Recurrent Neural Networks

Danna Gurari
University of Texas at Austin
Spring 2021

https://www.ischool.utexas.edu/~dannag/Courses/IntroToMachineLearning/CourseContent.html

Review

- Last week:
  - History of Convolutional Neural Networks (CNNs)
  - CNNs – Convolutional Layers
  - CNNs – Pooling Layers
  - Deep Features
- Assignments (Canvas):
  - Project proposal due tonight
  - Project outline due next week
- Questions?

Today’s Topics

- Machine Learning for Sequential Data
- Recurrent Neural Networks (RNNs)
- Training Deep Neural Networks: Hardware & Software

Today's Topics

- Machine Learning for Sequential Data
- Recurrent Neural Networks (RNNs)
- Training Deep Neural Networks: Hardware & Software

Sequence Definition: Data of Arbitrary Length

e.g., Document

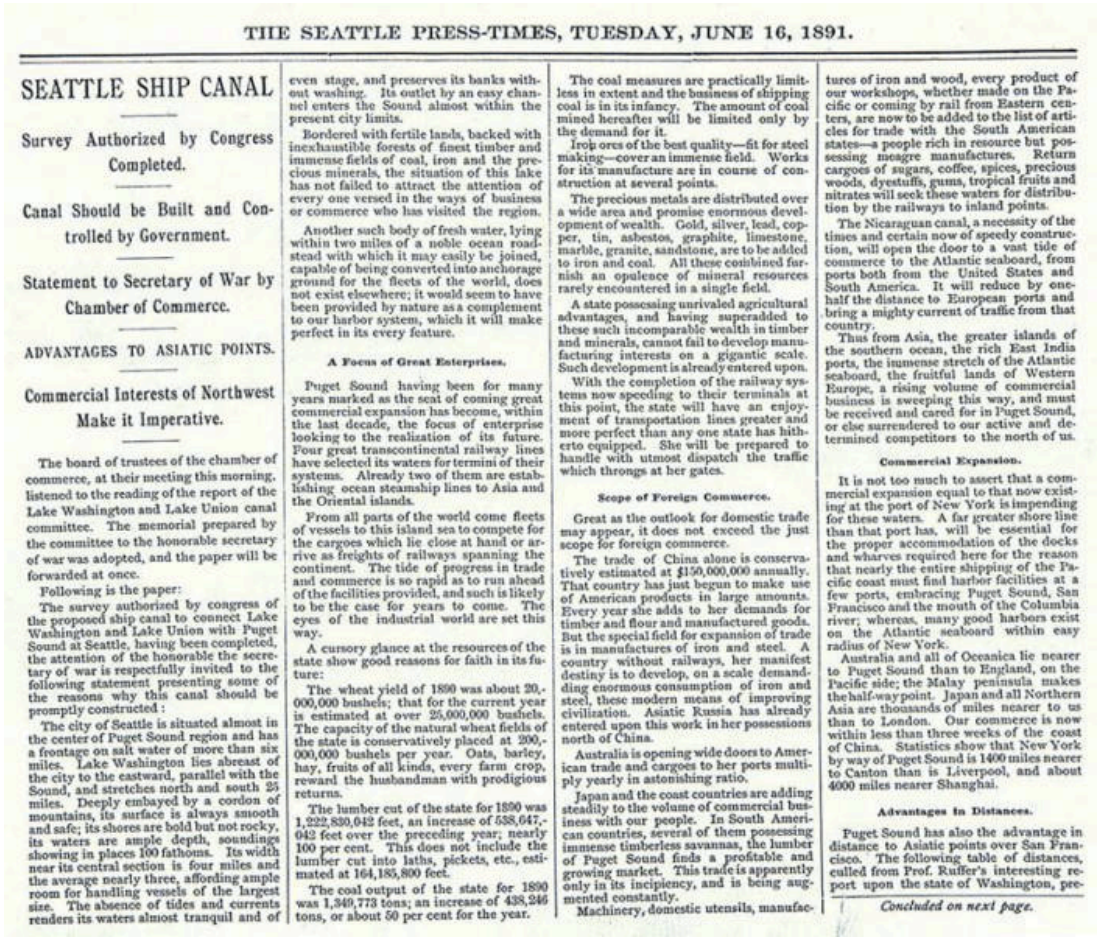

THE SEATTLE PRESS-TIMES, TUESDAY, JUNE 16, 1891.

SEATTLE SHIP CANAL

Survey Authorized by Congress Completed.

Canal Should be Built and Controlled by Government.

Statement to Secretary of War by Chamber of Commerce.

ADVANTAGES TO ASIATIC POINTS.

Commercial Interests of Northwest Make it Imperative.

The board of trustees of the chamber of commerce, at their meeting this morning, listened to the reading of the report of the Lake Washington and Lake Union canal committee. The memorial prepared by the committee to the honorable secretary of war was adopted, and the paper will be forwarded at once.

Following is the paper:

The survey authorized by congress of the proposed ship canal to connect Lake Washington and Lake Union with Puget Sound at Seattle, having been completed, the attention of the honorable the secretary of war is respectfully invited to the following statement presenting some of the reasons why this canal should be promptly constructed:

The city of Seattle is situated almost in the center of Puget Sound region and has a frontage on salt water of more than six miles. Lake Washington lies abreast of the city to the eastward, parallel with the Sound, and stretches north and south 25 miles. Deeply embayed by a cordon of mountains, its surface is always smooth and safe; its shores are bold but not rocky, its waters are ample depth, soundings showing in places 100 fathoms. Its width near its central section is four miles and the average nearly three, affording ample room for handling vessels of the largest size. The absence of tides and currents renders its waters almost tranquil and of even stage, and preserves its banks without washing. Its outlet by an easy channel enters the Sound almost within the present city limits.

Bordered with fertile lands, backed with inexhaustible forests of finest timber and immense fields of coal, iron and the precious minerals, the situation of this lake has not failed to attract the attention of every one versed in the ways of business or commerce who has visited the region.

Another such body of fresh water, lying within two miles of a noble ocean roadstead with which it may easily be joined, capable of being converted into anchorage ground for the fleets of the world, does not exist elsewhere; it would seem to have been provided by nature as a complement to our harbor system, which it will make perfect in its every feature.

A Focus of Great Enterprises.

Puget Sound having been for many years marked as the seat of coming great commercial expansion has become, within the last decade, the focus of enterprise looking to the realization of its future. Four great transcontinental railway lines have selected its waters for termini of their systems. Already two of them are establishing ocean steamship lines to Asia and the Oriental islands.

From all parts of the world come fleets of vessels to this island sea to compete for the cargoes which lie close at hand or arrive as freights of railways spanning the continent. The tide of progress in trade and commerce is so rapid as to run ahead of the facilities provided, and such is likely to be the case for years to come. The eyes of the industrial world are set this way.

A cursory glance at the resources of the state show good reasons for faith in its future:

The wheat yield of 1890 was about 20,000,000 bushels; that for the current year is estimated at over 25,000,000 bushels. The capacity of the natural wheat fields of the state is conservatively placed at 200,000,000 bushels per year. Oats, barley, hay, fruits of all kinds, every farm crop, reward the husbandman with prodigious returns.

The lumber cut of the state for 1890 was 1,222,830,042 feet, an increase of 538,647,042 feet over the preceding year; nearly 100 per cent. This does not include the lumber cut into laths, pickets, etc., estimated at 164,185,800 feet.

The coal output of the state for 1890 was 1,349,773 tons; an increase of 438,246 tons, or about 50 per cent for the year.

The coal measures are practically limitless in extent and the business of shipping coal is in its infancy. The amount of coal mined hereafter will be limited only by the demand for it.

Iron ores of the best quality—fit for steel making—cover an immense field. Works for its manufacture are in course of construction at several points.

The precious metals are distributed over a wide area and promise enormous development of wealth. Gold, silver, lead, copper, tin, asbestos, graphite, limestone, marble, granite, sandstone, are to be added to iron and coal. All these combined furnish an opulence of mineral resources rarely encountered in a single field.

A state possessing unrivaled agricultural advantages, and having superadded to these such incomparable wealth in timber and minerals, cannot fail to develop manufacturing interests on a gigantic scale. Such development is already entered upon.

With the completion of the railway systems now speeding to their terminals at this point, the state will have an enjoyment of transportation lines greater and more perfect than any one state has hitherto equipped. She will be prepared to handle with utmost dispatch the traffic which throngs at her gates.

Scope of Foreign Commerce.

Great as the outlook for domestic trade may appear, it does not exceed the just scope for foreign commerce.

The trade of China alone is conservatively estimated at $150,000,000 annually. That country has just begun to make use of American products in large amounts. Every year she adds to her demands for timber and flour and manufactured goods. But the special field for expansion of trade is in manufactures of iron and steel. A country without railways, her manifest destiny is to develop, on a scale demanding enormous consumption of iron and steel, these modern means of improving civilization. Asiatic Russia has already entered upon this work in her possessions north of China.

Australia is opening wide doors to American trade and cargoes to her ports multiply yearly in astonishing ratio.

Japan and the coast countries are adding steadily to the volume of commercial business with our people. In South American countries, several of them possessing immense timberless savannas, the lumber of Puget Sound finds a profitable and growing market. This trade is apparently only in its incipiency, and is being augmented constantly.

Machinery, domestic utensils, manufactures of iron and wood, every product of our workshops, whether made on the Pacific or coming by rail from Eastern centers, are now to be added to the list of articles for trade with the South American states—a people rich in resource but possessing meagre manufactures. Return cargoes of sugars, coffee, spices, precious woods, dyestuffs, gums, tropical fruits and nitrates will seek these waters for distribution by the railways to inland points.

The Nicaraguan canal, a necessity of the times and certain now of speedy construction, will open the door to a vast tide of commerce to the Atlantic seaboard, from ports both from the United States and South America. It will reduce by one-half the distance to European ports and bring a mighty current of traffic from that country.

Thus from Asia, the greater islands of the southern ocean, the rich East India ports, the immense stretch of the Atlantic seaboard, the fruitful lands of Western Europe, a rising volume of commercial business is sweeping this way, and must be received and cared for in Puget Sound, or else surrendered to our active and determined competitors to the north of us.

Commercial Expansion.

It is not too much to assert that a commercial expansion equal to that now existing at the port of New York is impending for these waters. A far greater shore line than that port has, will be essential for the proper accommodation of the docks and wharves required here for the reason that nearly the entire shipping of the Pacific coast must find harbor facilities at a few ports, embracing Puget Sound, San Francisco and the mouth of the Columbia river; whereas, many good harbors exist on the Atlantic seaboard within easy radius of New York.

Australia and all of Oceanica lie nearer to Puget Sound than to England, on the Pacific side; the Malay peninsula makes the half-way point. Japan and all Northern Asia are thousands of miles nearer to us than to London. Our commerce is now within less than three weeks of the coast of China. Statistics show that New York by way of Puget Sound is 1400 miles nearer to Canton than is Liverpool, and about 4000 miles nearer Shanghai.

Advantages in Distances.

Puget Sound has also the advantage in distance to Asiatic points over San Francisco. The following table of distances, culled from Prof. Ruffer's interesting report upon the state of Washington, pre-

Concluded on next page.

e.g., Images

e.g., Time-Series Data

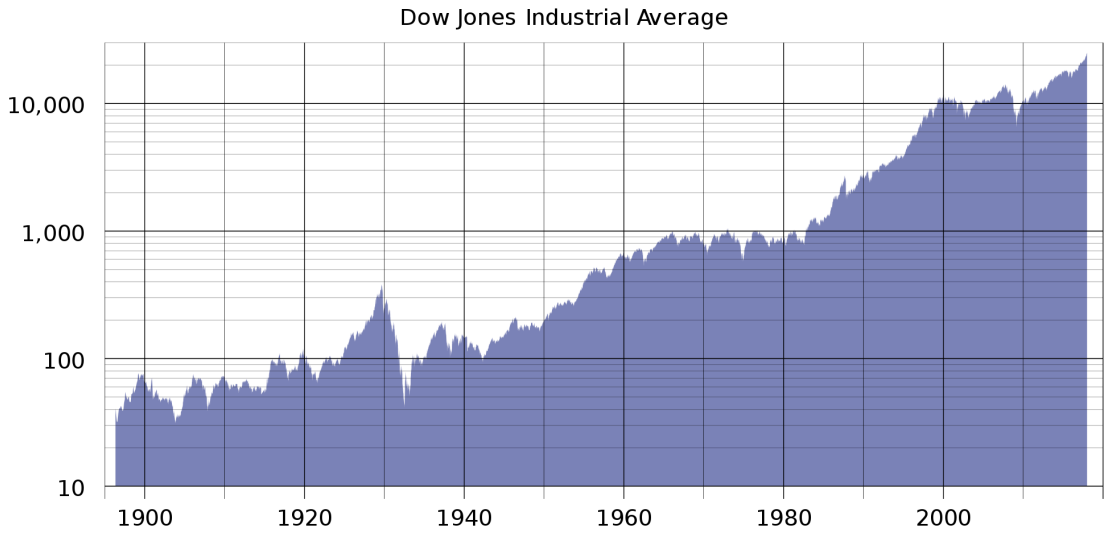

e.g., sentences, audio samples, brain waves, radio waves, air temperature

Properties of Sequences?

* Elements of a sequence occur in a certain order
* Elements depend on each other

e.g., Document

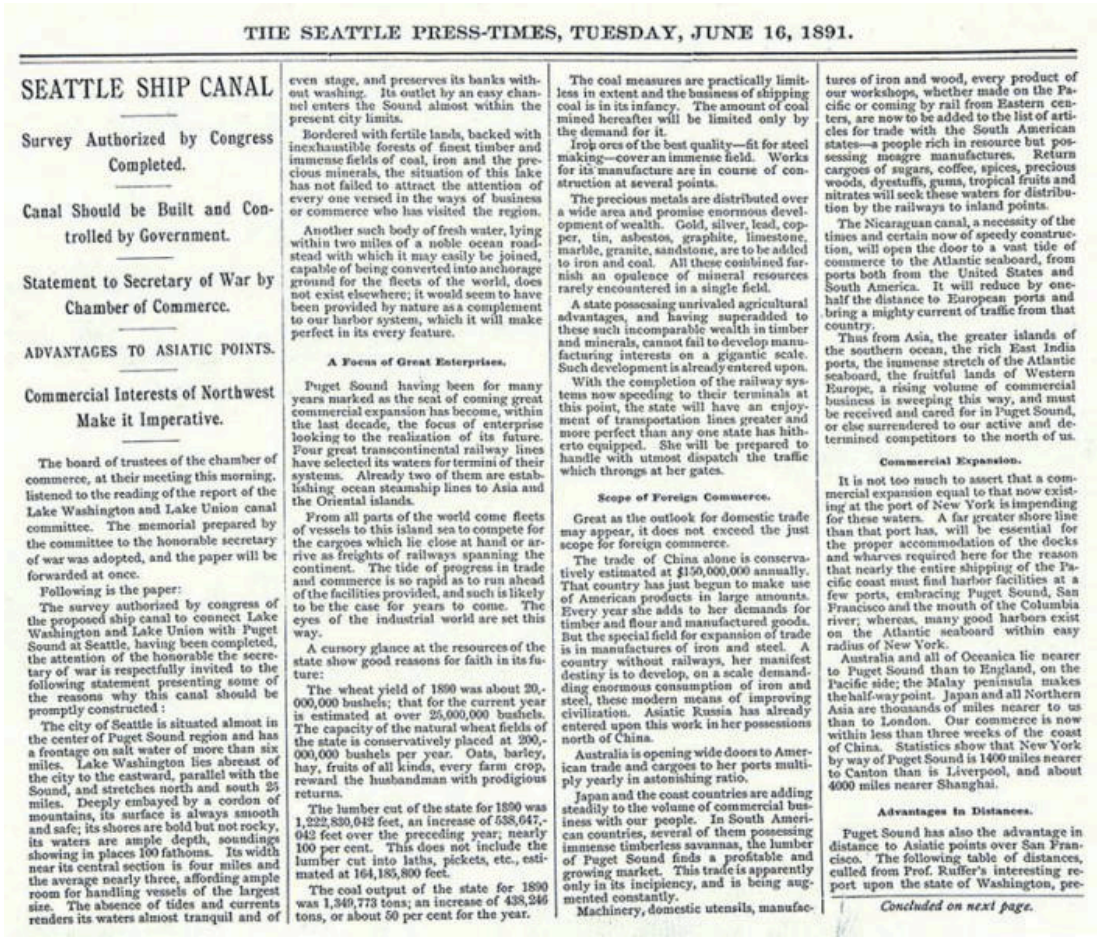

THE SEATTLE PRESS-TIMES, TUESDAY, JUNE 16, 1891.

SEATTLE SHIP CANAL

Survey Authorized by Congress Completed.

Canal Should be Built and Controlled by Government.

Statement to Secretary of War by Chamber of Commerce.

ADVANTAGES TO ASIATIC POINTS.

Commercial Interests of Northwest Make it Imperative.

The board of trustees of the chamber of commerce, at their meeting this morning, listened to the reading of the report of the Lake Washington and Lake Union canal committee. The memorial prepared by the committee to the honorable secretary of war was adopted, and the paper will be forwarded at once.

Following is the paper:

The survey authorized by congress of the proposed ship canal to connect Lake Washington and Lake Union with Puget Sound at Seattle, having been completed, the attention of the honorable the secretary of war is respectfully invited to the following statement presenting some of the reasons why this canal should be promptly constructed:

The city of Seattle is situated almost in the center of Puget Sound region and has a frontage on salt water of more than six miles. Lake Washington lies abreast of the city to the eastward, parallel with the Sound, and stretches north and south 25 miles. Deeply embayed by a cordon of mountains, its surface is always smooth and safe; its shores are bold but not rocky, its waters are ample depth, soundings showing in places 100 fathoms. Its width near its central section is four miles and the average nearly three, affording ample room for handling vessels of the largest size. The absence of tides and currents renders its waters almost tranquil and of even stage, and preserves its banks without washing. Its outlet by an easy channel enters the Sound almost within the present city limits.

Bordered with fertile lands, backed with inexhaustible forests of finest timber and immense fields of coal, iron and the precious minerals, the situation of this lake has not failed to attract the attention of every one versed in the ways of business or commerce who has visited the region.

Another such body of fresh water, lying within two miles of a noble ocean roadstead with which it may easily be joined, capable of being converted into anchorage ground for the fleets of the world, does not exist elsewhere; it would seem to have been provided by nature as a complement to our harbor system, which it will make perfect in its every feature.

A Focus of Great Enterprises.

Puget Sound having been for many years marked as the seat of coming great commercial expansion has become, within the last decade, the focus of enterprise looking to the realization of its future. Four great transcontinental railway lines have selected its waters for termini of their systems. Already two of them are establishing ocean steamship lines to Asia and the Oriental islands.

From all parts of the world come fleets of vessels to this island sea to compete for the cargoes which lie close at hand or arrive as freights of railways spanning the continent. The tide of progress in trade and commerce is so rapid as to run ahead of the facilities provided, and such is likely to be the case for years to come. The eyes of the industrial world are set this way.

A cursory glance at the resources of the state show good reasons for faith in its future:

The wheat yield of 1890 was about 20,000,000 bushels; that for the current year is estimated at over 25,000,000 bushels. The capacity of the natural wheat fields of the state is conservatively placed at 200,000,000 bushels per year. Oats, barley, hay, fruits of all kinds, every farm crop, reward the husbandman with prodigious returns.

The lumber cut of the state for 1890 was 1,222,830,042 feet, an increase of 538,647,042 feet over the preceding year; nearly 100 per cent. This does not include the lumber cut into laths, pickets, etc., estimated at 164,185,800 feet.

The coal output of the state for 1890 was 1,349,773 tons; an increase of 438,246 tons, or about 50 per cent for the year.

The coal measures are practically limitless in extent and the business of shipping coal is in its infancy. The amount of coal mined hereafter will be limited only by the demand for it.

Iron ores of the best quality—fit for steel making—cover an immense field. Works for its manufacture are in course of construction at several points.

The precious metals are distributed over a wide area and promise enormous development of wealth. Gold, silver, lead, copper, tin, asbestos, graphite, limestone, marble, granite, sandstone, are to be added to iron and coal. All these combined furnish an opulence of mineral resources rarely encountered in a single field.

A state possessing unrivaled agricultural advantages, and having superadded to these such incomparable wealth in timber and minerals, cannot fail to develop manufacturing interests on a gigantic scale. Such development is already entered upon.

With the completion of the railway systems now speeding to their terminals at this point, the state will have an enjoyment of transportation lines greater and more perfect than any one state has hitherto equipped. She will be prepared to handle with utmost dispatch the traffic which throngs at her gates.

Scope of Foreign Commerce.

Great as the outlook for domestic trade may appear, it does not exceed the just scope for foreign commerce.

The trade of China alone is conservatively estimated at $150,000,000 annually. That country has just begun to make use of American products in large amounts. Every year she adds to her demands for timber and flour and manufactured goods. But the special field for expansion of trade is in manufactures of iron and steel. A country without railways, her manifest destiny is to develop, on a scale demanding enormous consumption of iron and steel, these modern means of improving civilization. Asiatic Russia has already entered upon this work in her possessions north of China.

Australia is opening wide doors to American trade and cargoes to her ports multiply yearly in astonishing ratio.

Japan and the coast countries are adding steadily to the volume of commercial business with our people. In South American countries, several of them possessing immense timberless savannas, the lumber of Puget Sound finds a profitable and growing market. This trade is apparently only in its incipiency, and is being augmented constantly.

Machinery, domestic utensils, manufactures of iron and wood, every product of our workshops, whether made on the Pacific or coming by rail from Eastern centers, are now to be added to the list of articles for trade with the South American states—a people rich in resource but possessing meagre manufactures. Return cargoes of sugars, coffee, spices, precious woods, dyestuffs, gums, tropical fruits and nitrates will seek these waters for distribution by the railways to inland points.

The Nicaraguan canal, a necessity of the times and certain now of speedy construction, will open the door to a vast tide of commerce to the Atlantic seaboard, from ports both from the United States and South America. It will reduce by one-half the distance to European ports and bring a mighty current of traffic from that country.

Thus from Asia, the greater islands of the southern ocean, the rich East India ports, the immense stretch of the Atlantic seaboard, the fruitful lands of Western Europe, a rising volume of commercial business is sweeping this way, and must be received and cared for in Puget Sound, or else surrendered to our active and determined competitors to the north of us.

Commercial Expansion.

It is not too much to assert that a commercial expansion equal to that now existing at the port of New York is impending for these waters. A far greater shore line than that port has, will be essential for the proper accommodation of the docks and wharves required here for the reason that nearly the entire shipping of the Pacific coast must find harbor facilities at a few ports, embracing Puget Sound, San Francisco and the mouth of the Columbia river; whereas, many good harbors exist on the Atlantic seaboard within easy radius of New York.

Australia and all of Oceanica lie nearer to Puget Sound than to England, on the Pacific side; the Malay peninsula makes the half-way point. Japan and all Northern Asia are thousands of miles nearer to us than to London. Our commerce is now within less than three weeks of the coast of China. Statistics show that New York by way of Puget Sound is 1400 miles nearer to Canton than is Liverpool, and about 4000 miles nearer Shanghai.

Advantages in Distances.

Puget Sound has also the advantage in distance to Asiatic points over San Francisco. The following table of distances, culled from Prof. Ruffer's interesting report upon the state of Washington, pre-

Concluded on next page.

e.g., Images

e.g., Time-Series Data

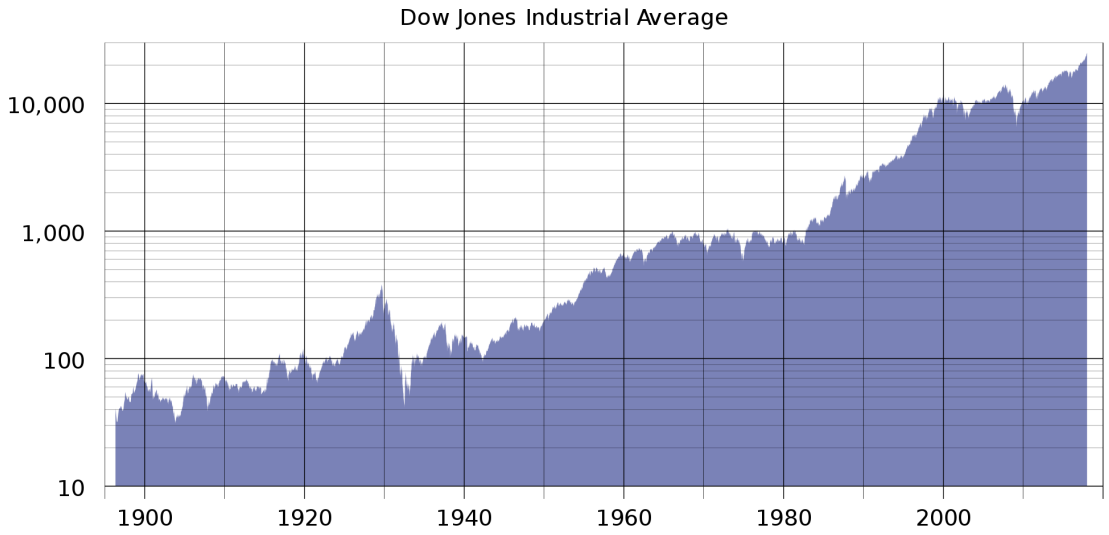

e.g., sentences, audio samples, brain waves, radio waves, air temperature

Sequence Sources

* Elements of a sequence occur in a certain order
* Elements depend on each other

AUDIO

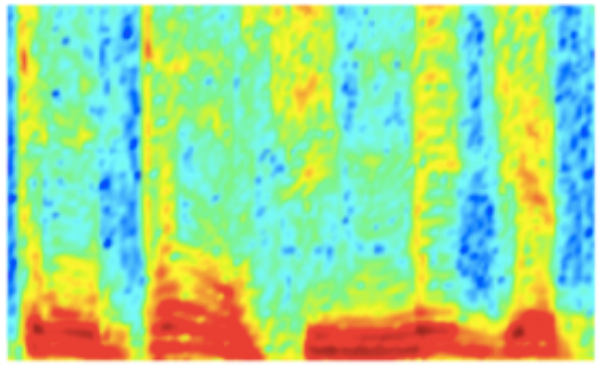

Audio Spectrogram

IMAGES

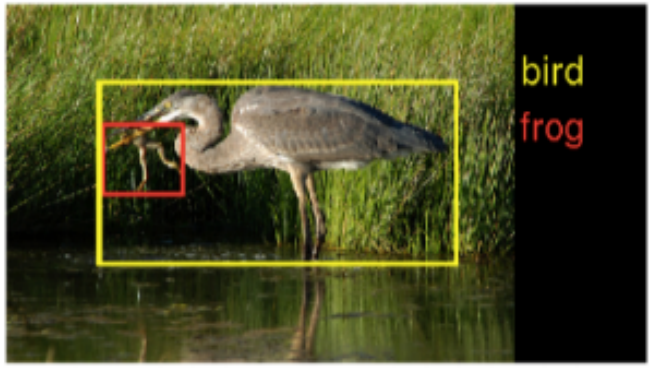

Image pixels

TEXT

0	0	0	0.2	0	0.7	0	0	0	...	...

Word, context, or document vectors

https://www.tensorflow.org/tutorials/representation/word2vec

Sequence Applications: One-to-Many

- **Input**: fixed-size
- **Output**: sequence
- e.g., image captioning

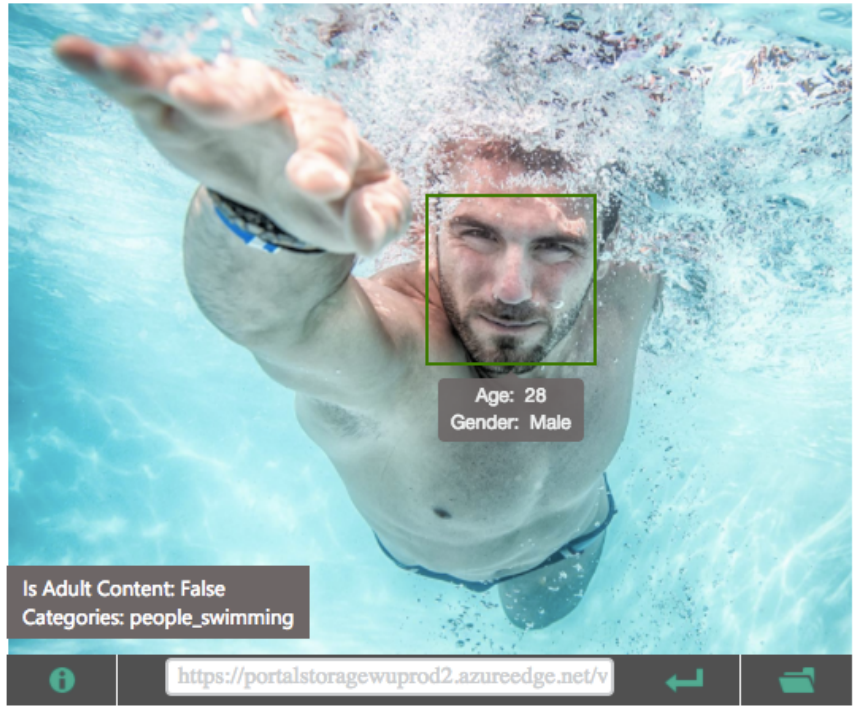

Features:

Feature Name	Value
Description	{ "type": 0, "captions": [{ "text": "a man swimming in a pool of water", "confidence": 0.7850108693093019 }] }
Tags	[{ "name": "water", "confidence": 0.9996442794799805 }, { "name": "sport", "confidence": 0.9504992365837097 }, { "name": "swimming", "confidence": 0.9062818288803101, "hint": "sport" }, { "name": "pool", "confidence": 0.8787588477134705 }, { "name": "water sport", "confidence": 0.631849467754364, "hint": "sport" }]
Image Format	jpeg
Image Dimensions	1500 x 1155
Clip Art Type	0 Non-clipart
Line Drawing Type	0 Non-LineDrawing
Black & White Image	False

Captions: https://www.microsoft.com/cognitive-services/en-us/computer-vision-api

Sequence Applications: Many-to-One

- **Input**: sequence
- **Output**: fixed-size
- e.g., sentiment analysis (hate? love?, etc)

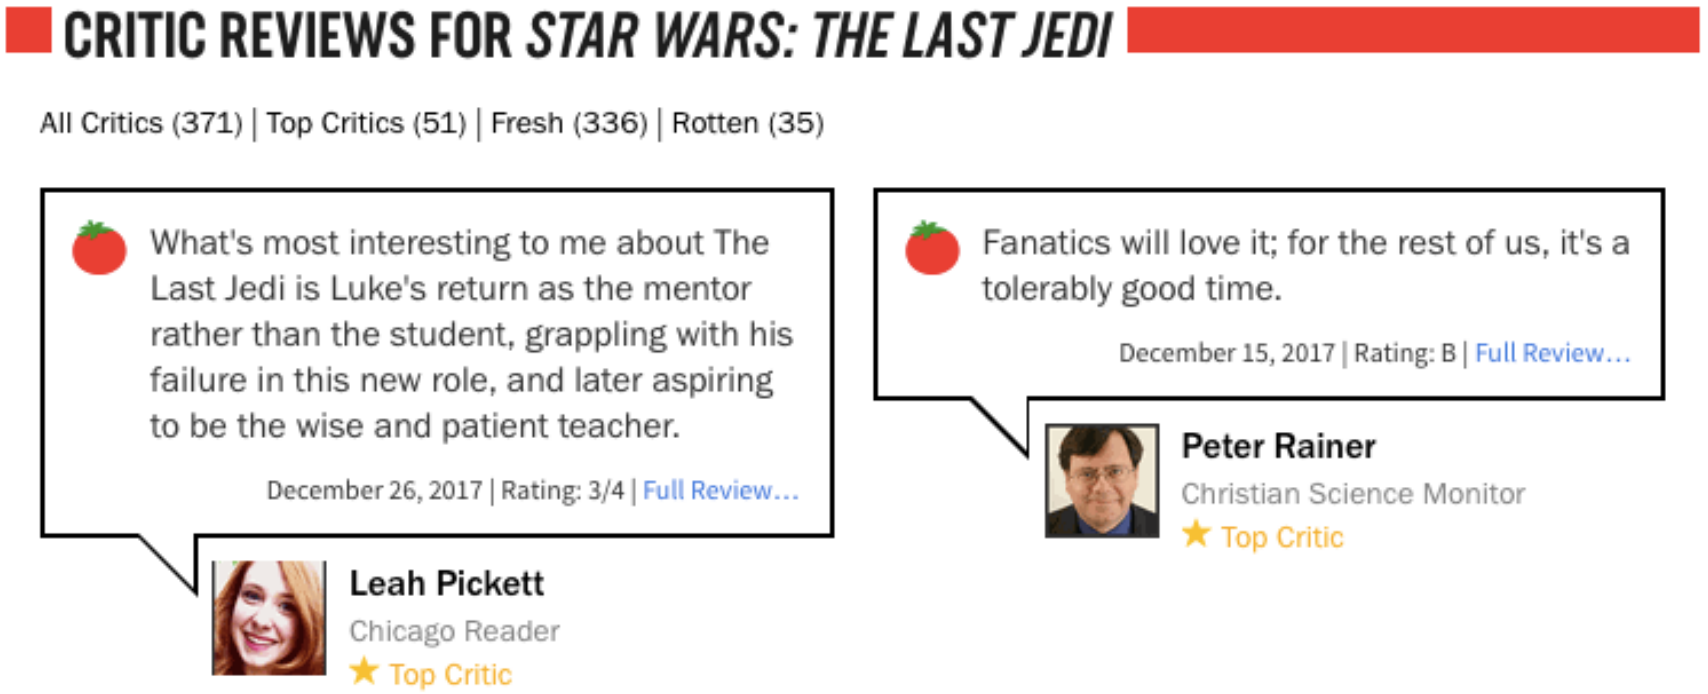

https://www.rottentomatoes.com/m/star_wars_the_last_jedi

Sequence Applications: Many-to-Many

- **Input**: sequence
- **Output**: sequence
- e.g., language translation

Sequence Applications

1. What are other examples of “one-to-many” applications?
2. What are other examples of “many-to-one” applications?
3. What are other examples of “many-to-many” applications?

Today's Topics

- Machine Learning for Sequential Data
- Recurrent Neural Networks (RNNs)
- Training Deep Neural Networks: Hardware & Software

Recall: Feedforward Neural Networks

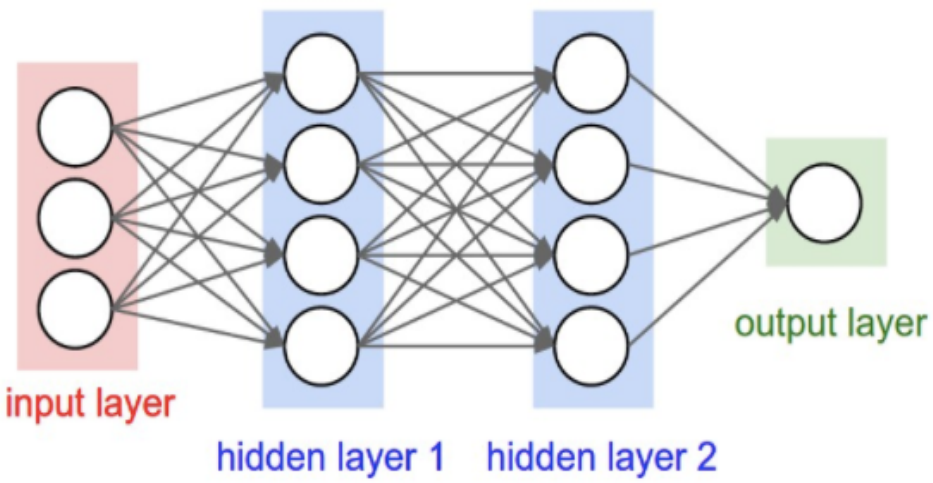

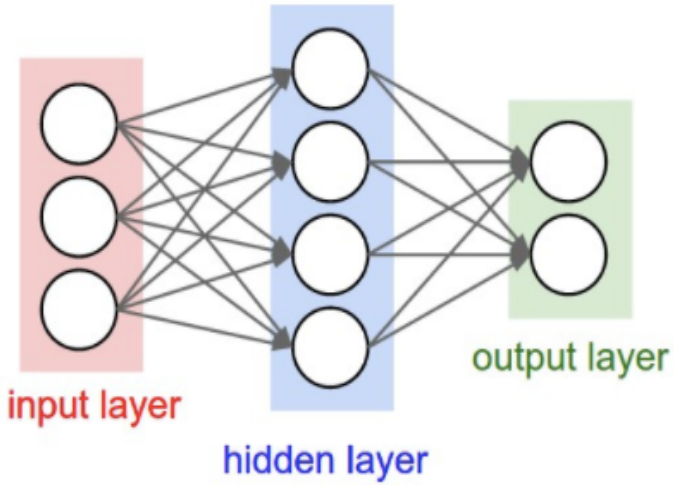

Problem: many model parameters!

Problem: no memory of past since weights learned independently

Each layer serves as input to the next layer with no loops

Figure Source: http://cs231n.github.io/neural-networks-1/

RNN History

Rumelhart, Hinton, & Williams.
"Learning Representations by back-propagating errors"

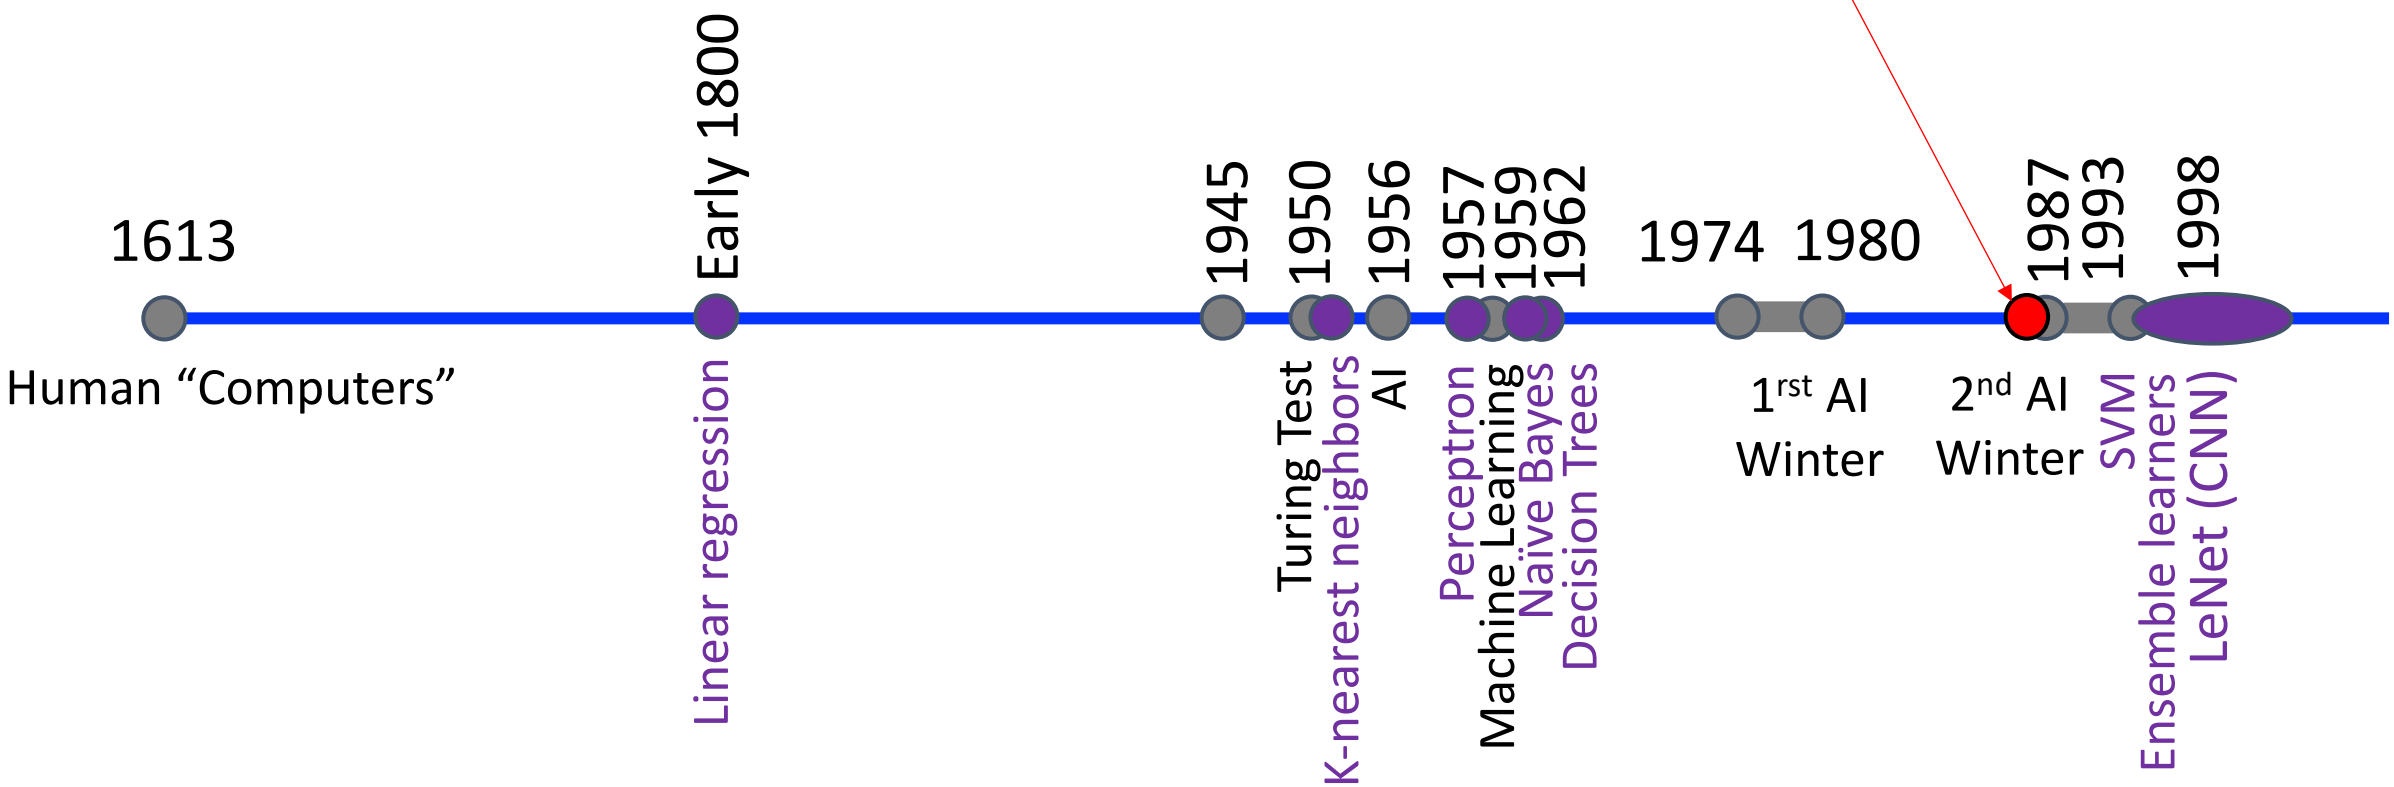

Recurrent Neural Networks (RNNs)

- Main idea: use hidden state to capture information about the past

Feedforward Network
Each layer receives input from the previous layer with no loops

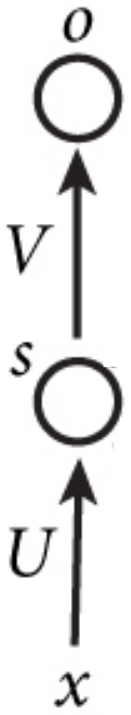

Recurrent Network
Each layer receives input from the previous layer and the output from the previous time step

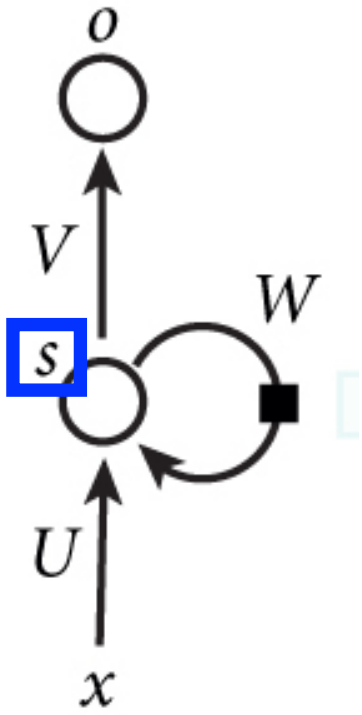

http://www.wildml.com/2015/09/recurrent-neural-networks-tutorial-part-1-introduction-to-rnns/

Recurrent Neural Networks (RNNs)

- Main idea: use hidden state to capture information about the past

Recurrence formula applied at every time step:

Model parameters

$$s_t = f_m(s_{t-1}, x_t)$$

New state | Old state | Input at time step

Recurrent Network

Each layer receives input from the previous layer and the output from the previous time step

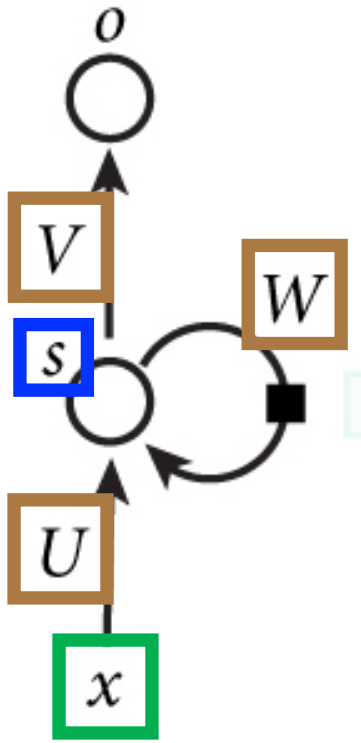

http://www.wildml.com/2015/09/recurrent-neural-networks-tutorial-part-1-introduction-to-rnns/

RNN: Time Step 1

- Main idea: use hidden state to capture information about the past

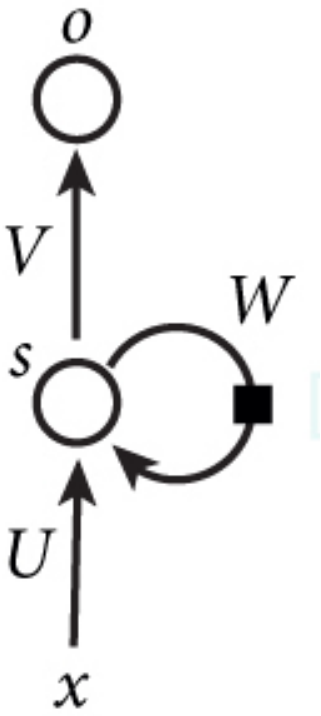

http://www.wildml.com/2015/09/recurrent-neural-networks-tutorial-part-1-introduction-to-rnns/

RNN: Time Step 1

- Main idea: use hidden state to capture information about the past

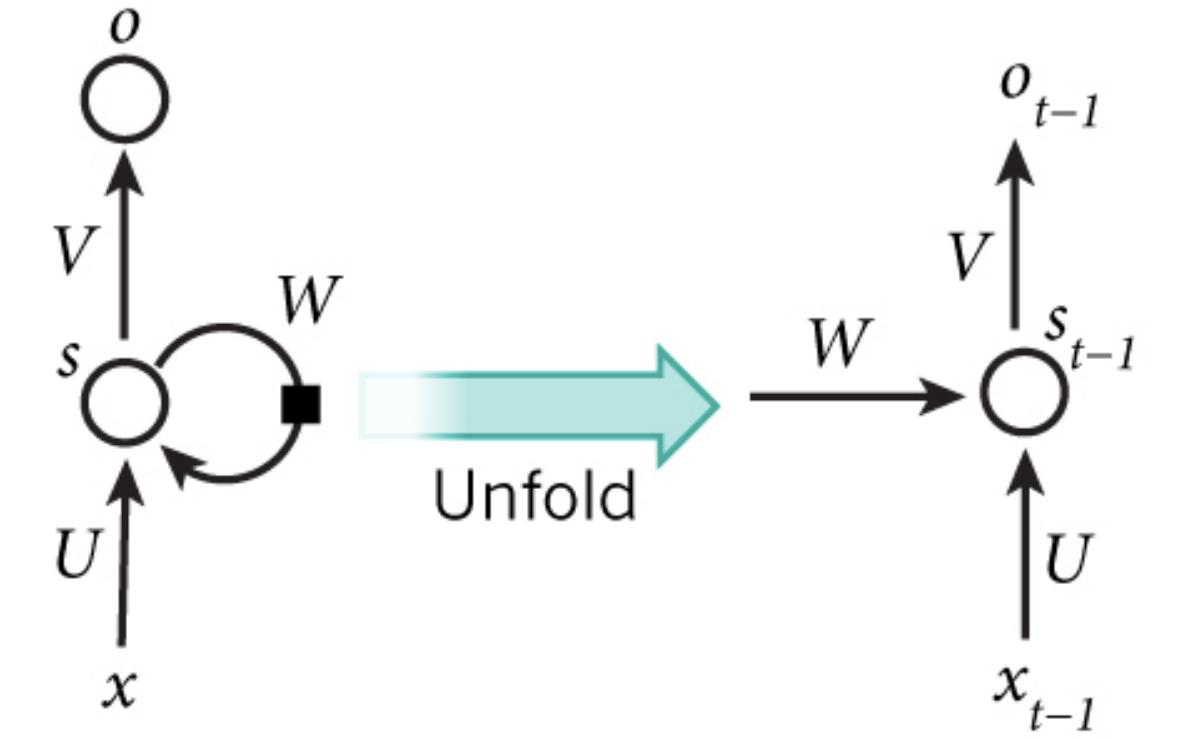

http://www.wildml.com/2015/09/recurrent-neural-networks-tutorial-part-1-introduction-to-rnns/

RNN: Time Step 2

- Main idea: use hidden state to capture information about the past

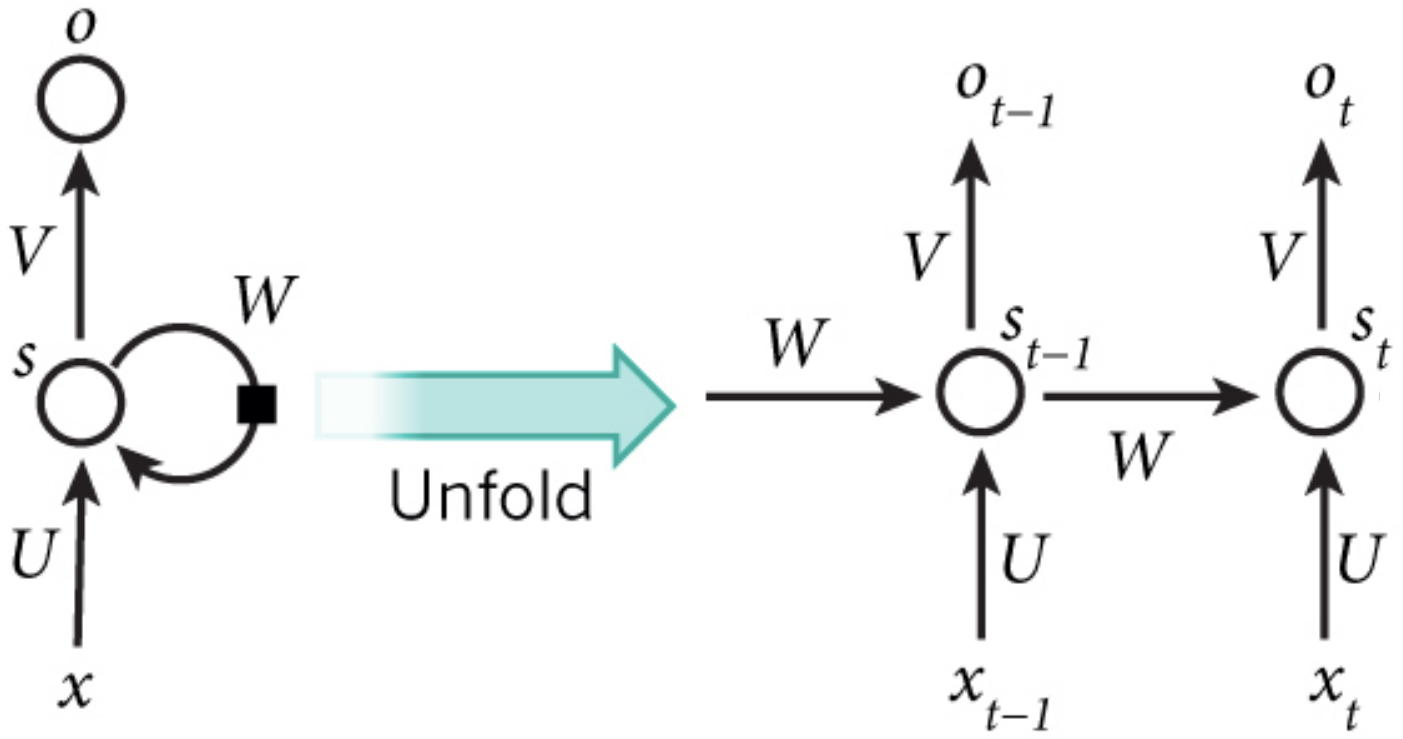

http://www.wildml.com/2015/09/recurrent-neural-networks-tutorial-part-1-introduction-to-rnns/

RNN: Time Step 3

- Main idea: use hidden state to capture information about the past

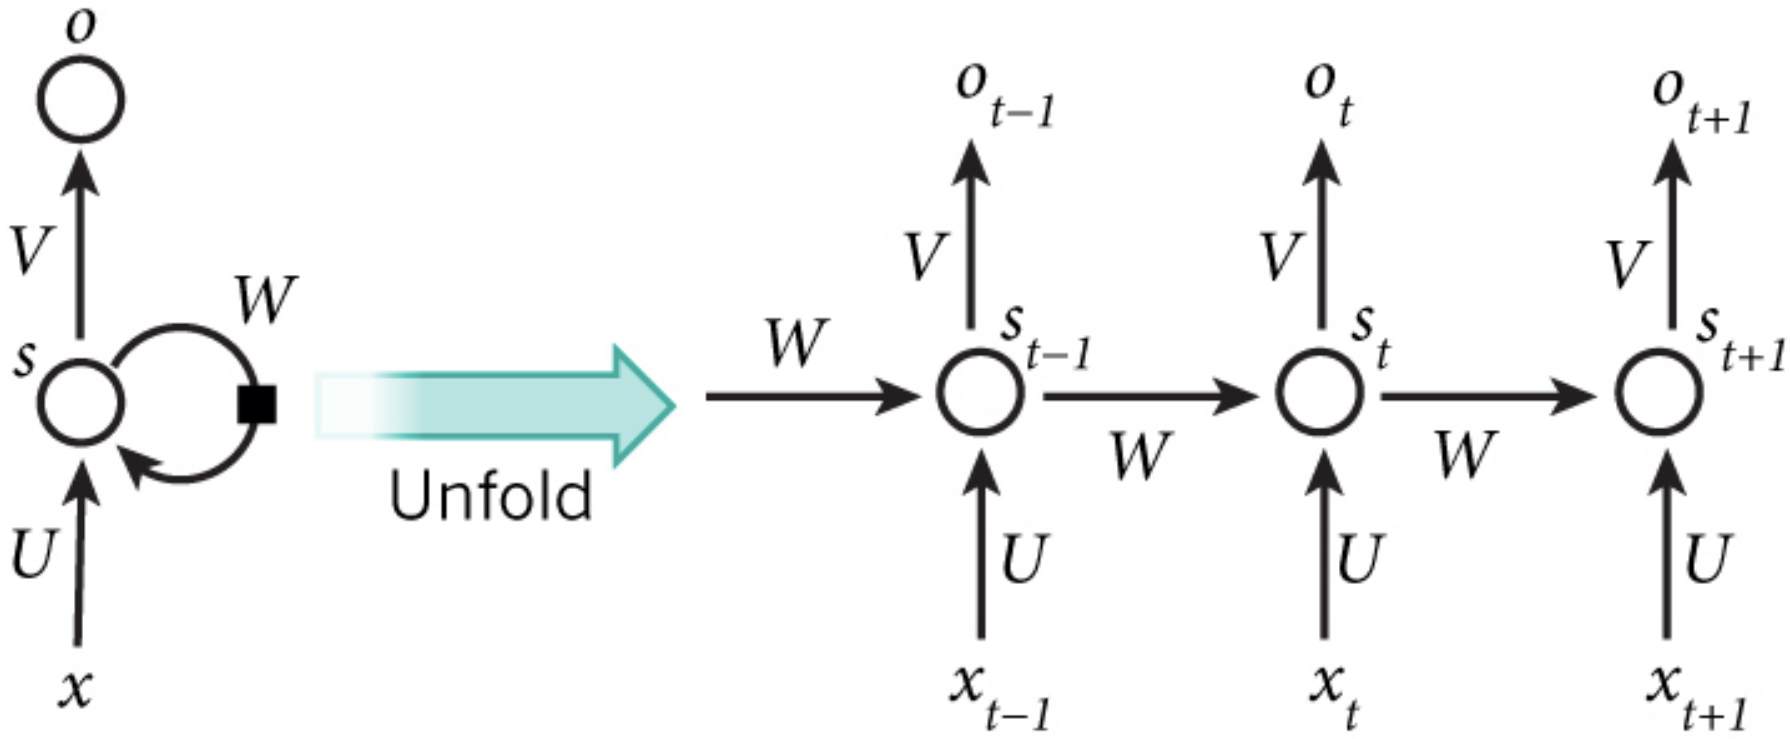

http://www.wildml.com/2015/09/recurrent-neural-networks-tutorial-part-1-introduction-to-rnns/

RNN: And So On…

- Main idea: use hidden state to capture information about the past

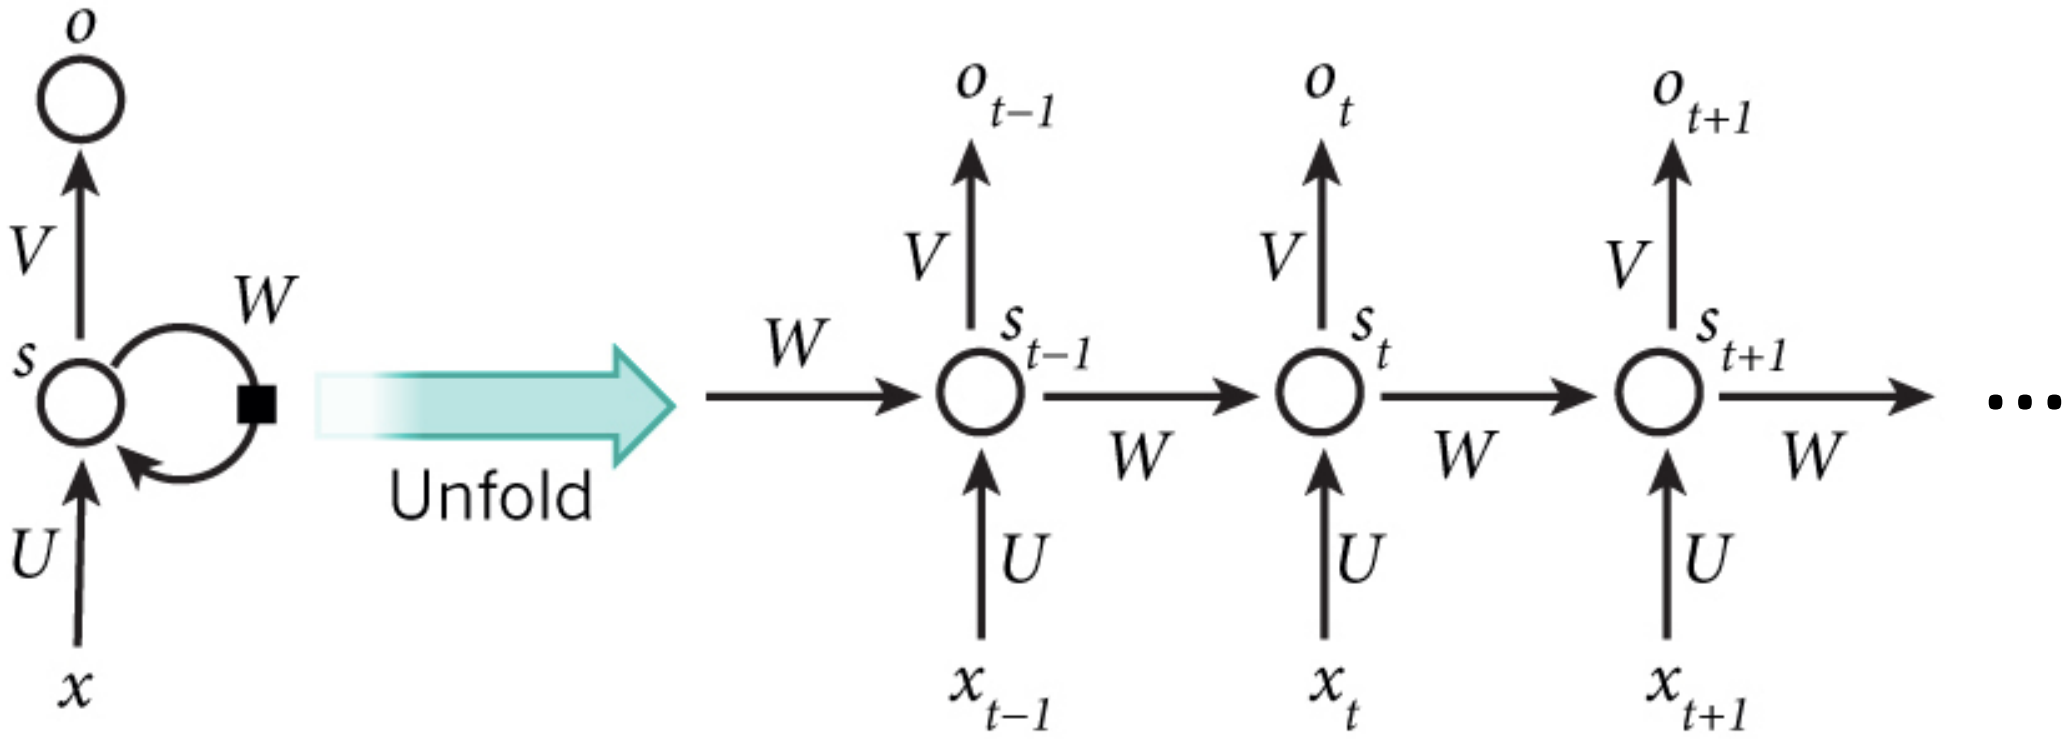

http://www.wildml.com/2015/09/recurrent-neural-networks-tutorial-part-1-introduction-to-rnns/

RNN: Model Parameters and Inputs

- Main idea: use hidden state to capture information about the past

Recurrence formula applied at every time step:

Model parameters

$s_t = f_m(s_{t-1}, x_t)$

New state — Old state — Input at time step

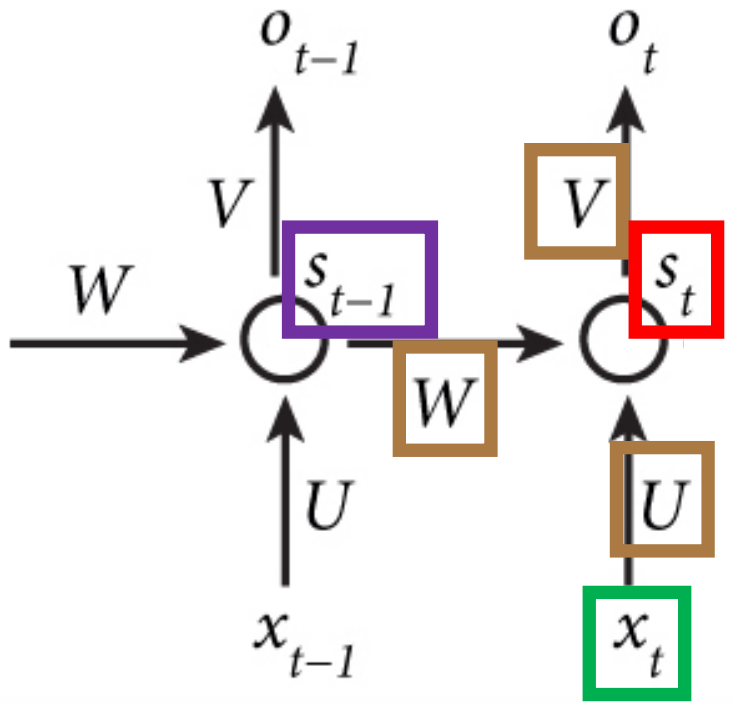

http://www.wildml.com/2015/09/recurrent-neural-networks-tutorial-part-1-introduction-to-rnns/

RNN: Model Parameters and Inputs

- Main idea: use hidden state to capture information about the past

Recurrence formula applied at every time step:

Model parameters

$s_t = f_m(s_{t-1}, x_t)$

New state | Old state | Input at time step

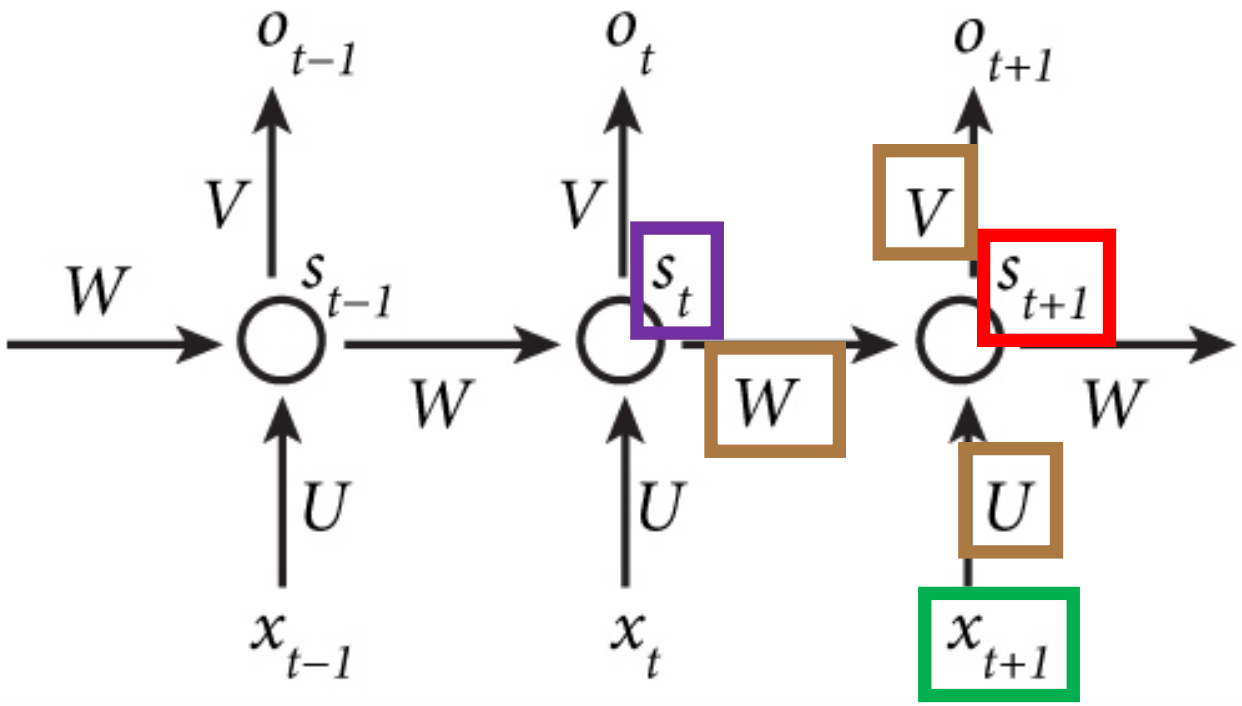

http://www.wildml.com/2015/09/recurrent-neural-networks-tutorial-part-1-introduction-to-rnns/

RNN: Model Parameters and Inputs

- All layers share the same model parameters (U, V, W)
  - What is different between the layers?

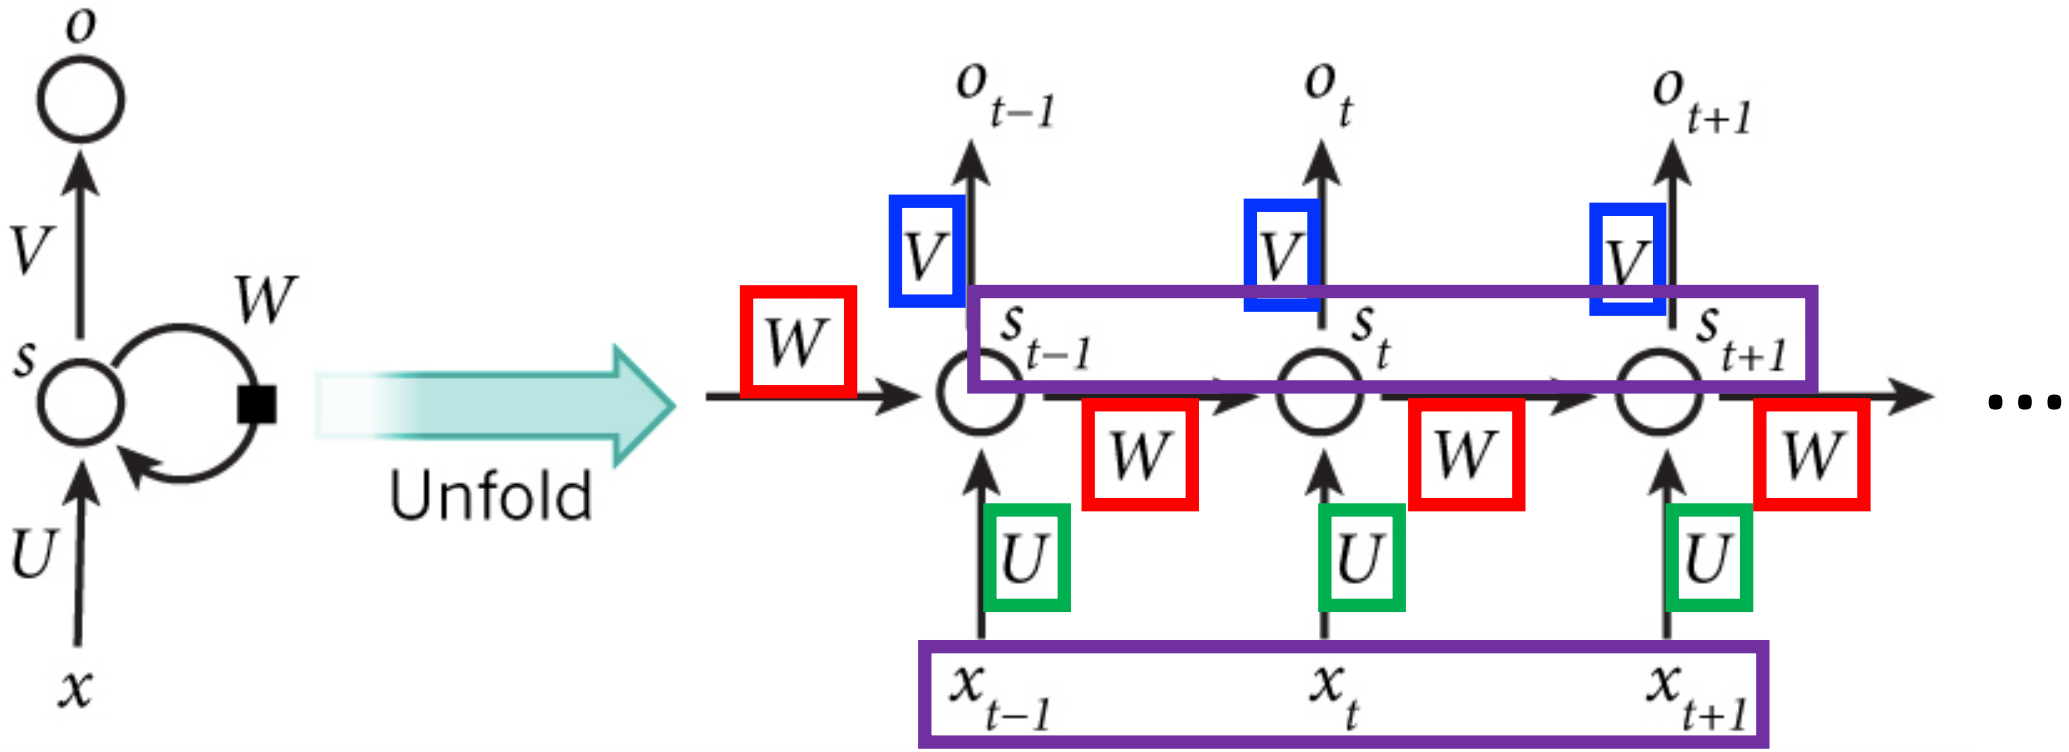

http://www.wildml.com/2015/09/recurrent-neural-networks-tutorial-part-1-introduction-to-rnns/

RNN: Model Parameters and Inputs

- When unfolded, a RNN is a deep feedforward network with shared weights!

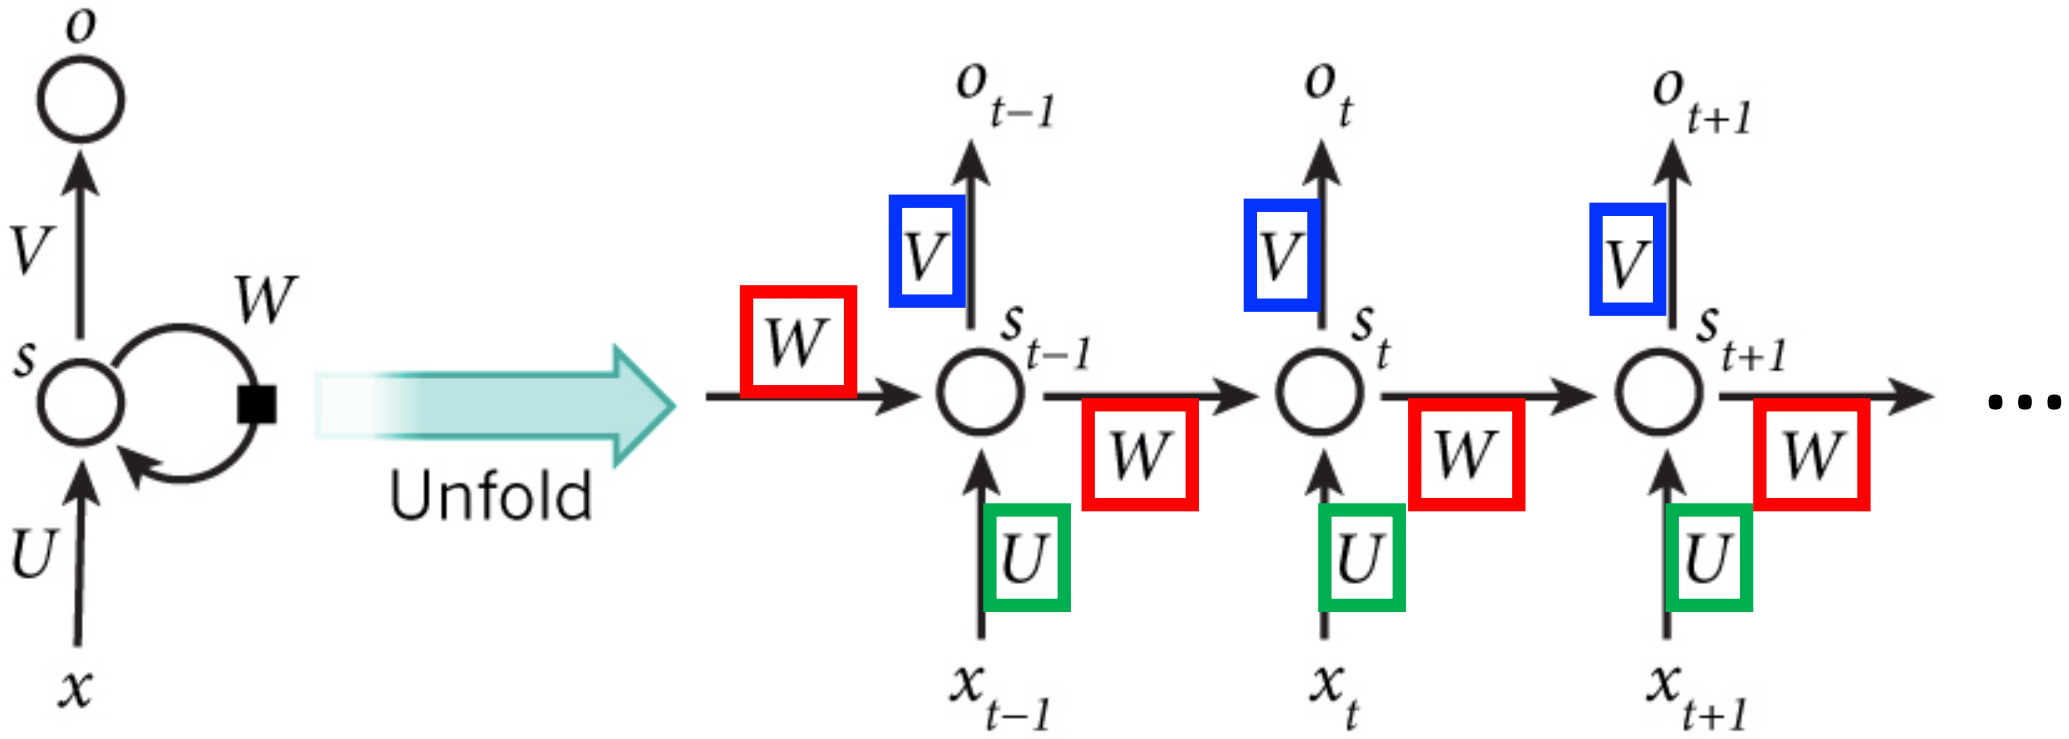

http://www.wildml.com/2015/09/recurrent-neural-networks-tutorial-part-1-introduction-to-rnns/

RNN: Advantages

- Overcomes problem that weights of each layer are learned **independently** by using previous hidden state
- Overcomes problem that model has many parameters since weights are shared across layers

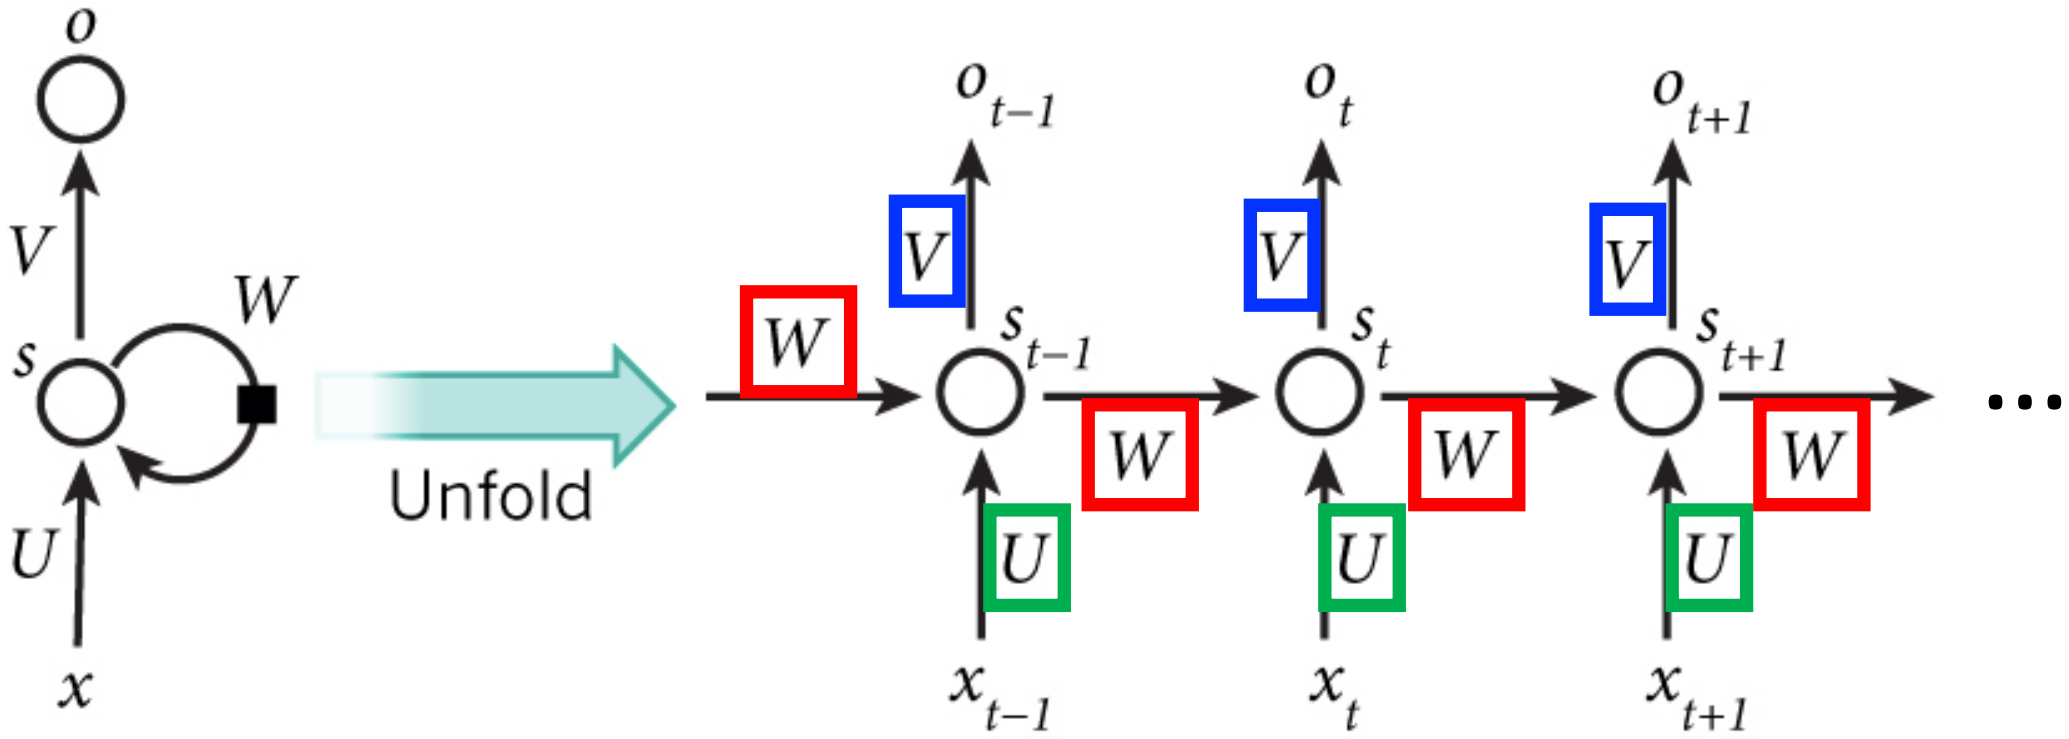

http://www.wildml.com/2015/09/recurrent-neural-networks-tutorial-part-1-introduction-to-rnns/

RNN: Advantages

- Retains information about past inputs for an amount of time that depends on the model's weights and input data rather than a fixed duration selected a priori

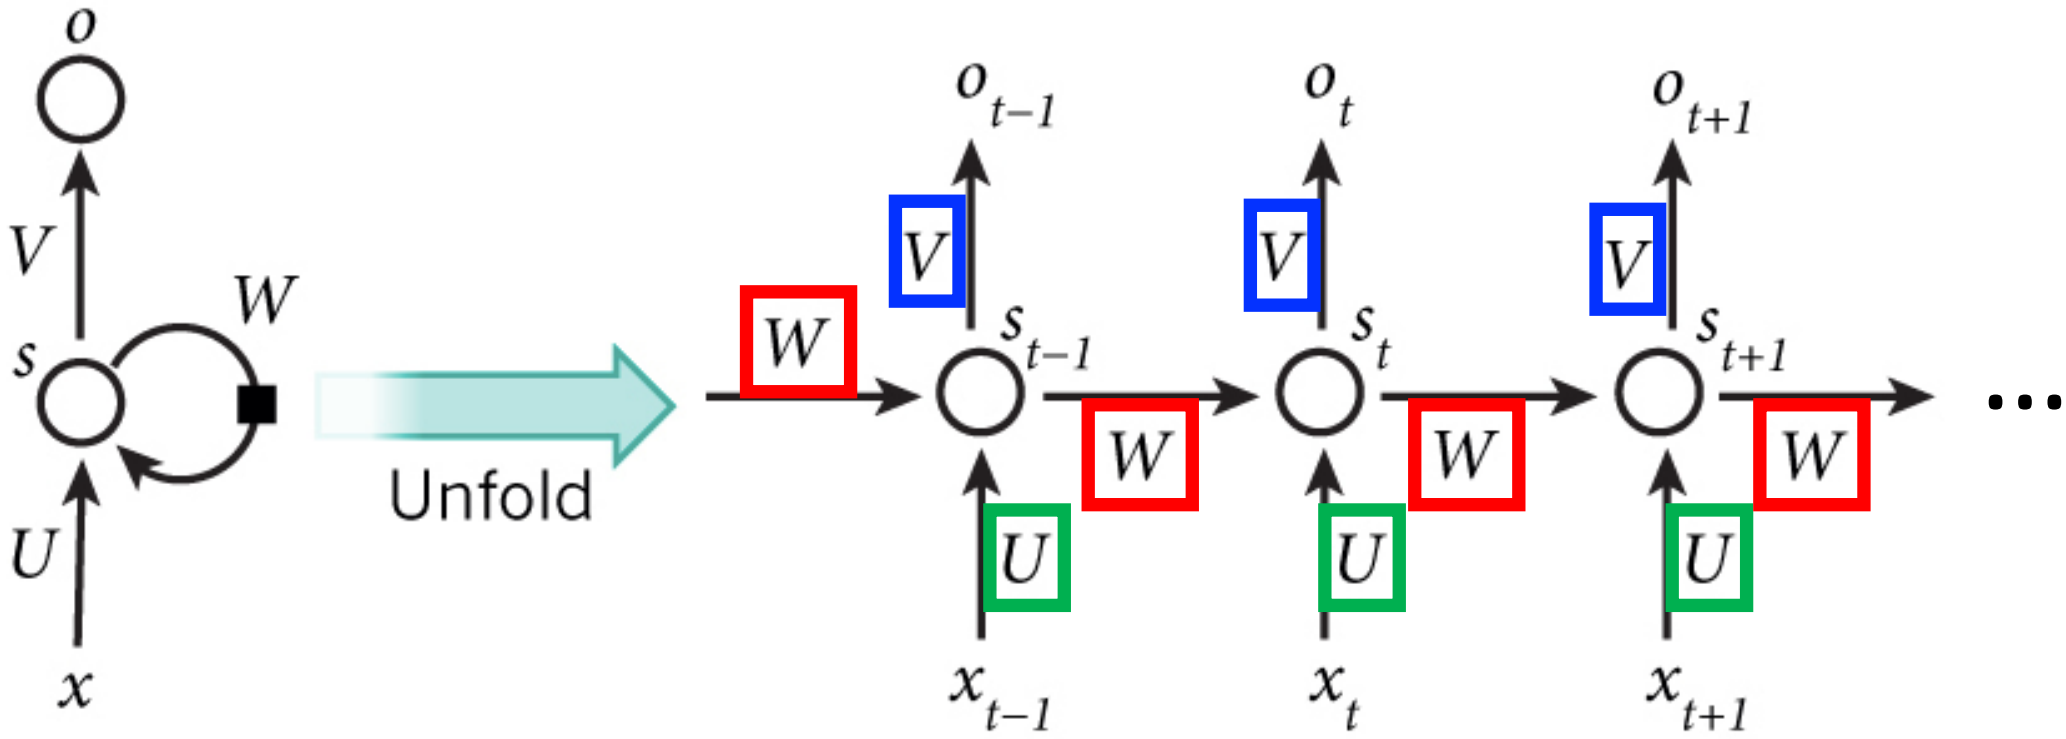

http://www.wildml.com/2015/09/recurrent-neural-networks-tutorial-part-1-introduction-to-rnns/

RNN Example: Predict Sequence of Characters

- Goal: predict next character in text
- Training Data: sequence of characters represented as one-hot vectors

RNN Example: Predict Sequence of Characters; e.g., To Write a Wikipedia Page

Training Input

Predicted Output

```
Naturalism and decision for the majority of Arab countries' capitalide was grounded
by the Irish language by [[John Clair]], [[An Imperial Japanese Revolt]], associated
with Guangzham's sovereignty. His generals were the powerful ruler of the Portugal
in the [[Protestant Immineners]], which could be said to be directly in Cantonese
Communication, which followed a ceremony and set inspired prison, training. The
emperor travelled back to [[Antioch, Perth, October 25|21]] to note, the Kingdom
of Costa Rica, unsuccessful fashioned the [[Thrales]], [[Cynth's Dajoard]], known
in western [[Scotland]], near Italy to the conquest of India with the conflict.
Copyright was the succession of independence in the slop of Syrian influence that
was a famous German movement based on a more popular servicious, non-doctrinal
and sexual power post. Many governments recognize the military housing of the
[[Civil Liberalization and Infantry Resolution 265 National Party in Hungary]],
that is sympathetic to be to the [[Punjab Resolution]]
(PJS)[http://www.humah.yahoo.com/guardian.
cfm/7754800786d17551963s89.htm Official economics Adjoint for the Nazism, Montgomery
was swear to advance to the resources for those Socialism's rule,
was starting to signing a major tripad of aid exile.]]
```

http://karpathy.github.io/2015/05/21/rnn-effectiveness/

RNN Example: Predict Sequence of Characters; e.g., To Write Like Shakespeare

Training Input (All Works of Shakespeare)

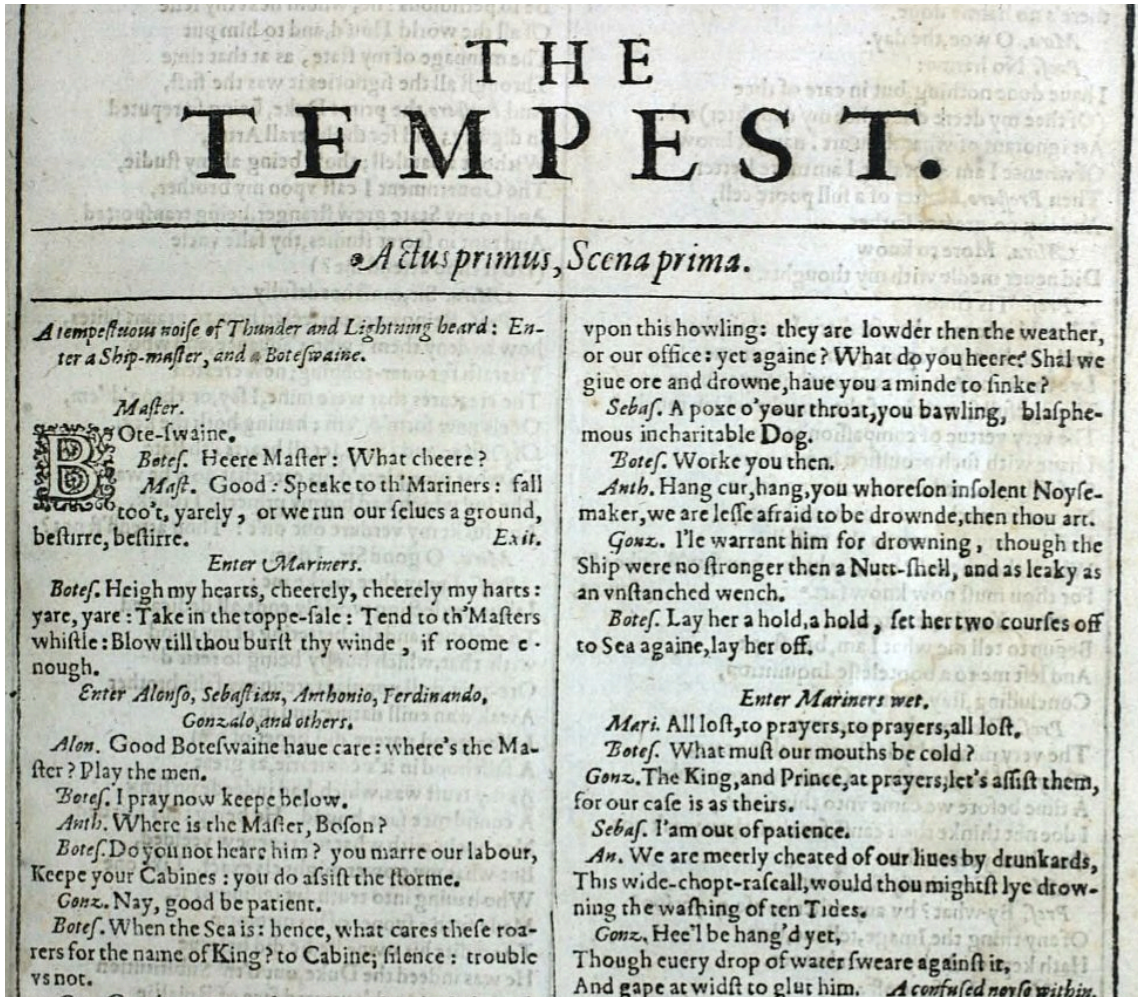

THE
TEMPEST.

Actus primus, Scena prima.

A tempestuous noise of Thunder and Lightning heard: Enter a Ship-master, and a Boteswaine.

Master.

BOte-swaine.

Bote. Heere Master: What cheere?

Mast. Good: Speake to th'Mariners: fall too't, yarely, or we run our selues a ground, bestirre, bestirre. *Exit.*

Enter Mariners.

Bote. Heigh my hearts, cheerely, cheerely my harts: yare, yare: Take in the toppe-sale: Tend to th'Masters whistle: Blow till thou burst thy winde, if roome enough.

Enter Alonso, Sebastian, Anthonio, Ferdinando, Gonzalo, and others.

Alon. Good Boteswaine haue care: where's the Master? Play the men.

Bote. I pray now keepe below.

Anth. Where is the Master, Boson?

Bote. Do you not heare him? you marre our labour, Keepe your Cabines: you do assist the storme.

Gonz. Nay, good be patient.

Bote. When the Sea is: hence, what cares these roarers for the name of King? to Cabine; silence: trouble vs not.

Gonz. Good, yet remember whom thou hast aboord.

Bote. None that I more loue then my selfe. You are a Counsellor, if you can command these Elements to silence, and worke the peace of the present, we will not hand a rope more; vse your authority: if you cannot, giue thankes you haue liu'd so long, and make your selfe ready in your Cabine for the mischance of the houre, if it so hap. Cheerely good hearts: Out of our way I say.

Gonz. I haue great comfort from this fellow: Methinks he hath no drowning marke vpon him; his complexion is perfect Gallowes. Stand fast good Fate to his hanging: make the rope of his destiny our cable, for our owne doth little aduantage. If he be not borne to be hang'd, our case is miserable.

Exeunt.

Enter Boteswaine.

Bote. Downe with the Top-mast: yare, lower, lower, bring her to try with main-course. A plague vpon this howling: they are louder then the weather, or our office: yet againe? What do you heere? Shal we giue ore and drowne, haue you a minde to sinke?

Sebast. A poxe o'your throat, you bawling, blasphemous incharitable Dog.

Bote. Worke you then.

Anth. Hang cur, hang, you whoreson insolent Noyse-maker, we are leſſe afraid to be drownde, then thou art.

Gonz. I'le warrant him for drowning, though the Ship were no stronger then a Nutt-shell, and as leaky as an vnstanched wench.

Bote. Lay her a hold, a hold, set her two courses off to Sea againe, lay her off.

Enter Mariners wet.

Mari. All lost, to prayers, to prayers, all lost.

Bote. What must our mouths be cold?

Gonz. The King, and Prince, at prayers; let's assist them, for our case is as theirs.

Sebast. I'am out of patience.

An. We are merely cheated of our liues by drunkards, This wide-chopt-rascall, would thou mightst lye drowning the washing of ten Tides.

Gonz. Hee'l be hang'd yet, Though euery drop of water sweare against it, And gape at widst to glut him. *A confused noyse within.*

Predicted Output

```
PANDARUS:
Alas, I think he shall be come approached and the day
When little srain would be attain'd into being never fe
And who is but a chain and subjects of his death,
I should not sleep.

Second Senator:
They are away this miseries, produced upon my soul,
Breaking and strongly should be buried, when I perish
The earth and thoughts of many states.

DUKE VINCENTIO:
Well, your wit is in the care of side and that.

Second Lord:
They would be ruled after this chamber, and
my fair nues begun out of the fact, to be conveyed,
Whose noble souls I'll have the heart of the wars.

Clown:
Come, sir, I will make did behold your worship.

VIOLA:
I'll drink it.
```

http://karpathy.github.io/2015/05/21/rnn-effectiveness/ https://github.com/martin-gorner/tensorflow-rnn-shakespeare

RNN Example: Predict Sequence of Characters; e.g., To Write Code

Training Input (C code on GitHub)

```c
/*
 * Bad block management
 *
 * - Heavily based on MD badblocks code from Neil Brown
 *
 * Copyright (c) 2015, Intel Corporation.
 *
 * This program is free software; you can redistribute it and/or modify it
 * under the terms and conditions of the GNU General Public License,
 * version 2, as published by the Free Software Foundation.
 *
 * This program is distributed in the hope it will be useful, but WITHOUT
 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
 * FITNESS FOR A PARTICULAR PURPOSE.  See the GNU General Public License for
 * more details.
 */

#include <linux/badblocks.h>
#include <linux/seqlock.h>
#include <linux/device.h>
#include <linux/kernel.h>
#include <linux/module.h>
#include <linux/stddef.h>
#include <linux/types.h>
#include <linux/slab.h>
```

Predicted Output

```c
 * Increment the size file of the new incorrect UI_FILTER group information
 * of the size generatively.
 */
static int indicate_policy(void)
{
  int error;
  if (fd == MARN_EPT) {
    /*
     * The kernel blank will coeld it to userspace.
     */
    if (ss->segment < mem_total)
      unblock_graph_and_set_blocked();
    else
      ret = 1;
    goto bail;
  }
  segaddr = in_SB(in.addr);
  selector = seg / 16;
  setup_works = true;
  for (i = 0; i < blocks; i++) {
    seq = buf[i++];
    bpf = bd->bd.next + i * search;
    if (fd) {
      current = blocked;
    }
  }
  rw->name = "Getjbbregs";
  bprm_self_clearl(&iv->version);
  regs->new = blocks[(BPF_STATS << info->historidac)] | PFMR_CLOBATHINC_SECONDS << 12;
  return segtable;
}
```

http://karpathy.github.io/2015/05/21/rnn-effectiveness/ https://github.com/martin-gorner/tensorflow-rnn-shakespeare

RNN Example: Predict Sequence of Characters; e.g., To Write Facebook Messages

Training Input

Facebook messages from…

Predicted Output

1. *The meaning of life is to* find them? Oh, I don't know if I would be able to publish a paper on that be climbing today, but it will definitely know what that makes sense. I'm sure they wanted to socialis that I am bringing or

2. *What a cruel twist of fate, that we should* be persuate that 😂 And cook :D I will think that's mean I think I need to go to the phoebe? That's awesome though Haha, sorry, I don't know if it was more time to clas for it's badass though I jus

3. *The fact of the matter is* just the world to invite your stuff? I don't know how to right it wouldn't be as offriving for anything, so that would be awesome, thanks :) I have no idea... She would get to worry about it :P And I

4. *At the very least, you should remember that* as a house of a perfect problems 😂 Yeah :D I wonder how perfect for this trank though So it's probably foltower before the bathers will be fine and haven't want to make it worse Thanks for one of

https://hjweide.github.io/char-rnn

Example: Predict Sequence of Characters

- Goal: predict next character in text
- Prediction: feed training sequence of one-hot encoded characters; e.g., “hello”
  - For simplicity, assume the following vocabulary (i.e., character set): {h, e, l, o}
  - What is our input at time step 1?
  - What is our input at time step 2?
  - What is our input at time step 3?
  - What is our input at time step 4?
  - And so on…

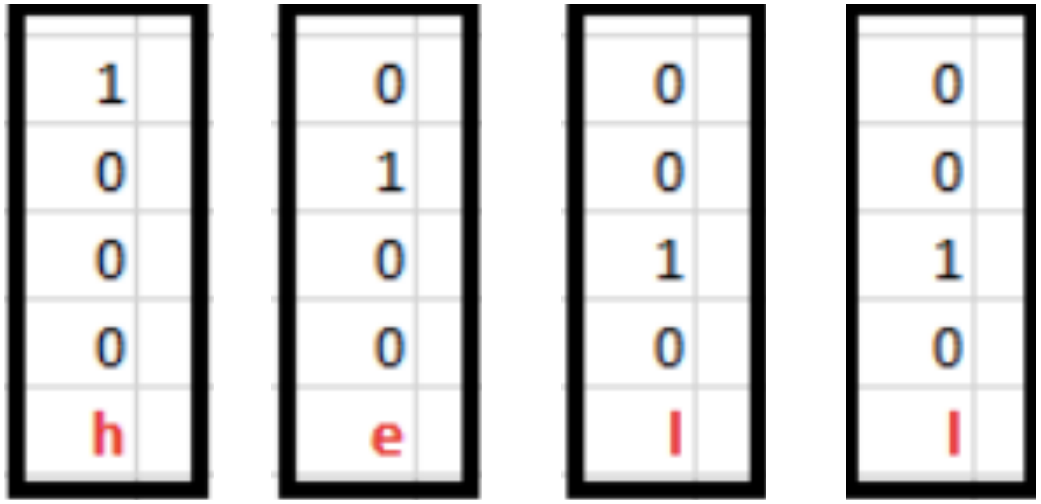

https://www.analyticsvidhya.com/blog/2017/12/introduction-to-recurrent-neural-networks/

Example: Predict Sequence of Characters

Recall activation functions: use tanh as activation function

Sigmoid

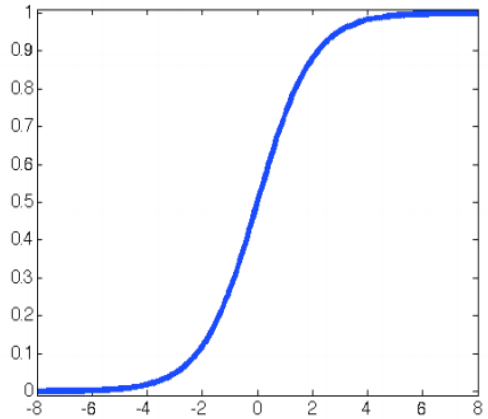

$$\sigma(z) = \frac{1}{1+\exp(-z)}$$

Tanh

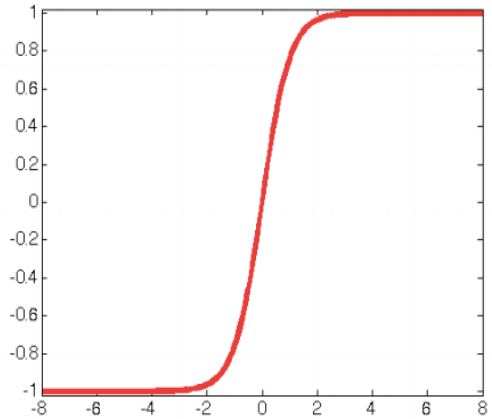

$$\tanh(z) = \frac{\exp(z)-\exp(-z)}{\exp(z)+\exp(-z)}$$

ReLU

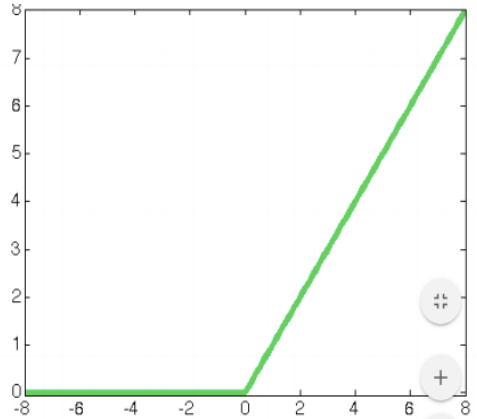

$$\mathrm{ReLU}(z) = \max(0, z)$$

https://www.analyticsvidhya.com/blog/2017/12/introduction-to-recurrent-neural-networks/

Example: Predict Sequence of Characters

Initialize to random value: 0.567001

$$h_t = \tanh(W_{hh} h_{t-1} + W_{xh} x_t + \text{bias})$$

Initialize to 0

wxh			
0.287027	0.84606	0.572392	0.486813
0.902874	0.871522	0.691079	0.18998
0.537524	0.09224	0.558159	0.491528

× | 1 | 0 | 0 | 0 | h | = | 0.287027 | 0.902874 | 0.537524 |

Input at next time step

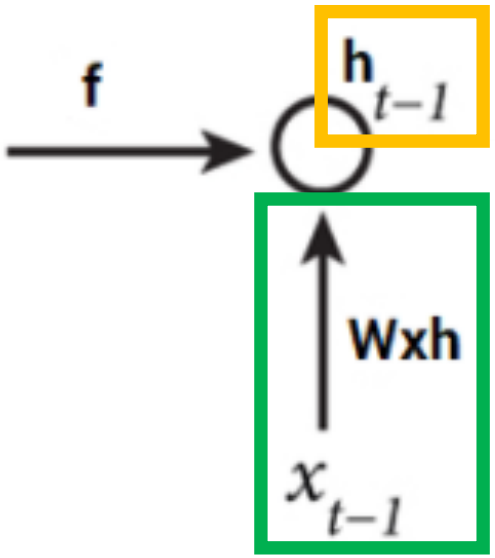

Ht = TANH { 0.287027359 | 0.902874425 | 0.537523791 } + { 0.567001 | 0.567001 | 0.567001 }

= 0.693168 | 0.899554 | 0.802118

https://www.analyticsvidhya.com/blog/2017/12/introduction-to-recurrent-neural-networks/

Example: Predict Sequence of Characters

Initialize to random value: 0.427043

Initialized to random value: 0.567001

$$h_t = \tanh(W_{hh} h_{t-1} + W_{xh} x_t + \text{bias})$$

Output at previous time step

wxh			
0.287027	0.84606	0.572392	0.486813
0.902874	0.871522	0.691079	0.18998
0.537524	0.09224	0.558159	0.491528

0
1
0
0
e

0.84606
0.871522
0.09224

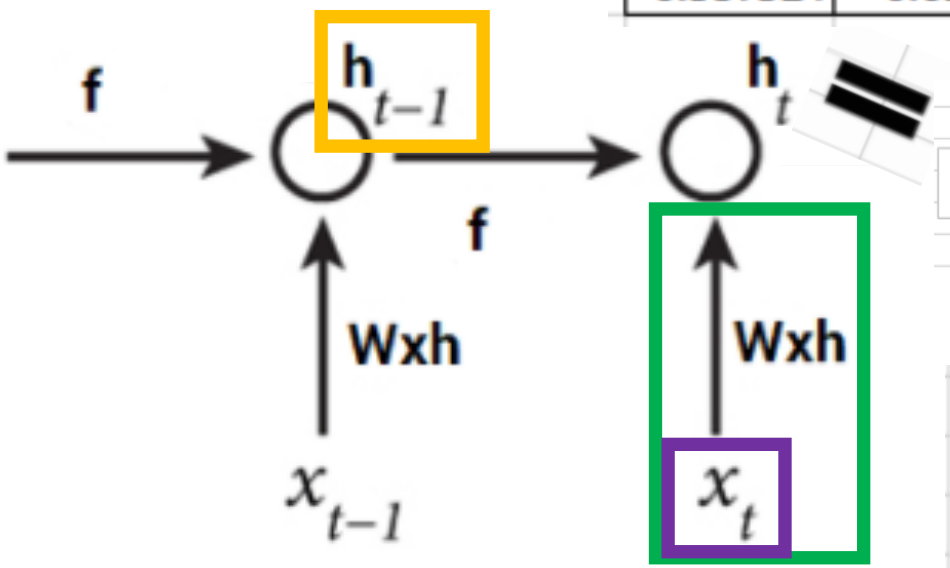

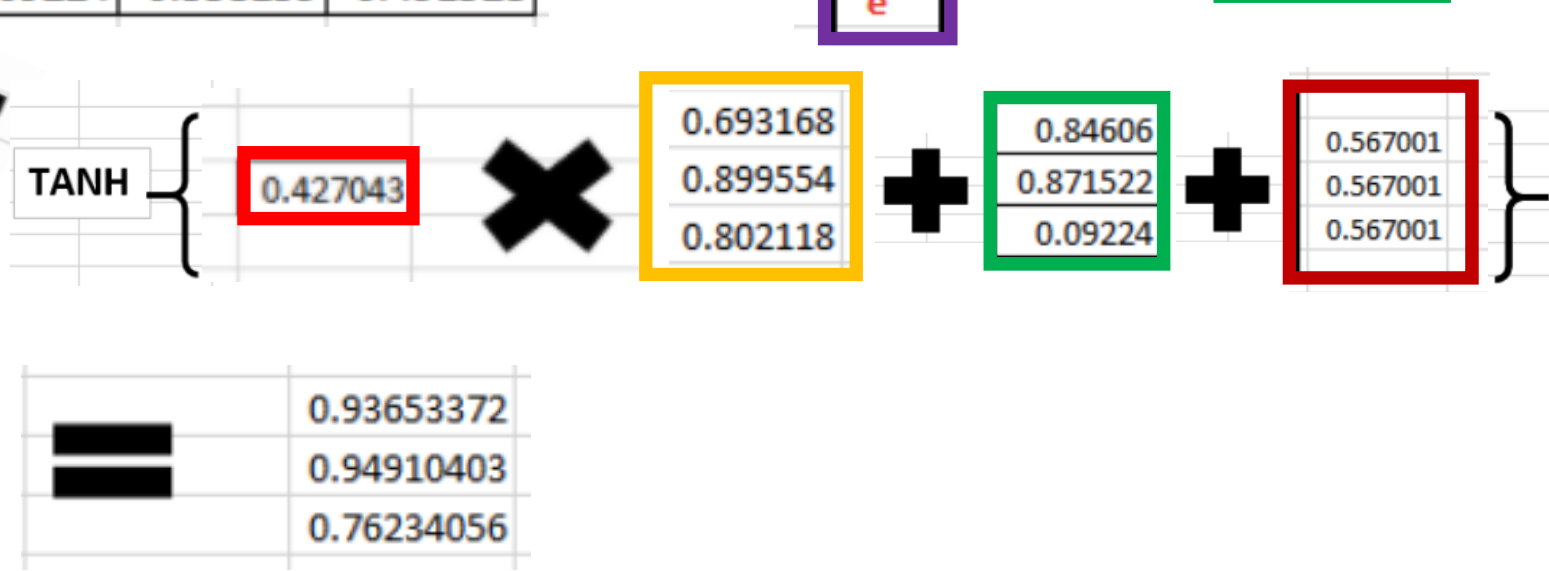

0.93653372
0.94910403
0.76234056

https://www.analyticsvidhya.com/blog/2017/12/introduction-to-recurrent-neural-networks/

Example: Predict Sequence of Characters

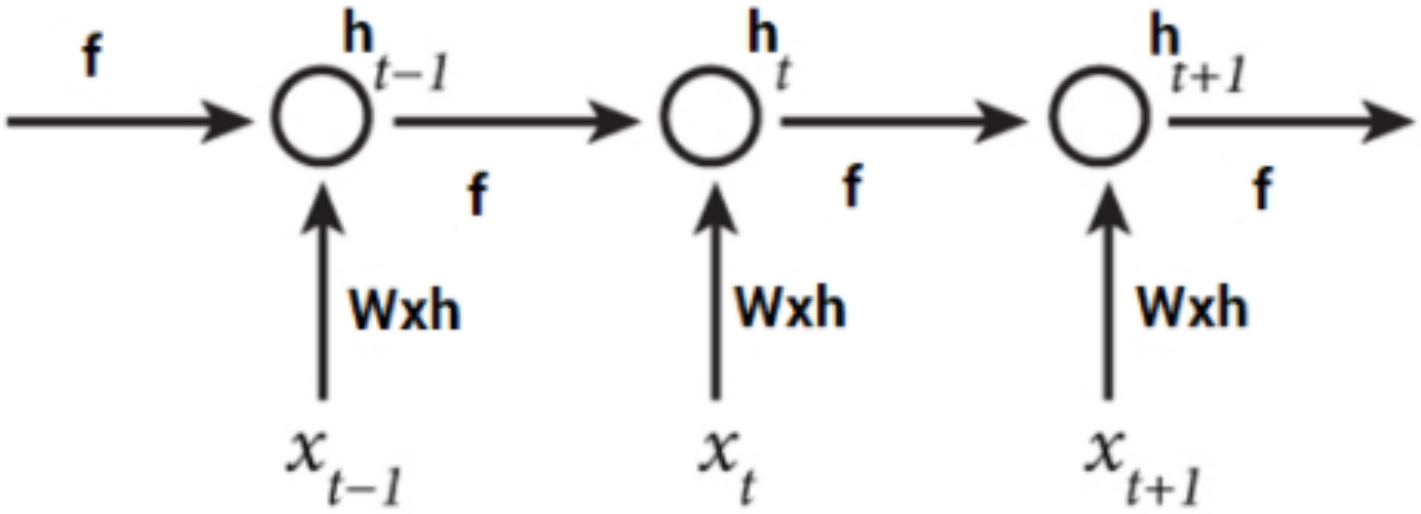

https://www.analyticsvidhya.com/blog/2017/12/introduction-to-recurrent-neural-networks/

Example: Prediction (Many-To-One)

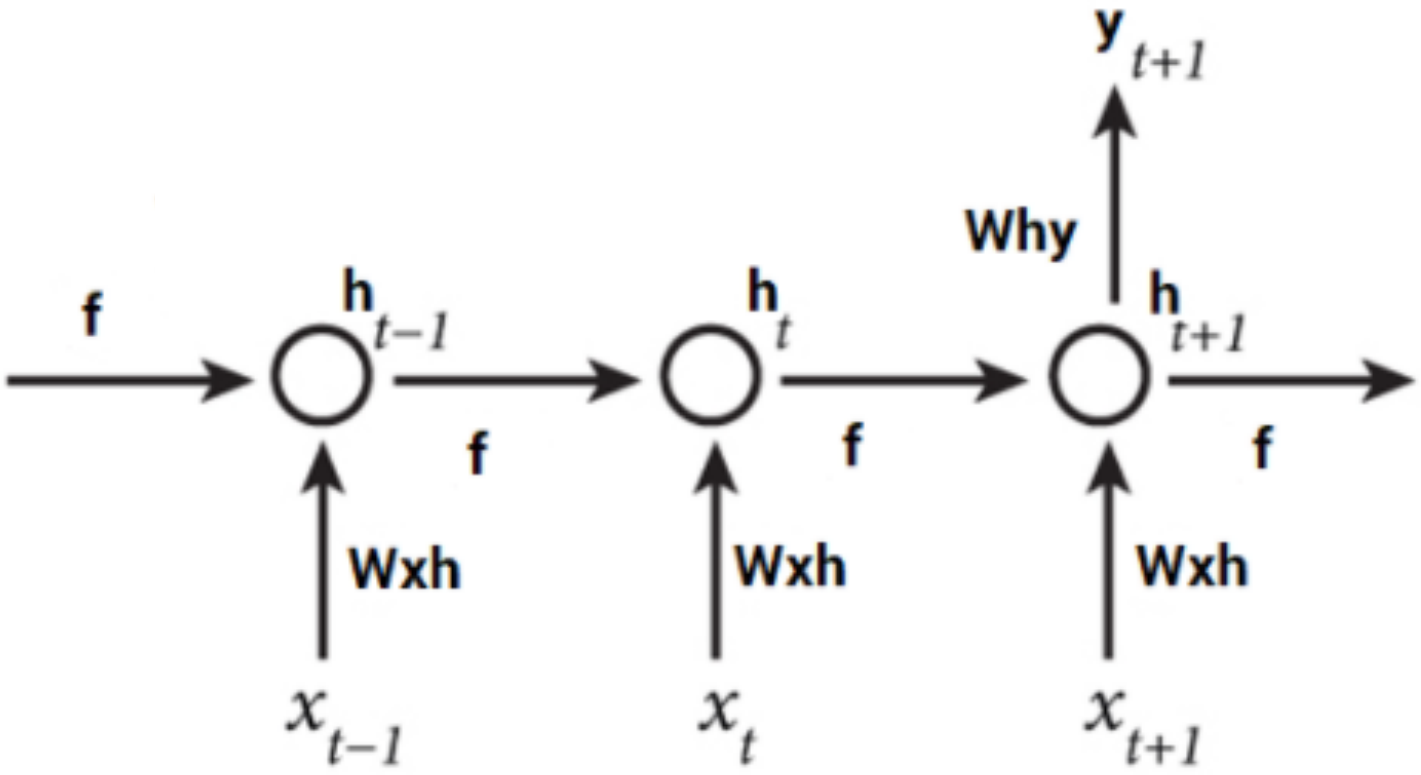

https://www.analyticsvidhya.com/blog/2017/12/introduction-to-recurrent-neural-networks/

Example: Prediction (Many-To-Many)

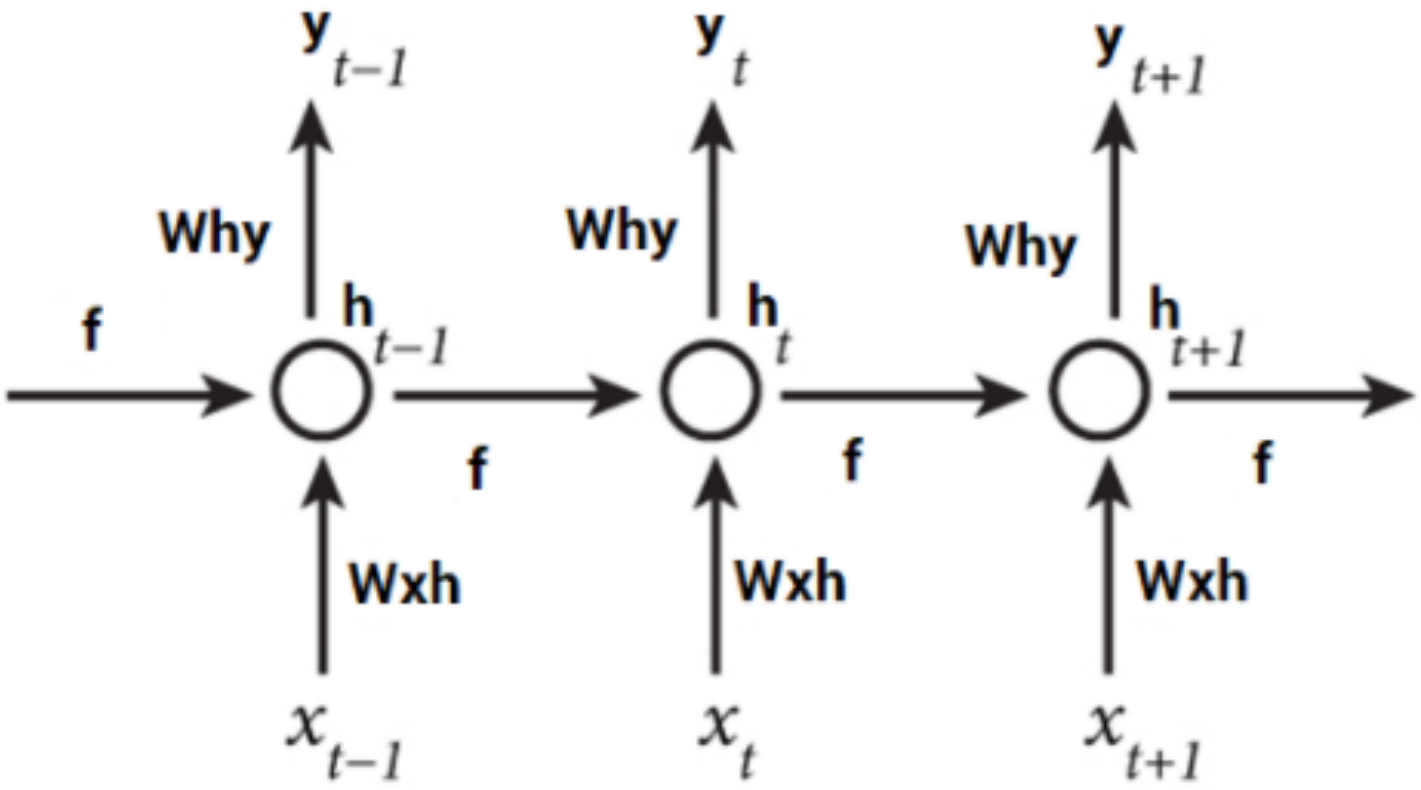

https://www.analyticsvidhya.com/blog/2017/12/introduction-to-recurrent-neural-networks/

Example: Prediction for Time Step 2

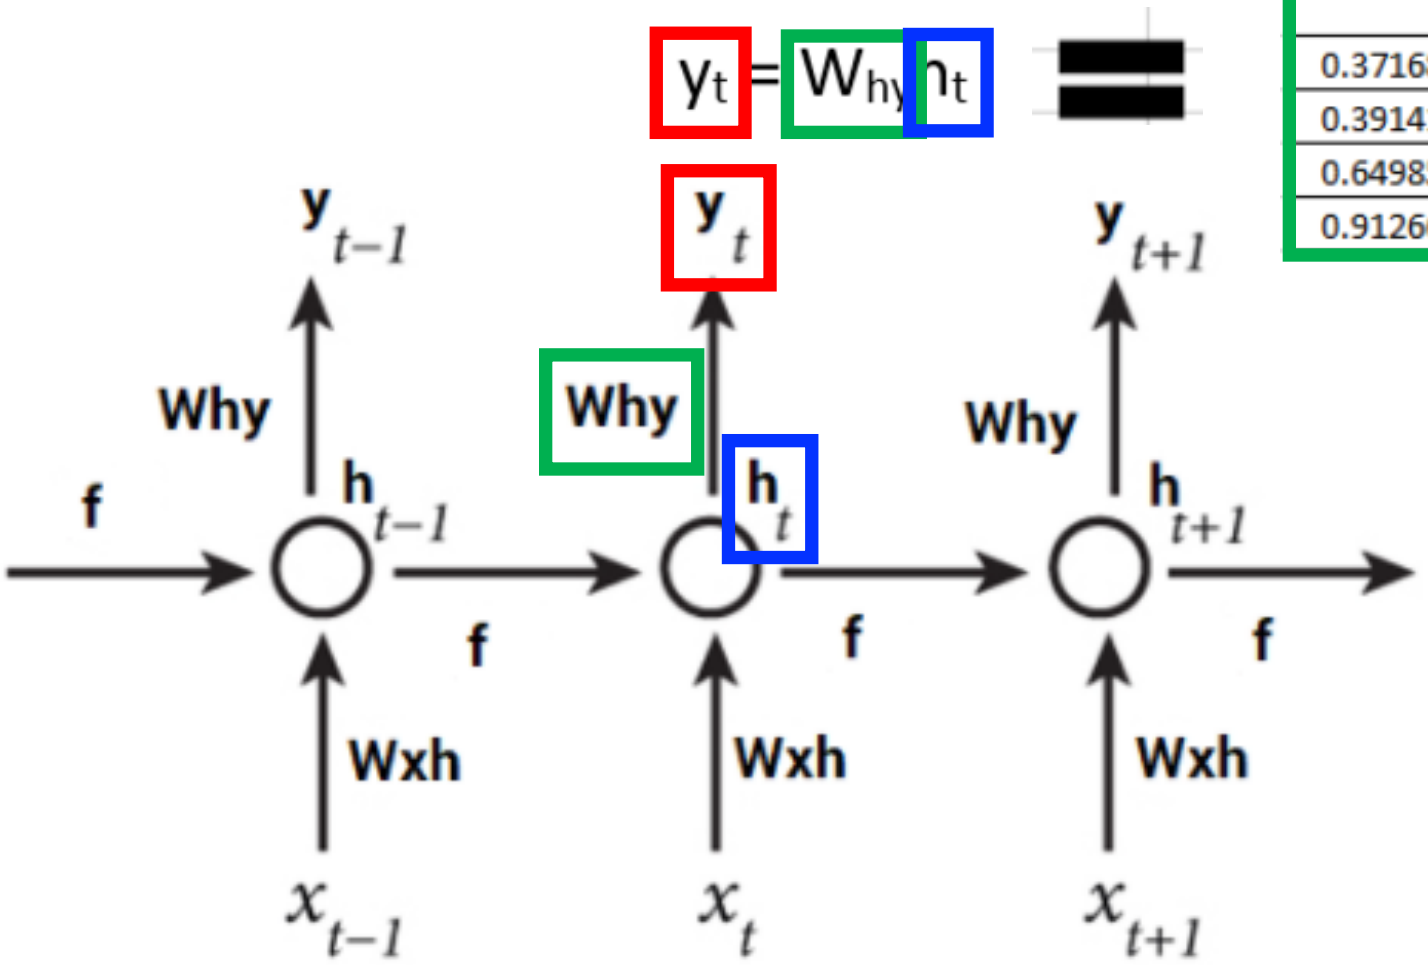

why		
0.37168	0.974829459	0.830034886
0.39141	0.282585823	0.659835709
0.64985	0.09821557	0.334287084
0.91266	0.32581642	0.144630018

×

Ht
0.936534
0.949104
0.762341

=

yt
1.90607732
1.13779113
0.95666016
1.27422602

Applying softmax, to compute letter probabilities:

0.419748
0.194682
0.162429
0.223141

https://www.analyticsvidhya.com/blog/2017/12/introduction-to-recurrent-neural-networks/

Example: Prediction for Time Step 2

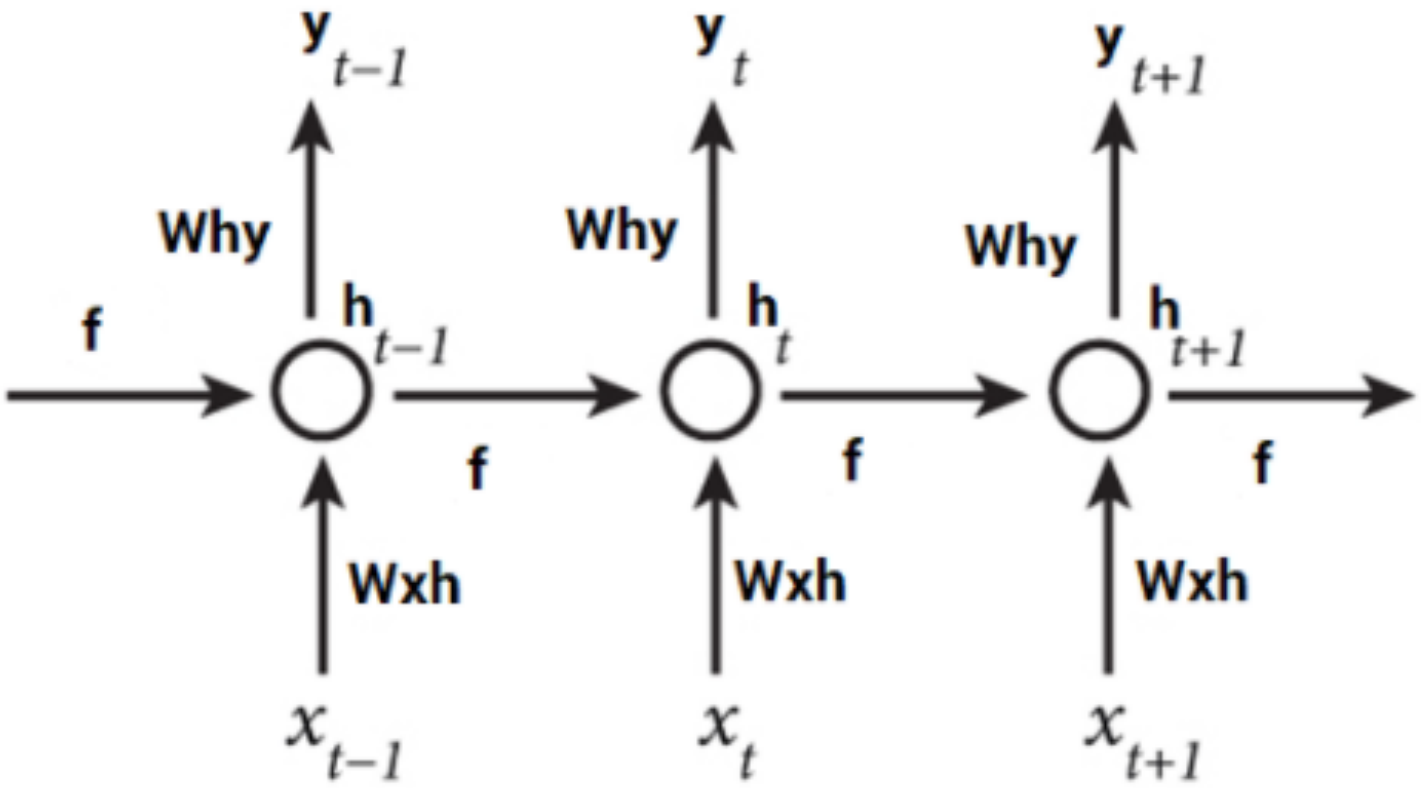

Given our vocabulary is {h, e, l, o}, what letter is predicted?

Applying softmax, to compute letter probabilities:

0.419748
0.194682
0.162429
0.223141

https://www.analyticsvidhya.com/blog/2017/12/introduction-to-recurrent-neural-networks/

RNN Variants: Different Number of Hidden Layers

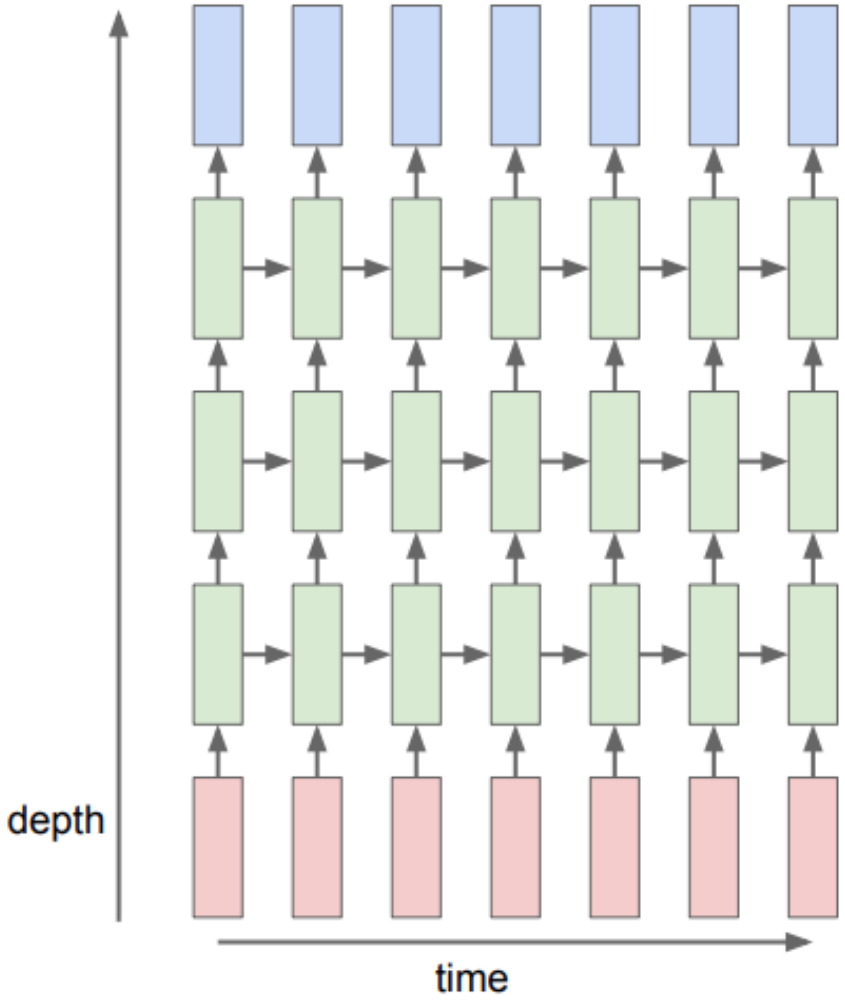

Experimental evidence suggests deeper models can perform better:

- Graves et al.; Speech Recognition with Deep Recurrent Neural Networks; 2013.
- Pascanu et al.; How to Construct Deep Recurrent Neural Networks; 2014.

http://cs231n.stanford.edu/slides/2016/winter1516_lecture10.pdf

RNN: Training

1. Forward pass: make prediction
2. Compute prediction error (with respect to a loss function)
3. Backpropagate error to all model parameters
   - Note: Since weight is same across all time steps, can combine gradients from all time steps
4. Update all model parameters

RNN: Vanishing Gradient Problem

- Problem: training to learn long-term dependencies
  - e.g., language: “In 2004, I started college” vs “I started college in 2004”

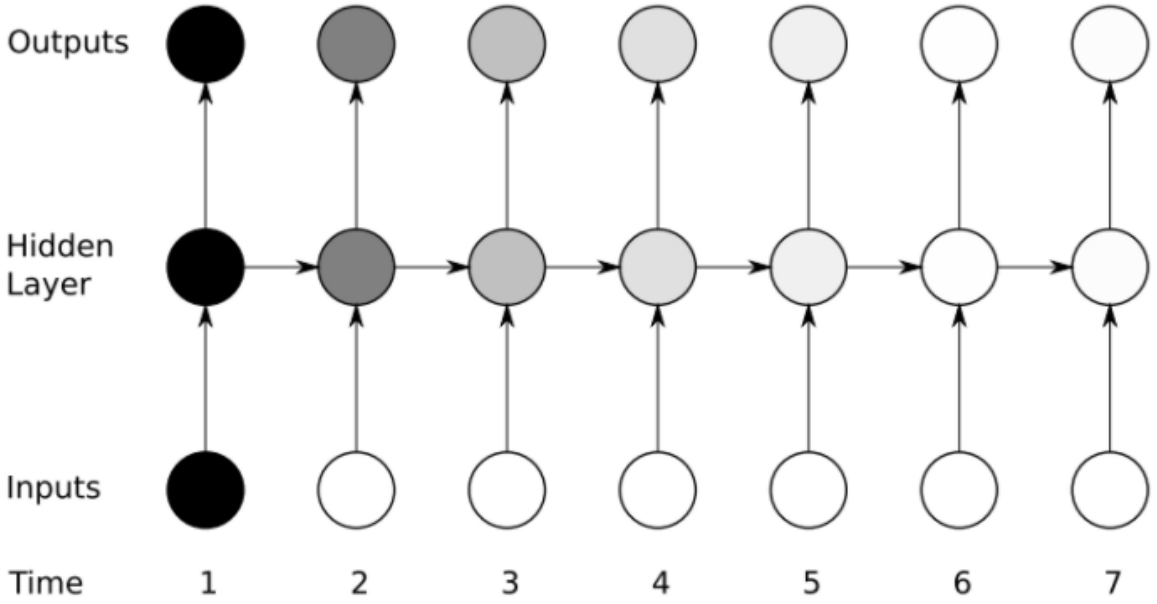

  - e.g., $\partial E/\partial W = \partial E/\partial y_3 * \partial y_3/\partial h_3 * \partial h_3/\partial y_2 * \partial y_2/\partial h_1$
- Vanishing gradient: a product of numbers less than 1 shrinks to zero
- Exploding gradient: a product of numbers greater than 1 explodes to infinity

https://www.analyticsvidhya.com/blog/2017/12/introduction-to-recurrent-neural-networks/

RNN History

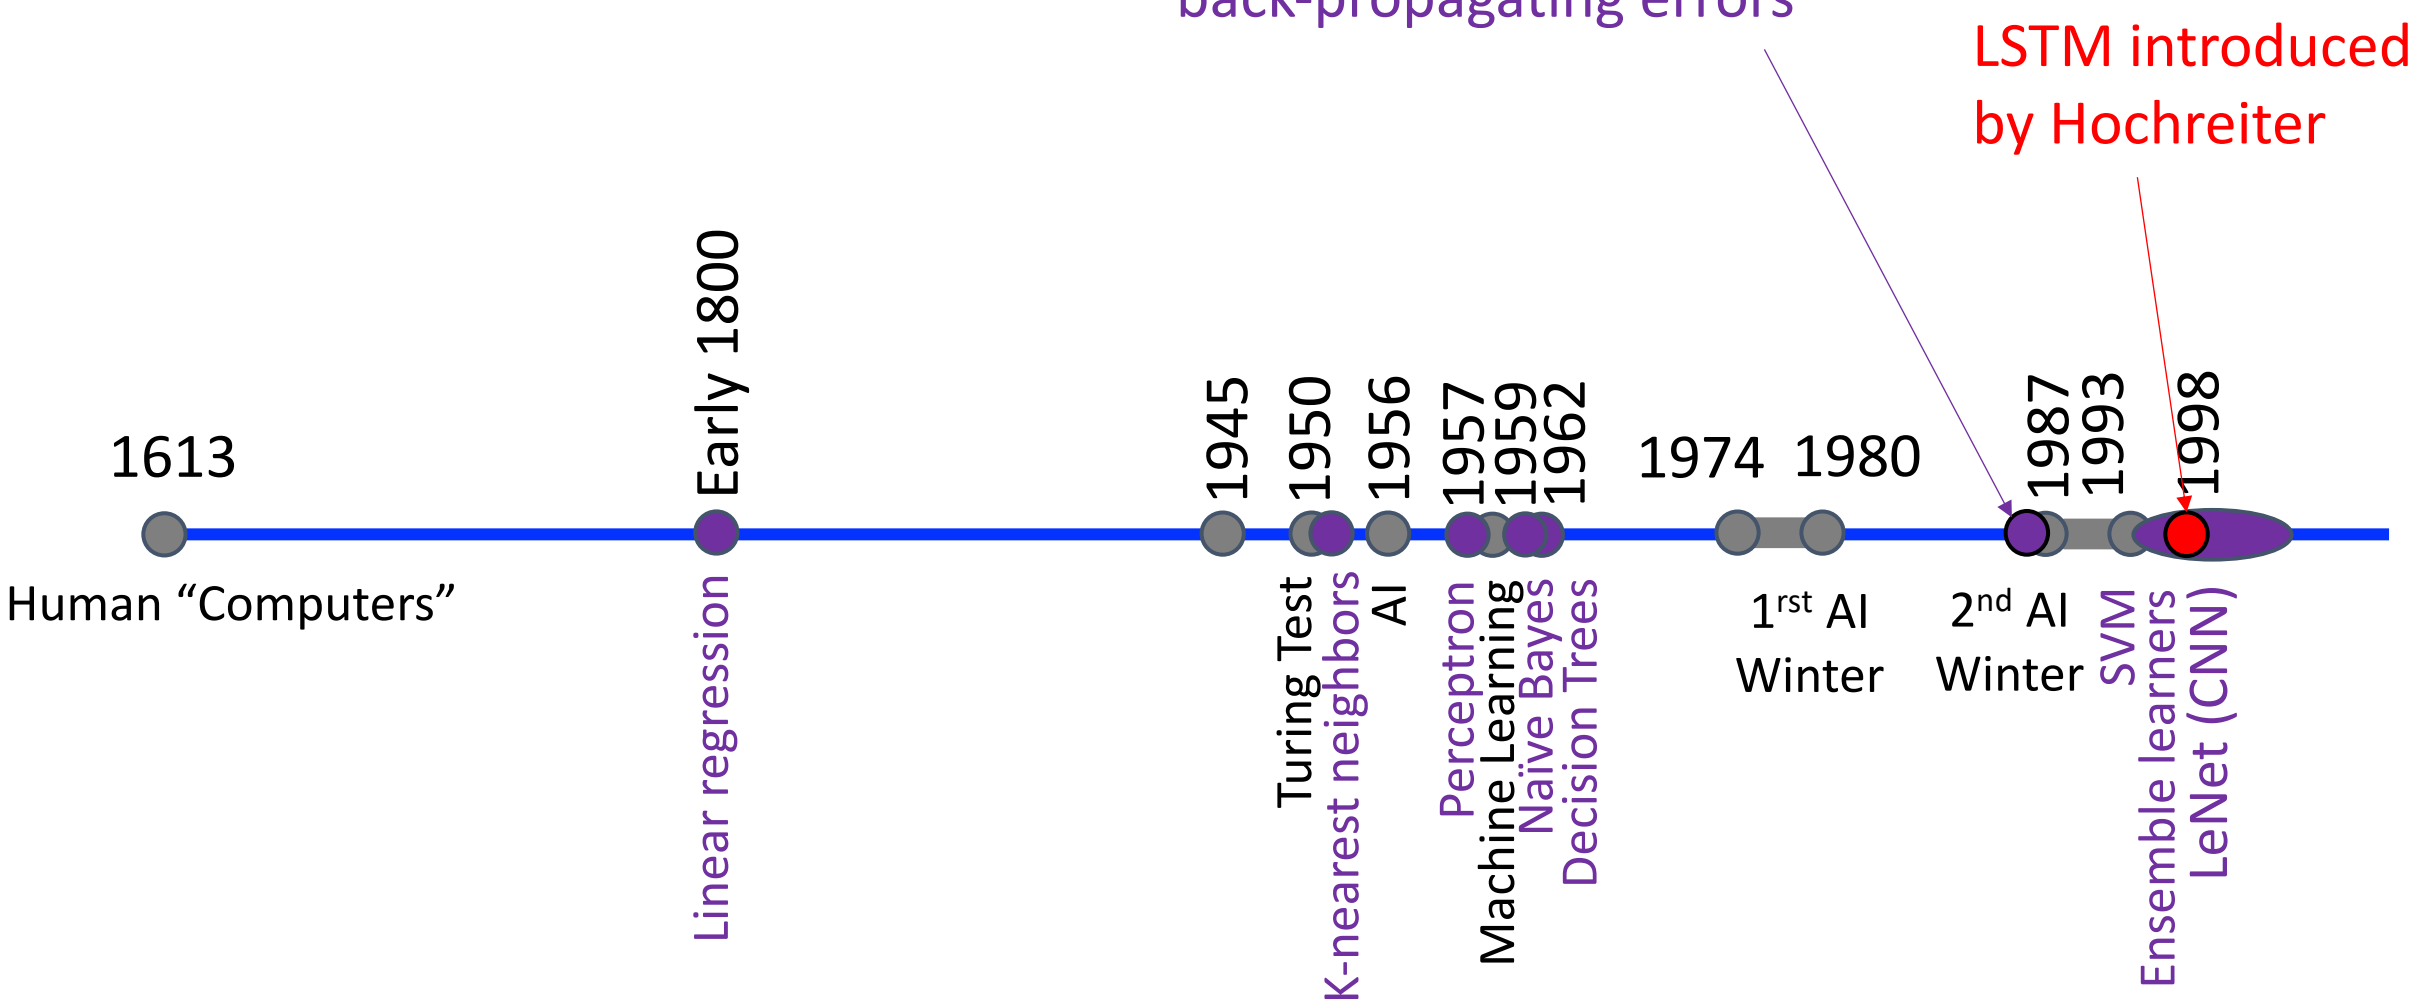
Rumelhart, Hinton, & Williams.
"Learning Representations by
back-propagating errors"
LSTM introduced
by Hochreiter
1613
Early 1800
1945
1950
1956
1957
1959
1962
1974
1980
1987
1993
1998
Human "Computers"
Linear regression
Turing Test
K-nearest neighbors
AI
Perceptron
Machine Learning
Naïve Bayes
Decision Trees
1rst AI
Winter
2nd AI
Winter
SVM
Ensemble learners
LeNet (CNN)

RNN Variants: Mitigate Vanishing Gradients

- Long Short Term Memory (LSTM): modification of basic RNN architecture
- Introduces preservation of memory over time

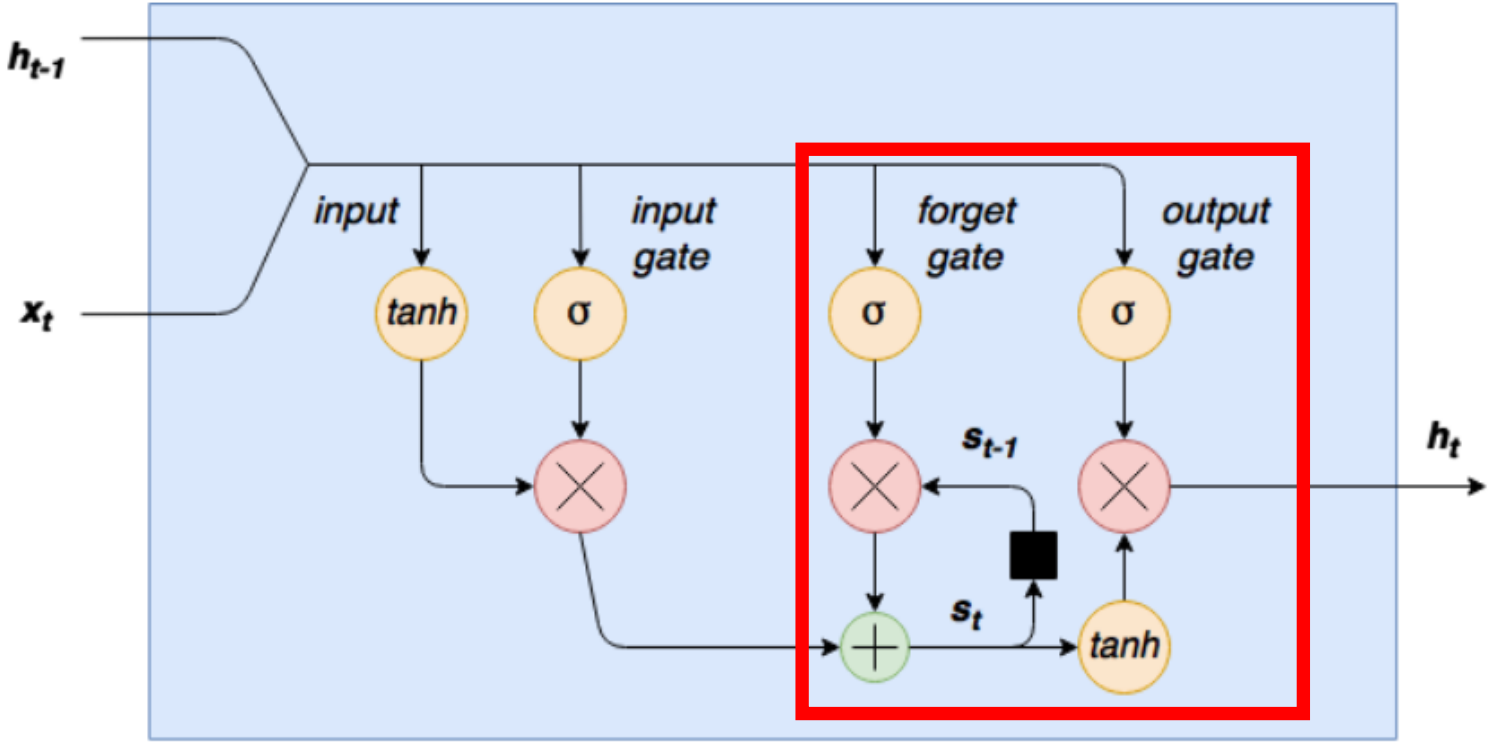

- Forget gate: determines which previous states to remember (i.e. forget gate output near 1) and which to forget (i.e. forget gate output near 0)

Figure Credit: http://adventuresinmachinelearning.com/recurrent-neural-networks-lstm-tutorial-tensorflow/

RNN Variants: Mitigate Vanishing Gradients

- Gated Recurrent Unit (GRU): simplification of LSTM unit to merge forget and input gates

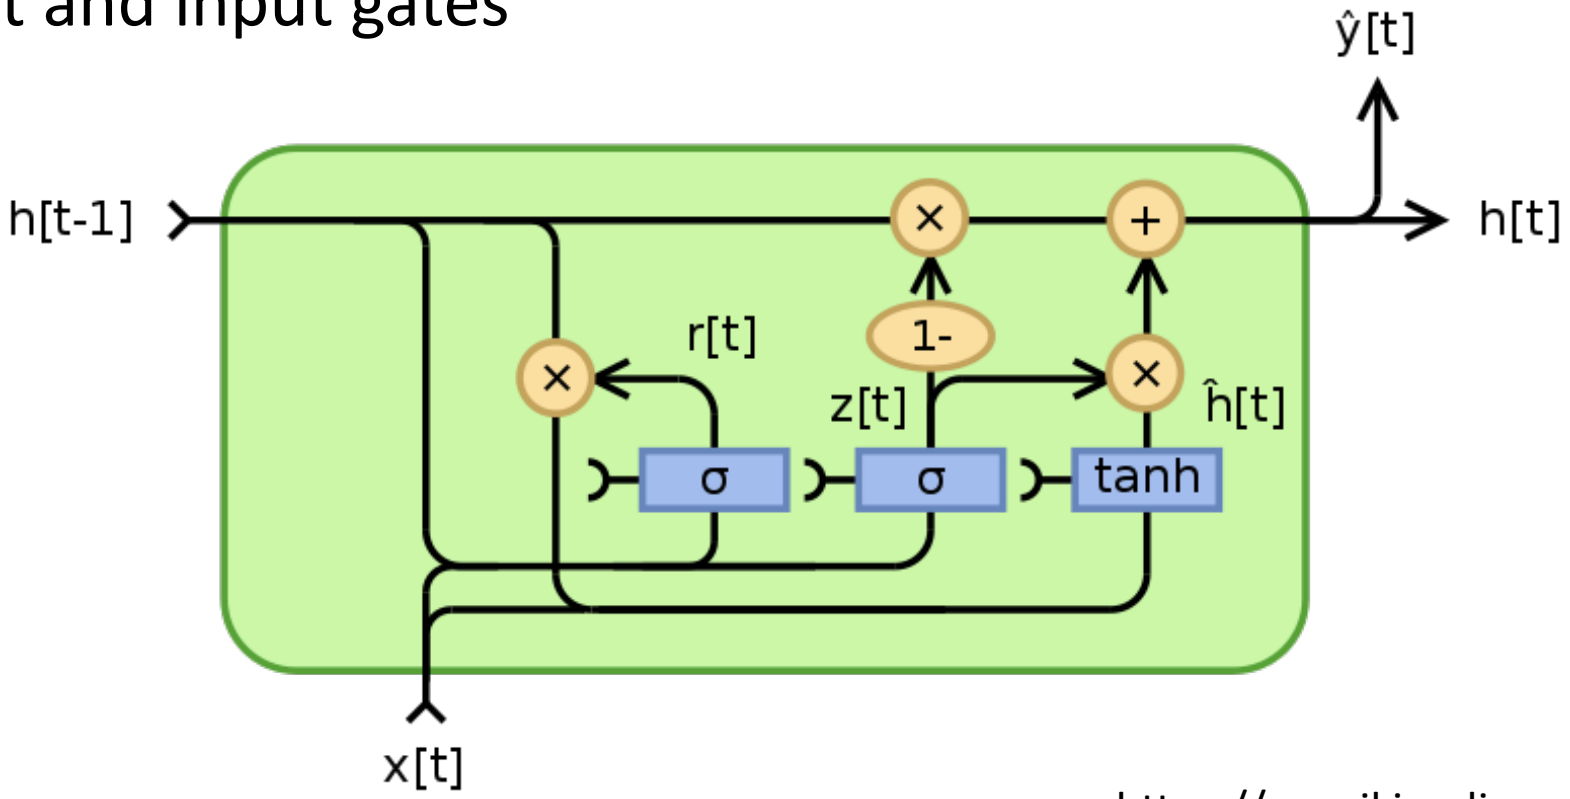

https://en.wikipedia.org/wiki/Gated_recurrent_unit

K. Chou et al; Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation; 2014

Group Discussion: Empirical Model Comparison

1. **What is the best model to use and why?**
2. **What is a good number of layers and why?**
3. **What is a good number of neurons and why?**

	LSTM			RNN			GRU		
Layers	1	2	3	1	2	3	1	2	3
Size	War and Peace Dataset								
64	1.449	1.442	1.540	1.446	1.401	1.396	1.398	**1.373**	1.472
128	1.277	1.227	1.279	1.417	1.286	1.277	1.230	**1.226**	1.253
256	1.189	**1.137**	1.141	1.342	1.256	1.239	1.198	1.164	1.138
512	1.161	1.092	1.082	-	-	-	1.170	1.201	**1.077**
Linux Kernel Dataset									
64	1.355	**1.331**	1.366	1.407	1.371	1.383	1.335	1.298	1.357
128	1.149	1.128	1.177	1.241	**1.120**	1.220	1.154	1.125	1.150
256	1.026	**0.972**	0.998	1.171	1.116	1.116	1.039	0.991	1.026
512	0.952	0.840	0.846	-	-	-	0.943	0.861	**0.829**

A. Karpathy, J. Johnson, and L. Fei-Fei; Visualizing and Understanding Recurrent Networks; 2016

Today's Topics

- Machine Learning for Sequential Data
- Recurrent Neural Networks (RNNs)
- Training Deep Neural Networks: Hardware & Software

Recall: Machine Learning Analogous to a Love Story of Partnering Up and Road Tripping Somewhere

An **algorithm** **learns** from **data** patterns that a final model will use to make a prediction

Recall: Machine Learning Analogous to a Love Story of Partnering Up and Road Tripping Somewhere

Key Issue: How Fast Will It Take to Get There?

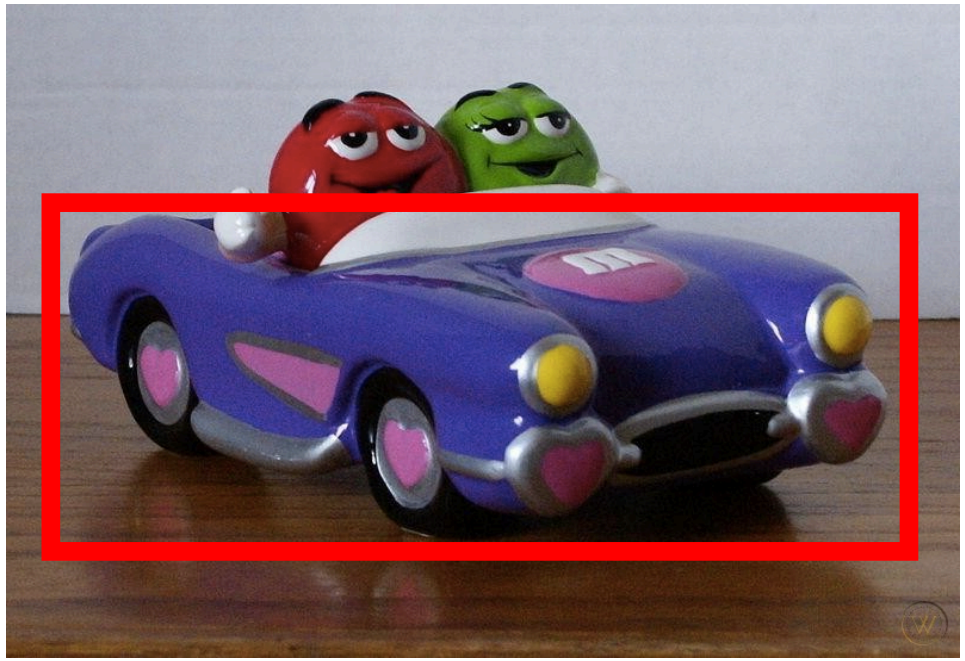

Challenge: Training Neural Network **Requires Many Computations** (e.g., millions of model parameters)

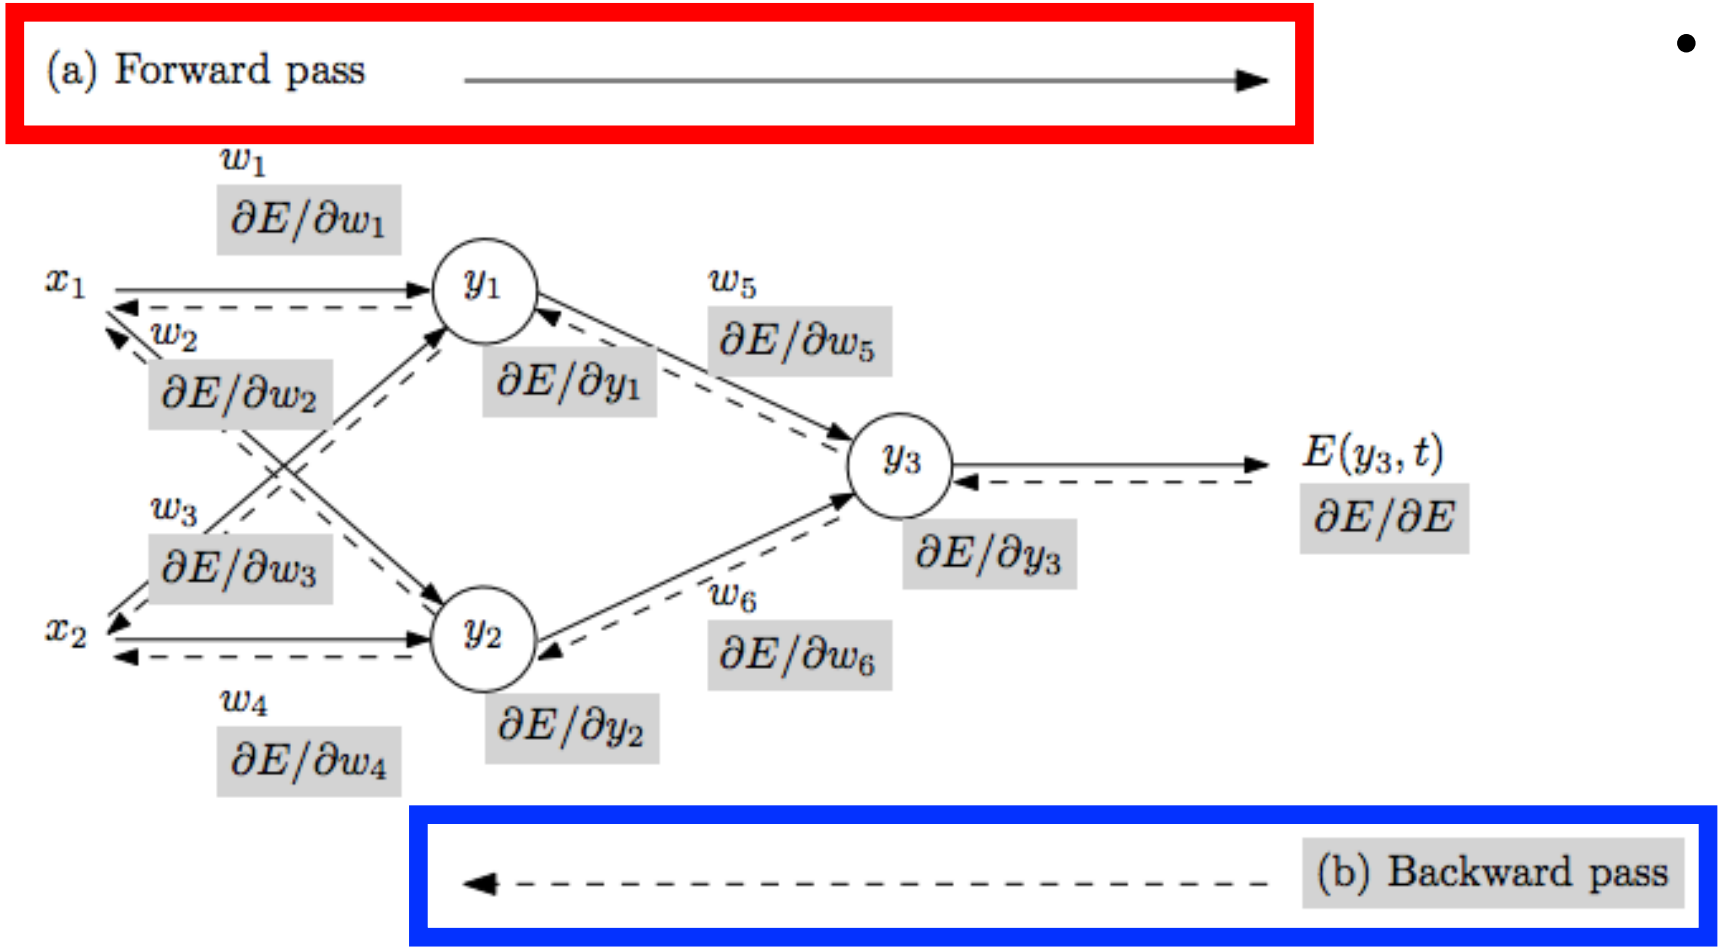

- **Repeat until stopping criterion met:**
  1. **Forward pass**: propagate training data through network to make prediction
  2. **Backward pass**: using predicted output, calculate gradients backward
  3. Update each weight using calculated gradients

Figure from: Atilim Gunes Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, Jeffrey Mark Siskind; Automatic Differentiation in Machine Learning: a Survey; 2018

Recall: Machine Learning Analogous to a Love Story of Partnering Up and Road Tripping Somewhere

Idea: Train Algorithms Using GPUs (think Porsche) Instead of CPUs (think Golf Cart)

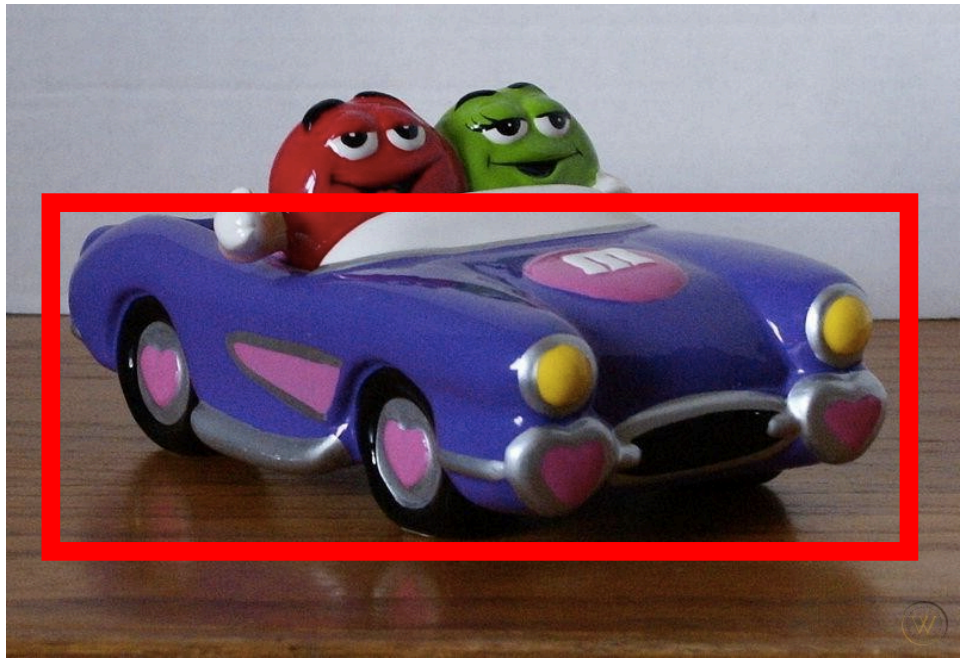

Hardware: CPU versus GPU

Spot the CPU!

(central processing unit)

http://cs231n.stanford.edu/slides/2018/cs231n_2018_lecture08.pdf

Hardware: CPU versus GPU

Spot the GPUs!

(graphics processing unit)

This image is in the public domain

http://cs231n.stanford.edu/slides/2018/cs231n_2018_lecture08.pdf

Hardware: CPU versus GPU

- Graphical Processing Units: accelerates computational workloads due to MANY more processing cores

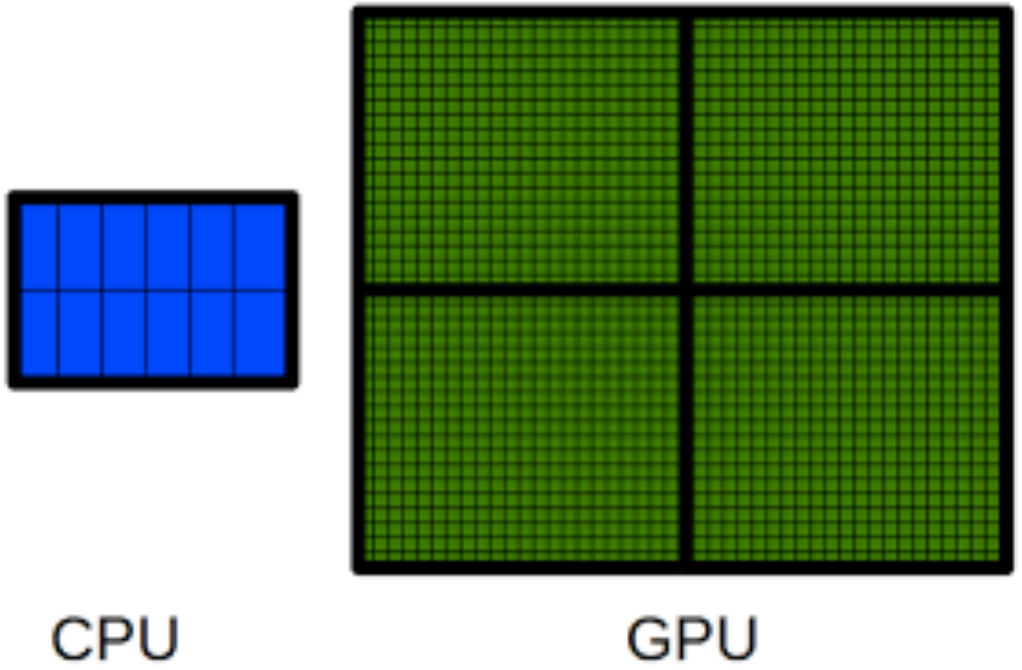

https://www.researchgate.net/figure/The-main-difference-between-CPUs-and-GPUs-is-related-to-the-number-of-available-cores-A_fig7_273383346

Hardware: Training Models with GPUs

Model is here

Data is here

If you aren't careful, training can bottleneck on reading data and transferring to GPU!

Solutions:

- Read all data into RAM
- Use SSD instead of HDD
- Use multiple CPU threads to prefetch data

(TPU = Tensor Processing Unit)

Hardware: CPU versus GPU

	Cores	Clock Speed	Memory	Price	Speed
CPU (Intel Core i7-7700k)	4 (8 threads with hyperthreading)	4.2 GHz	System RAM	$339	~540 GFLOPs FP32
GPU (NVIDIA GTX 1080 Ti)	3584	1.6 GHz	11 GB GDDR5 X	$699	~11.4 TFLOPs FP32
TPU NVIDIA TITAN V	5120 CUDA, 640 Tensor	1.5 GHz	12GB HBM2	$2999	~14 TFLOPs FP32 ~112 TFLOP FP16
TPU Google Cloud TPU	?	?	64 GB HBM	$6.50 per hour	~180 TFLOP

CPU: Fewer cores, but each core is much faster and much more capable; great at sequential tasks

GPU: More cores, but each core is much slower and "dumber"; great for parallel tasks

TPU: Specialized hardware for deep learning

http://cs231n.stanford.edu/slides/2018/cs231n_2018_lecture08.pdf

GPU Clusters (Google Cloud's TPU Servers)

https://www.extremetech.com/extreme/249499-google-takes-swipe-nvidia-powerful-new-learning-capable-cloud-tpu

GPU Machines: Rent Versus Buy?

Rent from Cloud (Microsoft Azure):

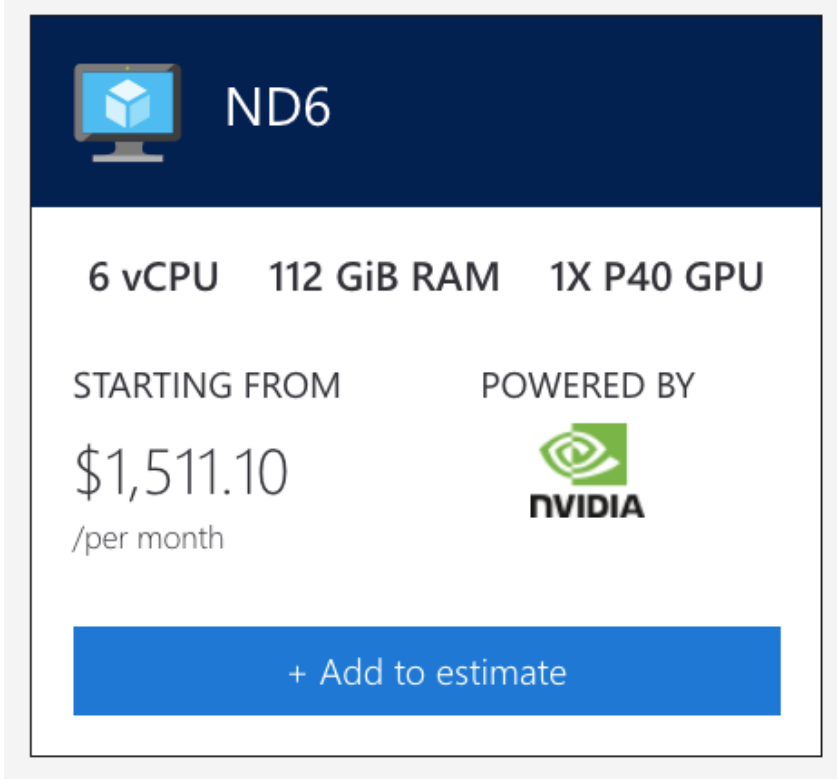

Buy:

Basic

2x RTX 2080 Ti

2-Way NVLink

Intel i9-9820X (10 cores, 3.30 GHz)

2x RTX 2080 Ti (11 GB VRAM)

64 GB RAM

2 TB SSD

4 TB HDD

Starting at

$7,059

Customize

Rise of “Deep Learning” Open Source Platforms

Motivation:

Can run on GPUs:

OpenMP support	OpenCL support	CUDA support	Automatic differentiation[1]

Simplifies using popular neural network architectures:

Has pretrained models	Recurrent nets	Convolutional nets	RBM/DBNs	Parallel execution (multi node)

https://en.wikipedia.org/wiki/Comparison_of_deep_learning_software

Rise of “Deep Learning” Open Source Platforms

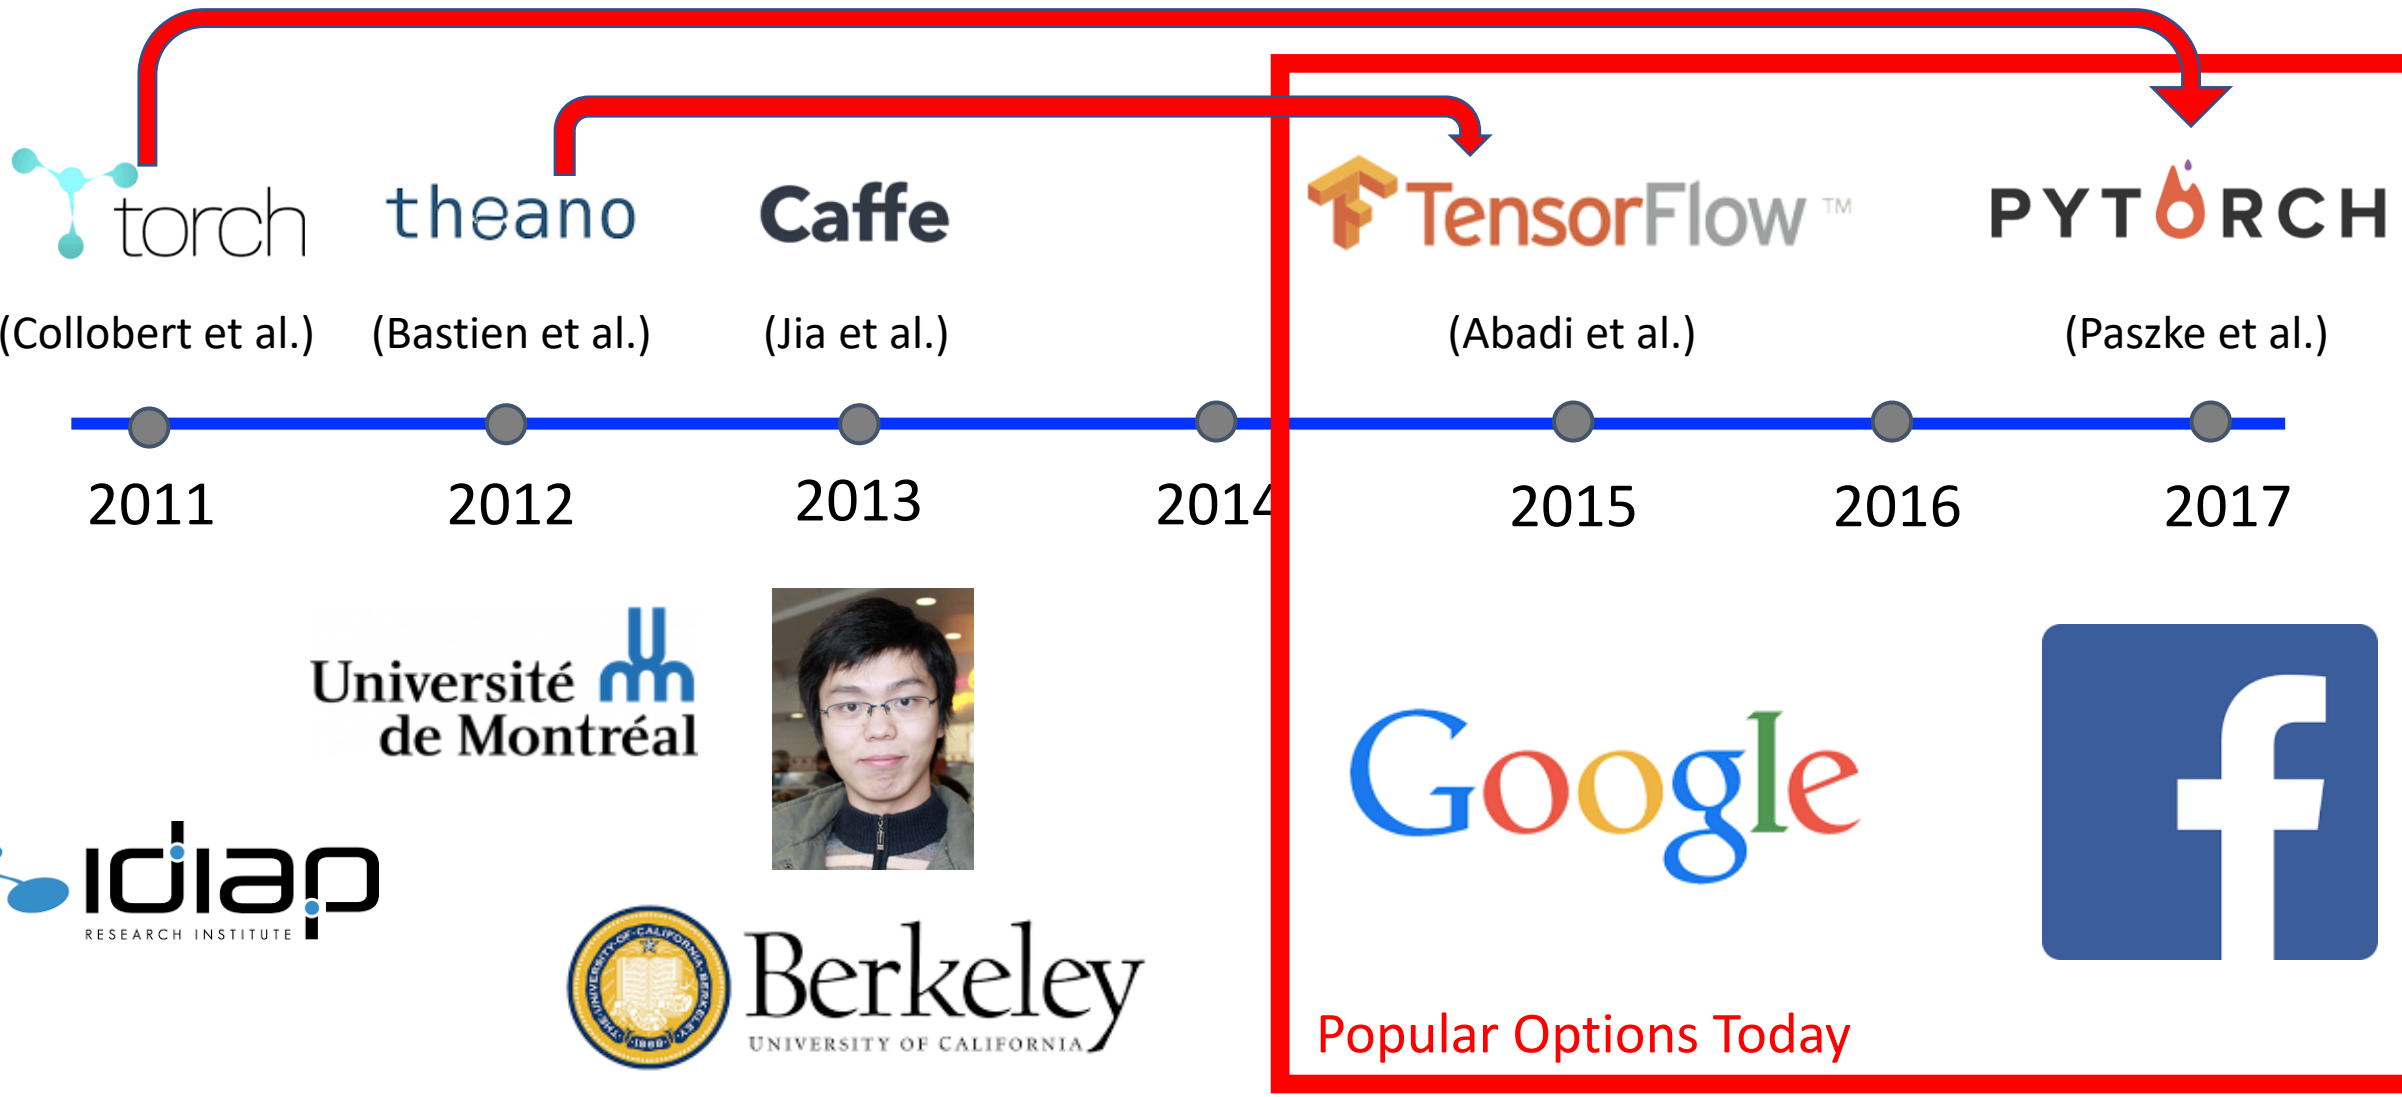

Rise of “Deep Learning” Open Source Platforms

Software	Creator	Software license[a]	Open source	Platform	Written in	Interface	OpenMP support	OpenCL support	CUDA support	Automatic differentiation[1]	Has pretrained models	Recurrent nets	Convolutional nets	RBM/DBNs	Parallel execution (multi node)	Actively Developed
roNNie.ai	Kevin Lok	MIT license	Yes	Linux, macOS, Windows	Python	Python			Yes		Yes	Yes	Yes			
BigDL	Jason Dai	Apache 2.0	Yes	Apache Spark	Scala	Scala, Python			No		Yes	Yes	Yes			
Caffe	Berkeley Vision and Learning Center	BSD	Yes	Linux, macOS, Windows[2]	C++	Python, MATLAB, C++	Yes	Under development[3]	Yes	Yes	Yes[4]	Yes	Yes	No	?	
Deeplearning4j	Skymind engineering team; Deeplearning4j community; originally Adam Gibson	Apache 2.0	Yes	Linux, macOS, Windows, Android (Cross-platform)	C++, Java	Java, Scala, Clojure, Python (Keras), Kotlin	Yes	On roadmap[5]	Yes[6][7]	Computational Graph	Yes[8]	Yes	Yes	Yes	Yes[9]	
Chainer	Preferred Networks	MIT license	Yes	Linux, macOS, Windows		Python	No	No[10][11]	Yes	Yes	Yes	Yes	Yes			
Darknet	Joseph Redmon	Public Domain	Yes	Cross-Platform	C	C, Python	Yes	No[12]	Yes	Yes						
Dlib	Davis King	Boost Software License	Yes	Cross-Platform	C++	C++	Yes	No	Yes	Yes	Yes	No	Yes	Yes	Yes	
DataMelt (DMelt)	S.Chekanov	Freemium	Yes	Cross-Platform	Java	Java	No	No	No	No	No	No	No	No	No	
DyNet	Carnegie Mellon University	Apache 2.0	Yes	Linux, macOS, Windows		C++, Python		No[13]	Yes	Yes	Yes					
Intel Data Analytics Acceleration Library	Intel	Apache License 2.0	Yes	Linux, macOS, Windows on Intel CPU[14]	C++, Python, Java	C++, Python, Java[14]	Yes	No	No	Yes	No		Yes		Yes	
Intel Math Kernel Library	Intel	Proprietary	No	Linux, macOS, Windows on Intel CPU[15]		C[16]	Yes[17]	No	No	Yes	No	Yes[18]	Yes[18]		No	
Keras	François Chollet	MIT license	Yes	Linux, macOS, Windows	Python	Python, R	Only if using Theano as backend	Can use Theano or Tensorflow as backends	Yes	Yes	Yes[19]	Yes	Yes	Yes	Yes[20]	
MATLAB + Neural Network Toolbox	MathWorks	Proprietary	No	Linux, macOS, Windows	C, C++, Java, MATLAB	MATLAB	No	No	Train with Parallel Computing Toolbox and generate CUDA code with GPU Coder[21]	No	Yes[22][23]	Yes[22]	Yes[22]	No	With Parallel Computing Toolbox[24]	
Microsoft Cognitive Toolkit	Microsoft Research	MIT license[25]	Yes	Windows, Linux[26] (macOS via Docker on roadmap)	C++	Python (Keras), C++, Command line,[27] BrainScript[28] (.NET on roadmap[29])	Yes[30]	No	Yes	Yes	Yes[31]	Yes[32]	Yes[32]	No[33]	Yes[34]	
Apache MXNet	Apache Software Foundation	Apache 2.0	Yes	Linux, macOS, Windows,[35][36] AWS, Android,[37] iOS, JavaScript[38]	Small C++ core library	C++, Python, Julia, Matlab, JavaScript, Go, R, Scala, Perl	Yes	On roadmap[39]	Yes	Yes[40]	Yes[41]	Yes	Yes	Yes	Yes[42]	
Neural Designer	Artelnics	Proprietary	No	Linux, macOS, Windows	C++	Graphical user interface	Yes	No	No	?	?	No	No	No	?	
OpenNN	Artelnics	GNU LGPL	Yes	Cross-platform	C++	C++	Yes	No	Yes	?	?	No	No	No	?	
PaddlePaddle	Baidu	Apache License	Yes	Linux, macOS, Windows	C++, Python	Python	No	Yes	Yes	Yes	Yes	Yes	Yes	?	Yes	
PlaidML	Vertex.AI	AGPL3	Yes	Linux, macOS, Windows	C++, Python	Keras, Python, C++, C	No	Yes	Yes	Yes		Yes	Yes	?	Yes	
PyTorch	Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan	BSD	Yes	Linux, macOS, Windows	Python, C, CUDA	Python	Yes	Via separately maintained package[43][44][45]	Yes	Yes	Yes	Yes	Yes		Yes	
Apache SINGA	Apache Incubator	Apache 2.0	Yes	Linux, macOS, Windows	C++	Python, C++, Java	No	No	Yes	?	Yes	Yes	Yes	Yes	Yes	
TensorFlow	Google Brain team	Apache 2.0	Yes	Linux, macOS, Windows,[46] Android	C++, Python, CUDA	Python (Keras), C/C++, Java, Go, R[47], Julia, Swift	No	On roadmap[48] but already with SYCL[49] support	Yes	Yes[50]	Yes[51]	Yes	Yes	Yes	Yes	
TensorLayer	Hao Dong	Apache 2.0	Yes	Linux, macOS, Windows,[52] Android	C++, Python,	Python	No	On roadmap[48] but already with SYCL[49] support	Yes	Yes[53]	Yes[54]	Yes	Yes	Yes	Yes	
Theano	Université de Montréal	BSD	Yes	Cross-platform	Python	Python (Keras)	Yes	Under development[55]	Yes	Yes[56][57]	Through Lasagne's model zoo[58]	Yes	Yes	Yes	Yes[59]	No
Torch	Ronan Collobert, Koray Kavukcuoglu, Clement Farabet	BSD	Yes	Linux, macOS, Windows,[60] Android,[61] iOS	C, Lua	Lua, LuaJIT,[62] C, utility library for C++/OpenCL[63]	Yes	Third party implementations[64][65]	Yes[66][67]	Through Twitter's Autograd[68]	Yes[69]	Yes	Yes	Yes	Yes[70]	
Wolfram Mathematica	Wolfram Research	Proprietary	No	Windows, macOS, Linux, Cloud computing	C++, Wolfram Language, CUDA	Wolfram Language	Yes	No	Yes	Yes	Yes[71]	Yes	Yes	Yes	Under Development	
VerAI	VerAI	Proprietary	No	Linux, Web-based	C++,Python, Go, Angular	Graphical user interface, cli	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	

Excellent comparison: https://skymind.ai/wiki/comparison-frameworks-dl4j-tensorflow-pytorch

Excellent comparison: https://arxiv.org/pdf/1511.06435.pdf

https://en.wikipedia.org/wiki/Comparison_of_deep_learning_software

Microsoft Azure: Supported Platforms

- Caffe: A deep learning framework built for speed, expressivity, and modularity
- Caffe2: A cross-platform version of Caffe
- Microsoft Cognitive Toolkit: A deep learning software toolkit from Microsoft Research
- H2O: An open-source big data platform and graphical user interface
- Keras: A high-level neural network API in Python for Theano and TensorFlow
- MXNet: A flexible, efficient deep learning library with many language bindings
- NVIDIA DIGITS: A graphical system that simplifies common deep learning tasks
- PyTorch: A high-level Python library with support for dynamic networks
- TensorFlow: An open-source library for machine intelligence from Google
- Theano: A Python library for defining, optimizing, and efficiently evaluating mathematical expressions involving multi-dimensional arrays
- Torch: A scientific computing framework with wide support for machine learning algorithms
- CUDA, cuDNN, and the NVIDIA driver
- Many sample Jupyter notebooks

https://docs.microsoft.com/en-us/azure/machine-learning/data-science-virtual-machine/dsvm-ubuntu-intro

GPU vs CPU Demo

- Using Keras in interactive Python notebooks

Today’s Topics

- Machine Learning for Sequential Data
- Recurrent Neural Networks (RNNs)
- Training Deep Neural Networks: Hardware & Software

Google Form: Guest Speaker

- Guest: Dr. Cheryl Martin, Chief Data Scientist at Alegion (https://www.alegion.com/company/leadership)
  - Share one question for her for tomorrow's visit