

10-701: Introduction to Machine Learning Lecture 12 – RNNs

Henry Chai


2/26/24

Front Matter


- Announcements
 - HW3 released 2/19, due 2/28 (Wednesday) at 11:59 PM
 - HW4 released 2/28 (Wednesday), due 3/15 (after break) at 11:59 PM
 - Project details will be released 3/1 (Friday)
 - **You must work in groups of 2 or 3 on the project**
- Recommended Readings
 - Zhang, Lipton, Li & Smola, [Chapters 9 & 10](#)

Convolutional Neural Networks


- Neural networks are frequently applied to inputs with some inherent spatial structure, e.g., images
- Idea: use the first few layers to identify relevant macro-features, e.g., edges
- Insight: for spatially-structured inputs, many useful macro-features are shift or location-invariant, e.g., an edge in the upper left corner of a picture looks like an edge in the center
- Strategy: learn a *filter* for macro-feature detection in a small window and apply it over the entire image






LeNet (LeCun et al., 1998)




- One of the earliest, most famous deep learning models – achieved remarkable performance at handwritten digit recognition (< 1% test error rate on MNIST)
- Used sigmoid (or logistic) activation functions between layers and mean-pooling, both of which are pretty uncommon in modern architectures



Wait how did we go from 6 to 16?



Channels



4	1	2	16	3	6
1	7	5	8	19	27
5	2	5	12	17	8
0	4	9	9	6	11

5	2	6	14	15	8
26	3	6	8	4	9
0	15	24	6	1	8
7	4	9	5	24	17

4	6	8	9	5	3
16	5	2	8	2	1
5	2	14	11	7	8
15	2	5	0	9	8

- An image can be represented as the sum of red, green and blue pixel intensities
- Each color corresponds to a *channel*




Example:
 $3 \times 4 \times 6$ tensor

4	1	2	16	3	6							
1						5	2	6	14	15	8	
5						26						
0						0						
						7						
							4	6	8	9	5	3
							16	5	2	8	2	1
							5	2	14	11	7	8
							15	2	5	0	9	8

- An image can be represented as a *tensor* or multidimensional array

Convolutions on Multiple Input Channels

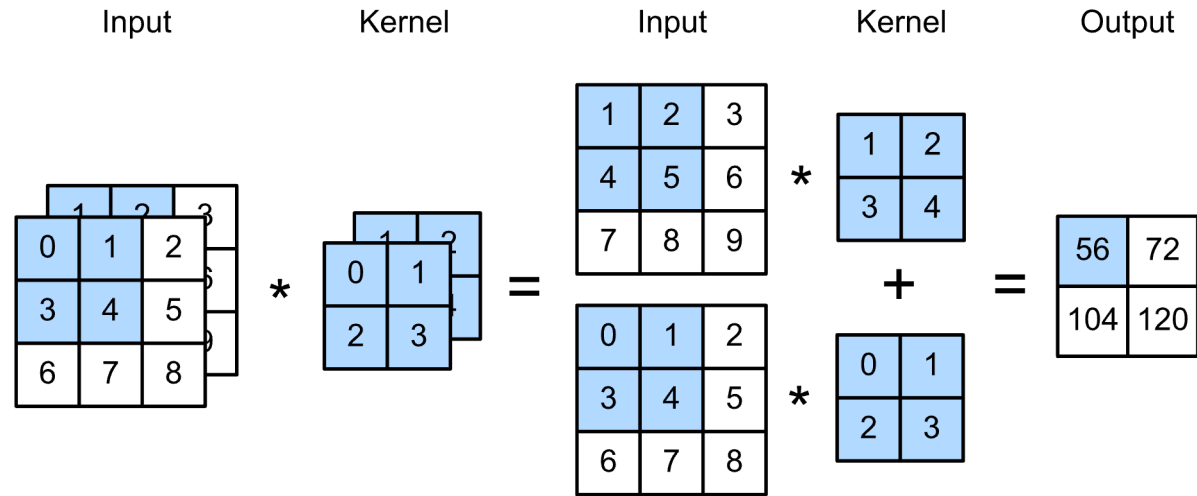
- Given multiple input channels, we can specify a filter for each one and sum the results to get a 2-D output tensor




- For c channels and $h \times w$ filters, we have $chw + c$ learnable parameters (each filter has a bias term)

Convolutions on Multiple Input Channels


- Given multiple input channels, we can specify a filter for each one and sum the results to get a 2-D output tensor



- Questions:
 - Why might we want a different filter for each input?
 - Why do we combine them together into a single output channel?



- Channels in hidden layers correspond to different macro-features, which we might want to manipulate differently → one filter per channel




	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	X				X	X	X			X	X	X	X		X	X
1	X	X				X	X	X			X	X	X	X	X	
2	X	X	X				X	X	X			X		X	X	X
3		X	X	X			X	X	X	X			X		X	X
4			X	X	X			X	X	X	X		X	X		X
5				X	X	X			X	X	X	X		X	X	X

TABLE I
EACH COLUMN INDICATES WHICH FEATURE MAP IN S2 ARE COMBINED BY THE UNITS IN A PARTICULAR FEATURE MAP OF C3.

- We can combine these macro-features into a new, interesting, “higher-level” feature
 - But we don’t always need to combine all of them!
 - Different combinations → multiple output channels
 - Common architecture: more output channels and smaller outputs in deeper layers

C3: f. maps 16@10x10

S2: f. maps
6@14x14



	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	X				X	X	X			X	X	X	X		X	X
1	X	X				X	X	X			X	X	X	X	X	
2	X	X	X				X	X	X			X		X	X	X
3		X	X	X			X	X	X	X			X		X	X
4			X	X	X			X	X	X	X		X	X		X
5				X	X	X			X	X	X	X		X	X	X


TABLE I

EACH COLUMN INDICATES WHICH FEATURE MAP IN S2 ARE COMBINED BY THE UNITS IN A PARTICULAR FEATURE MAP OF C3.

Okay, but what if our layers become too big in the channel dimension?

Downsampling: 1×1 Convolutions

- Convolutional layers can be represented as 4-D tensors of size $c_o \times c_i \times h \times w$ where c_o is the number of output channels and c_i is the number of input channels
- Layers of size $c_o \times c_i \times 1 \times 1$ can condense many input channels into fewer output channels (if $c_o < c_i$)



- Practical note: 1×1 convolutions are typically followed by a nonlinear activation function; otherwise, they could simply be folded into other convolutions




Key Takeaways

- The loss function for neural networks is non-convex!
 - Momentum can help break out of local minima
 - Adaptive gradients help when parameters behave differently w.r.t. step sizes
 - Random restarts can improve the chances of finding better local minima
 - Jitter & dropout act like regularization for neural networks by preventing them fitting the training dataset perfectly
- MLPs and neural networks of sufficient depth are universal approximators

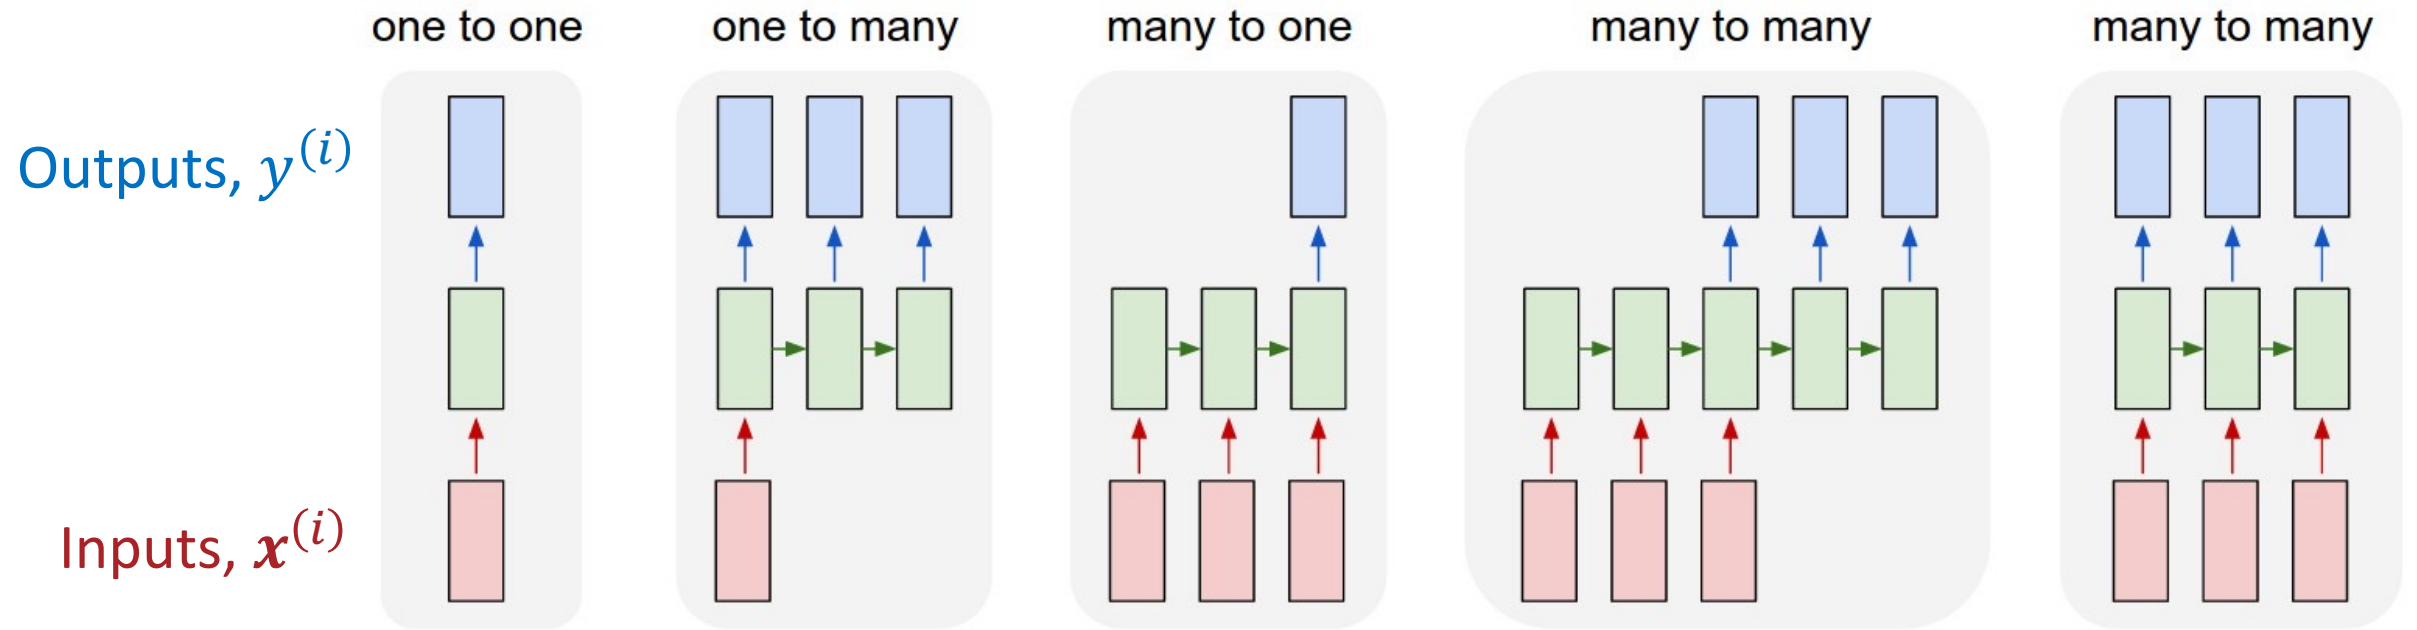
Key Takeaways

- Convolutional neural networks use convolutions to learn macro-features
 - Can be thought of as slight modifications to the fully-connected feed-forward neural network
 - Can still be learned using SGD
 - Padding is used to preserve spatial dimensions while pooling, stride and 1×1 convolutions are used to downsample intermediate representations

Example: Handwriting Recognition



$$\mathbf{y}^{(i)} = [\mathbf{y}_1^{(i)}, \mathbf{y}_2^{(i)}, \dots, \mathbf{y}_{T_i}^{(i)}]$$



$$\mathbf{x}^{(i)} = [\mathbf{x}_1^{(i)}, \mathbf{x}_2^{(i)}, \dots, \mathbf{x}_{T_i}^{(i)}]$$


Sequential Data

Recurrent Neural Networks

- Neural networks are frequently applied to inputs with some inherent temporal or sequential structure (e.g., text or video) of variable length
- Idea: use the information from previous parts of the input to inform subsequent predictions
- Insight: the hidden layers learn a useful representation (relative to the task)
- Approach: incorporate the output from earlier hidden layers into later ones.

Recurrent Neural Networks

$$\mathbf{h}_t = \left[1, \theta \left(W^{(1)} \mathbf{x}_t^{(i)} + W_h \mathbf{h}_{t-1} \right) \right]^T \text{ and } \mathbf{o}_t = \hat{y}_t^{(i)} = \theta \left(W^{(2)} \mathbf{h}_t \right)$$



- Training dataset consists of (input **sequence**, label **sequence**) pairs, potentially of varying lengths

$$\mathcal{D} = \left\{ \left(\mathbf{x}^{(n)}, \mathbf{y}^{(n)} \right) \right\}_{n=1}^N$$


$$\mathbf{x}^{(n)} = \left[\mathbf{x}_1^{(n)}, \dots, \mathbf{x}_{T_n}^{(n)} \right]$$

$$\mathbf{y}^{(n)} = \left[\mathbf{y}_1^{(n)}, \dots, \mathbf{y}_{T_n}^{(n)} \right]$$

- This model requires an initial value for the hidden representation, \mathbf{h}_0 , typically a vector of all zeros


Unrolling Recurrent Neural Networks

$$\mathbf{h}_t = \left[1, \theta \left(W^{(1)} \mathbf{x}_t^{(i)} + W_h \mathbf{h}_{t-1} \right) \right]^T \text{ and } \mathbf{o}_t = \hat{y}_t^{(i)} = \theta \left(W^{(2)} \mathbf{h}_t \right)$$




Deep Recurrent Neural Networks

$$\mathbf{h}_t^{(l)} = \left[1, \theta \left(W^{(l)} \mathbf{h}_t^{(l-1)} + W_h^{(l)} \mathbf{h}_{t-1}^{(l)} \right) \right]^T \text{ and } \mathbf{o}_t = \hat{y}_t^{(i)} = \theta \left(W^{(L)} \mathbf{h}_t^{(L-1)} \right)$$




Deep Recurrent Neural Networks

$$\mathbf{h}_t^{(l)} = \left[1, \theta \left(W^{(l)} \mathbf{h}_t^{(l-1)} + W_h^{(l)} \mathbf{h}_{t-1}^{(l)} \right) \right]^T \text{ and } \mathbf{o}_t = \hat{y}_t^{(i)} = \theta \left(W^{(L)} \mathbf{h}_t^{(L-1)} \right)$$




But why do we only pass information forward?
What if later time steps have useful information as well?

$$\mathbf{h}_t^{(l)} = \left[1, \theta \left(W^{(l)} \mathbf{h}_t^{(l-1)} + W_h^{(l)} \mathbf{h}_{t-1}^{(l)} \right) \right]^T \text{ and } \mathbf{o}_t = \hat{y}_t^{(i)} = \theta \left(W^{(L)} \mathbf{h}_t^{(L-1)} \right)$$




But why do we only pass information forward?
What if later time steps have useful information as well?

$$\mathbf{h}_t = \left[1, \theta \left(W^{(1)} \mathbf{x}_t^{(i)} + W_h \mathbf{h}_{t-1} \right) \right]^T \text{ and } \mathbf{o}_t = \hat{y}_t^{(i)} = \theta \left(W^{(2)} \mathbf{h}_t \right)$$



But why do we only pass information forward?
 What if later time steps have useful information as well?


$$\mathbf{h}_t = \left[1, \theta \left(W^{(1)} \mathbf{x}_t^{(i)} + W_h \mathbf{h}_{t-1} \right) \right]^T \text{ and } \mathbf{o}_t = \hat{y}_t^{(i)} = \theta \left(W^{(2)} \mathbf{h}_t \right)$$



Bidirectional Recurrent Neural Networks

$$\mathbf{h}_t^{(f)} = \left[1, \theta \left(W_f^{(1)} \mathbf{x}_t^{(i)} + W_f \mathbf{h}_{t-1}\right)\right]^T \text{ and } \mathbf{h}_t^{(b)} = \left[1, \theta \left(W_b^{(1)} \mathbf{x}_t^{(i)} + W_b \mathbf{h}_{t+1}\right)\right]^T$$


$$\mathbf{o}_t = \hat{y}_t^{(i)} = \theta \left(W_f^{(2)} \mathbf{h}_t^{(f)} + W_b^{(2)} \mathbf{h}_t^{(b)}\right)$$



Unrolling Bidirectional Recurrent Neural Networks

$$\mathbf{o}_t = \hat{y}_t^{(i)} = \theta \left(W_f^{(2)} \mathbf{h}_t^{(f)} + W_b^{(2)} \mathbf{h}_t^{(b)} \right)$$

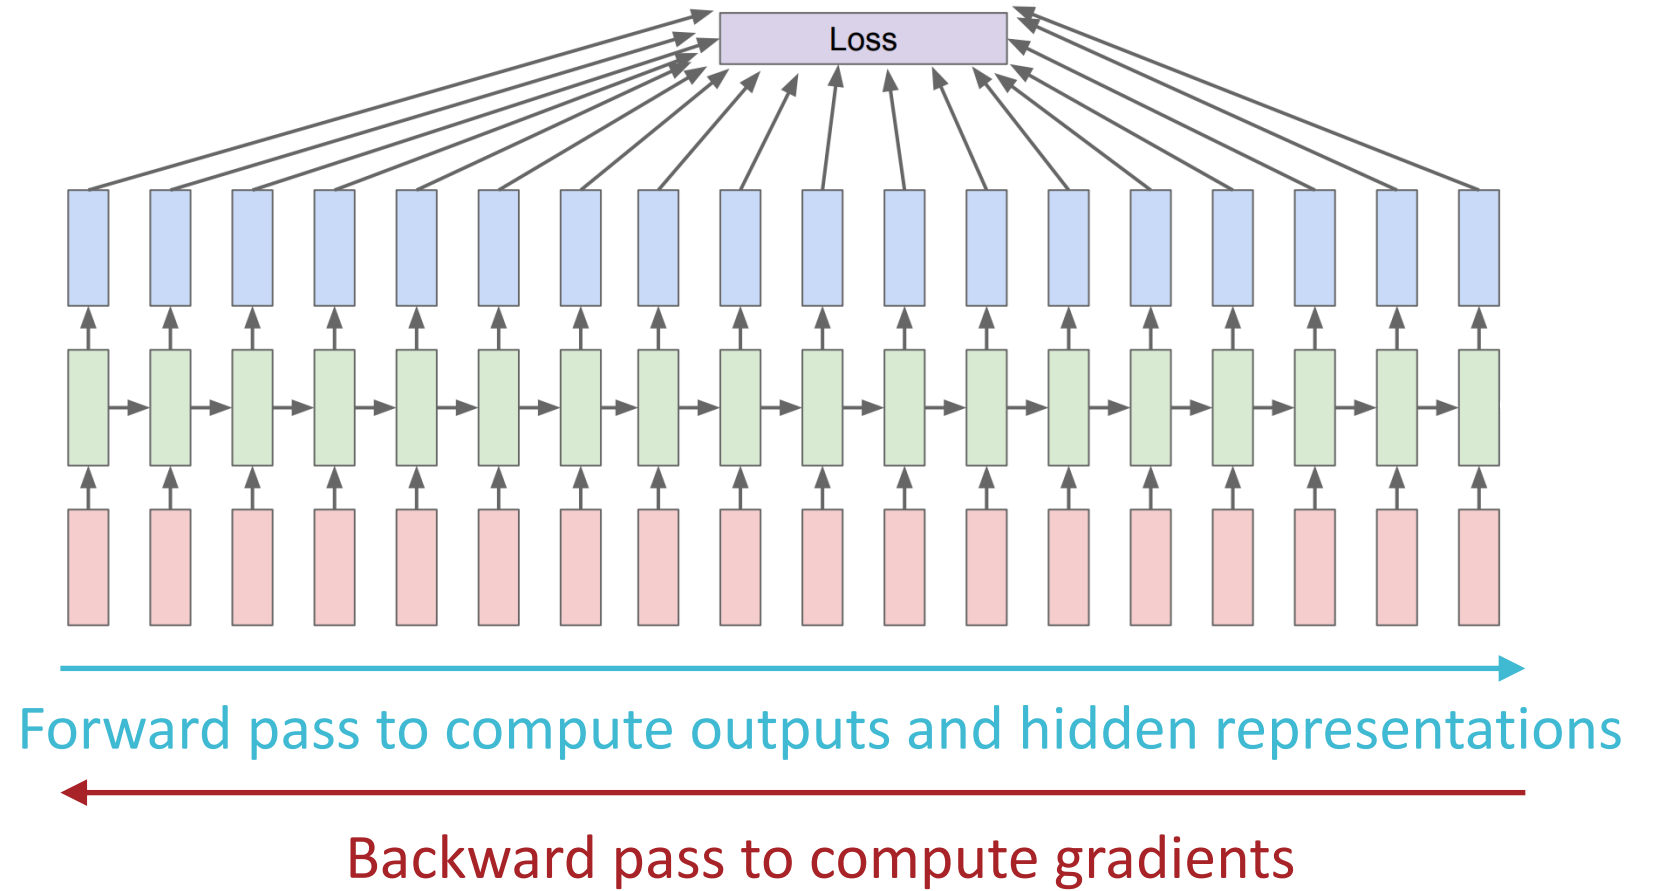
$$\mathbf{h}_t^{(f)} = \left[1, \theta \left(W_f^{(1)} \mathbf{x}_t^{(i)} + W_f \mathbf{h}_{t-1} \right) \right]^T \text{ and } \mathbf{h}_t^{(b)} = \left[1, \theta \left(W_b^{(1)} \mathbf{x}_t^{(i)} + W_b \mathbf{h}_{t+1} \right) \right]^T$$



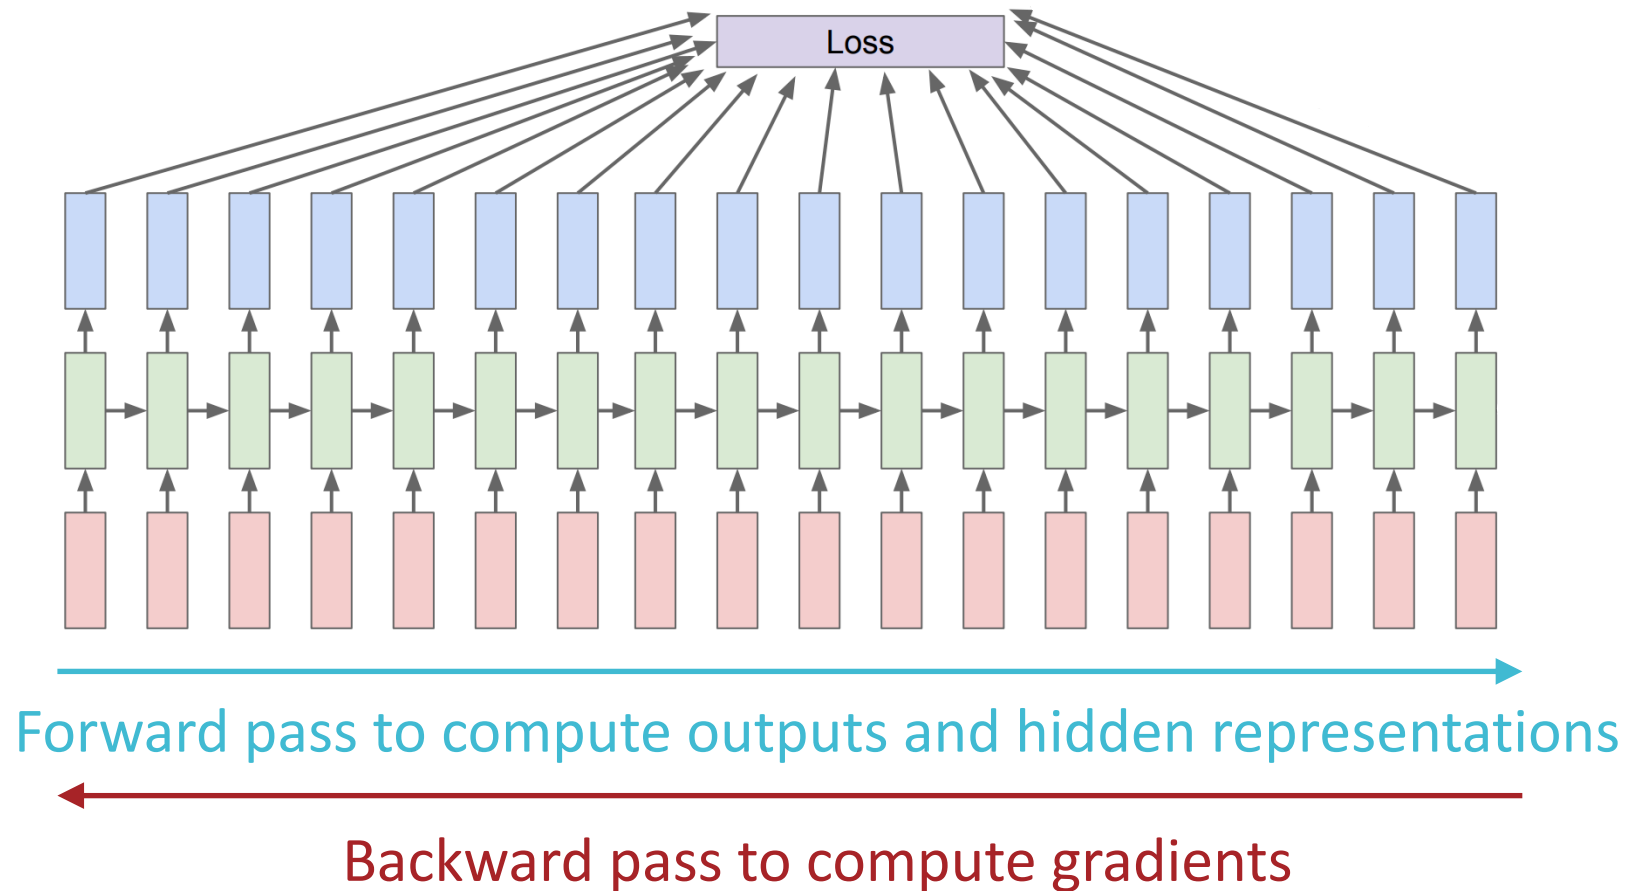
Training RNNs

- A (deep/bidirectional) RNN simply represents a (somewhat complicated) computation graph
 - Weights are shared between different timesteps, significantly reducing the number of parameters to be learned!
- Can be trained using (stochastic) gradient descent/backpropagation → “backpropagation through time”

Training RNNs



Training RNNs: Challenges



- Issue: as the sequence length grows, the gradient is more likely to explode or vanish

Recall: Vanishing Gradients

Insight: $s_b^{(l)}$ only affects $\ell^{(i)}$ via $o_b^{(l)}$

$$\text{Chain rule: } \delta_b^{(l)} = \frac{\partial \ell^{(i)}}{\partial o_b^{(l)}} \left(\frac{\partial o_b^{(l)}}{\partial s_b^{(l)}} \right)$$


$$\begin{aligned} o_b^{(l)} = \theta \left(s_b^{(l)} \right) &\rightarrow \frac{\partial o_b^{(l)}}{\partial s_b^{(l)}} = \frac{\partial \theta \left(s_b^{(l)} \right)}{\partial s_b^{(l)}} \\ &= 1 - \left(\tanh \left(s_b^{(l)} \right) \right)^2 \leq 1 \end{aligned}$$

when $\theta(\cdot) = \tanh(\cdot)$

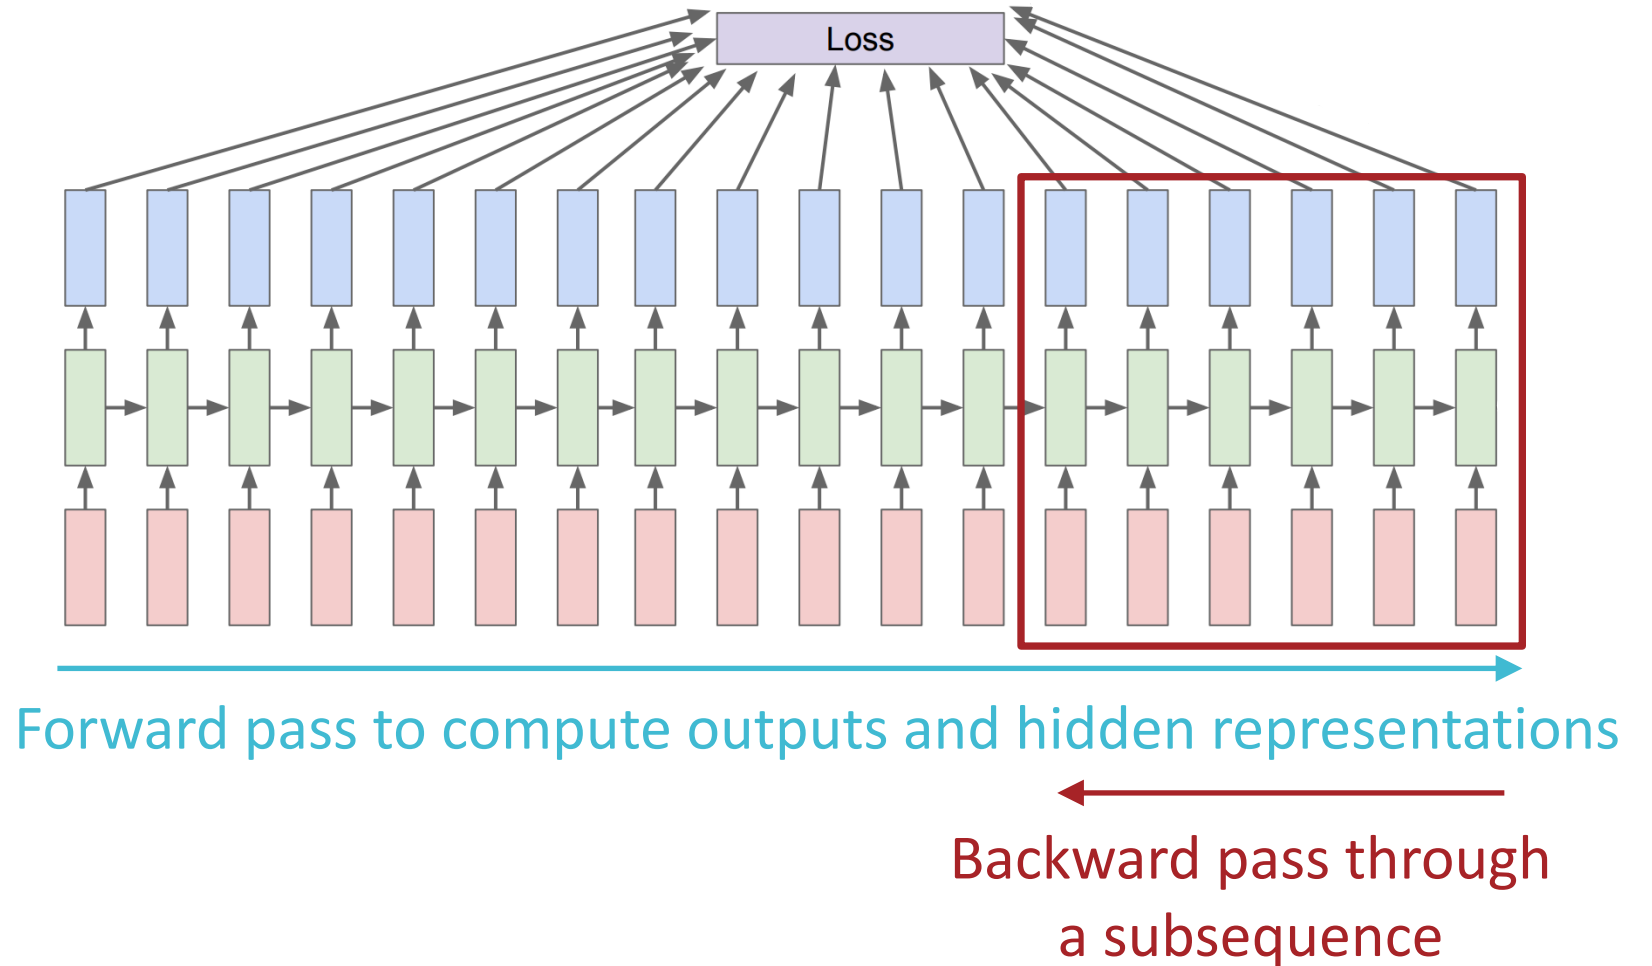
Gradient Clipping (Pascanu et al., 2013)

- Common strategy to deal with exploding gradients: if the magnitude of the gradient ever exceeds some threshold, simply scale it down to the threshold

$$G = \begin{cases} \nabla_W \ell^{(i)} & \text{if } \|\nabla_W \ell^{(i)}\|_2 \leq \tau \\ \left(\frac{\tau}{\|\nabla_W \ell^{(i)}\|_2} \right) \nabla_W \ell^{(i)} & \text{otherwise} \end{cases}$$



Truncated Backpropagation Through Time



- Idea: limit the number of time steps to backprop through

Long Short-Term Memory (Hochreiter & Schmidhuber, 1997)

- LSTM networks address the vanishing gradient problem by replacing hidden layers with *memory cells*
- Each cell still computes a hidden representation but also maintains a separate internal *state*, C_t
- The flow of information through a cell is manipulated by three *gates*:
 - An input gate, I_t , that controls how much the state looks like the normal RNN hidden layer
 - An output gate, O_t , that “releases” the hidden representation to later timesteps
 - A forget gate, F_t , that determines if the previous memory cell’s state affects the current internal state

Long Short-Term Memory (Hochreiter & Schmidhuber, 1997)

- LSTM networks address the vanishing gradient problem by replacing hidden layers with *memory cells*
- Each cell still computes a hidden representation but also maintains a separate internal *state*, C_t
- Gates are implemented as sigmoids: a value of 0 would be a fully closed gate and 1 would be fully open

$$I_t = \sigma \left(W_{ix} \mathbf{x}_t^{(i)} + W_{ih} \mathbf{h}_{t-1} \right)$$

$$O_t = \sigma \left(W_{ox} \mathbf{x}_t^{(i)} + W_{oh} \mathbf{h}_{t-1} \right)$$


$$F_t = \sigma \left(W_{fx} \mathbf{x}_t^{(i)} + W_{fh} \mathbf{h}_{t-1} \right)$$

$$C_t = F_t \odot C_{t-1} + I_t \odot \theta \left(W^{(1)} \mathbf{x}_t^{(i)} + W_h \mathbf{h}_{t-1} \right)$$

$$\mathbf{h}_t = C_t \odot O_t$$


Long Short-Term Memory (Hochreiter & Schmidhuber, 1997)

- LSTM networks address the vanishing gradient problem by replacing hidden layers with *memory cells*
- Each cell still computes a hidden representation but also maintains a separate internal *state*, C_t



Long Short-Term Memory (Hochreiter & Schmidhuber, 1997)

- LSTM networks address the vanishing gradient problem by replacing hidden layers with *memory cells*
- Each cell still computes a hidden representation but also maintains a separate internal state, C_t



- The internal state allows information to move through time without needing to affect the hidden representations!