

CS685 Quiz 3: *parameter-efficient adaptation*

Released 3/8, due 3/15 on Gradescope (please upload a PDF!)

Please answer both questions in 2-4 sentences each.

1. Explain how the [prompt tuning](#) method discussed in class allows us to solve multiple different NLP tasks within a single batch.
2. You decide to set **A** and **B** in [LoRA](#) to *full-rank* matrices instead of low-rank matrices. Is the resulting approach equivalent to normal fine-tuning? Why or why not?