

```
# Aim: Get LSTM to learn addition mod n
# Requirements: sequence modeling
import numpy as np
import matplotlib.pyplot as plt
import torch.nn as nn
import torch.nn.functional as F
Iteration: 10
In [3]: # mock test of overriding as would be expected in recurrent structure
# this overwriting works only for pytorch
# pytorch keeps versions of stuff
mock_in = torch.tensor([1.0, requires_grad=True])
mock_rec = torch.tensor(2.0, requires_grad=True)
mock_rec2 = torch.tensor(2.0, requires_grad=True)
mock_in = mock_in + mock_rec
mock_in = mock_in + mock_rec2
print(mock_rec.backward())
# if you do not detach here (e.g some sort of multihanded loss) you get error as graph got garbaged when you backprop
mock_cpy = mock_in.detach()
mock_cpy = mock_cpy + mock_rec2
mock_cpy.backward()
# trying to access mock_cpy's grad gives error
# trying to
print(mock_rec2.grad) # makes sense, doesnt include anything from mock_in
Iteration: 10
In [4]: # class PrefixSumDataset(Dataset):
def __init__(self, n, length, num_samples):
    self.x = np.random.randint(0, n, size=(num_samples, length))
    self.y = np.cumsum(self.x, axis=1) % n
    self.n = n
def __len__(self):
    return len(self.x)
def __getitem__(self, idx):
    # onehot encode everything, onehot requires long vector
    return F.one_hot(torch.tensor(self.x[idx], dtype=torch.long), num_classes=self.n), \
           F.one_hot(torch.tensor(self.y[idx], dtype=torch.long), num_classes=self.n)
Iteration: 10
In [5]: # testing
PS = PrefixSumDataset(3, 5, 10)
PS.__getitem__(1)
Out[5]: (tensor([[1, 0, 0],
          [1, 0, 0],
          [0, 1, 0],
          [1, 0, 0],
          [0, 1, 0]]),
        tensor([[1, 0, 0],
          [1, 0, 0],
          [0, 1, 0],
          [0, 1, 0],
          [1, 0, 0]]))
Iteration: 10
In [90]: # Demo of LSTMs
Can't do teacher forcing as output isn't fed back into input
Will train as is
class SeqModel(nn.Module):
    def __init__(self, n, hidden_dim):
        """ LSTM, initC -> data -> LSTM -> output layer, hidden, cell
            output layer -> (dropout) -> decoder -> price prediction """
        super(SeqModel, self).__init__()
        # if num_layers = 1 then dropout is not used, it is not applied to final output layer
        self.lstm = LSTM(input_size = n, hidden_size=hidden_dim, num_layers = 2, batch_first=True, dropout=0)
        self.decoder = nn.Sequential(torch.nn.Dropout(p=0.1),
                                     nn.Linear(hidden_dim, n),
                                     nn.Softmax())
        self.init_hidden = torch.zeros(1, hidden_dim, dtype=torch.float)
        self.init_cell = torch.zeros(1, hidden_dim, dtype=torch.float)
    def forward(self, x):
        # x is a one hot encoded array to prefix sum shape ('', len, n)
        if len(x.shape) == 2: # no batch
            lstm_outputs, (final_hidden, final_cell) = self.lstm(x.float(), (
                self.init_hidden, self.init_cell))
        else:
            lstm_outputs, (final_hidden, final_cell) = self.lstm(x.float(), ( # broadcast to batch
                self.init_hidden.repeat(1, len(x), 1), self.init_cell.repeat(1, len(x), 1)))
        return self.decoder(lstm_outputs)
# Pool operator style model
Can do teacher forcing
Bit different to NAR in that we dont work in latent space
In NAR it would be
new lstm encodes, concat to previous latent (representing prev sum) -> new latent -> decode -> new sum
and we can still do teacher forcing by taking correct prev sum and encoding it as latent
class AggregatorModel(nn.Module):
    def __init__(self, n, hidden_dim):
        """ LSTM, initC -> prev_sum, cur_val -> concat -> 2 layer NN -> new_sum
            new_sum is fed back into prev_sum """
        super(AggregatorModel, self).__init__()
        # if num_layers = 1 then dropout is not used, it is not applied to final output layer
        self.processor = nn.Sequential(nn.Linear(2*n, hidden_dim),
                                       nn.ReLU(),
                                       nn.Dropout(p=0.1),
                                       nn.Linear(hidden_dim, n),
                                       nn.Softmax())
    def forward(self, cur_val, prev_sum):
        # x is a one hot encoded array to prefix sum shape ('', len, n)
        return self.processor(torch.concat((cur_val, prev_sum), dim=-1))
Iteration: 10
In [94]: # from tqdm import tqdm
# train aggregator model
def trainMatch(batchX, batchY, mModel, mLoss, mOptimizer, parameters_dict):
    l = batchX.shape[1]
    batchSize = batchX.shape[0]
    trunc_size = parameters_dict['trunc_size']
    mOptimizer.zero_grad()
    loss = 0
    if teacher_forced:
        # feed in correct 'prev' sums
        for i in range(1, l):
            true_prev_sum = torch.zeros(batchSize, n, dtype=torch.float) if i == 0 else batchY[:,i-1,:].float()
            preds = mModel(torch.cat([true_prev_sum, batchX[:,i,:]])
                           .float(), parameters_dict)
            loss += mLoss_fn(preds, batchY[:,i,:].float())
    else:
        # dont feed in correct 'prev' sums
        prev_sum = torch.zeros(batchSize, n, dtype=torch.float)
        for i in range(1, l):
            if ltrunc_size == 0:
                prev_sum = prev_sum.detach()
                print(preds, batchY[:,i,:])
                prev_sum = mModel(batchX[:,i,:], prev_sum)
                loss += mLoss_fn(prev_sum.squeeze(), batchY[:,i,:].squeeze().float())
            else:
                loss.backward()
                mOptimizer.step()
                return loss.item()
def trainBatchLSTM(batchX, batchY, mModel, mLoss, mOptimizer, parameters_dict):
    # standard training protocol
    # no teacher forcing no batching
    preds = mModel(batchX)
    loss = mLoss_fn(pred, torch.squeeze(batchY).float())
    mOptimizer.zero_grad()
    loss.backward()
    mOptimizer.step()
    return loss.item()
def trainLoop(train_dataset, mModel, mLoss, mOptimizer, prob_teacher_force, trunc_size=5):
    # execute one train pass over data
    # return total training loss
    mDataloader = DataLoader(train_dataset, batch_size = 20, shuffle=True)
    total_loss = 0
    use_lstm = False
    if mModel.__class__.__name__ == 'SeqModel':
        use_lstm = True
    for batchSize, (X, Y) in tqdm(enumerate(mDataloader)):
        if use_lstm:
            total_loss += trainBatchLSTM(X, Y, mModel, mLoss, mOptimizer, parameters_dict)
        else:
            use_teacher_force = np.random.random() < prob_teacher_force
            parameters_dict = {'teacher_forced': use_teacher_force, 'trunc_size': trunc_size}
            total_loss += trainMatch(X, Y, mModel, mLoss, mOptimizer, parameters_dict)
    return total_loss
def evaluate_testset(test_dataset, model):
    # returns accuracy and plots confusion matrix
    test_loader = DataLoader(test_dataset, batch_size = test_dataset.__len__() // vectorize)
    [(textX, texty)] = [c for c in enumerate(test_loader)]
    predictions = torch.argmax(model(textX), -1)
    print('Test Accuracy: (torch.sum(predictions == texty))/(test_dataset.__len__() correct)')
    sns.heatmap(confusion_matrix(texty, predictions))
def evaluate_testset(test_dataset, mModel, print_result = True):
    # returns accuracy and plots confusion matrix
    with torch.inference_mode():
        mModel.eval() # turn off dropout
        test_loader = DataLoader(test_dataset, batch_size = test_dataset.__len__() // vectorize)
        [(textX, texty)] = [c for c in enumerate(test_loader)]
        n = testX.shape[-1]
        correct = 0
        total = len(testy.flatten()) // testy.shape[-1] # total number of tests
        use_lstm = False
        if mModel.__class__.__name__ == 'SeqModel':
            use_lstm = True
        if use_lstm:
            predictions = torch.argmax(model(textX), -1)
            actual = torch.argmax(testy, -1)
            correct = torch.sum(predictions == actual)
        else:
            prev_sum = torch.zeros(test_dataset.__len__(), n, dtype=torch.float) # probabilities
            for i in range(testy.shape[0]):
                prev_sum = mModel(textX[:,i:], prev_sum)
                predictions = torch.argmax(prev_sum, -1)
                actual = torch.argmax(testy[:,i:], -1)
                print(predictions.shape, actual.shape)
                correct += torch.sum(predictions == actual)
        # I wonder if later values are harder to predict
        if print_result:
            print('Evaluate on test set: (correct) out of (total) values correct')
            mModel.train() # turn dropout back on
            return correct.item()/total
Iteration: 10
In [81]: N = 3
seq_len = 20
num_train = 5000
num_test = 1000
PS_train = PrefixSumDataset(N, seq_len, num_train)
PS_test = PrefixSumDataset(N, seq_len, num_test)
Iteration: 10
In [107]: optimizer_fn = torch.optim.Adam
LR = 3e-3
L2REG = 1e-6
model_agg = AggregatorModel(N, 20) # hidden dim size 20
teacher_probability = 0.5
loss_fn = nn.BCELoss()
Iteration: 10
In [108]: optimizer_agg = optimizer_fn(model_agg.parameters(), lr=LR, weight_decay = L2REG)
epochs = 0
for i in range(epochs):
    print(f'Iteration {i}:')
    print(f'training loss: {trainLoop(PS_train, model_agg, loss_fn, optimizer_agg, teacher_probability)}')
    evaluate_testset(PS_test, model_agg)
Iteration 0:
250it [00:04, 57.491t/s]
training loss: 3070.76281836487
Evaluate on test set: 7445 out of 20000 values correct
Iteration 1:
250it [00:04, 59.821t/s]
training loss: 2609.64158476457
Evaluate on test set: 7442 out of 20000 values correct
Iteration 2:
250it [00:04, 60.701t/s]
training loss: 2247.48628682014
Evaluate on test set: 9431 out of 20000 values correct
Iteration 3:
250it [00:03, 62.941t/s]
training loss: 2112.8623947906494
Evaluate on test set: 12509 out of 20000 values correct
Iteration 4:
250it [00:04, 58.321t/s]
training loss: 1999.3551097106934
Evaluate on test set: 12639 out of 20000 values correct
Iteration 5:
250it [00:03, 62.531t/s]
training loss: 1535.3752093315125
Evaluate on test set: 20903 out of 20000 values correct
Iteration 6:
250it [00:04, 62.051t/s]
training loss: 1213.2252872294378
Evaluate on test set: 20000 out of 20000 values correct
Iteration 7:
250it [00:04, 62.221t/s]
training loss: 1277.7338566446669
Evaluate on test set: 20000 out of 20000 values correct
Iteration: 10
In [109]: # with torch.inference_mode(): # cuz I got dropout going
model_agg.eval()
for i in range(N):
    for j in range(N):
        V1 = F.one_hot(torch.tensor(1, dtype=torch.long), num_classes=N).float()
        V2 = F.one_hot(torch.tensor(j, dtype=torch.long), num_classes=N).float()
        print(model_agg.eval(V1, V2), i, j)
        model_agg.train()
        tensor([0.9745, 0.0105, 0.0150]) 0 0
        tensor([0.0315, 0.9435, 0.0252]) 0 1
        tensor([0.0346, 0.0640, 0.9014]) 0 2
        tensor([0.0343, 0.9385, 0.0272]) 1 0
        tensor([0.0323, 0.0436, 0.9240]) 1 1
        tensor([0.0289, 0.0980, 0.8731]) 1 2
        tensor([0.0277, 0.0205, 0.9519]) 2 0
        tensor([0.0605, 0.0278, 0.9882]) 2 1
        tensor([0.0262, 0.0643, 0.9685]) 2 2
Iteration: 10
In [111]: results = []
lengths = [n*500 for n in range(1,10)]
# # #
exceptional generalisation
for length in lengths:
    results.append(evaluate_testset(PrefixSumDataset(N, length, num_test), model_agg, False))
print(lengths, results)
plt.scatter(lengths, results)
[500, 1000, 1500, 2000, 2500, 3000, 3500, 4000, 4500] [1.0, 1.0, 1.0, 1.0, 1.0, 1.0, 1.0, 1.0, 1.0]
Iteration: 10
In [112]: 
Iteration: 10
In [18]: # torch.save(model_agg.state_dict(), 'model_aggregator_1.pth')
Out[18]: <All keys matched successfully>
Iteration: 10
In [19]: model_agg.load_state_dict(torch.load('model_aggregator_1.pth'))
Out[19]: <All keys matched successfully>
Iteration: 10
In [22]: # no teacher forcing
# it gets there pretty quick but a bit slower than with teacher forcing
model_agg_no_tf = AggregatorModel(N, 20)
optimizer_agg_no_tf = optimizer_fn(model_agg_no_tf.parameters(), lr=LR, weight_decay = L2REG)
epochs = 10
for i in range(epochs):
    print(f'Iteration {i}:')
    print(f'training loss: {trainLoop(PS_train, model_agg_no_tf, loss_fn, optimizer_agg_no_tf, 0)}')
    evaluate_testset(PS_test, model_agg_no_tf)
Iteration 0:
250it [00:04, 59.151t/s]
training loss: 3179.372380256653
Evaluate on test set: 7375 out of 20000 values correct
Iteration 1:
250it [00:09, 27.451t/s]
training loss: 3183.36835239800146
Evaluate on test set: 7375 out of 20000 values correct
Iteration 2:
250it [00:04, 51.751t/s]
training loss: 3191.0283526265593
Evaluate on test set: 7484 out of 20000 values correct
Iteration 3:
250it [00:07, 25.131t/s]
training loss: 3104.084663391133
Evaluate on test set: 7526 out of 20000 values correct
Iteration 4:
250it [00:04, 56.331t/s]
training loss: 3059.462224006653
Evaluate on test set: 7709 out of 20000 values correct
Iteration 5:
250it [00:05, 46.281t/s]
training loss: 3017.5321502515149
Evaluate on test set: 8488 out of 20000 values correct
Iteration 6:
250it [00:05, 42.781t/s]
training loss: 2979.76588410645
Evaluate on test set: 8927 out of 20000 values correct
Iteration 7:
250it [00:06, 49.351t/s]
training loss: 2936.732525252926
Evaluate on test set: 9381 out of 20000 values correct
Iteration 8:
250it [00:04, 55.871t/s]
training loss: 2827.7850449386597
Evaluate on test set: 9413 out of 20000 values correct
Iteration 9:
250it [00:04, 51.801t/s]
training loss: 2637.06275587106
Evaluate on test set: 10997 out of 20000 values correct
Iteration 10:
250it [00:05, 46.231t/s]
training loss: 1976.331780944824
Evaluate on test set: 20000 out of 20000 values correct
Iteration 11:
250it [00:04, 51.531t/s]
training loss: 1544.2978539466858
Evaluate on test set: 20000 out of 20000 values correct
Iteration 12:
250it [00:04, 54.531t/s]
training loss: 1376.1043495522837
Evaluate on test set: 20000 out of 20000 values correct
Iteration 13:
250it [00:04, 56.361t/s]
training loss: 1292.206323795165
Evaluate on test set: 20000 out of 20000 values correct
Iteration 14:
250it [00:04, 56.381t/s]
training loss: 1159.765312075615
Evaluate on test set: 20000 out of 20000 values correct
Iteration: 10
In [72]: PS_train_long = PrefixSumDataset(N, 500, 1000)
Iteration: 10
In [73]: # does no truncation still work if we have very long sequences?
model_agg_no_tf_no_truncate = AggregatorModel(N, 20)
optimizer_agg_no_tf_no_truncate = optimizer_fn(model_agg_no_tf_no_truncate.parameters(), lr=3e-3, weight_decay = 0)
epochs = 10
for i in range(epochs):
    print(f'Iteration {i}:')
    print(f'training loss: {trainLoop(PS_train, model_agg_no_tf_no_truncate, loss_fn, optimizer_agg_no_tf_no_truncate, 0, 1000)}')
    evaluate_testset(PS_test, model_agg_no_tf_no_truncate)
Iteration 0:
50it [00:22, 2.211t/s]
training loss: 18920.43191282820
Evaluate on test set: 7089 out of 20000 values correct
Iteration 1:
50it [00:18, 2.761t/s]
training loss: 15923.708034667959
Evaluate on test set: 7421 out of 20000 values correct
Iteration 2:
50it [00:16, 2.981t/s]
training loss: 15914.55371109375
Evaluate on test set: 7420 out of 20000 values correct
Iteration 3:
50it [00:19, 2.561t/s]
training loss: 15913.419342041016
Evaluate on test set: 7274 out of 20000 values correct
Iteration 4:
50it [00:17, 2.791t/s]
training loss: 15913.23832939062
Evaluate on test set: 7238 out of 20000 values correct
Iteration 5:
50it [00:17, 2.881t/s]
training loss: 15912.35232222266
Evaluate on test set: 7286 out of 20000 values correct
Iteration 6:
50it [00:16, 2.761t/s]
training loss: 15912.246948242188
Evaluate on test set: 7071 out of 20000 values correct
Iteration 7:
50it [00:16, 3.051t/s]
training loss: 15911.862049505647
Evaluate on test set: 7184 out of 20000 values correct
Iteration 8:
50it [00:16, 2.971t/s]
training loss: 15911.733001709884
Evaluate on test set: 7421 out of 20000 values correct
Iteration 9:
50it [00:16, 3.121t/s]
training loss: 15911.045011230459
Evaluate on test set: 7421 out of 20000 values correct
Iteration: 10
In [79]: # how about with truncated training on size 20 samples
# works about same
model_lstm = SeqModel(N, 20)
optimizer_lstm = optimizer_fn(model_lstm.parameters(), lr=LR, weight_decay = L2REG)
epochs = 8
for i in range(epochs):
    print(f'Iteration {i}:')
    print(f'training loss: {trainLoop(PS_train, model_lstm, loss_fn, optimizer_lstm, 0, 0)}')
    evaluate_testset(PS_test, model_lstm)
Iteration 0:
250it [00:05, 48.031t/s]
training loss: 258.072349190712
Evaluate on test set: 7375 out of 20000 values correct
Iteration 1:
250it [00:05, 49.381t/s]
training loss: 257.70398897256355
Evaluate on test set: 7383 out of 20000 values correct
Iteration 2:
250it [00:04, 53.201t/s]
training loss: 254.9025342464447
Evaluate on test set: 7897 out of 20000 values correct
Iteration 3:
250it [00:05, 44.601t/s]
training loss: 251.589487718521
Evaluate on test set: 8733 out of 20000 values correct
Iteration 4:
250it [00:05, 47.541t/s]
training loss: 234.269292571282
Evaluate on test set: 19999 out of 20000 values correct
Iteration 5:
250it [00:04, 52.191t/s]
training loss: 177.022140117825
Evaluate on test set: 20000 out of 20000 values correct
Iteration 6:
250it [00:04, 50.061t/s]
training loss: 172.67809599637985
Evaluate on test set: 20000 out of 20000 values correct
Iteration 7:
250it [00:04, 51.061t/s]
training loss: 171.262413031578
Evaluate on test set: 20000 out of 20000 values correct
Iteration: 10
In [ ]: evaluate_testset(PrefixSumDataset(N, 2000, 1000), model_lstm)
Iteration: 10
In [103]: # lstm do well training on size 20 samples
# but it cant learn on more than that
model_lstm2 = SeqModel(N, 20)
optimizer_lstm2 = optimizer_fn(model_lstm2.parameters(), lr=1e-2, weight_decay = 0)
epochs = 8
for i in range(epochs):
    print(f'Iteration {i}:')
    print(f'training loss: {trainLoop(PS_train_long, model_lstm2, loss_fn, optimizer_lstm2, 0, 0)}')
    evaluate_testset(PS_test, model_lstm2)
Iteration 0:
50it [00:16, 3.401t/s]
training loss: 6.10475924910604
Evaluate on test set: 6739 out of 20000 values correct
Iteration 1:
50it [00:18, 2.861t/s]
training loss: 51.641952252197
Evaluate on test set: 6744 out of 20000 values correct
Iteration 2:
50it [00:16, 2.361t/s]
training loss: 51.64139246940613
Evaluate on test set: 6756 out of 20000 values correct
Iteration 3:
50it [00:19, 2.601t/s]
training loss: 51.64142048358917
Evaluate on test set: 6757 out of 20000 values correct
Iteration 4:
50it [00:18, 2.681t/s]
training loss: 51.641587197917
Evaluate on test set: 6748 out of 20000 values correct
Iteration 5:
50it [00:16, 3.111t/s]
training loss: 51.6418632224579
Evaluate on test set: 6756 out of 20000 values correct
Iteration 6:
50it [00:16, 3.041t/s]
training loss: 51.64198422431946
Evaluate on test set: 6759 out of 20000 values correct
Iteration 7:
50it [00:19, 2.591t/s]
training loss: 51.64186716079712
Evaluate on test set: 6759 out of 20000 values correct
Iteration: 10
In [ ]: 
```