

Chapter 1

Recurrent neural nets

Draft chapter from Torralba, Isola, Freeman


RNNs are a neural net architecture for modeling sequences – they perform *sequential* computation (rather than parallel computation).

So far we have seen **feedforward** neural net architectures. These architectures can be represented by directed acyclical computation graphs. **Recurrent neural networks**, or **RNNs**, are neural nets with feedback connections. Their computation graphs have cycles. That means that the outputs of one layer can send a message back to earlier in the computational pipeline. To make this well-defined, RNNs introduce the notion of *timestep*. One timestep of running an RNN corresponds to one functional call to each layer (or computation node) in the RNN, mapping that layer’s inputs to outputs.

A feedback connection defines a **recurrence** in the computational graph.

RNNs are all about adaptation. They solve the problem: what if we want our computations in the past to influence our computations in the future? If we have a stream of data, a feedforward net has no way to adapting its behavior based on what outputs it previously produced.

To motivate RNNs, we will first consider a deficiency of CNNs. Consider that we are processing a temporal signal with a CNN, like so:




The filter w slides over the input signal and produces outputs. In this example, imagine the input is a video of your dog, Rufus. The output produced for the first set of frames is entirely independent from the output produced for a later frame, as long as the receptive fields do not overlap. Maybe it got dark out as the video progressed, and a clear image of Rufus became hard to discern as the light dimmed. Then the convolutional response, and the net’s prediction, for later in the video may be of low quality, maybe the net now predicts this is your neighbor’s dog, Douglas!

Wouldn’t it be nice if we could inform the CNN’s later predictions that earlier in the day we had a brighter view and this was clearly Rufus? With a CNN, the only way to do this is to increase the size of the convolutional filters so that the receptive fields can see sufficiently far back in time. But what if we last saw Rufus a year ago, it would be infeasible to extend our receptive fields a year back. How can we humans solve this problem and recognize Rufus after a year of not seeing him?

The answer is memory. The recurrent feedback loops in an RNN are a kind of memory. They propagate information from timestep t to timestep $t + 1$. In other words, they remember information about timestep t when processing new information at timestep $t + 1$. By induction, then, an RNN can, in theory, remember information over arbitrarily long time intervals. In fact, RNNs are Turing Complete: you can think of them as a little computer. They can do anything a computer can do.

Here’s how RNNs arrange their connections to perform this propagation:



The key idea of an RNN is that there are lateral connections, between hidden units, through time. The inputs, outputs, and hidden units may be multidimensional tensors; to denote this we will use \square , as \circ was reserved for denoting a single (scalar) neuron. Here is a simple 1-layer RNN:

The idea of a “time” in an RNN does not necessarily mean time as it is measured on a clock.

Time just refers to *sequential computation*; the timestep t is an index into the sequence. The sequential computation could progress over a spatial dimension (processing an image pixel by pixel) or over the temporal dimension (processing a video frame by frame), or even over more abstracted computational modules.




Diagram (a), to the left, is the **unrolled computation graph**. Diagram (b), to the right, is the computation graph rolled up, with a feedback connection. Typically we will work with unrolled graphs because they are DAGs, and therefore all the tools we have developed for DAGs, including backprop, will extend naturally to them. Note however, that RNNs can run for unbounded time, and process signals that have unbounded temporal extent, while a DAG can only represent a finite graph. The DAG depicted above is therefore a *truncated* version of the RNN’s true computation graph.

1.1 Recurrent layer

A **recurrent layer** is defined by the following equations:

$$\mathbf{h}_t = f(\mathbf{h}_{t-1}, \mathbf{x}_{\text{in}_t}) \quad \triangleleft \quad \text{update state based on input and previous state} \quad (1.1)$$

$$\mathbf{x}_{\text{out}_t} = g(\mathbf{h}_t) \quad \triangleleft \quad \text{produce output based on state} \quad (1.2)$$

This is a layer that includes a state variable, \mathbf{h} . The operation of the layer depends on its state. The layer also updates its state every time it is called, therefore the state is a kind of memory.

f and g can be arbitrary functions, and we can imagine a wide variety of recurrent layers defined by different choices for the form of f and g .

One common kind of recurrent layer is the **Simple Recurrent Layer**, or **SRN**, which was defined in “Elman Networks” [Elman 1990]. For this layer, f is a linear layer followed by a pointwise nonlinearity, and g is another linear layer. The full computation is defined by the following equations:

$$\mathbf{h}_t = \sigma_1(\mathbf{W}\mathbf{h}_{t-1} + \mathbf{U}\mathbf{x}_{\text{in}_t} + \mathbf{b}) \quad (1.3)$$

$$\mathbf{x}_{\text{out}_t} = \sigma_2(\mathbf{V}\mathbf{h}_t + \mathbf{c}) \quad (1.4)$$


where σ_1 and σ_2 are pointwise nonlinearities. In SRNs, `tanh` is the common choice for the nonlinearity, but `relu` and other choices may also be used.

1.2 Backprop through time

Backprop is only defined for DAGs and RNNs are not DAGs. To apply backprop to RNNs, we unroll the net for T timesteps, which creates a DAG representation of the truncated RNN, and then do backprop through that DAG. This is known as **backpropagation through time** or **BPTT**.

This approach yields exact gradients of the loss w.r.t. the parameters when T is equal to total number of timesteps the RNN was run for. Otherwise, BPTT is an approximation that truncates the true computation graph. Essentially, BPTT ignores any dependencies in the computation graph greater than length T .


As an example of BPTT, suppose we want to compute the gradient of an output neuron at timestep five with respect to an input at timestep zero, i.e. $\frac{\partial x_{\text{out}_5}}{\partial x_{\text{in}_0}}$. The forward pass (black arrows) and backward pass (red arrows) would look like this, where we have only drawn arrows for the subpart of the computation graph that is relevant to the calculation of this particular gradient:



What about the gradient of the total cost w.r.t. the parameters, $\frac{\partial J}{\partial [\mathbf{W}, \mathbf{U}, \mathbf{V}]}$? Suppose the RNN is being applied to a video and predicts a class label for every frame of the video. Then x_{out_t} is a prediction for each frame and a loss can be applied to each x_{out_t} . The total cost is simply the summation of the losses at each timestep:

$$\frac{\partial J}{\partial [\mathbf{W}, \mathbf{U}, \mathbf{V}]} = \sum_{t=0}^T \frac{\partial \mathcal{L}(x_{\text{out}_t}, y_t)}{\partial [\mathbf{W}, \mathbf{U}, \mathbf{V}]} \quad (1.5)$$

The graph for backprop has branching structure, which we saw how to deal with in Chapter ??:



1.2.1 The problem of exploding and vanishing gradients

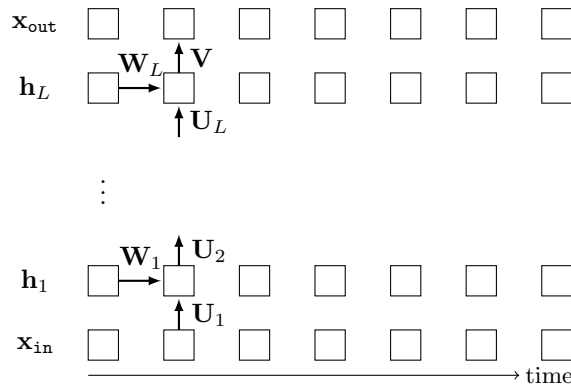
Notice that the we can get very large computation graphs when we run an RNN for many timesteps. The gradient computation involves a series of matrix multiplies that is $\mathcal{O}(T)$. For example, to compute our gradient $\frac{\partial \mathbf{x}_{\text{out}_5}}{\partial \mathbf{x}_{\text{in}_0}}$ in the example above involves the following product:

$$\frac{\partial \mathbf{x}_{\text{out}_t}}{\partial \mathbf{x}_{\text{in}_0}} = \frac{\partial \mathbf{x}_{\text{out}_t}}{\partial \mathbf{h}_T} \frac{\partial \mathbf{h}_T}{\partial \mathbf{h}_{T-1}} \dots \frac{\partial \mathbf{h}_1}{\partial \mathbf{h}_0} \frac{\partial \mathbf{h}_0}{\partial \mathbf{x}_{\text{in}_0}} \quad (1.6)$$

Suppose we use an RNN with `relu` nonlinearities. Then each term $\frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_{t-1}}$ equals $\mathbf{R}\mathbf{W}$, where \mathbf{R} is a diagonal matrix with 0's or 1's on the i -th diagonal element depending on whether the i -th neuron is above or below zero at the current operating point. This gives a product $\mathbf{R}_T \mathbf{W} \dots \mathbf{R}_1 \mathbf{W}$. In a worst case (for numerics), suppose all the `relu` are active; then we have a product of T matrices \mathbf{W} , which means the gradient is \mathbf{W}^T . If values in \mathbf{W} are large, the gradient will explode. If they are small, the gradient will vanish – basically, this system is not numerically well-behaved. One solution to this problem is to use RNN units that have better numerics, an example of which we will give in Section 1.4 below.

1.3 Stacking recurrent layers

A deep RNN can be constructed by stacking recurrent layers. Here is an example:



To emphasize that the parameters are shared, we only show them once. Here is code for a 2-layer version of this RNN:

```
# in_dim, out_dim: input and output dimensionality
# x : input data sequence
# h_init: initial conditions of hidden state
# T: number of timesteps to run for

# first define parameterized layers
U1 = nn.linear(in_dim, hid_dim, bias=False)
U2 = nn.linear(hid_dim, hid_dim, bias=False)
W1 = nn.linear(hid_dim, hid_dim)
W2 = nn.linear(hid_dim, hid_dim)
V = nn.linear(hid_dim, out_dim)

# then run data through network
h = h_init
for t in range(T):
    h[1] = nn.tanh(W1(h)+U1(x[t]))
    h[2] = nn.tanh(W2(h)+U2(x[t]))
    x_out[t] = V(h[2])
```

We set the bias to be “False” for U1 and U2 since the recurrent layer already gets a bias term from W1 and W2.

1.4 LSTMs

Long Short Term Memory layers, or **LSTMs** are a special kind of recurrent layer, that is designed to avoid the vanishing and exploding gradients problem described in Section ?? [Hochreiter and Schmidhuber 1997].

An LSTM is like a little memory controller. It has a **cell state**, \mathbf{c} , which is a vector it can read from and write to. The cell state can record memories and retrieve them as needed. The cell state plays a similar role as the hidden state, \mathbf{h} , from regular RNNs – both record memories – but the cell state is built to store *persistent* memories that can last “long” (hence the name long short term memory). We will see how next.

Like a normal recurrent layer, an LSTM takes as input a hidden state \mathbf{h}_{t-1} and an observation \mathbf{x}_t and produces as output an updated hidden state \mathbf{h}_t . Internally, the LSTM uses its cell state to decide how to update \mathbf{h}_{t-1} .

First, the cell state is updated based on the incoming signals, \mathbf{h}_{t-1} and \mathbf{x}_t . Then the cell state is used to determine the hidden state \mathbf{h}_t to output. Different subcomponents of the LSTM layer decide what to *delete* from the cell state, what to *write* to the cell state, and what to *read* off the cell state.

One component picks the indices of the cell state to *delete*:

$$\mathbf{f}_t = \sigma(\mathbf{W}_f \mathbf{h}_{t-1} + \mathbf{U}_f \mathbf{x}_t + \mathbf{b}_f) \quad \triangleleft \quad \text{decide which indices to forget} \quad (1.7)$$

\mathbf{f}_t is a vector of length equal to the length of the cell state \mathbf{c}_t . σ is the sigmoid function, which outputs values in the range $[0, 1]$. The idea is that most values in \mathbf{f}_t will be near either 1 (remember) or 0 (forget). Later \mathbf{f}_t will be pointwise multiplied by the cell state to “forget” the values where \mathbf{f}_t is zero.

Another component chooses what values to *write* to the cell state, and where to write them:

$$\mathbf{i}_t = \sigma(\mathbf{W}_i \mathbf{h}_{t-1} + \mathbf{U}_i \mathbf{x}_t + \mathbf{b}_i) \quad \triangleleft \quad \text{which indices to write to} \quad (1.8)$$

$$\tilde{\mathbf{c}}_t = \mathbf{tanh}(\mathbf{W}_c \mathbf{h}_{t-1} + \mathbf{U}_c \mathbf{x}_t + \mathbf{b}_c) \quad \triangleleft \quad \text{what to write to those indices} \quad (1.9)$$

\mathbf{i}_t is similar to \mathbf{f}_t – it’s a vector values that are nearly 0 or 1 of length equal to the length of the cell state – and it determines which indices of $\tilde{\mathbf{c}}_t$ will be written to the cell state. Next we use these vectors to actually update the cell state:

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tilde{\mathbf{c}}_t \quad \triangleleft \quad \text{update the cell state} \quad (1.10)$$

Finally, given the updated cell state, we *read* from it and use the values to determine the hidden state to output:

$$\mathbf{o}_t = \sigma(\mathbf{W}_o \mathbf{h}_{t-1} + \mathbf{U}_o \mathbf{x}_t + \mathbf{b}_o) \quad \triangleleft \quad \text{which indices of cell state to use} \quad (1.11)$$

$$\mathbf{h}_t = \mathbf{o}_t \odot \mathbf{tanh}(\mathbf{C}_t) \quad \triangleleft \quad \text{use these to determine next hidden state} \quad (1.12)$$

The basic idea is it should be easy, by default, not to forget. The cell state controls what is remembered. It gets modulated, but \mathbf{f}_t and $\tilde{\mathbf{c}}_t$ can learn to be zero so no change occurs and information propagates forward through an identity connection. This idea is similar to the idea of skip connections and residual connections in ResNets (Chapter ??).

The presentation of LSTMs in this section draws inspiration from a great blog post by Chris Olah [Olah 2015]

Bibliography

J. L. Elman. Finding structure in time. *Cognitive science*, 14(2):179–211, 1990.

S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

C. Olah. Understanding lstm networks, 2015. URL <http://colah.github.io/posts/2015-08-Understanding-LSTMs/>.