2022-01-31: Questions and answers

Made with a warm hug

CHEERYTEXT3944 JAN 20, 2022 09:24PM

What is the GPU performance on Colab like in comparison to using a normal computer with GPU?

I do not quite now what you consider to be "a normal computer" 😉, but with Colab you get a single Tesla K80 with 12GB RAM. (You can check by running the command "!nvidia-smi" from a notebook cell.) — CHEERYTEXT3944

Lab1 A optimization

so we got 28 min running for training the model for one epoch in lab 1 A, and I noticed in the function 'training examples', we used normal arrays instead of tensors to form our target-word vectors and such. I know tensor operations (like sum) are much faster than iterating through normal vectors, so could that improve the running time by any considerable amount?

Whether you would see an instant speed-up would depend on your code, but the general answer is "yes". Our reference implementation runs one epoch of training in ca. 8 minutes on Colab CPU and ca. 2 minutes on Colab GPU., so I guess there is some room for performance improvement. 😉 — CHEERYTEXT3944

Smoothing, it could be viewed somewhat like what Google PageRank tries to solve - spreading some "weight" to sites with very little connection to other websites?

I addressed this question in the session ("cream cheese"). — CHEERYTEXT3944

Gates in LSTM-Cell

In lecture 2.4 you motivated the gating mechanism of a RNN with elementwise multiplication with the gating vectors g and 1-g.
How does this translate to the next slide where the weights of the forget and input layer in the LSTM-cell are independent?

I addressed this question during the session. — CHEERYTEXT3944

prediction

when predicting the t+1 word, does the models compute conditional probabilities for every dictionary word and choose the word with the highest probability?

I addressed this question during the session. — CHEERYTEXT3944

2022-01-31: Questions and answers

Post your questions about the course content on this padlet.
Btw, Marco is cheerytext3944.