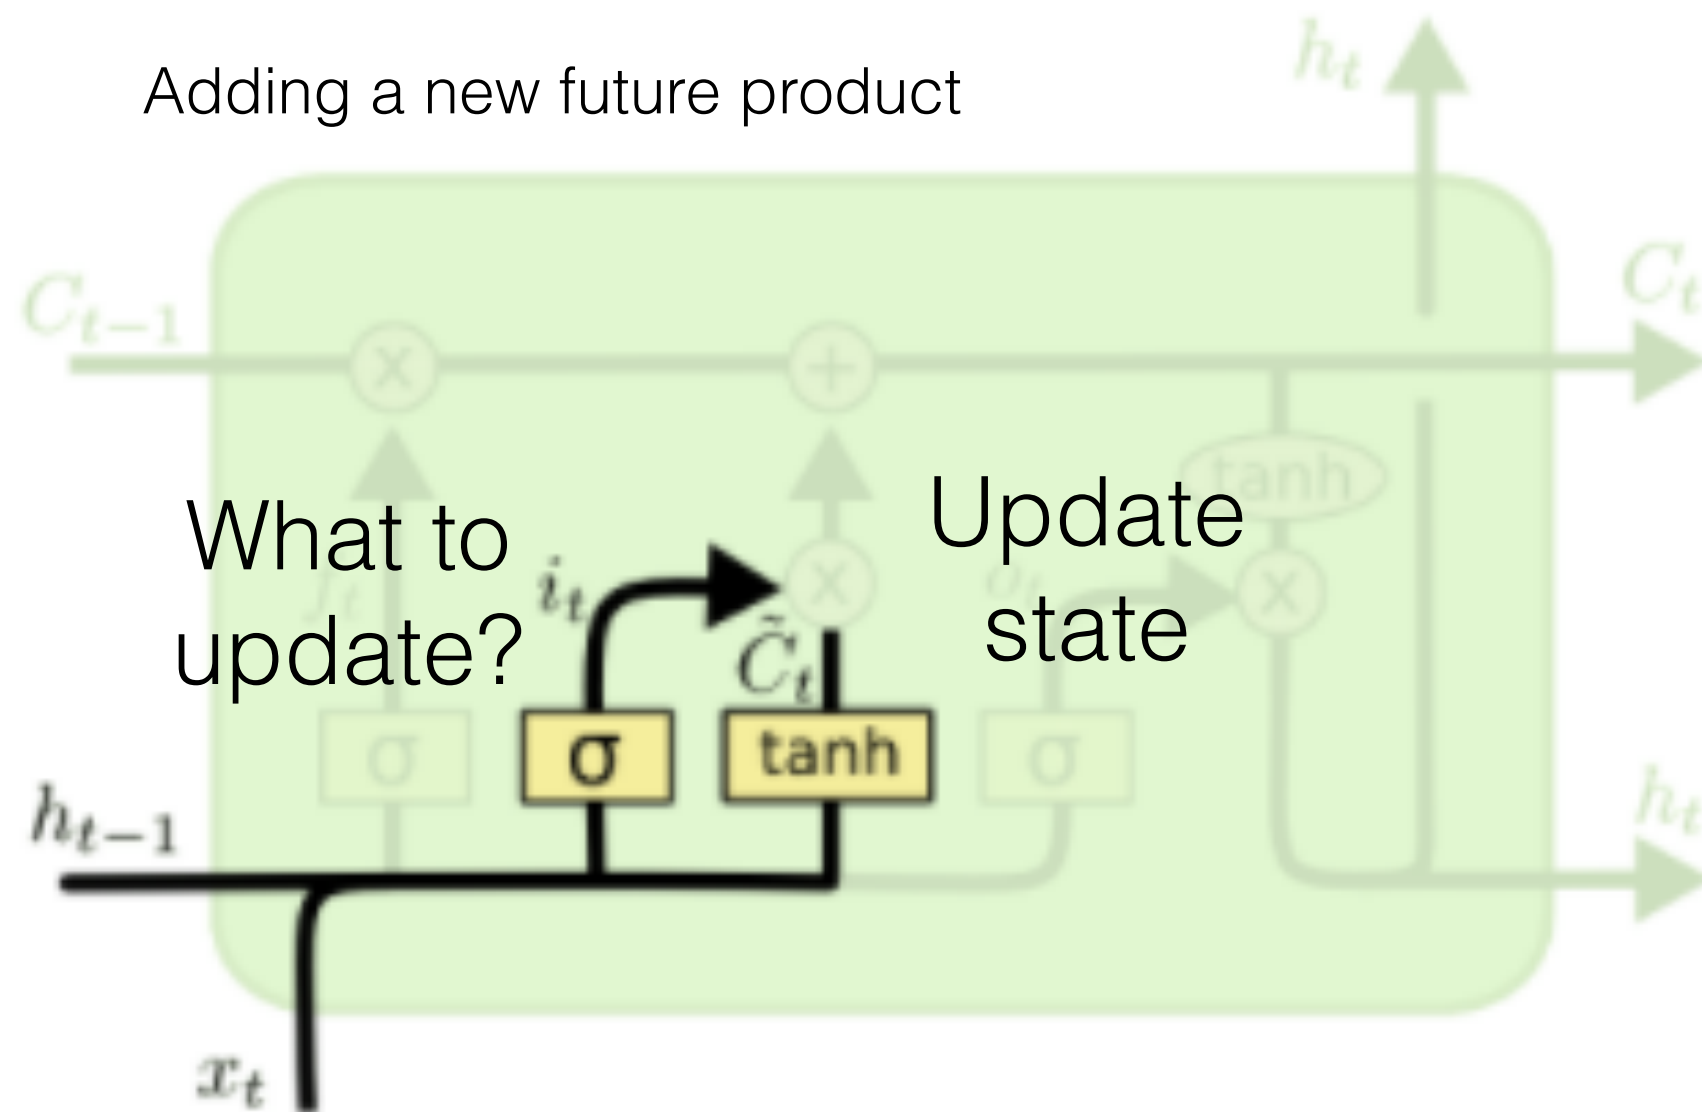


Long Short Term Memory cell - Forget gate

Forgetting a past product



Long Short Term Memory cell - Input gate



Neural Network
Layer

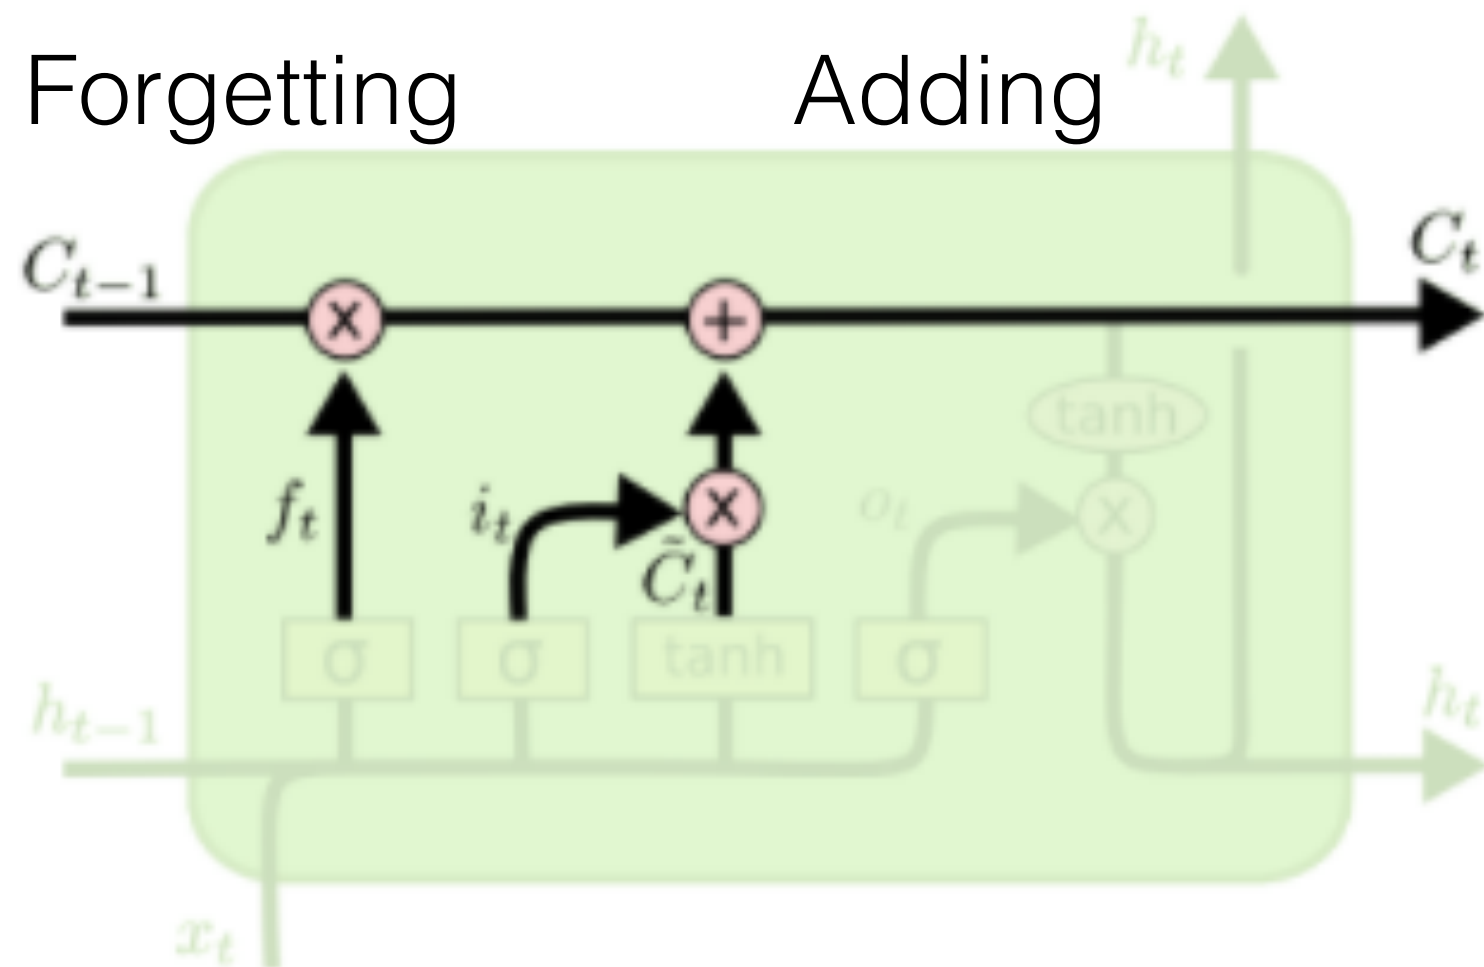
Pointwise
Operation


Vector
Transfer

Concatenate


Copy

Long Short Term Memory cell - New internal state




Neural Network
Layer

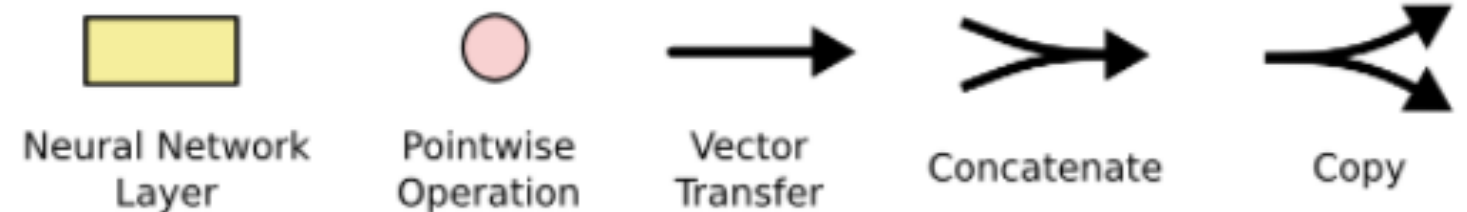
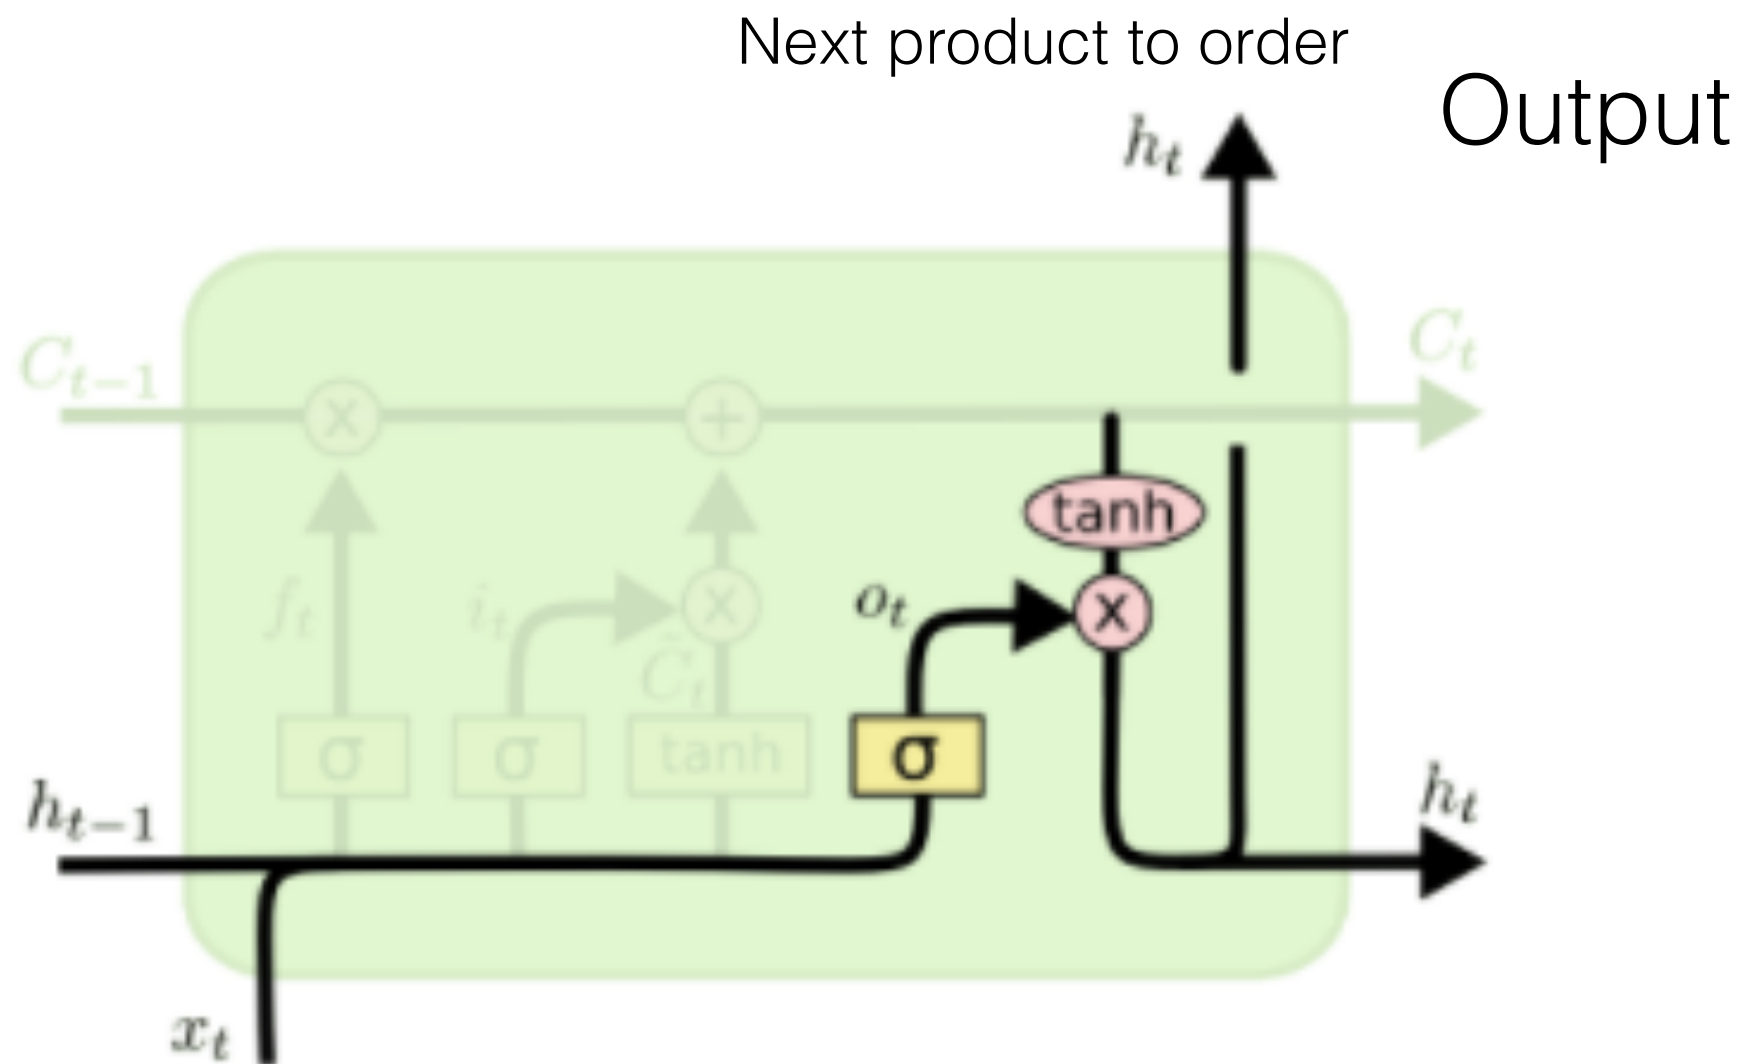

Pointwise
Operation

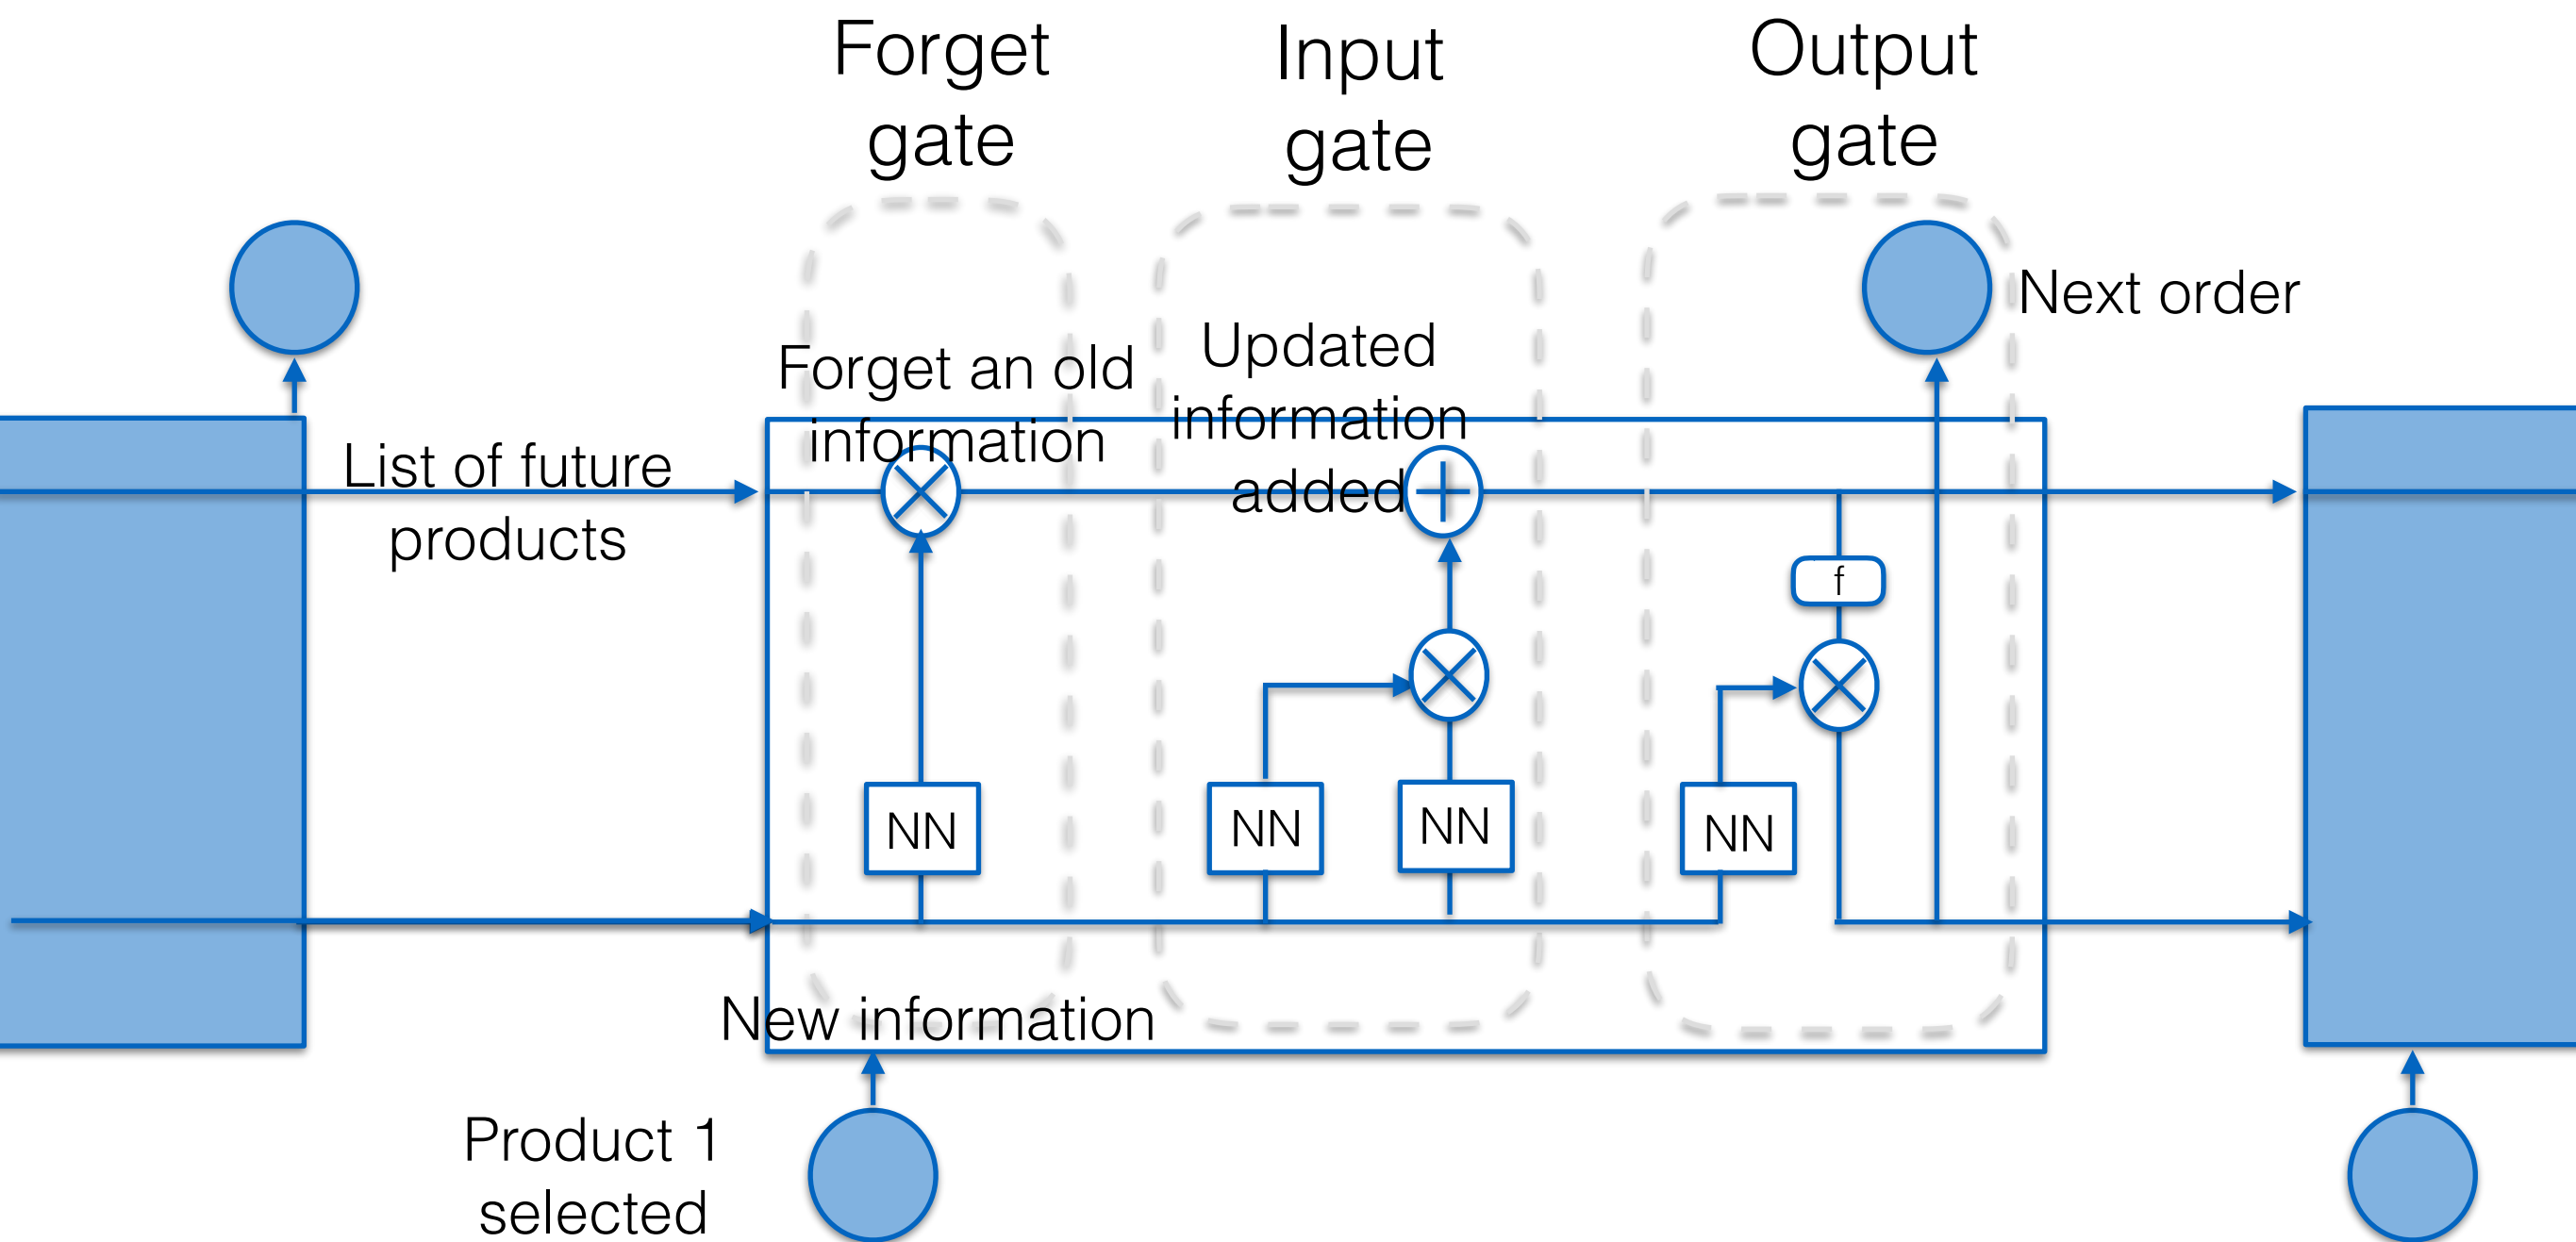

Vector
Transfer


Concatenate

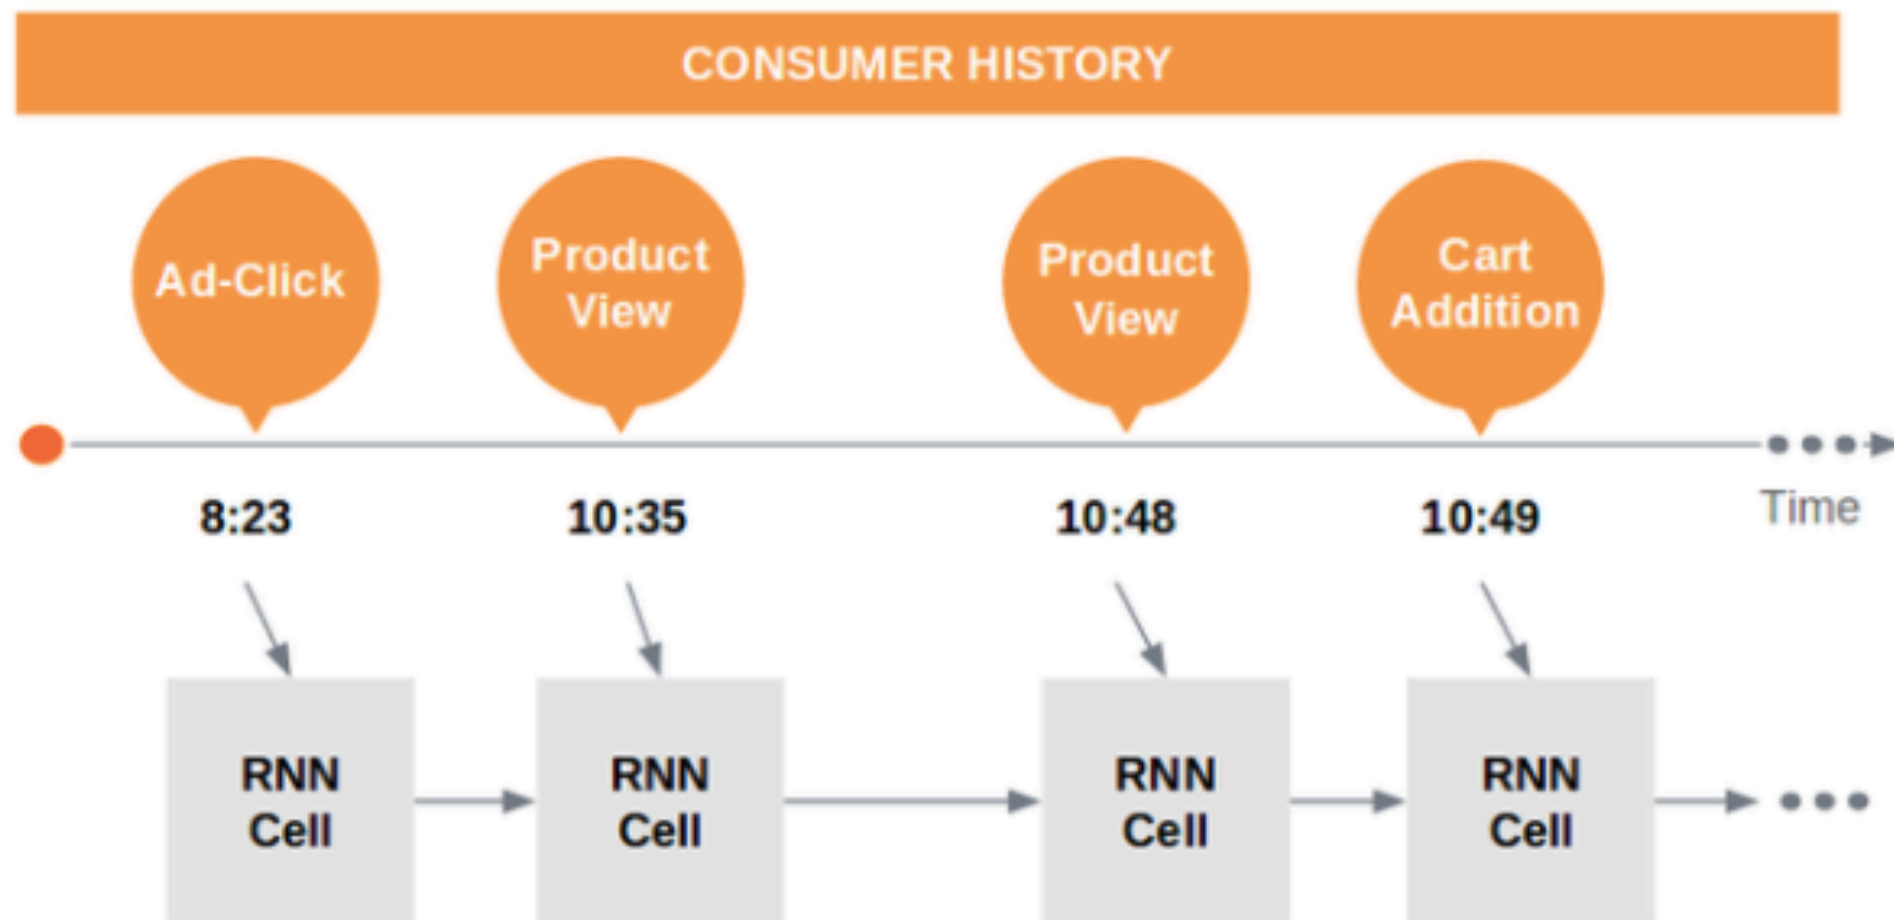

Copy

Long Short Term Memory cell - Output gate





Zolando's approach to Prediction



Future Products

Amazon Knows What You Want Before You Buy It

4.7k
SHARES



Share



Share



THE WALL STREET JOURNAL
Home World U.S. Politics Economy **Business** Tech Markets Opinion Arts Life Real Estate

MULTIPLIERS OF PROSPERITY



MetLife Foundation | WSI SYSTEMS

A nudge, not a nag,
toward savings

LEARN MORE

BUSINESS

Wal-Mart Wants to Know When Your Milk Is About to Expire

The retailer is thinking about using sensor technologies to trigger automatic delivery or suggestions for related products to buy, according to a patent application

By *Sarah Nassauer*

May 4, 2017 6:32 p.m. ET

Wal-Mart Stores Inc. is thinking about using sensor technologies to



Tensorflow LSTM API

TensorFlow™

Install

Develop

API r1.4

Deploy

Extend

Community

Versions

IFRC

Search

API r1.4

tf.contrib.reduce_slice_ops.python.
ops

tf.contrib.remote_fused_graph

tf.contrib.resampler

tf.contrib.rnn

Overview

AttentionCellWrapper

BasicLSTMCell

BasicRNNCell

BidirectionalGridLSTMCell

CompiledWrapper

Conv1DLSTMCell

Conv2DLSTMCell

Conv3DLSTMCell

ConvLSTMCell

CoupledInputForgetGateLSTM...

DeviceWrapper

DropoutWrapper

EmbeddingWrapper

FusedRNNCell

FusedRNNCellAdaptor

tf.contrib.rnn.BasicLSTMCell

Class **BasicLSTMCell**

Inherits From: [RNNCell](#)

Aliases:

- Class `tf.contrib.rnn.BasicLSTMCell`
- Class `tf.nn.rnn_cell.BasicLSTMCell`

Defined in [tensorflow/python/ops/rnn_cell_impl.py](#).

See the guide: [RNN and Cells \(contrib\) > Core RNN Cells for use with TensorFlow's core RNN methods](#)

Basic LSTM recurrent network cell.

The implementation is based on: <http://arxiv.org/abs/1409.2329>.

We add `forget_bias` (default: 1) to the biases of the forget gate in order to reduce the scale of forgetting in the beginning of the training.

Use-case feature for using LSTM

- **Repeated patterns of human action done over longer duration**

Any other use case?

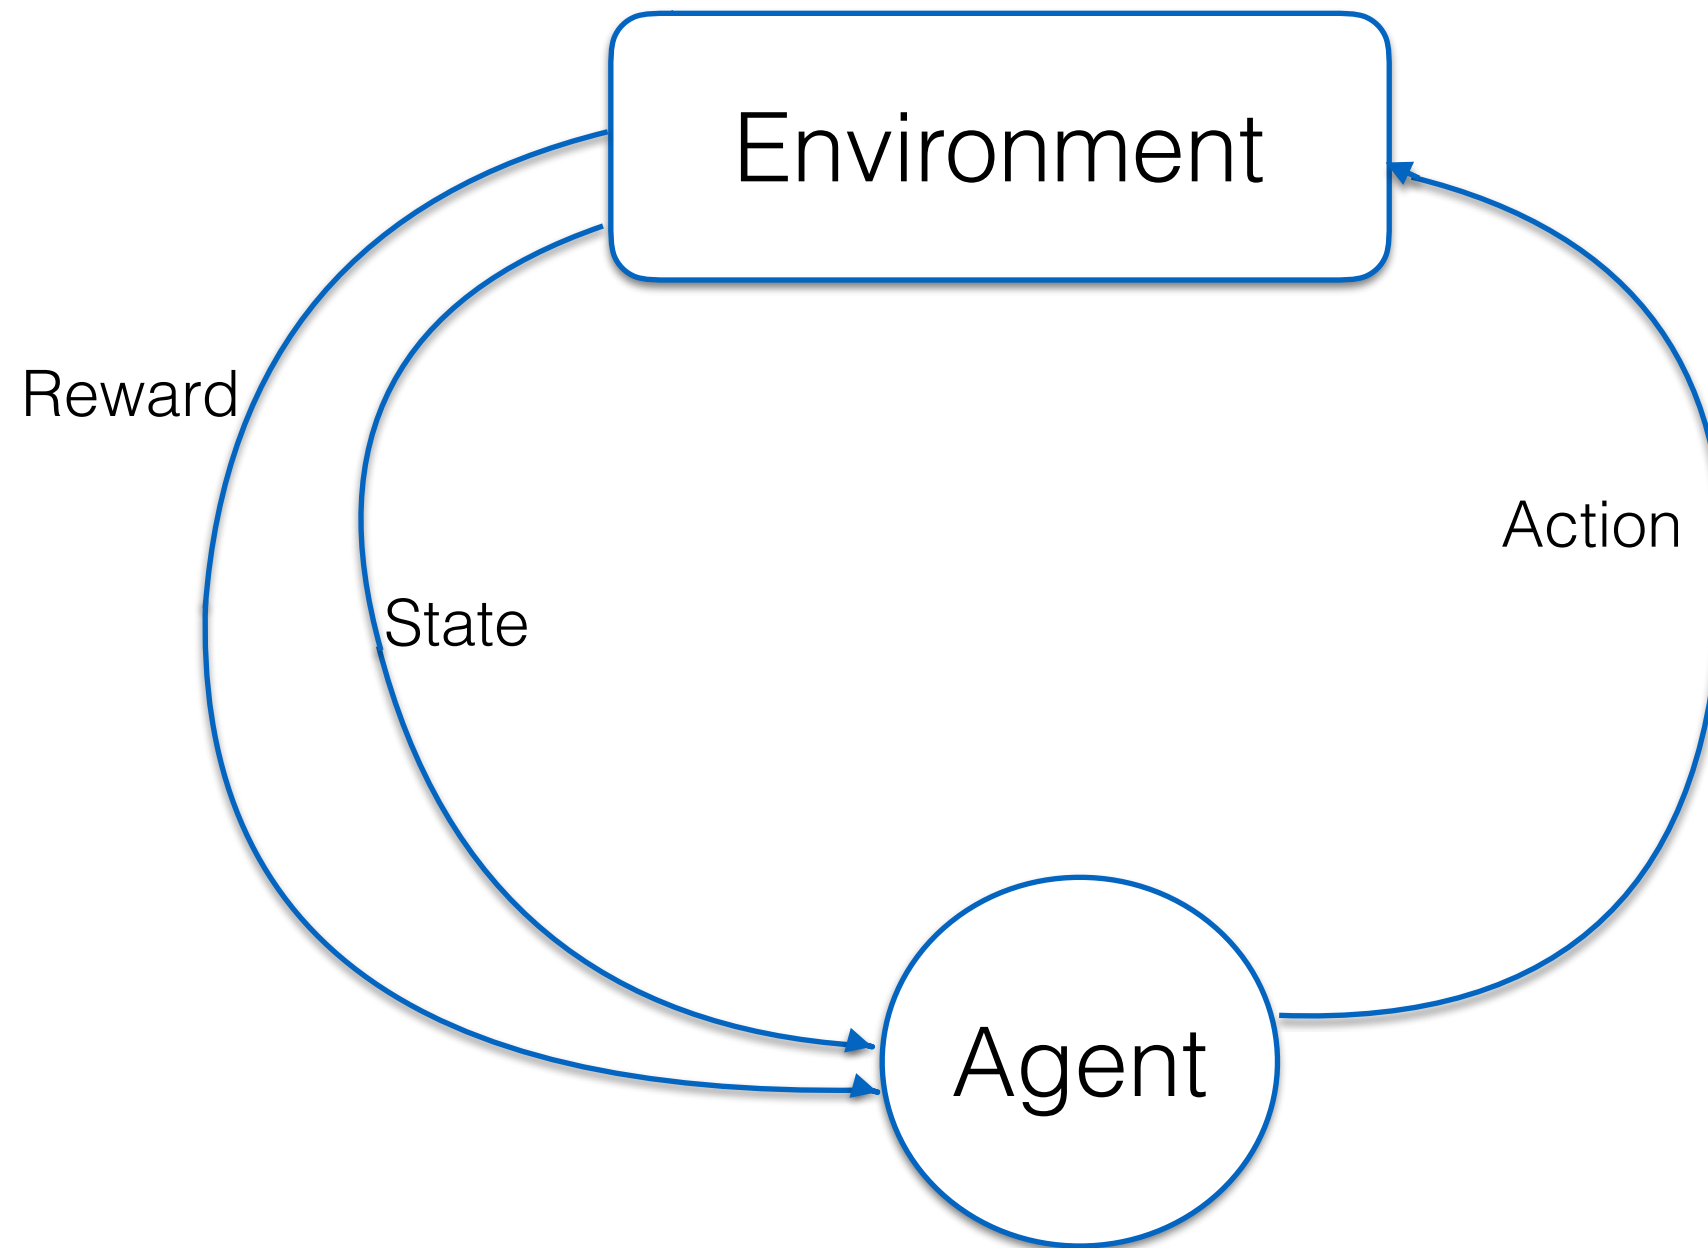
What is missing?

Training Data



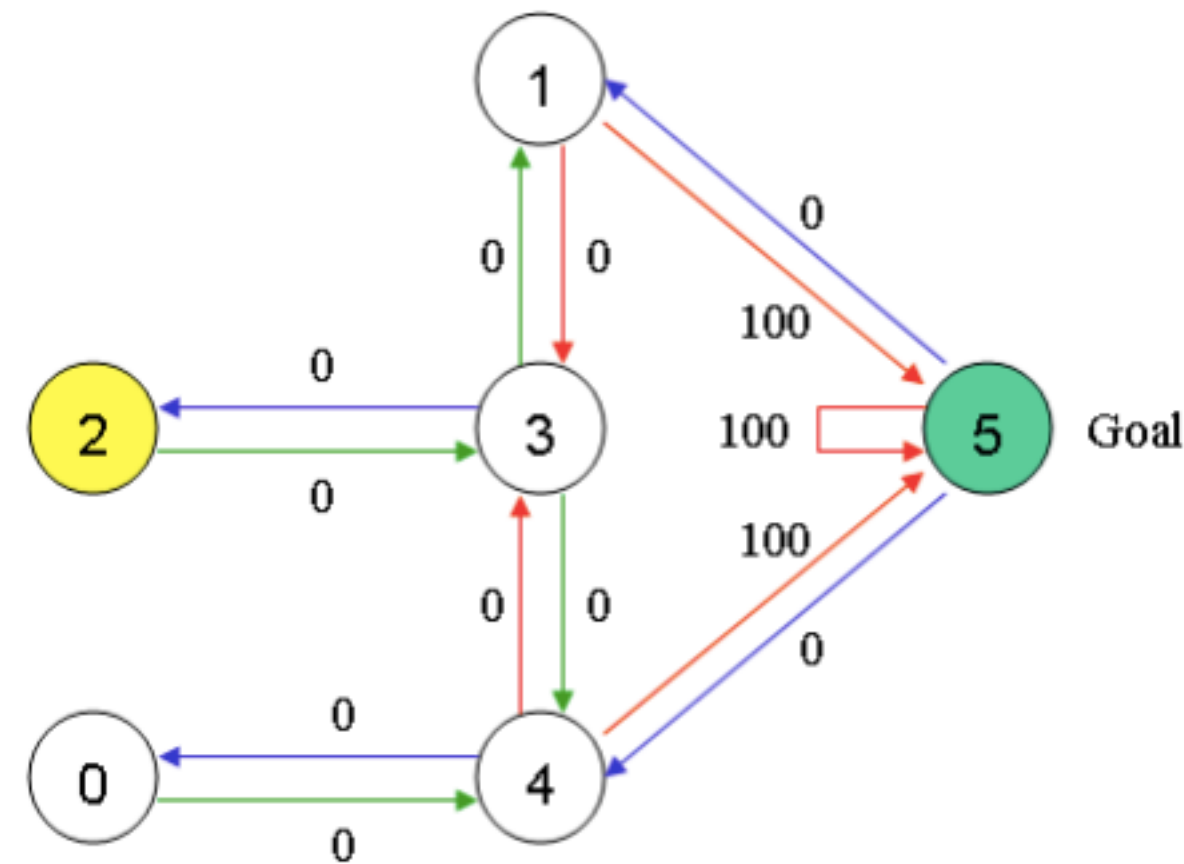
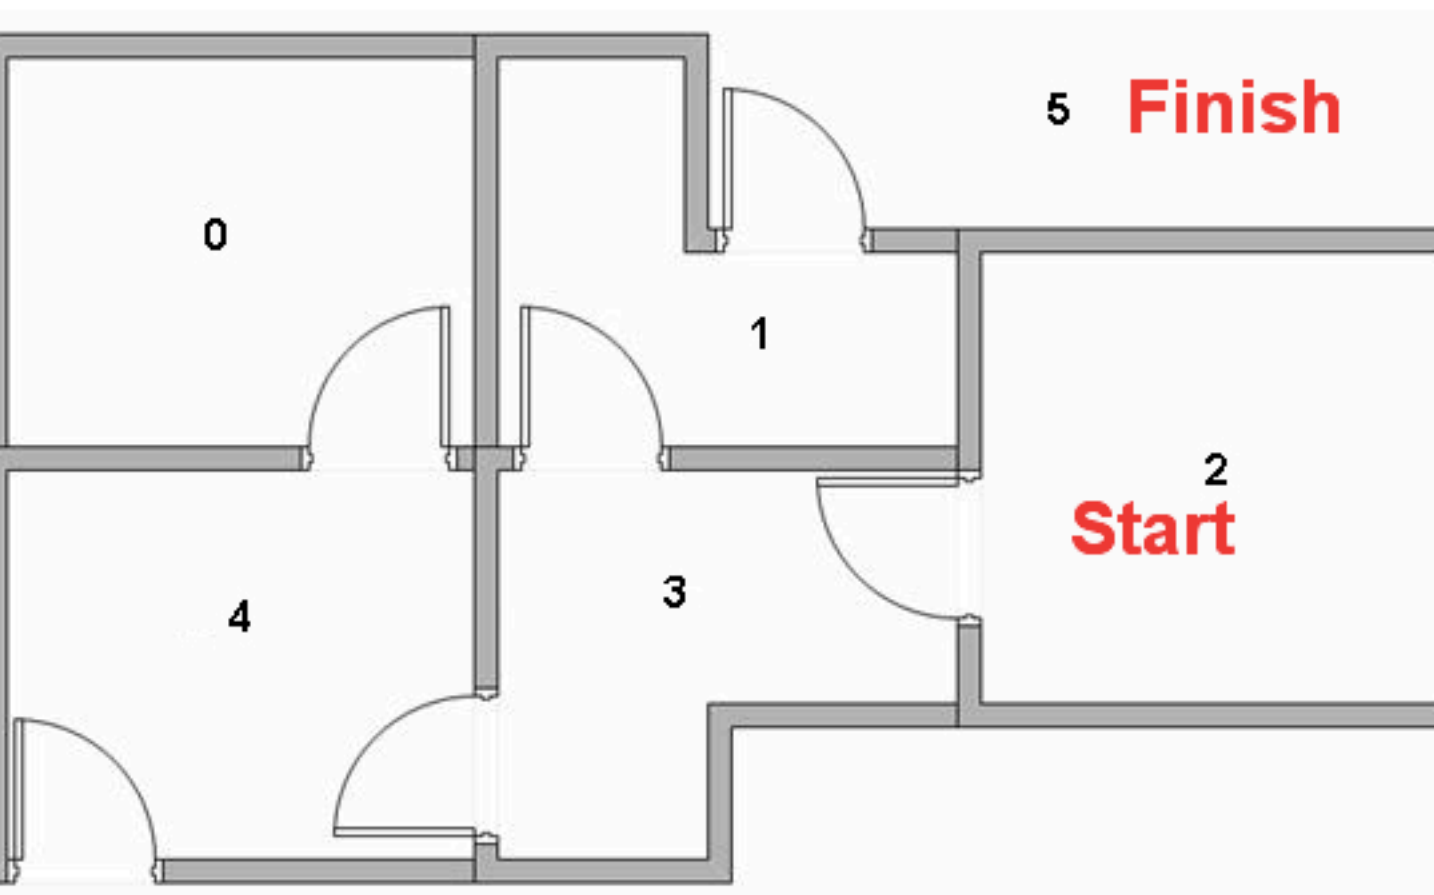
No Data Available

Reinforcement learning



Q-learning

- Reward function for each action in a given state
- $Q(s,a)=R(s,a) + \gamma * \max [Q(\text{next states}, \text{all actions})]$
- Initial Q is 0
- Value of γ is (0,1) depends how much you want future actions to influence current learning.



$$Q(s,a)=R(s,a) + \gamma * \max [Q(\text{next states, all actions})]$$

If $\gamma = 0.8$

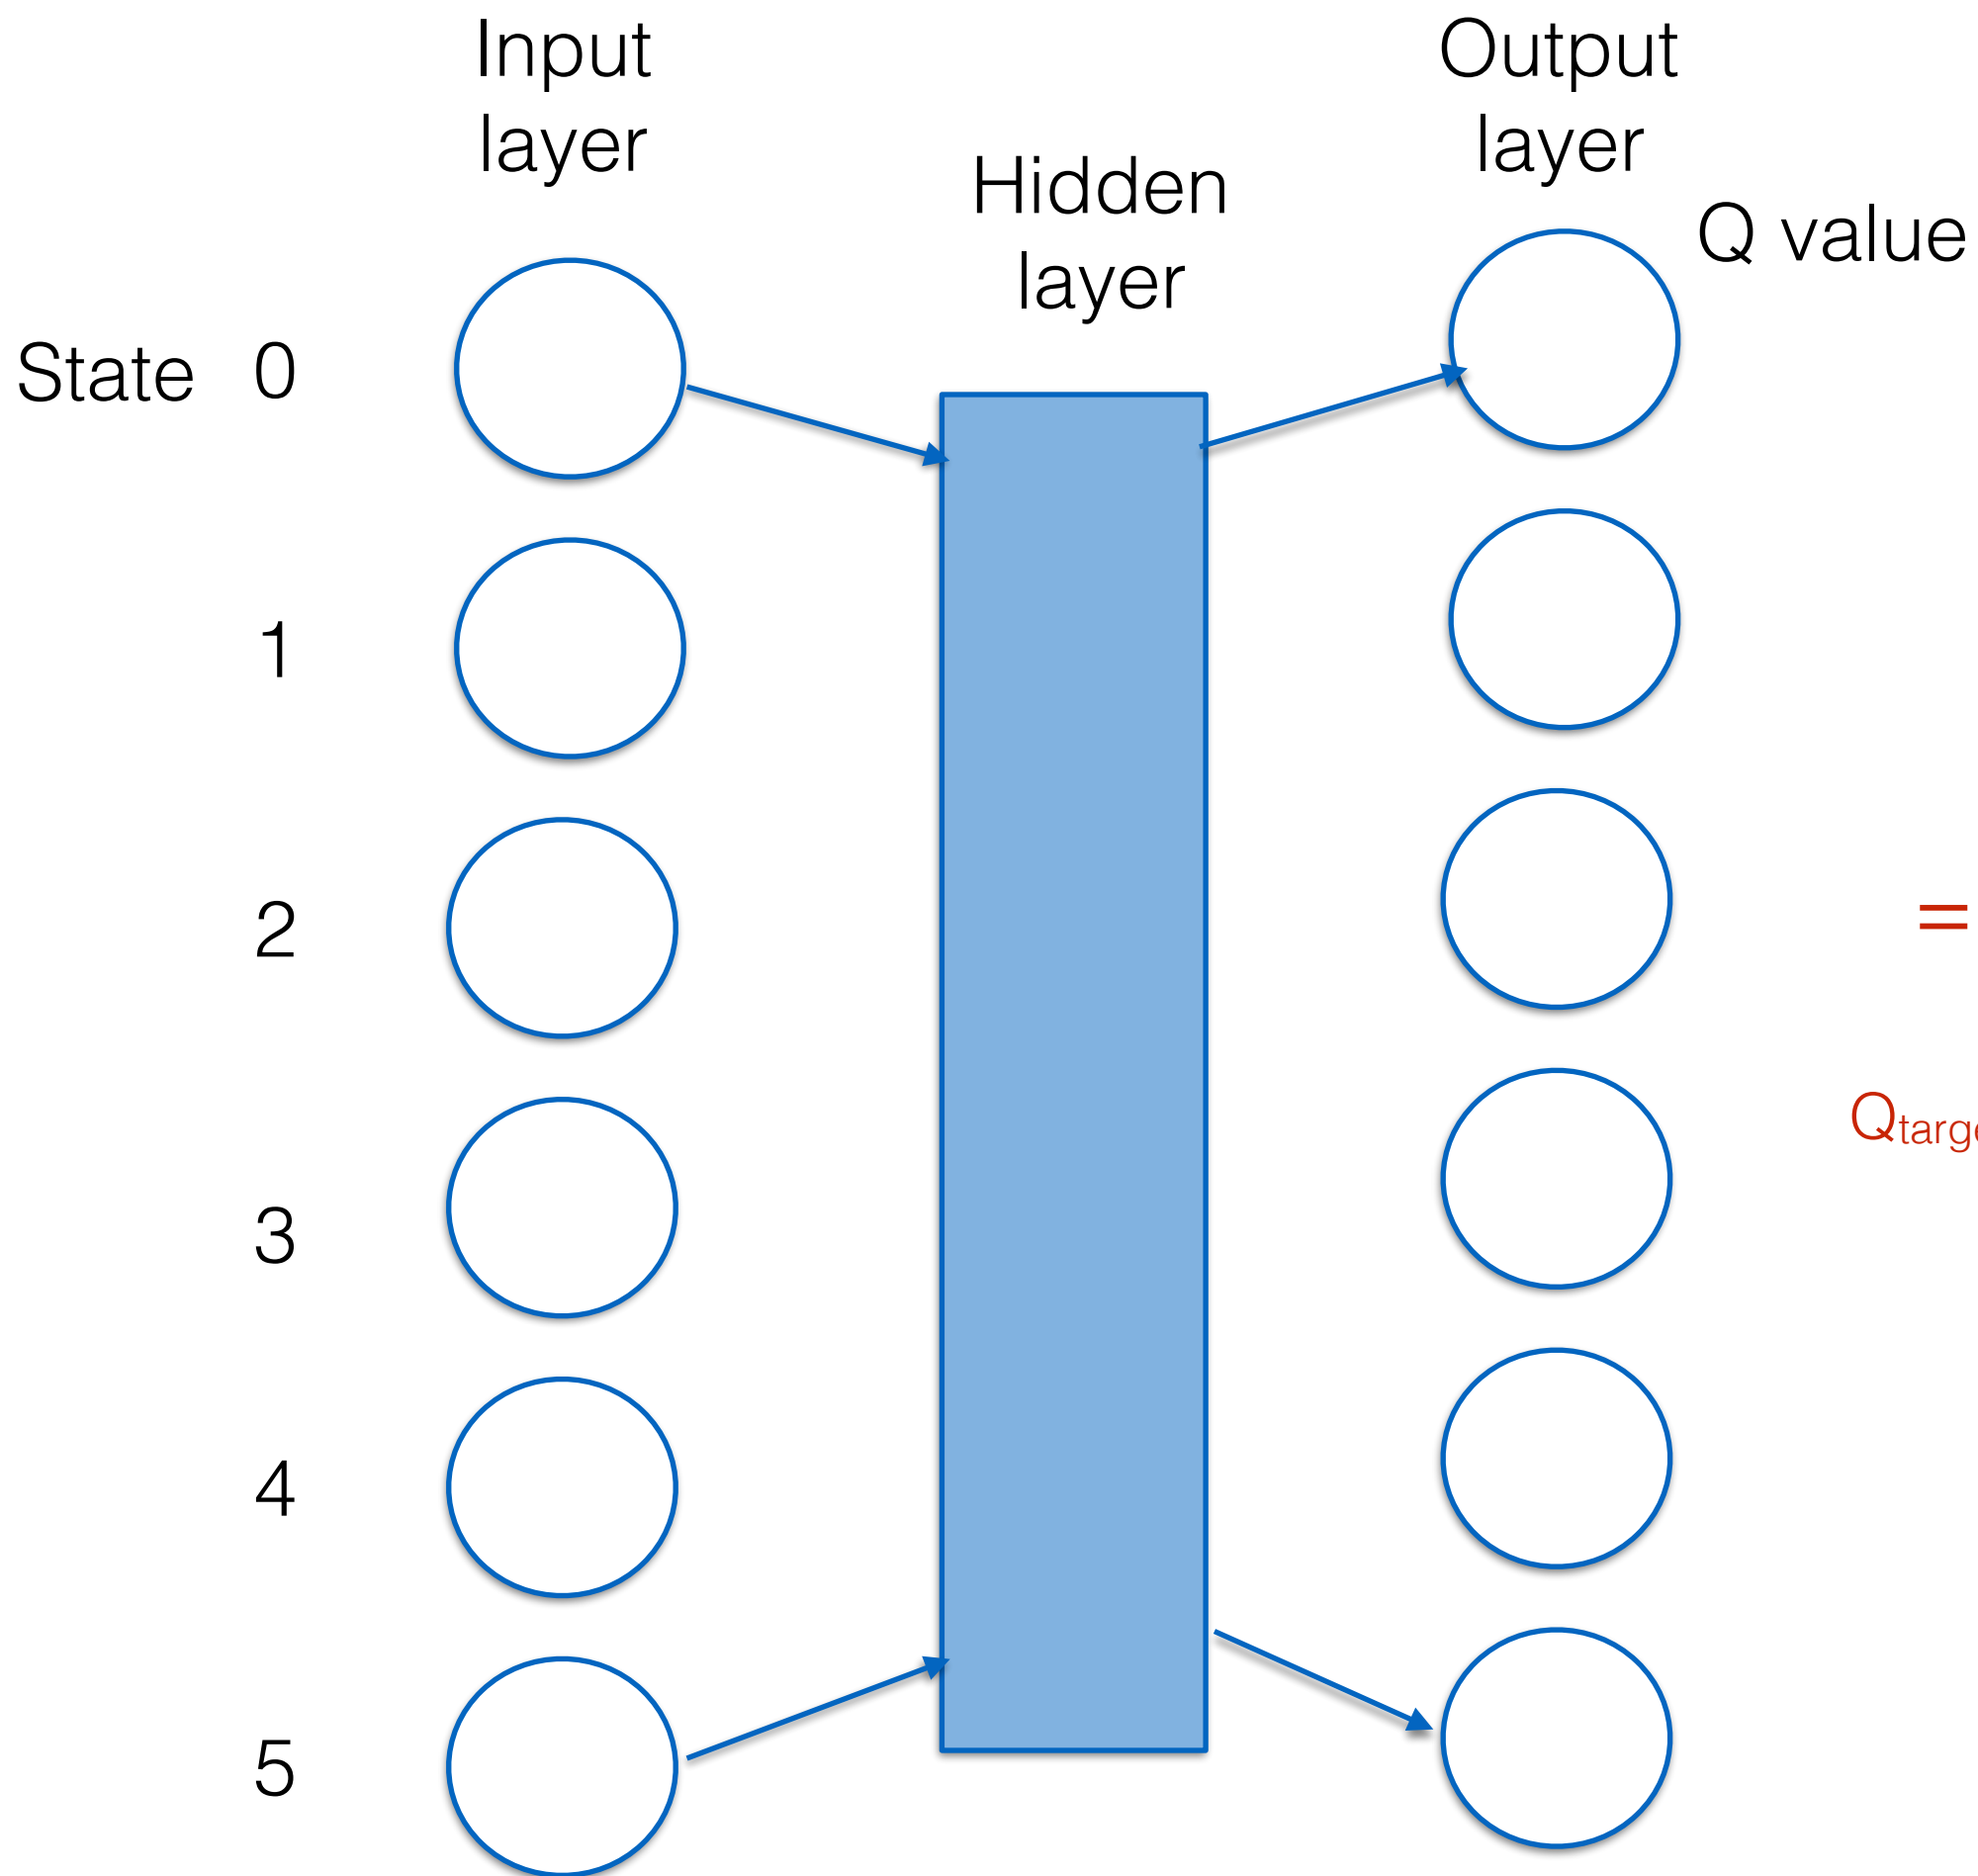
$$Q(s,a)=R(s,a) + 0.8 * \max [Q(\text{next states, actions})]$$

eg

$$Q(3,1)=0 + 0.8 [\max(Q(1,3),Q(1,5))] = 400$$

	Action					
State	0	1	2	3	4	5
0	-1	-1	-1	-1	0	-1
1	-1	-1	-1	0	-1	100
2	-1	-1	-1	0	-1	-1
3	-1	0	0	-1	0	-1
4	0	-1	-1	0	-1	100
5	-1	0	-1	-1	0	100

	0	1	2	3	4	5
0	0	0	0	0	400	0
1	0	0	0	320	0	500
2	0	0	0	320	0	0
3	0	400	256	0	400	0
4	320	0	0	320	0	500
5	0	400	0	0	400	500

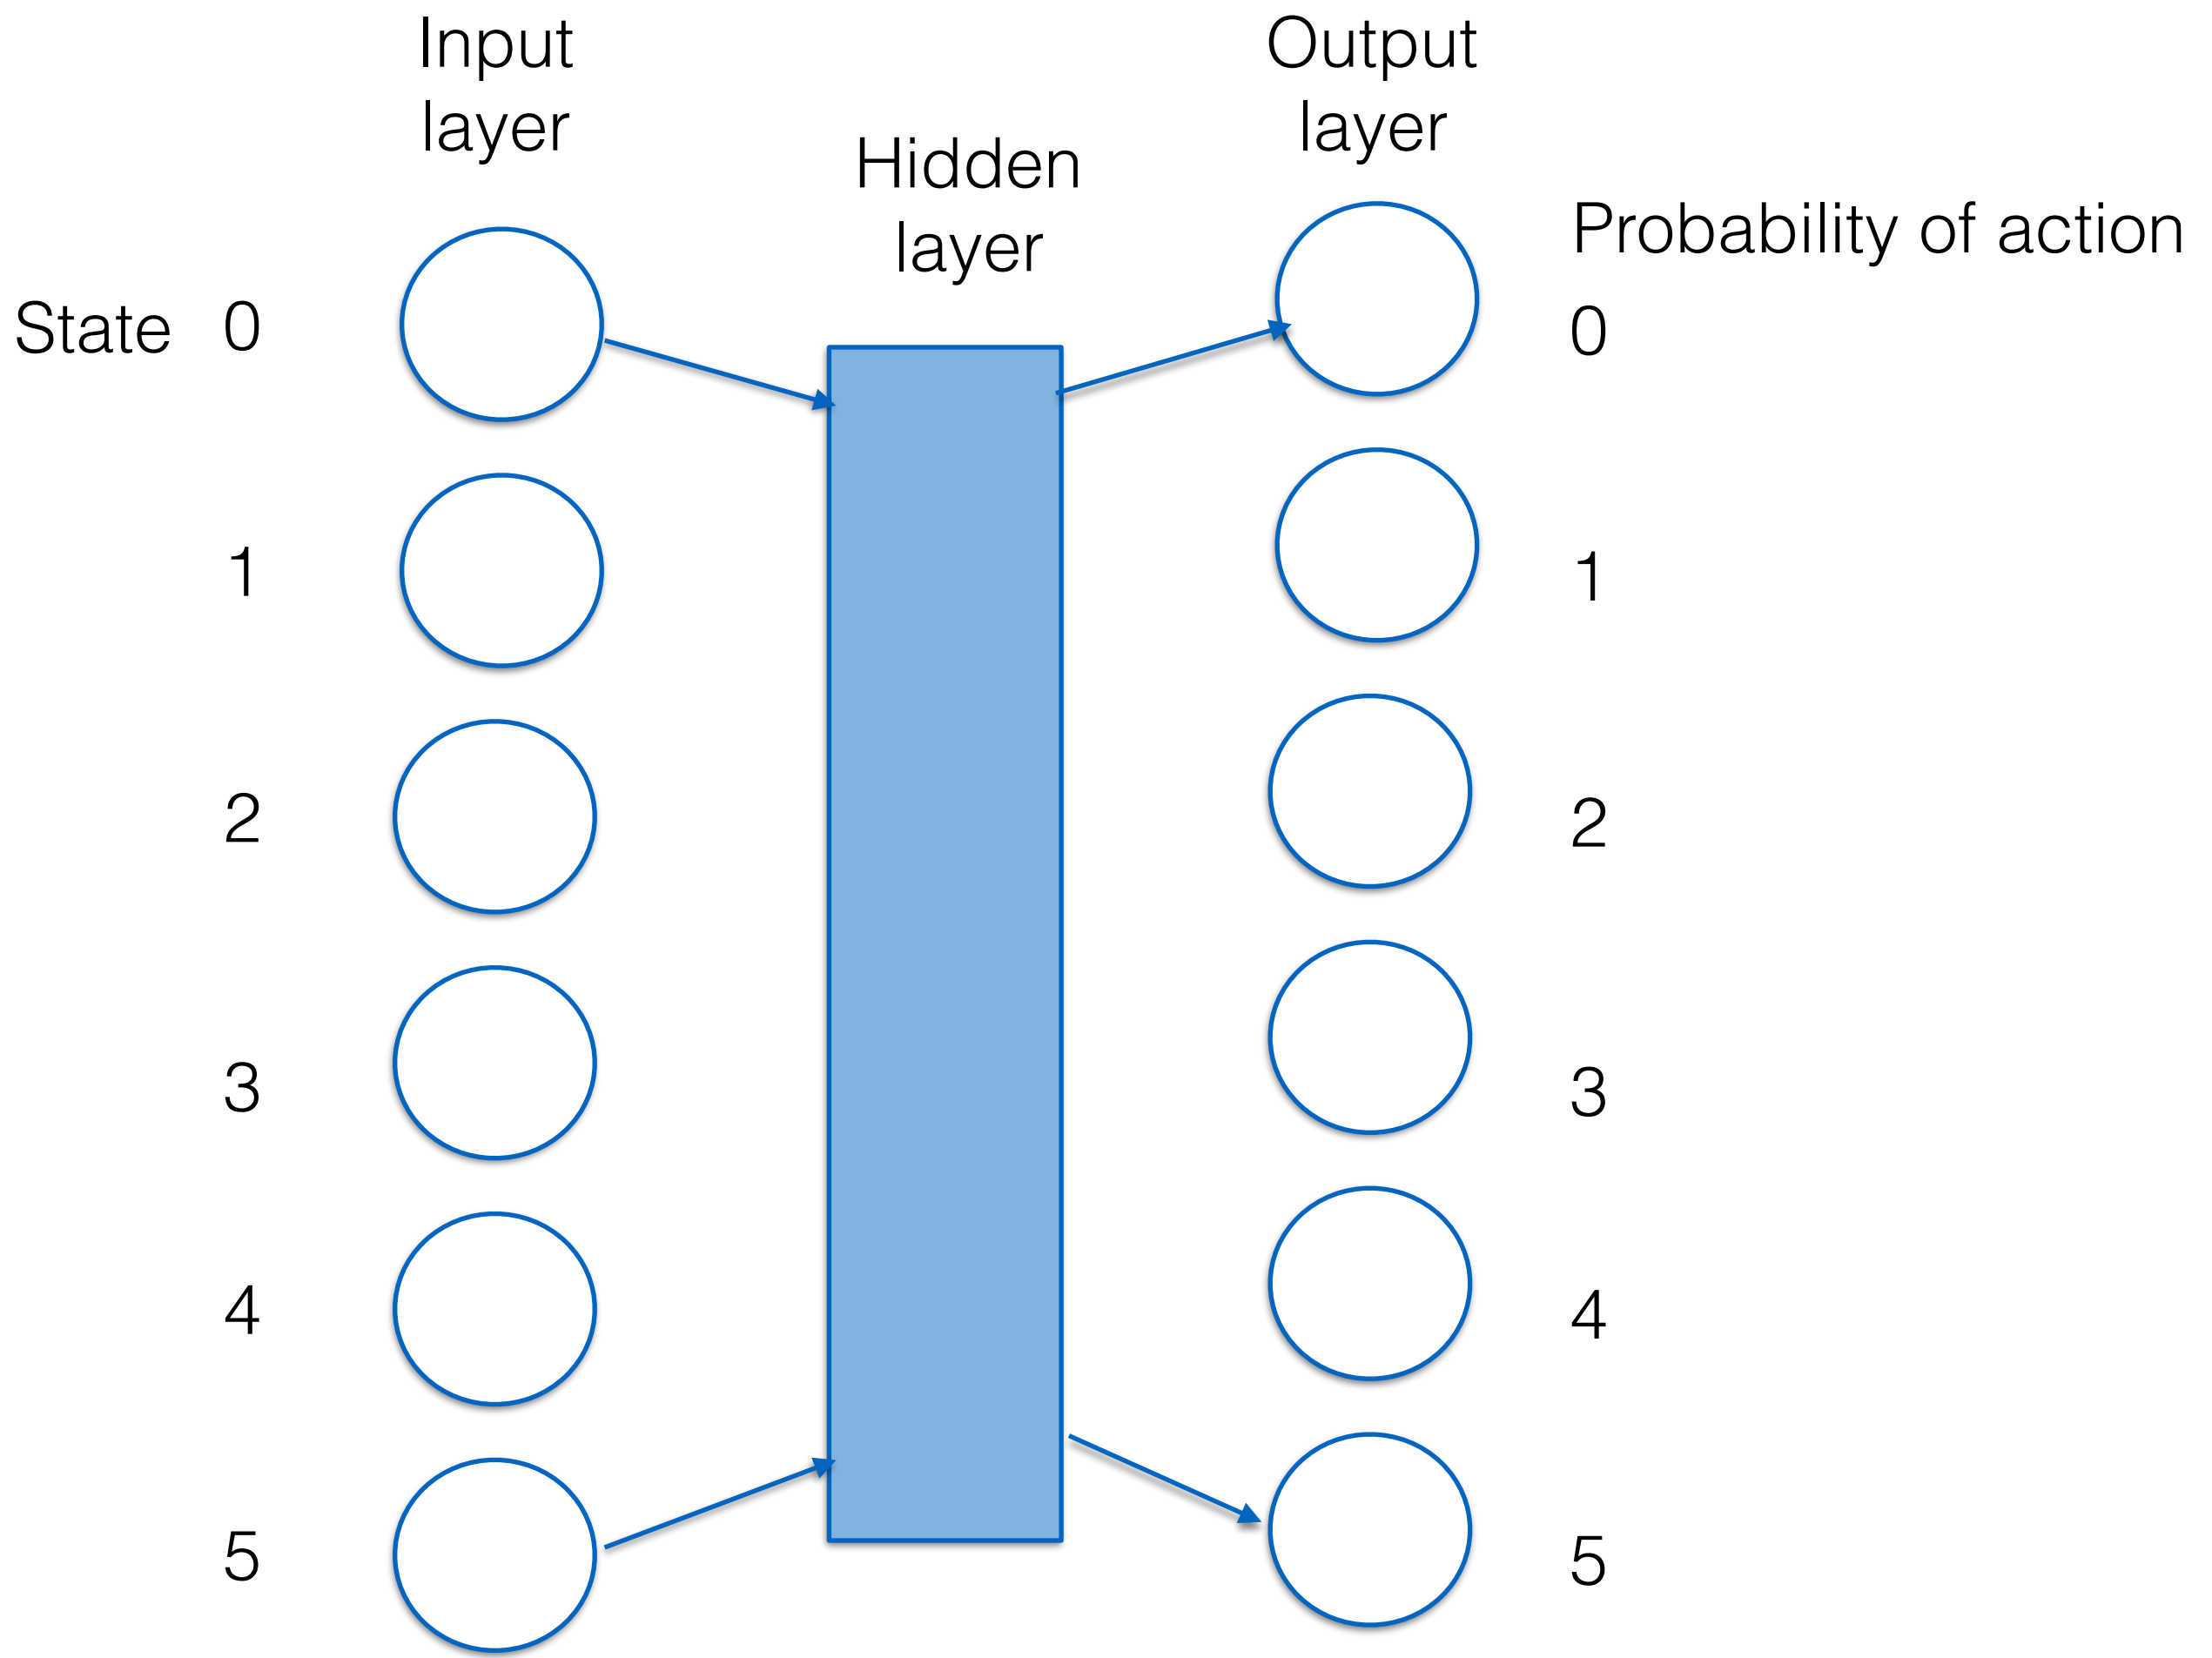


$$\text{Loss} = \sum (Q_{\text{target}i} - Q_i)$$

$$Q_{\text{target}i} = r + \gamma * \max(Q(s', a'))$$

Policy Gradient

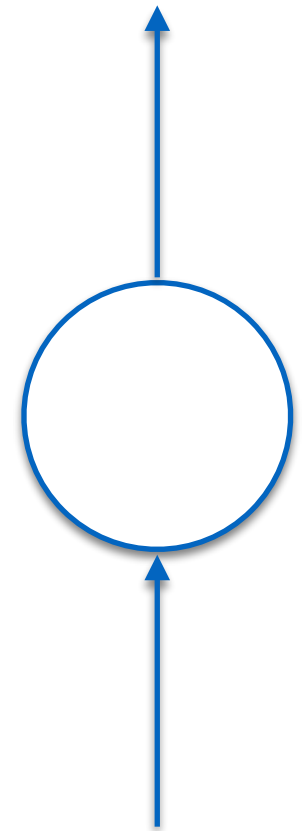
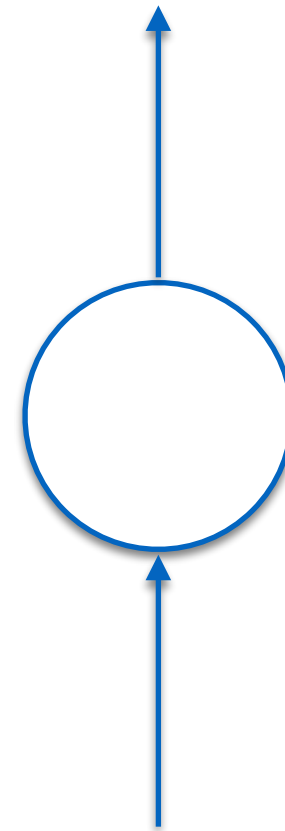
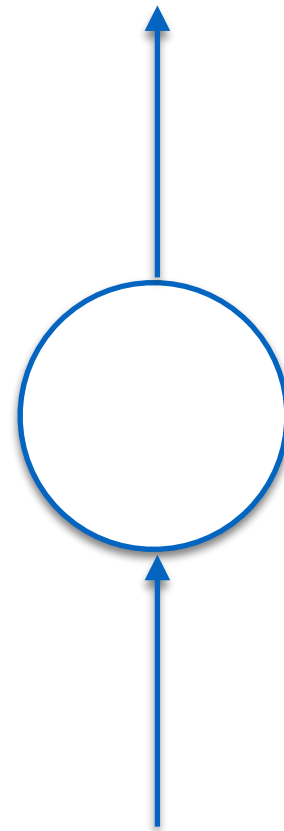
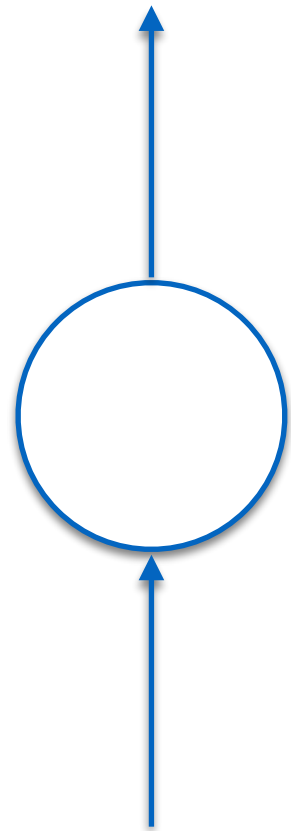
- Optimizes the policy space so the Neural network directly models the action space
- More complex problems esp. continuous action space



Predicting next order

Action

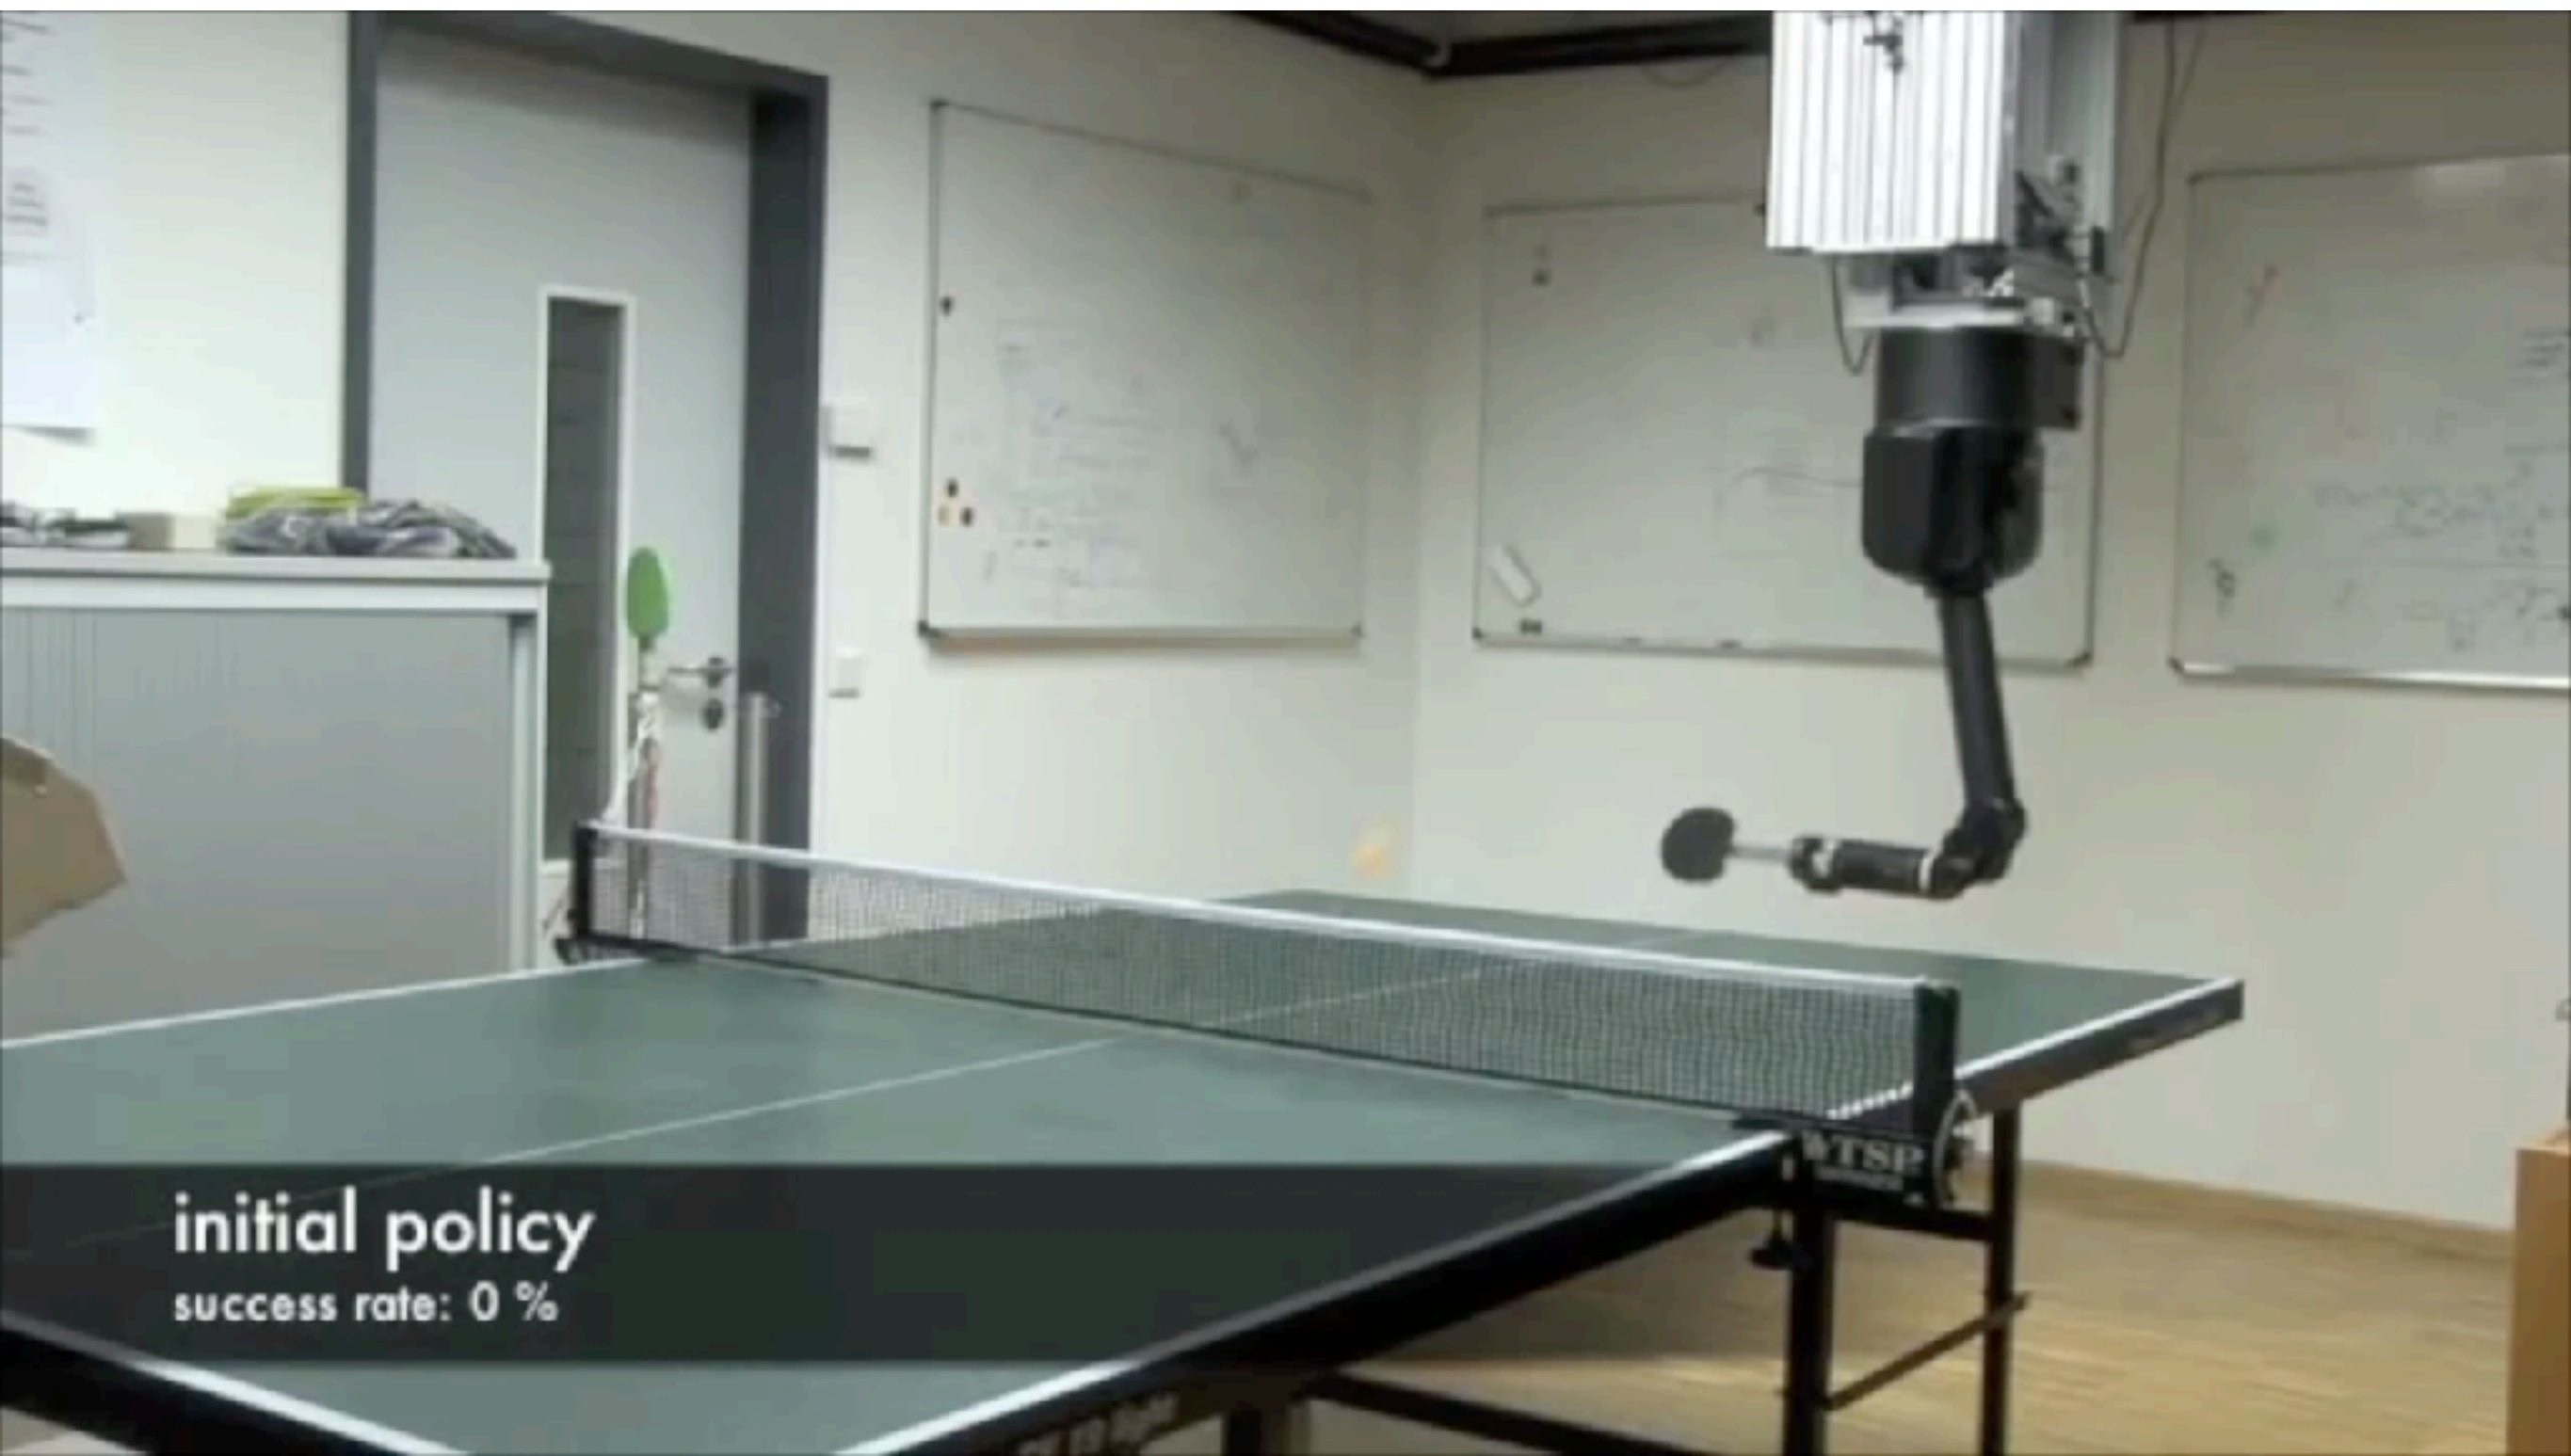
Product 5



State

Product 1

Product 3



initial policy
success rate: 0 %

Use-case feature for using RL

- **Hard to get past data**

Any other use case?

Implementation

📖 README.md

Tensorflow-Reinforce

A collection of [Tensorflow](#) implementations of reinforcement learning models. Models are evaluated in [OpenAI Gym](#) environments. Any contribution/feedback is more than welcome. **Disclaimer:** These implementations are used for educational purposes only (i.e., to learn deep RL myself). There is no guarantee that the exact models will work on any of your particular RL problems without changes.

Environments

This codebase works in both Python 2.7 and 3.5. The models are implemented in [Tensorflow 1.0](#).

Models

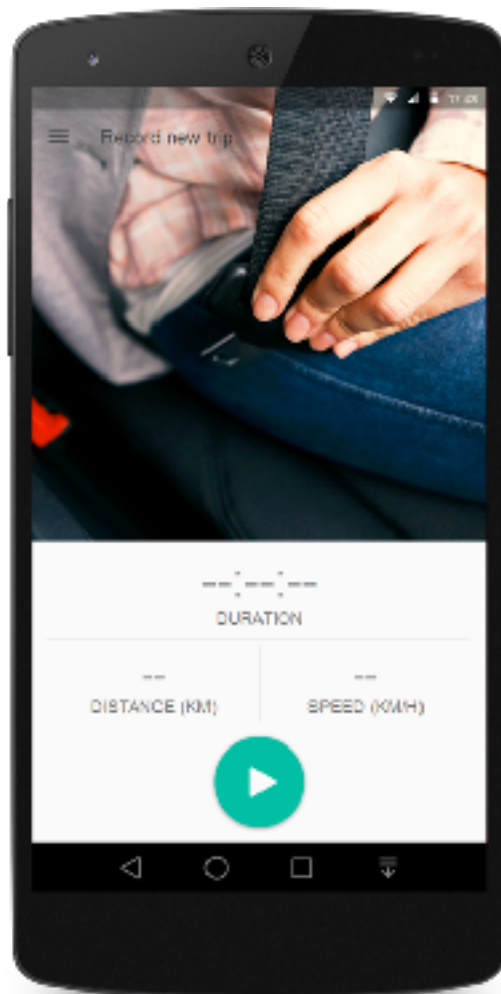
Model	Code	References
Cross-Entropy Method	run_cem_cartpole	Cross-entropy method
Tabular Q Learning	rl/tabular_q_learner	Sutton and Barto, Chapter 8
Deep Q Network	rl/neural_q_learner	Mnih et al.
Double Deep Q Network	rl/neural_q_learner	van Hasselt et al.
REINFORCE Policy Gradient	rl/pg_reinforce	Sutton et al.
Actor-critic Policy Gradient	rl/pg_actor_critic	Minh et al.

Other applications?

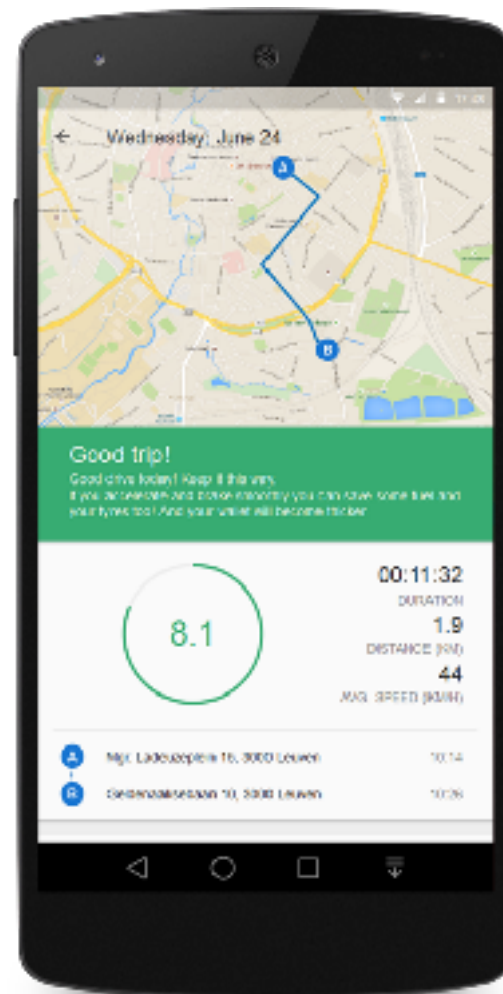
Finance



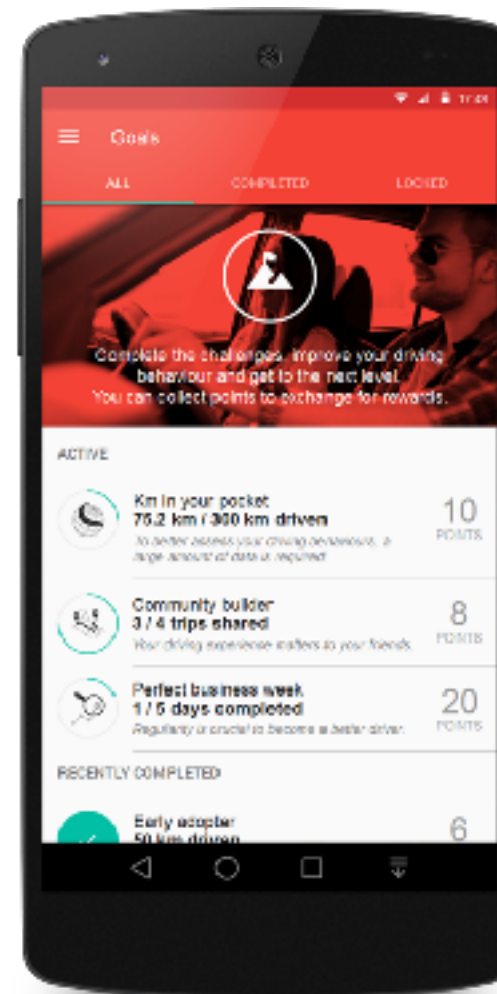
Insurance



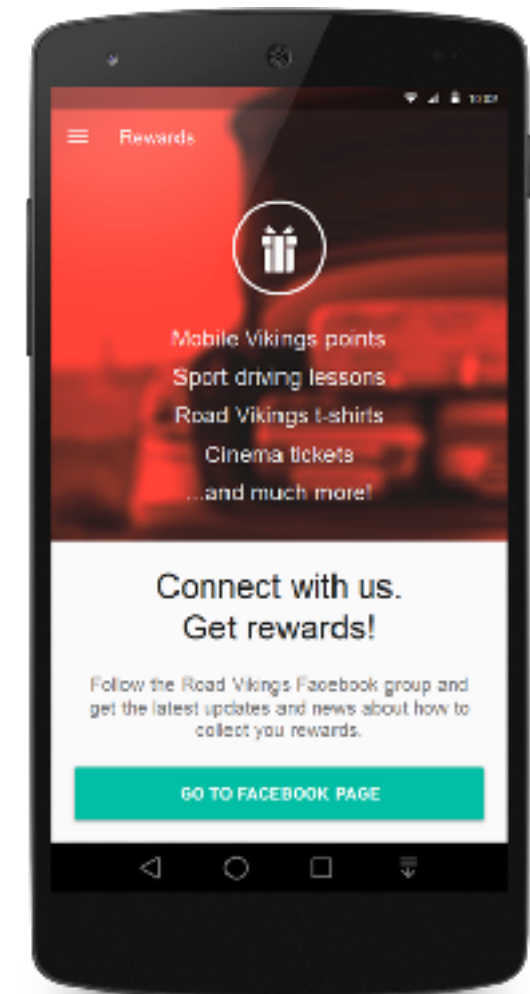
Record a trip



Trip feedback



Goals & challenges



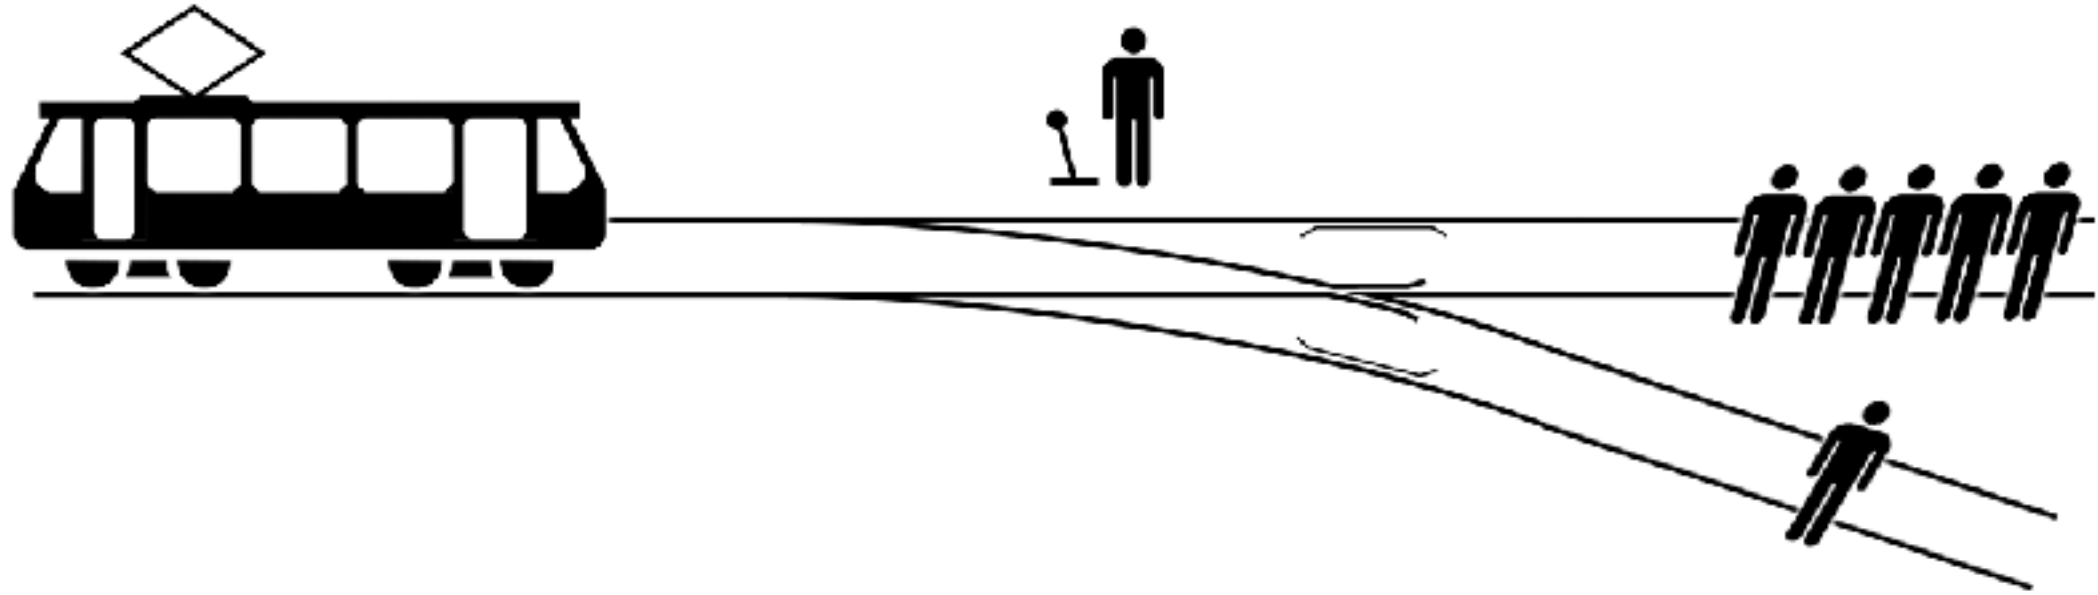
Rewards

Healthcare



Pitfalls?

Trolley problem



Ethical dilemma

Guidelines



- AI systems should be designed so that they **always are able to show the process which led to their actions** (Government decision-making)

Ethically Aligned design

- How can we assure that AI/AS are **accountable**?
- How can we ensure that AI/AS are **transparent**?
- How can we extend the benefits and **minimize** the risks of AI/AS technology being **misused**?

Thank you for listening!
Questions?

Feel free to add me at

<https://www.linkedin.com/in/mitrar/>

An example

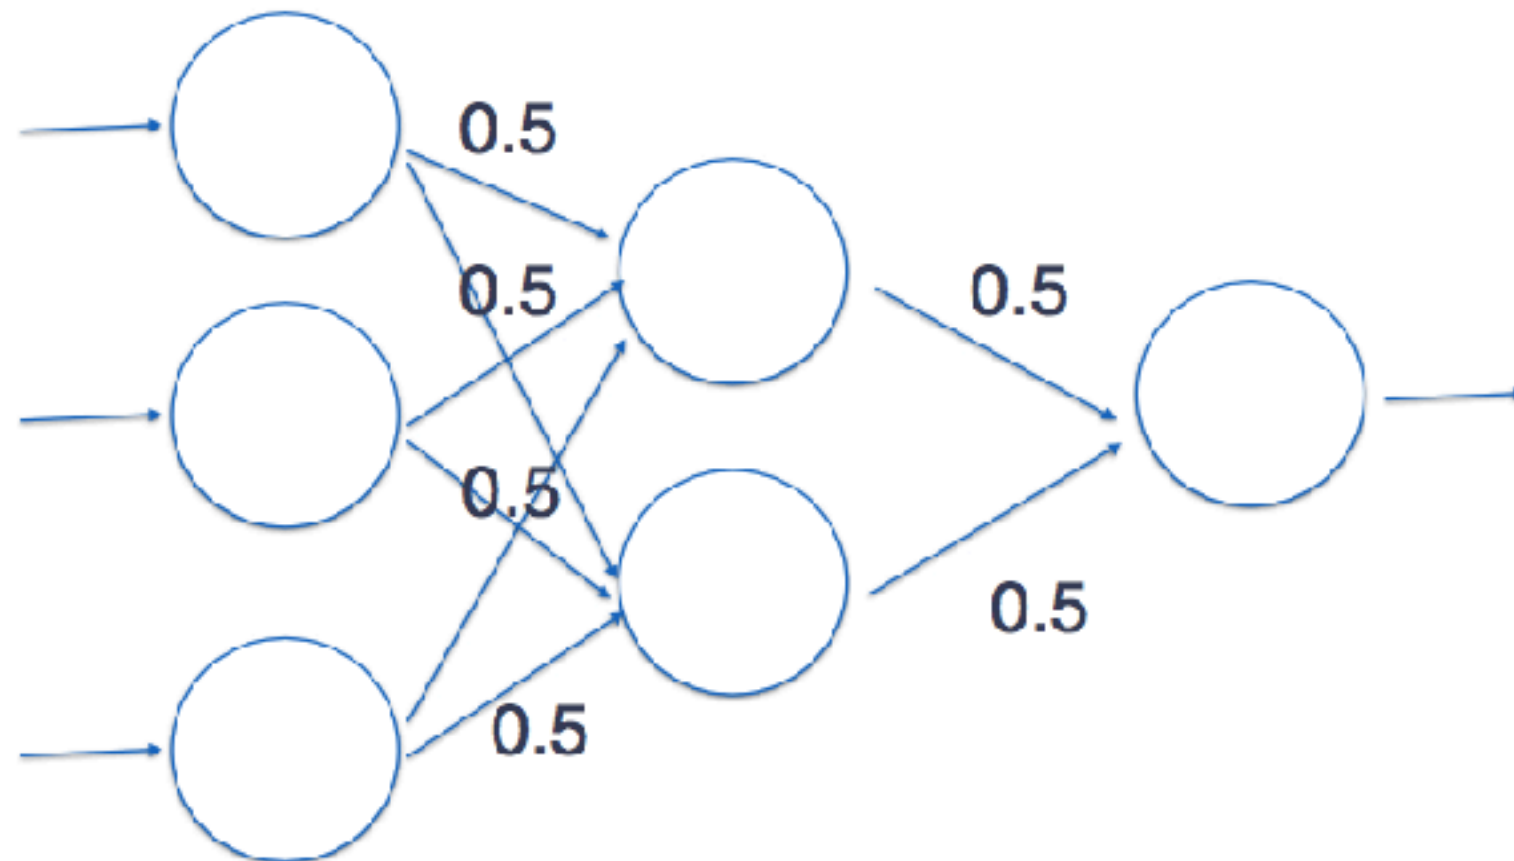
Training

Input: 0 1 0
Output: 0.75

Input layer

Hidden layer

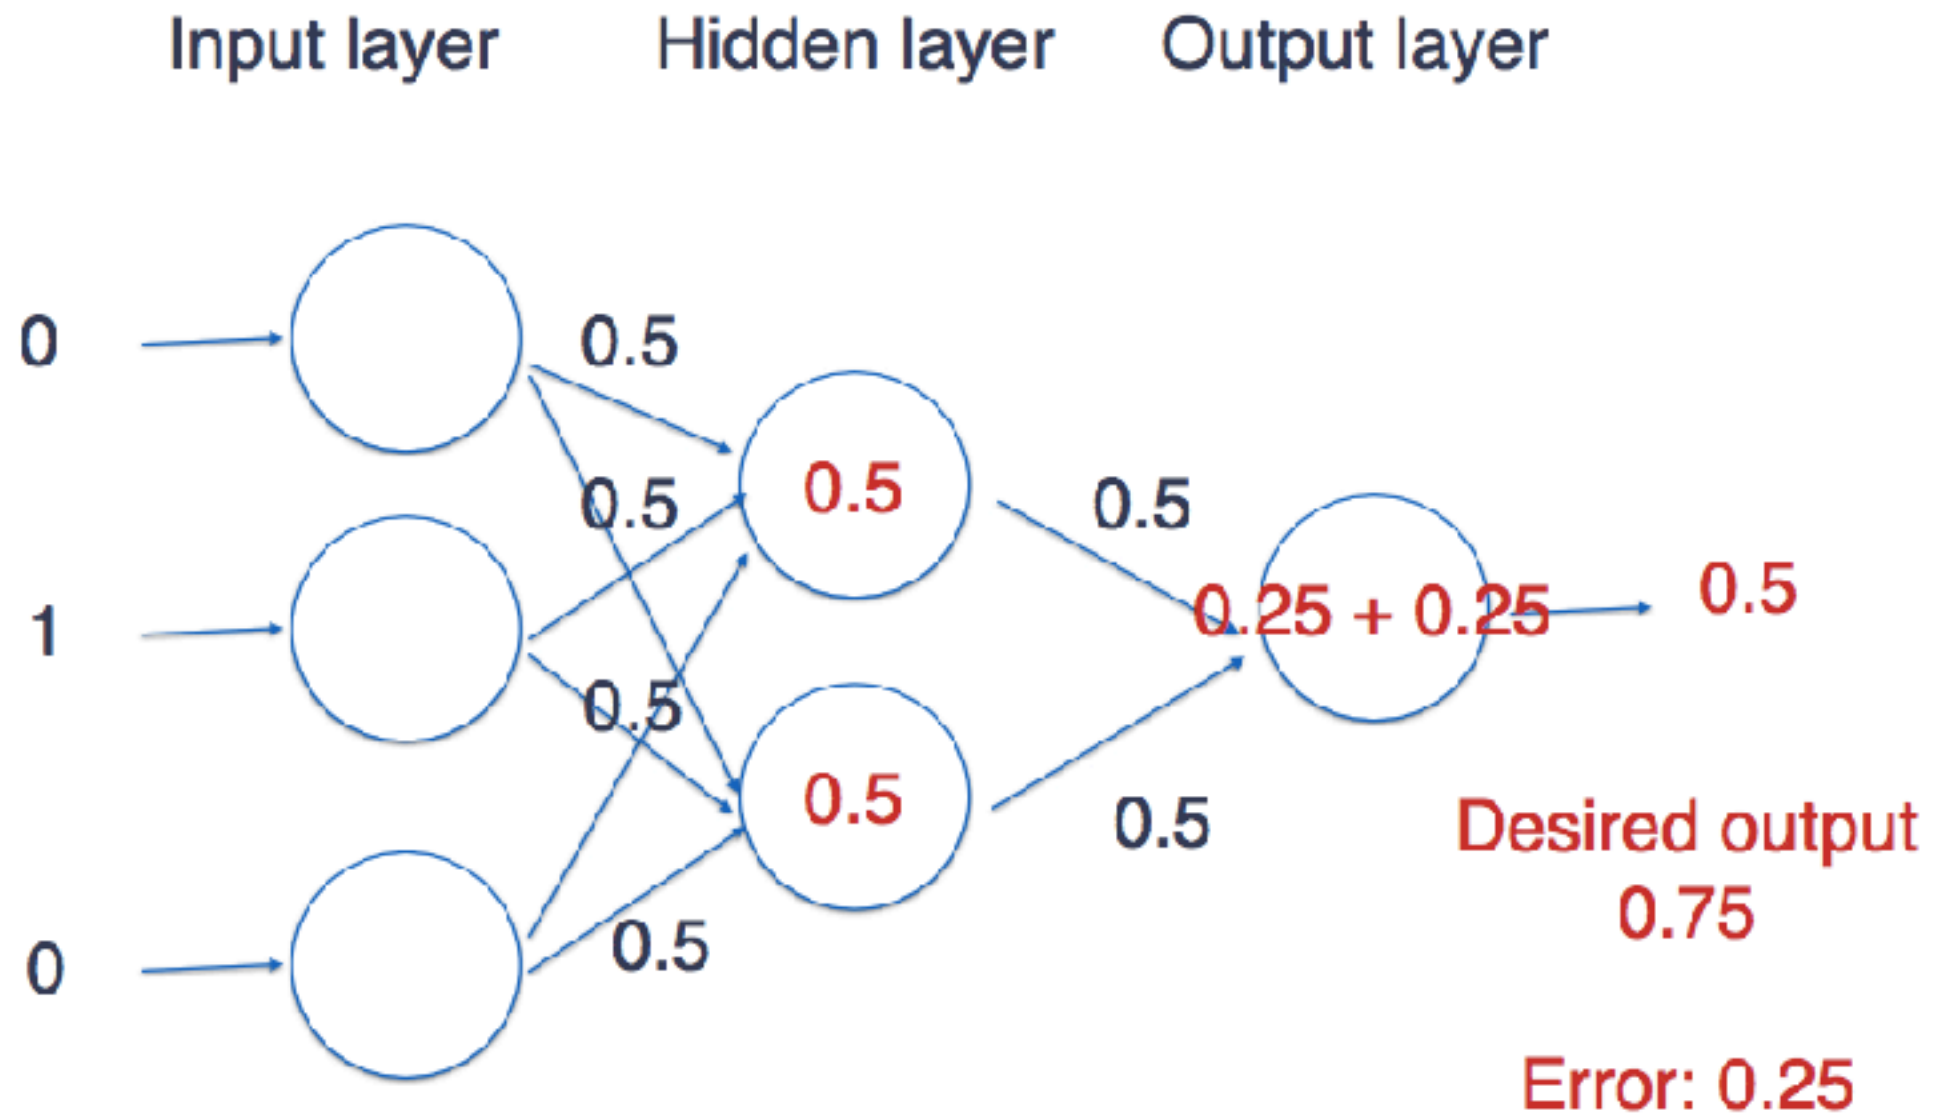
Output layer



An example

Training

Input: 0 1 0
Output: 0.75



An example

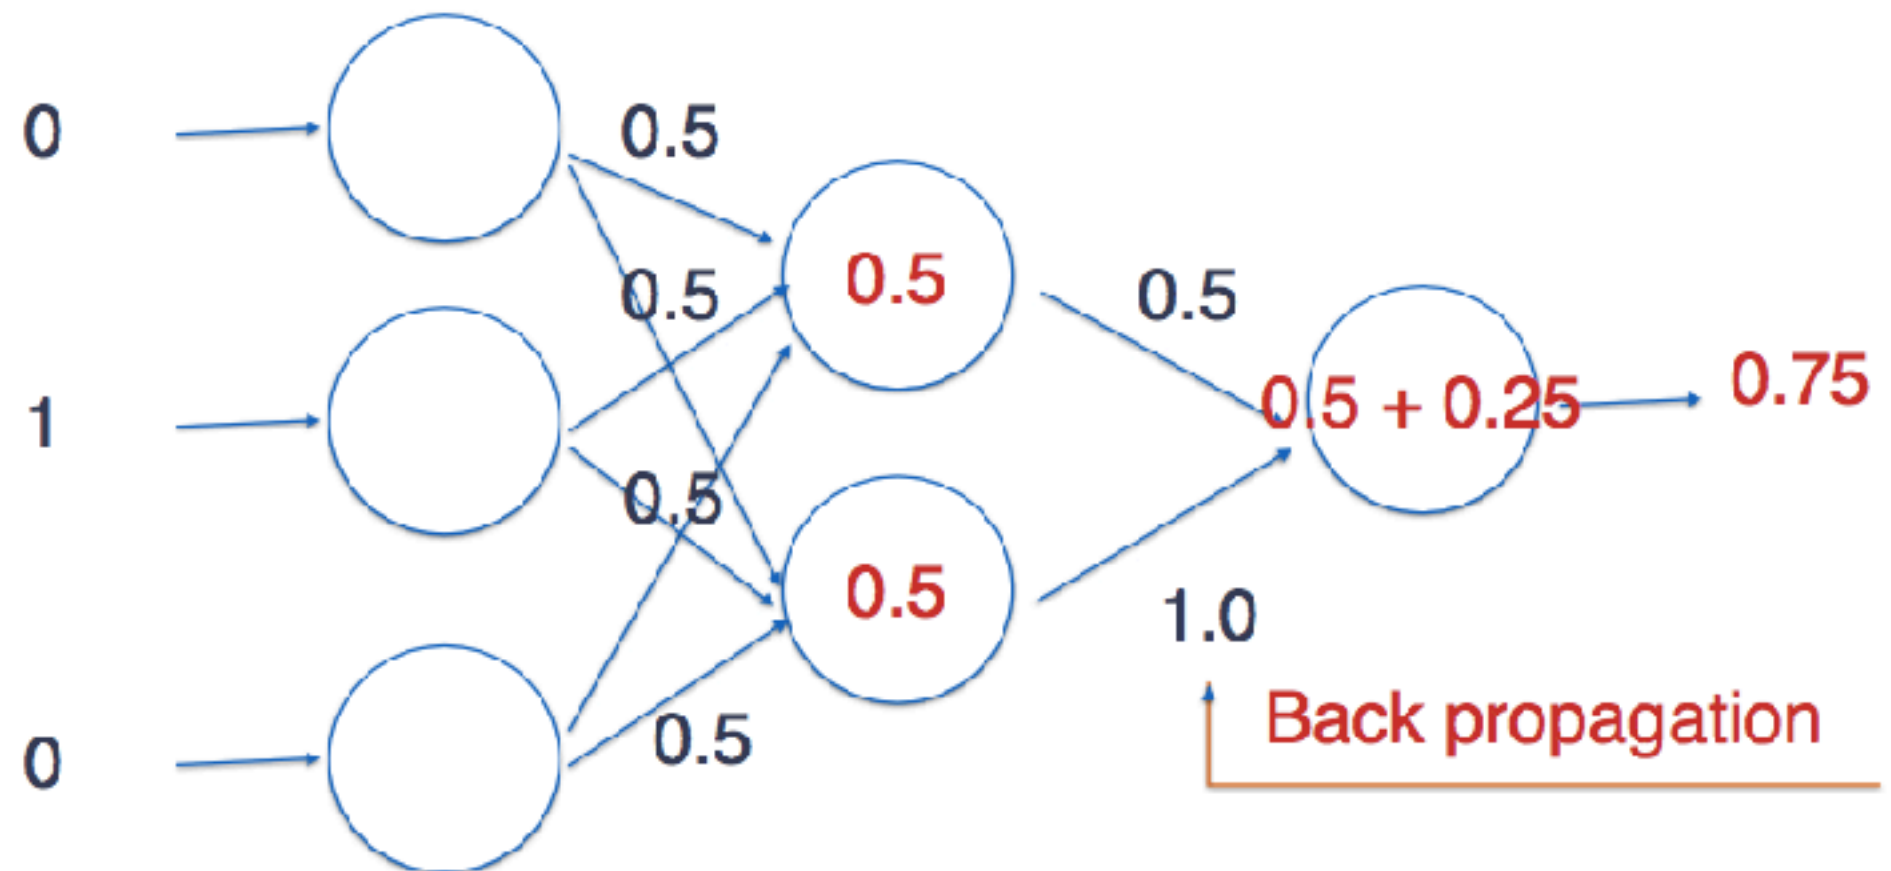
Training

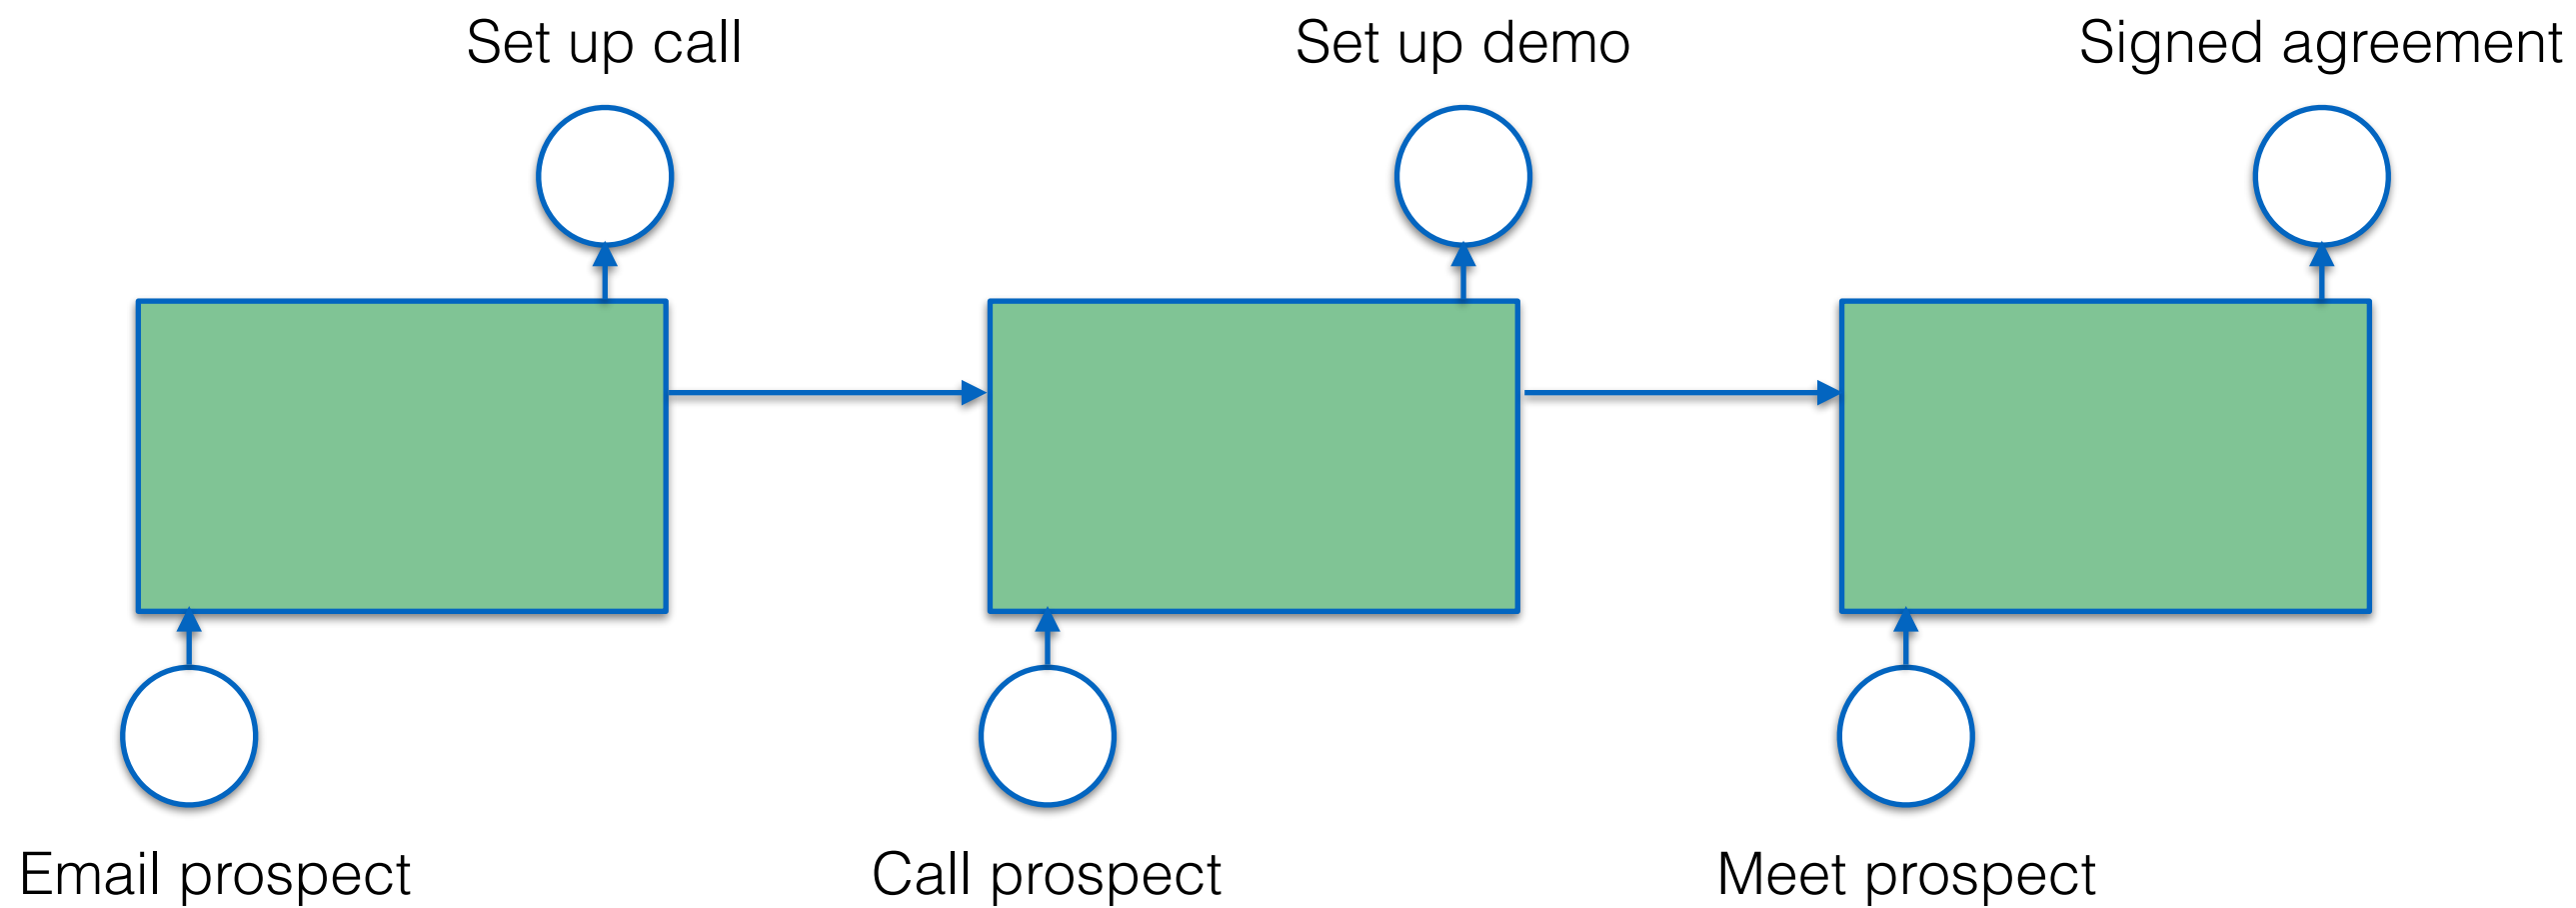
Input layer

Hidden layer

Output layer

Input: 0 1 0
Output: 0.75





Training LSTM