# Recurrent Neural Networks

# Outline

- Recap convolutional neural networks
- Compare conventional and recurrent neural networks (RNNs)
- Text translation example
- Backpropagation through time
- Code comparison
- Parts-of-speech example
- RNN paradigms
- LSTMs
- A mention of attention

### Recap: Convolutional Neural Networks

• Take five minutes to draw

#### Image Dataset

- We need an image dataset for next week's lecture on inference/deploying.
- I want you all to take photos of whatever we're classifying.
- From previous semesters
  - Frank or Frary
  - Pine or Palm
  - Spoon or Fork
  - Cup or Bowl
  - Salt or Sugar

# Conventional Neural Networks (including CNNs)

- Input: fixed sized tensor
  - Though the batch size can be any value due to broadcasting
- Output: fixed sized tensor
  - Though the batch size can be any value due to broadcasting
- Functionally deterministic (always produce the same output for a given input)
  - When might you want different outputs on the same input?

# Motivating Example: Text Translation

- Input: I love purple cats. Cats are neat.
- Output: J'adore les chats violets. Les chats sont soignés.

### Recurrent Neural Networks

Operate over sequences (data with temporal dependencies).

## Linear vs (Elman) Recurrent Neurons

```
class Neuron(torch.Module):
                                                                         class RecurrentNeuron():
   def init (self, input size, output size):
                                                                             def init (self, input size, output size):
       self.W = torch.randn(output size, input size) * 0.01
                                                                                 self.Wx = torch.randn(output size, input size) * 0.01
       self.b = torch.randn(output size, 1) * 0.01
                                                                                 self.Wh = torch.randn(output size, output size) * 0.01
                                                                                 self.bh = torch.zeros(output size, 1)
                                                                                 self.output size = output size
   def forward(self, X):
       linear = X @ self.W.T + self.b.T
       return F.sigmoid(linear)
                                                                             def forward(self, X, state=None):
                                                                                 L, N, input size = X.shape
                                                                                 if not state:
                                                                                     state = torch.zeros(N, self.output size)
                                                                                 output sequence = []
                                                                                 for x t in X:
                                                                                    state = F.tanh(x t @ self.Wx + state @ self.Wh + self.bh)
                                                                                    output sequence.append(state)
                                                                                 return torch.tensor(output sequence), state
• Input shape:
                                                                         • Input shape:
                                                                         • Output shape:
  Output shape:
•
```

#### Unrolled Visualization

# Motivating Example: Text Translation

- Input: I love purple cats. Cats are neat.
- Output: J'adore les chats violets. Les chats sont soignés.

# Backpropgation Through Time (BPTT)

# Backpropgation Through Time (BPTT)

#### Some BPTT Math

# RNN Paradigm: One to One (RNN unneeded)

No need for recurrent connections.

#### RNN Paradigm: One to Many

#### RNN Paradigm: Many to One

# RNN Paradigm: Many to Many (Synced)

# RNN Paradigm: Many to Many (Encoder/Decoder)

# Parts-Of-Speech Example

You all must submit sentences for the dataset.

https://docs.google.com/spreadsheets/d/1HJmlehaYhGWclDo1t0k6i1VHxN15zr8ZmJj7Rf VEal/edit#gid=1031300490

"I is a teeth"

How do we pass this into a neural network?

## Processing Natural Language with an NN

Here's one way to convert text into numbers

- 1. Assign every word a unique number (e.g., 1 .. vocab\_size)
- 2. Assign every part-of-speech a unique number (e.g., 1 .. num\_classes)
- 3. Convert sentences into index tensors (using mapping from step 1)
- 4. Pass index tensors into an <u>embedding</u> layer (i.e., a simple lookup table)
- 5. Pass embedding outputs into the recurrent neural network (RNN)
- 6. Pass the RNN output into a fully-connected (FC) classification network
- 7. Convert the FC output into a part-of-speech (one-hot)

class POS\_LSTM(torch.nn.Module):
 """Parts-of-speech LSTM model."""

```
def __init__ (self, vocab_size, embed_dim, hidden_dim, num_layers, parts_size):
    super().__init__()
    self.embed = torch.nn.Embedding(vocab_size, embed_dim)
    self.lstm = torch.nn.LSTM(embed_dim, hidden_dim, num_layers=num_layers)
    self.linear = torch.nn.Linear(hidden_dim, parts_size)
```

```
def forward(self, X):
    X = self.embed(X)
    X, _ = self.lstm(X.unsqueeze(1))
    return self.linear(X)
```

Output activation function handled by <u>torch.nn.CrossEntropyLoss</u>

#### LSTMs



https://colah.github.io/posts/2015-08-Understanding-LSTMs/

#### RNNs

#### Advantages

- Process varying input length
- Model size remains constant
- Maintains historical information

Input: I love purple cats. Cats are neat. Output: J'adore les chats violets. Les chats sont soignés.

#### Disadvantages

## RNNs

#### Advantages

- Process varying input length
- Model size remains constant
- Maintains historical information

Input: I love purple cats. Cats are neat. Output: J'adore les chats violets. Les chats sont soignés.

#### Disadvantages

- Slower to computer
- Poor handling of long-term dependencies
- Does not consider future inputs to produce current state
- Largely replaced by transformers

We propose a new simple network architecture, the Transformer, based solely on attention mechanisms, dispensing with recurrence and convolutions entirely.



https://ai.googleblog.com/2016/09/a-neural-network-for-machine.html



### Summary

- Recurrent neural networks maintain an internal state (memory)
- This internal state is useful when data has a temporal component
- They were frequently used in translation and audio processing
- We don't see them as much over the last few years, but the concepts are still worthwhile to know