


# 11-695: Competitive Engineering Recurrent Neural Networks


Spring 2018

- 1 Regularization in Recurrent Neural Networks
- 2 Coding an RNN with TF Dynamic Graph

# General Regularization Strategy: Dropout



- Each colored arrowed can be dropped using *the same mask*.
  - Word embeddings dropout mean to remove *the whole word*

Other Strategies:  $l_p$ 

- $l_2$  norm of all or some parameters
- $l_2$  norm of all or some hidden states:  $\sum_i \|\mathbf{e}_i\|^2, \sum_j \|\mathbf{f}_j\|^2$
- $l_2$  difference of all or some hidden states:  $\sum_i \|\mathbf{e}_i - \mathbf{e}_{i-1}\|^2, \sum_j \|\mathbf{f}_j - \mathbf{f}_{j-1}\|^2$

- ① Regularization in Recurrent Neural Networks
- ② Coding an RNN with TF Dynamic Graph


## `tf_while_loop.py`

```
1 def build_tf_graph:
2     def condition(i, *args): return tf.less(i, 10)
3
4     def body(i, a, b): return i+1, b, a+b
5
6     loop_vars = [tf.constant(0, dtype=tf.int32), tf.constant(1, dtype=tf.int32),
7                 tf.constant(1, dtype=tf.int32)]
8
9     loop_outputs = tf.while_loop(condition, body, loop_vars)
```

- What does this do?
  - Computes the Fibonacci numbers.

## `tf_while_loop.py`

```
1 def build_tf_graph:
2     def condition(i, *args): return tf.less(i, 10)
3
4     def body(i, a, b): return i+1, b, a+b
5
6     loop_vars = [tf.constant(0, dtype=tf.int32), tf.constant(1, dtype=tf.int32),
7                 tf.constant(1, dtype=tf.int32)]
8
9     loop_outputs = tf.while_loop(condition, body, loop_vars)
```



## `tf_while_loop_outputs.py`

```
1 def build_tf_graph:
2     def condition(i, *args): return tf.less(i, 10)
3     def body(i, a, b): return i+1, b, a+b
4     loop_vars = [tf.constant(0, dtype=tf.int32), tf.constant(1, dtype=tf.int32),
5                 tf.constant(1, dtype=tf.int32)]
6
7     loop_outputs = tf.while_loop(condition, body, loop_vars)
8     print(type(loop_outputs)) # <type 'list'>
9     for loop_output in loop_outputs:
10        print(type(loop_output)) # TF Tensor
```

- `loop_outputs` is a *nested structure*
  - the same structure with `loop_vars`
  - Each *inner-most* element is **the TF ops** that **triggers the loop** and returns the corresponding output.




Using `tf.while_loop` to Build Machine Translation

## tf\_rnn.py


```

1 def build_tf_graph:
2     # en_sent, fr_sent: Tensors with unknown shapes [1, E], [1, F].
3     en_sent, fr_sent = read_data()
4
5     # encoder_states: [1, E, hidden_size]
6     encoder_states = encoder(en_sent)
7
8     # decoder_states: [1, F, hidden_size]
9     decoder_states = decoder(encoder_states, fr_sent)

```



- We will use *2 while loops*
  - one for encoder; one for decoder

Coding the Encoder with `tf.while_loop``tf_seq2seq_encoder.py`

```

1 def encoder(sent, vocab_size=10000, hidden_size=128):
2     # sent: Tensors with unknown shape [1, E]
3     with tf.variable_scope("encoder"):
4         w_emb = tf.get_variable("w_emb", [vocab_size, hidden_size]) # "encoder/w_emb"
5
6         def condition(i, sent, *args): return tf.less(i, tf.shape(E)[-1])
7
8         def body(i, sent, *args): # later

```

- Use the *unknown* length of `sent` to stop the loop
- Problem: we **do not** know how much memory to allocate

# Coding the Encoder with `tf.while_loop`

## `tf_seq2seq_encoder.py`

```
1 def encoder(sent, vocab_size=10000, hidden_size=128):
2     # sent: Tensors with unknown shape [1, E]
3     with tf.variable_scope("encoder"):
4         w_emb = tf.get_variable("w_emb", [vocab_size, hidden_size]) # "encoder/w_emb"
5
6         hidden_states = tf.TensorArray(tf.float32, size=tf.shape(E)[-1],
7                                       clear_after_read=False)
8         def condition(i, sent, *args): return tf.less(i, tf.shape(E)[-1])
9
10        def body(i, sent, hidden_states):
11            # do the RNN computations; write the new information to hidden_states
12            return i+1, sent, hidden_states
13
14        loop_vars = [tf.constant(0, dtype=tf.int32), sent, hidden_states]
15        loop_outputs = tf.while_loop(condition, body, loop_vars)
16        hidden_states = loop_outputs[-1].stack() # [E, 1, hidden_states]
17        return tf.transpose(hidden_states, [1, 0, 2])
```

- Problem: we **do not** know how much memory to allocate
  - Solution: `tf.TensorArray` supports *dynamic* memory allocation.

# Writing to tf.TensorArray

## tf\_seq2seq\_encoder.py

```

1 def encoder(sent, vocab_size=10000, hidden_size=128):
2     # create variables: w_emb, ...
3     hidden_states = tf.TensorArray(tf.float32, size=tf.shape(E)[-1],
4                                   clear_after_read=False)
5     def condition(i, sent, *args): return tf.less(i, tf.shape(E)[-1])
6
7     def body(i, sent, prev_state, hidden_states):
8         # sent[:, i]: [1, 1] --> emb: [1, 1, hidden_size]
9         emb = tf.nn.embedding_lookup(w_emb, sent[:, i])
10
11        # do the RNN computations; write the new information to hidden_states
12        next_state = prev_state + emb
13        hidden_states = hidden_states.write(i, next_state)
14        return i+1, sent, next_state, hidden_states
15
16    loop_vars = [tf.constant(0, dtype=tf.int32), sent,
17                tf.zeros([1, 1, hidden_size], dtype=tf.float32), hidden_states]
18    loop_outputs = tf.while_loop(condition, body, loop_vars)
19    hidden_states = loop_outputs[-1].stack() # [|E|, 1, hidden_states]
20    return tf.transpose(hidden_states, [1, 0, 2])

```

- A dummy encoder network

## tf\_seq2seq\_encoder.py

```
1 def encoder(sent, vocab_size=10000, hidden_size=128):
2     with tf.variable_scope("encoder"):
3         w_emb = tf.get_variable("w_emb", [vocab_size, hidden_size])
4         w_rnn = tf.get_variable("w_rnn", [hidden_size, hidden_size])
5         w_inp = tf.get_variable("w_inp", [hidden_size, hidden_size])
6
7     def condition(i, sent, *args): return tf.less(i, tf.shape(E)[-1])
8     def body(i, sent, prev_state, hidden_states):
9         # sent[:, i]: [1, 1] --> emb: [1, 1, hidden_size]
10        emb = tf.nn.embedding_lookup(w_emb, sent[:, i])
11
12        # do the RNN computations; write the new information to hidden_states
13        next_state = tf.tanh(tf.matmul(prev_state, w_rnn) + tf.matmul(emb, w_inp))
14        hidden_states = hidden_states.write(i, next_state)
15        return i+1, sent, next_state, hidden_states
16
17    # create loop_vars, calling tf.While\loop, stack, transpose, return
```

- A less dummy encoder network:

$$f(\mathbf{x}_t, \mathbf{h}_{t-1}) = \tanh(\mathbf{h}_{t-1} \cdot \mathbf{W}_h + \mathbf{x}_t \cdot \mathbf{W}_x) \quad (1)$$

- Want Long Short-Term Memory (LSTM)? No problem!

$$\begin{aligned} \mathbf{i}_t &= \text{Sigmoid}(\mathbf{x}_t \cdot \mathbf{W}_{xi} + \mathbf{h}_{t-1} \cdot \mathbf{W}_{hi}) \\ \mathbf{f}_t &= \text{Sigmoid}(\mathbf{x}_t \cdot \mathbf{W}_{xf} + \mathbf{h}_{t-1} \cdot \mathbf{W}_{hf}) \\ \mathbf{o}_t &= \text{Sigmoid}(\mathbf{x}_t \cdot \mathbf{W}_{xo} + \mathbf{h}_{t-1} \cdot \mathbf{W}_{ho}) \\ \mathbf{g}_t &= \tanh(\mathbf{x}_t \cdot \mathbf{W}_{xg} + \mathbf{h}_{t-1} \cdot \mathbf{W}_{hg}) \\ \mathbf{c}_t &= \mathbf{f}_t \otimes \mathbf{c}_{t-1} + \mathbf{i}_t \otimes \mathbf{g}_t \\ \mathbf{h}_t &= \mathbf{o}_t \otimes \tanh(\mathbf{c}_t) \end{aligned} \tag{2}$$

- Want Long Short-Term Memory (LSTM)? No problem!

$$\begin{aligned} \mathbf{i}_t &= \text{Sigmoid}(\mathbf{x}_t \cdot \mathbf{W}_{xi} + \mathbf{h}_{t-1} \cdot \mathbf{W}_{hi}) \\ \mathbf{f}_t &= \text{Sigmoid}(\mathbf{x}_t \cdot \mathbf{W}_{xf} + \mathbf{h}_{t-1} \cdot \mathbf{W}_{hf}) \\ \mathbf{o}_t &= \text{Sigmoid}(\mathbf{x}_t \cdot \mathbf{W}_{xo} + \mathbf{h}_{t-1} \cdot \mathbf{W}_{ho}) \\ \mathbf{g}_t &= \tanh(\mathbf{x}_t \cdot \mathbf{W}_{xg} + \mathbf{h}_{t-1} \cdot \mathbf{W}_{hg}) \\ \mathbf{c}_t &= \mathbf{f}_t \otimes \mathbf{c}_{t-1} + \mathbf{i}_t \otimes \mathbf{g}_t \\ \mathbf{h}_t &= \mathbf{o}_t \otimes \tanh(\mathbf{c}_t) \end{aligned} \tag{2}$$

## tf\_lstm.py

```
1 def lstm(x, prev_c, prev_h, w_lstm):
2     ifog = tf.matmul(tf.concat([x, prev_h], axis=2))
3     i, f, o, g = tf.split(ifog, 4, axis=1)
4     i, f, o, g = tf.sigmoid(i), tf.sigmoid(f), tf.sigmoid(o), tf.tanh(g)
5     next_c = f * prev_c + i * g
6     next_h = o * tf.tanh(next_c)
7     return next_c, next_h
```


# Changing The “Footprint” of body

---

## tf\_seq2seq\_encoder.py

```
1 def lstm(x, prev_c, prev_h, w_lstm): # previous slide
2 def encoder(sent, vocab_size=10000, hidden_size=128):
3     with tf.variable_scope("encoder"):
4         w_emb = tf.get_variable("w_emb", [vocab_size, hidden_size])
5         w_lstm = tf.get_variable("w_lstm", [2 * hidden_size, 4 * hidden_size])
6
7     def condition(i, sent, *args): return tf.less(i, tf.shape(E)[-1])
8     def body(i, sent, prev_c, prev_h, hidden_states):
9         emb = tf.nn.embedding_lookup(w_emb, sent[:, i])
10        next_c, next_h = lstm(emb, prev_c, prev_h, w_lstm)
11        hidden_states = hidden_states.write(i, next_h)
12        return i+1, sent, next_c, next_h, hidden_states
13
14 # create loop_vars, calling tf.While\loop, stack, transpose, return
```



Coding the Decoder with `tf.while_loop``tf_seq2seq_decoder.py`


```

1 def build_tf_graph:
2     # en_sent, fr_sent: Tensors with unknown shapes [1, E], [1, F].
3     en_sent, fr_sent = read_data()
4
5     # encoder_states: [1, E, hidden_size]
6     encoder_states = encoder(en_sent)
7
8     # decoder_states: [1, F, hidden_size]
9     decoder_states = decoder(encoder_states, fr_sent)

```

- Idea: almost the same with Encoder...

## Handle encoder\_states from encoder



## tf\_seq2seq\_decoder.py

```

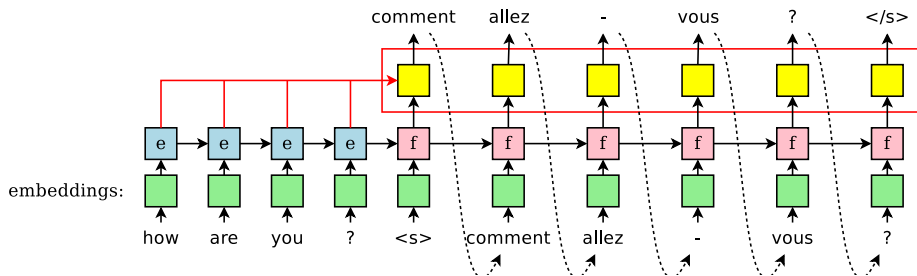
1 def decoder(sent, encoder_states, vocab_size=10000, hidden_size=128):
2     # sent: Tensor of size [1, F]
3     with tf.variable_scope("decoder"):
4         w_emb = tf.get_variable("w_emb", [vocab_size, hidden_size])
5         w_rnn = tf.get_variable("w_rnn", [hidden_size, hidden_size])
6         w_inp = tf.get_variable("w_inp", [hidden_size, hidden_size])
7
8     def condition(i, sent, *args): return tf.less(i, tf.shape(E)[-1])
9     def body(i, sent, prev_state, hidden_states):
10        emb = tf.nn.embedding_lookup(w_emb, sent[:, i])
11        next_state = tf.tanh(tf.matmul(prev_state, w_rnn) + tf.matmul(emb, w_inp))
12        hidden_states = hidden_states.write(i, next_state)
13        return i+1, sent, next_state, hidden_states
14
15    # In encoder: tf.zeros([1, hidden_size]). Now: prev_state
16    loop_vars = [tf.constant(0, dtype=tf.int32), sent, encoder_state, hidden_states]

```

# Dealing with hidden\_states

## tf\_seq2seq\_decoder.py

```
1 def decoder(sent, encoder_states, vocab_size=10000, hidden_size=128):
2     # sent: [1, F]. define other variables, define while_loop, etc.
3     loop_outputs = tf.while_loop(condition, body, loop_vars)
4
5     hidden_states = loop_outputs[-1].stack()           # [F, 1, hidden_size]
6     hidden_states = tf.transpose(hidden_states, [1, 0, 2]) # [1, F, hidden_size]
7
8     return hidden_states
9
10 def build_tf_graph():
11     # calling read_data, calling encoder, etc.
12     decoder_states = decoder(encoder_states, fr_sent[1, :-1])
13
14     # matmul and softmax
15     with tf.variable_scope("softmax"):
16         w_soft = tf.get_variable("w_soft", [hidden_size, vocab_size])
17         decoder_states = tf.reshape(decoder_states, [-1, hidden_size])
18
19         logits = tf.matmul(decoder_states, w_soft) # [1 * F, vocab_size]
20         logits = tf.reshape(logits, [1, tf.shape(sent)[1], vocab_size])
21         labels = tf.reshape(fr_sent[:, 1:], [-1])
22
23         loss = tf.nn.sparse_softmax_cross_entropy_with_logits(logits=logits, labels=labels)
```



- **Attention:** how is  $\mathbf{a}(\mathbf{f}, \mathbf{e}_{1..|x|})$  computed?

$$\alpha_i = g(\mathbf{f}, \mathbf{e}_i); a_i = \text{Softmax}(\alpha_{1..|x|}); \mathbf{a}(\mathbf{f}, \mathbf{e}_{1..|x|}) = \sum_{i=1}^{|x|} a_i \mathbf{e}_i \quad (3)$$

- Luong attention:  $g(\mathbf{f}, \mathbf{e}_i) = \mathbf{f} \cdot \mathbf{e}_i^\top$ . Quite easy to implement.

## tf\_seq2seq\_decoder.py

```
1 def build_tf_graph():
2     encoder_states = encoder(en_sent)
3     decoder_states = decoder(encoder_states, fr_sent[1, :-1])
4     attn_states = attention(encoder_states, decoder_states)
```

- **Attention:** how is  $\mathbf{a}(\mathbf{f}, \mathbf{e}_{1 \dots |x|})$  computed?

$$\alpha_i = g(\mathbf{f}, \mathbf{e}_i); a_i = \text{Softmax}(\alpha_{1 \dots |x|}); \mathbf{a}(\mathbf{f}, \mathbf{e}_{1 \dots |x|}) = \sum_{i=1}^{|x|} a_i \mathbf{e}_i \quad (4)$$

- Luong attention:  $g(\mathbf{f}, \mathbf{e}_i) = \mathbf{f} \cdot \mathbf{e}_i^\top$ . Quite easy to implement.
- Bahdanau attention:  $g(\mathbf{f}, \mathbf{e}_i) = \tanh(\mathbf{f} \cdot \mathbf{w}_f + \mathbf{e}_i \cdot \mathbf{w}_e) \cdot \mathbf{v}$ , where  $\mathbf{w}_f, \mathbf{w}_e \in \mathbf{R}^{H \times H}$  and  $\mathbf{v} \in \mathbf{R}^{H \times 1}$  are trainable parameters

## tf\_seq2seq\_decoder.py

```
1 def attention(enc_states, dec_states):
2     # enc_states, dec_states: [N, E, hidden_size], [N, F, hidden_size]
3     enc_len = tf.shape(enc_states)[1]
4     dec_len = tf.shape(dec_states)[1]
5     hidden_size = tf.shape(enc_states)[-1]
6
7     # auto batch matmul. attn_logits: [N, F, E]
8     attn_logits = tf.matmul(dec_states, enc_states, transpose_b=True)
9     attn_logits = tf.reshape(attn_logits, [-1, E])
10    attn_weights = tf.nn.softmax(attn_logits)
11
12    # weighted sum
13    attn_weights = tf.reshape(attn_weights, [-1, dec_len, enc_len]) # [N, F, E]
14    attn_outputs = tf.matmul(attn_weights, enc_states)
15    return attn_outputs
16
17 def build_tf_graph():
18     encoder_states = encoder(en_sent)
19     decoder_states = decoder(encoder_states, fr_sent[1, :-1])
20     attn_states = attention(encoder_states, decoder_states)
```