

Image captioning

1 Introduction

Image captioning is a term used in machine learning for the task of giving a textual description of an image. In this assignment we will see how we can use RNNs for this task. In addition we will use a convolutional neural network to extract features from the image. The convolutional neural network will extract a feature map of size $7 \times 7 \times 1024$ for an image of shape $224 \times 224 \times 3$. This feature extraction will be fixed throughout. We will then learn a projection of the 1024 - dimensional feature-vectors down to 256 - dimensional vectors, mainly due to computational considerations. Furthermore we will use *content-based soft attention* so that at each time step we transform the $7 \times 7 \times 256$ feature map into a 256 - dimensional vector that is a weighted average of the 49 feature vectors. Note that when using content-based attention, the vectors are just treated as a flat array without considerations to their spatial positions or relationships. If we want to use that kind of information we should use some spatial encoding scheme to embed that information into the feature vectors themselves. The idea is that at each time step we can attend to the location in the image that contains the relevant information for what we are currently describing.

There are **two** files included:

- **image_captioning.py** - Python script
- **image_captioning.ipynb** - Ipython notebook

The notebook may be opened locally using Jupyter or in the cloud with Google Colab (navigate to File \rightarrow Upload notebook). The notebook format may sometimes be easier to read, and work with interactively, so give it a try if you like. If you use Colab you may also use a GPU for free! Navigate to 'Edit \rightarrow Notebook settings' and then choose GPU as hardware accelerator.

Note that some summaries are already being logged into TensorBoard, so you probably want to check those out, and maybe add your own.

2 Implementation tasks

Our baseline attention mechanism (implemented in the class **UniformAttention**) attends equally to all parts of the image at every time step, i.e.

$$p_{t,j} = \frac{1}{49} \tag{1}$$

$$z^t = \sum_{j=1}^{49} p_{t,j} x_j \tag{2}$$

where z^t is the context vector given as input to the RNN at timestep t .

Your task is to improve upon this by implementing two other *trainable* attention schemes that estimate the attention weights $p_{t,j}$ in a way that we can pay more attention to relevant parts of the image. More specifically we shall try to learn a function f so that

$$\alpha_{t,j} = f(s^t, x_j) \tag{3}$$

$$p_{t,j} = \frac{e^{\alpha_{t,j}}}{\sum_k e^{\alpha_{t,k}}} \tag{4}$$

f should thus hopefully be able to figure out how relevant location j is at time step t . Note that the attention weights are just the softmax activation functions applied to the logits $\alpha_{t,j}$ at each time step.

2.1 Additive attention

The first method is an example of what is sometimes called *additive* attention. Here we replace f above with a feedforward neural network. In this problem you will implement a very simple feedforward network, following Neural Machine Translation by Jointly Learning to Align and Translate, which we shall denote *Bahdanau*-attention. The logits $\alpha_{t,j}$ are calculated as

$$\alpha_{t,j} = W \tanh(U s^t + V x_j + b) + c \tag{5}$$

for learnable matrices U, V, W and bias vectors b and c . Implement the attention mechanism given in (5) through a class, e.g. called **BahdanauAttention**, that inherits from `tf.keras.layers.Layer`. Train a model and try the `evaluate` function for different images and see if the attention weights seem to make sense. The `call` method of the class should have the same signature as that of **UniformAttention**, it should take two inputs:

- `feature_vectors`: tensor of shape `[batch_size, 49, 256]`

- `state_output`: tensor of shape `[batch_size, 512]`

and return

- `context_vector`: tensor of shape `[batch_size, 256]`, which is the weighted average of the feature vectors.
- `attention_weights`: tensor of shape `[batch_size, 49]`, which are the weights used in the averaging. These are returned for visualization purposes only.

2.1.1 Solution

2.2 Multiplicative/dot-product - attention

The second approach uses a more tailored to the estimation of attention weights. We learn a query matrix Q and key matrix K so that

$$\alpha_{t,j} = (Qs^t)^T Kx_j \tag{6}$$

Implement the attention mechanism in (6) and try to train a model with it.

2.2.1 Solution

3 Questions

3.1 Pretrained Convnet

We are using a convolutional neural network trained on imagenet to extract features, and their parameters are not adjusted during the optimization. What are some advantages and disadvantages to this approach compared to also updating these weights during optimization (or even starting with random weights)?

3.2 Pretrained word embedding

The embedding layer is initialized with random weights. Would it be possible to somehow use a pretrained embedding layer?

3.3 Supervised learning

One can not in general say that there is a *correct* caption for a given image. Is *supervised* learning still an appropriate framework?

3.4 Baseline loss

Having a simple baseline loss during training is often useful for a sanity check that things behave as expected. Can you think of any models you should be able to beat?

3.5 Validation loss

What are some possible issues with our validation loss? Why do this validation loss not tell the whole story?

3.6 RNN cell vs layer

In TensorFlow, what is the difference between the LSTM and LSTMCell layers (or in general between RNN and RNNCell layers)?