

Deep Learning (IST, 2022-23)

Homework 2

André Martins, Francisco Melo, Gonçalo Correia, João Santinha

Deadline: Friday, January 13, 2023.

Please turn in the answers to the questions below in a PDF file, together with the code you implemented to solve them (when applicable).

IMPORTANT: Please write 1 paragraph indicating clearly what was the contribution of each member of the group in this project. A penalization of 10 points will be applied if this information is missing.

Please submit a **single zip file** in Fenix under your group's name.

Question 1 (30 points)

- (30 points) Consider the convolutional neural network in Fig. 1, used to classify images $\mathbf{x} \in \mathbb{R}^{H \times W}$ into 3 possible classes. The network has a convolutional layer with a single $M \times N$ filter, a stride of 1 and no padding, and a ReLU non linearity. This layer is followed by a 2×2 max pooling layer with a stride of 2 and no padding, and an output layer with softmax activation.

Let \mathbf{W} denote the filter weights, \mathbf{z} the result of the convolution (i.e., $\mathbf{z} = \text{conv}(\mathbf{W}, \mathbf{x})$) and $\mathbf{h} = \text{ReLU}(\mathbf{z})$. Let also $\mathbf{x}' = \text{vec}(\mathbf{x})$ and $\mathbf{z}' = \text{vec}(\mathbf{z})$ denote the flattened versions of \mathbf{x} and \mathbf{z} , respectively.

- (5 points) What is the dimension of \mathbf{z}' ?
 - (10 points) Show that there is a matrix $\mathbf{M} \in \mathbb{R}^{H'W' \times HW}$, with $H' = H - M + 1$ and $W' = W - N + 1$, such that $\mathbf{z}' = \mathbf{M}\mathbf{x}'$. What is the general expression for element (i, j) of \mathbf{M} ?
 - (10 points) Indicate the number of parameters in the network. Compare that number with the number of parameters in a network where the convolutional and max pooling layers are replaced by a fully connected layer yielding an output with the same dimension as \mathbf{h}_2 . Ignore the bias terms.
- (5 points) Assume now that, instead of using convolutions, the same $M \times N$ image \mathbf{x} is flattened into a sequence $\mathbf{x}' = \text{vec}(\mathbf{x})$ of length MN and goes through a single-head self-attention layer with 1×1 projection matrices $\mathbf{W}_Q = \mathbf{W}_K = \mathbf{W}_V = 1$, without any positional encodings. Write the expression for the attention probabilities and attention output as a function of \mathbf{x}' .




Figure 1: Architecture of a simple convolutional neural network.

Question 2 (35 points)

Image classification with CNNs. In this exercise, you will implement a convolutional neural network to perform classification using the Kuzushiji-MNIST dataset.

As previously done in Homework 1, you will need to download the Kuzushiji-MNIST dataset. You can do this by running the following command in the homework directory:

```
python download_kuzushiji_mnist.py
```

Python skeleton code is provided (`hw2-q2.py`). You will now try out convolutional networks. For this exercise, we recommend you use a deep learning framework with automatic differentiation (suggested: Pytorch).

1. (4 points) Why does a CNN have fewer free parameters than a fully-connected network with the same input size and the same number of classes (assuming it is a network for classification)?
2. (4 points) Despite having fewer free parameters, a CNN usually achieves better generalization on images and patterns representing letters and numbers than a fully-connected network. Can you justify this?
3. (4 points) If the input is from a source composed of independent sensors (with no spatial structure), is a CNN still expected to achieve better generalization than a fully connected network? Justify why.

4. (20 points) Implement a simple convolutional network with the following structure:
- A convolution layer with 8 output channels, a kernel of size 5x5, stride of 1, and padding chosen to preserve the original image size. Note: Assign convolution layer initialization to a variable called `self.conv1` to
 - A rectified linear unit activation function.
 - A max pooling with kernel size 2x2 and stride of 2.
 - A convolution layer with 16 output channels, a kernel of size 3x3, stride of 1, and padding of zero.
 - A rectified linear unit activation function.
 - A max pooling with kernel size 2x2 and stride of 2.
 - An affine transformation with 600 output features (to determine the number of input features use the number of channels, width and height of the output of the second block. Hint: The number of input features = *number of output channels* × *output width* × *output height*).
 - A rectified linear unit activation function.
 - A dropout layer with a dropout probability of 0.3.
 - An affine transformation with 120 output features.
 - A rectified linear unit activation function.
 - An affine transformation with the number of classes followed by an output LogSoftmax layer.

Hint: use the functions `nn.Conv2d` and `nn.MaxPool2d`.

Train your model for 20 epochs using Adam tuning only the learning rate on your validation data, using the following values: 0.00001, 0.0005, 0.01. Report the learning rate of best configuration and plot two things: the training loss and the validation accuracy, both as a function of the epoch number.

5. (3 points) Plot the activation maps of the first convolutional layer and the original training example. From the comparison of the activation maps and the original image, what appears to be highlighted in the activation maps?

NOTE: the skeleton code already performs the extraction of the activation maps (saved as `activation_maps.pdf`) and the original image (saved as `original_example.pdf`). No coding is required.

Question 3 (35 points)

Character-level Machine Translation. Machine translation is the problem of automatically translating a sequence of tokens from a source language into a target language. This task is usually addressed with an encoder-decoder model. Typically, the encoder is based on LSTMs or self-attentions that process and summarises the source sentence into relevant feature representations. The representations are then passed to a target language decoder, commonly an LSTM model or a masked self-attention, which generates the respective translation of the source sentence token by token given the previous tokens and the encoder representations.

1. In the following, you will implement a character-level machine translation model from Spanish to English, that is, a model that translates a sequence of characters (a sentence in Spanish) into another sequence of characters (a sentence in English). The evaluation metric

is the mean error rate (which calculates the Levenshtein distance between the prediction and the true target, divided by the true length of the sequence).

- (a) (10 points) The dataset is provided in the folder `data`. Implement a vanilla character-level machine translation model using an encoder-decoder architecture with an auto-regressive LSTM as the decoder and a Bidirectional LSTM in the encoder. Specifically, in the skeleton code, you will need to implement the method `forward()` of both the `Encoder` and the `Decoder` in `models.py` and then run (`python hw2-q3.py`). Train the model over 50 epochs, using a learning rate of 0.003, a dropout rate of 0.3, a hidden size of 128, and a batch size of 64. The dropout layer should only be applied to the embeddings and the final outputs (in both the encoder and the decoder).

Plot the validation error rate over epochs. Also, report the final error rate in the test set.

You can run the code with the command `python hw2-q3.py`.

Hint: At each time-step, the LSTM decoder receives as input the embedding of the current word (`self.embedding`), then each output at each timestep has dropout applied (`self.dropout`) and the resulting decoder state is used in the next timestep. In the end, all decoder outputs are concatenated to form an output sequence.

- (b) (20 points) Add an attention mechanism to the decoder (bilinear attention, explained below), which weights the contribution of the different source characters, according to relevance for the current prediction.

Bilinear Attention. This attention mechanism is defined as follows:

$$\mathbf{p} = \text{softmax}(\mathbf{s}), \quad (1)$$

where \mathbf{s}_i is the score between the i -th source character and the target character at the current time step, and the softmax is applied in the dimension of the size of the source sequence. This is equivalent to applying `softmax` to the score vectors. The score is defined as:

$$s_i = \mathbf{z}^\top \mathbf{h}_i, \quad (2)$$

where $\mathbf{z} = \mathbf{W}_q^\top \mathbf{q}$ in which \mathbf{q} is the hidden state of the decoder at the current time step (the query), \mathbf{W}_q is a weight matrix (`linear_in` in the skeleton code), and \mathbf{h}_i is the hidden state of the encoder at the i -th time step. After computing the attention weights \mathbf{p} , we compute the context vector \mathbf{c} as:

$$\mathbf{c} = \sum_{i=1}^{T_x} p_i \mathbf{h}_i. \quad (3)$$

Finally, the attention layer output will be computed as:

$$\mathbf{a}_o = \tanh(\mathbf{W}_o^T [\mathbf{q}; \mathbf{c}]), \quad (4)$$

where \mathbf{W}_o is a weight matrix (`linear_out` in the skeleton code) and $[\mathbf{h}_q; \mathbf{c}]$ is the concatenation of the query and the context vector. \mathbf{a}_o is the attention layer output and will be used as the final representation at the current time step in the decoder.

Specifically, in the skeleton code, you will need to implement the `forward()` methods of the `Attention` class and slightly modify your `Decoder` implementation to account for the attention mechanism (it should modify the output of the LSTM at each timestep).

Plot the validation error rate over epochs. Also, report the final error rate in the test set.

You can run the code with the command `python hw2-q3.py -use_attn`. Use the same hyperparameters as in the previous exercise.

- (c) (5 points) What could we do to improve results without changing the model architecture? Explain why this change would improve results. Hint: think about the decoding process in `test()` implemented on the `hw2-q3.py` file.