Assume we have a pre-trained model M_0.

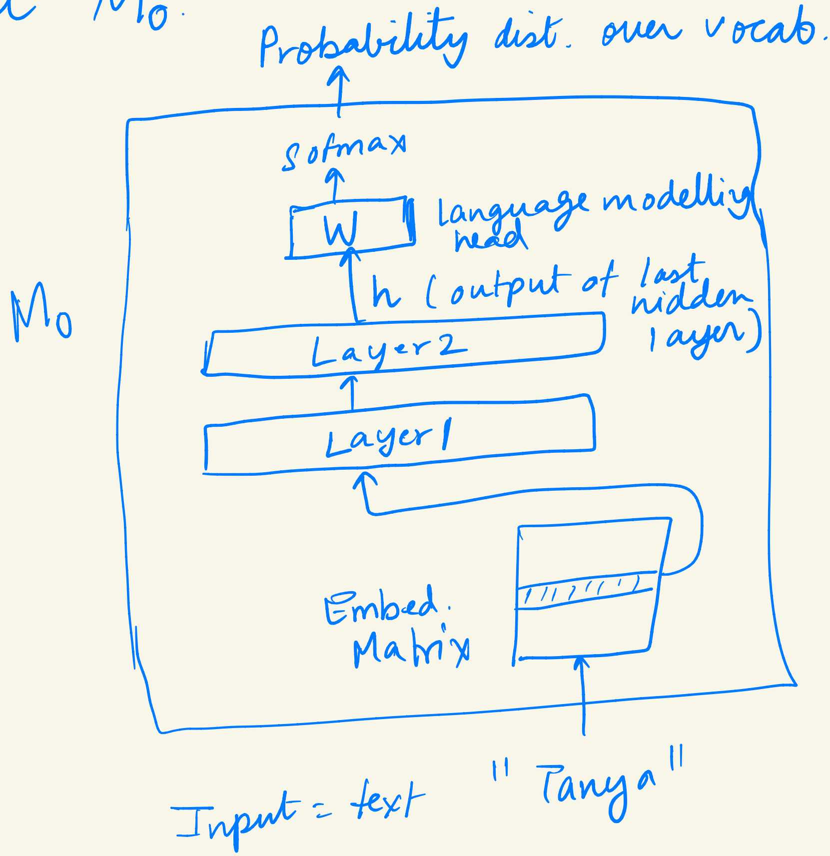

Question from slides: Fine-tune this model for Question-Answering. i.e, given a question q, the fine-tuned model must generate a.

1. Data — $\{(q,a)\}_{i=1}^{N}$ to fine-tune on.

2

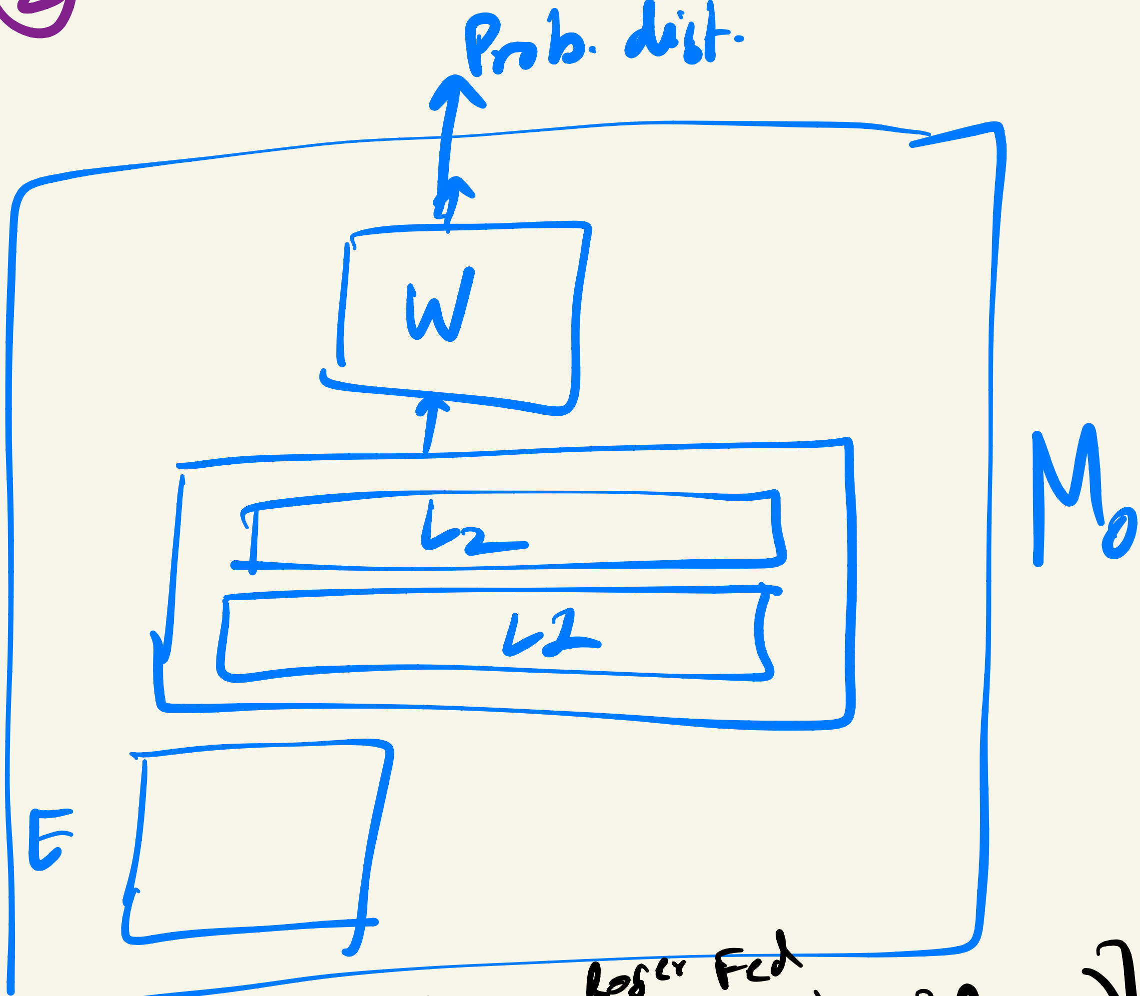

$$-\log \left[P(\text{Roger} \mid q?) \times P(\text{Federer} \mid q? \text{ Roger}) \right]$$

Roger Fed

minize loss

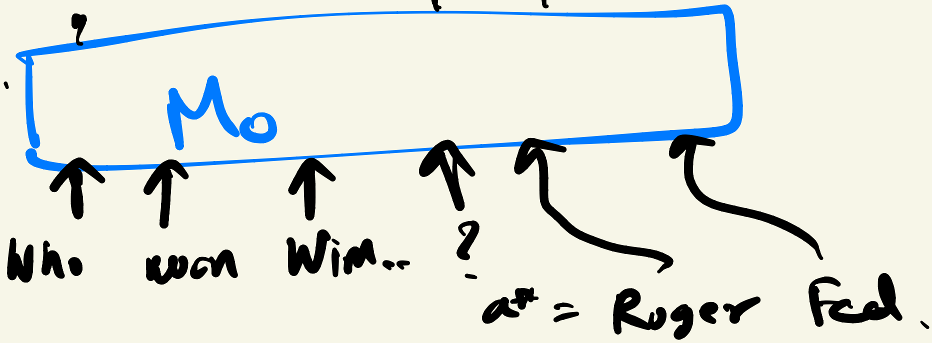

Usually predict a <EOS> as the last token

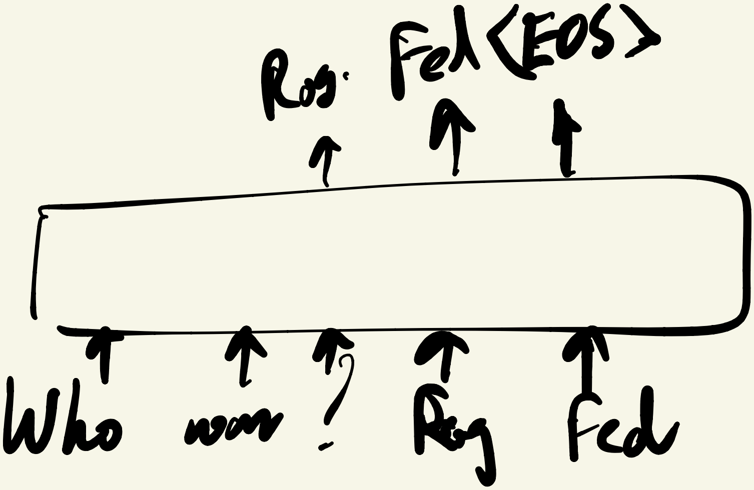

Pre-trained M_0

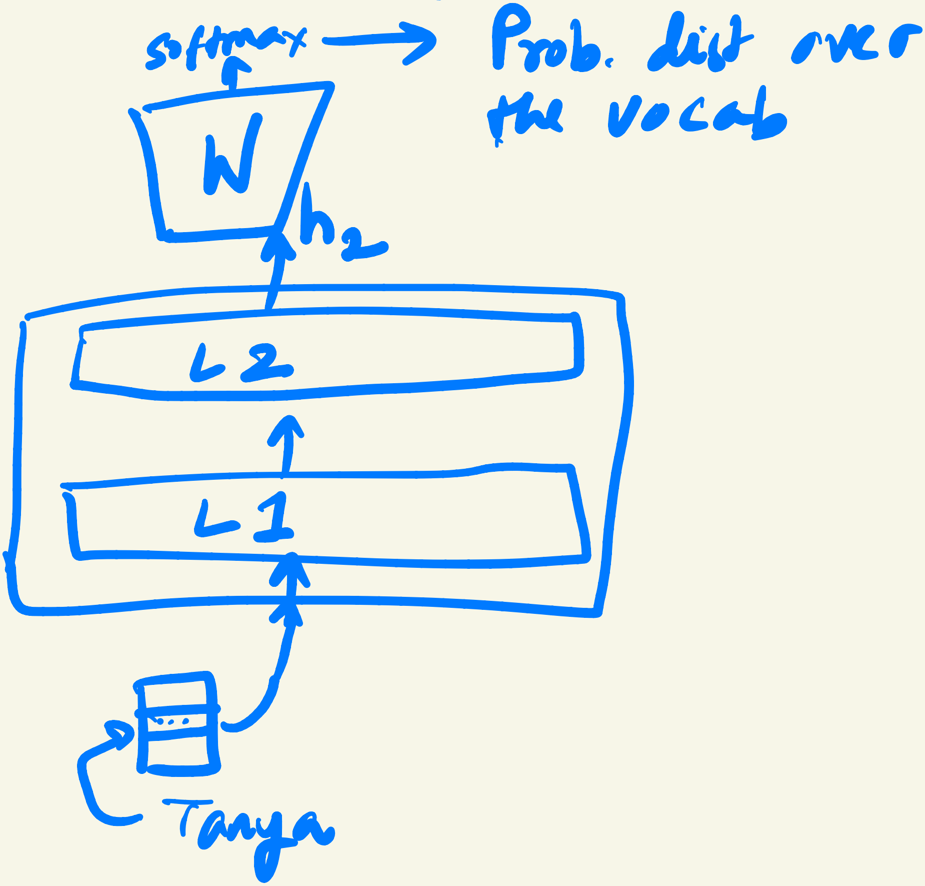

Fine-tune this model for sentiment classification.
given x, predict $y \in \{0, 1, 2\}$

① Collect data $\{(x, y)\}_{i=1}^{N}$

① Get rid of W from pre-training, use a diff head

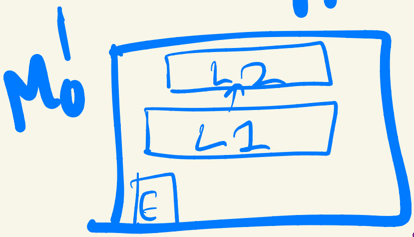

$$Loss = -\log P(y^*|x)$$

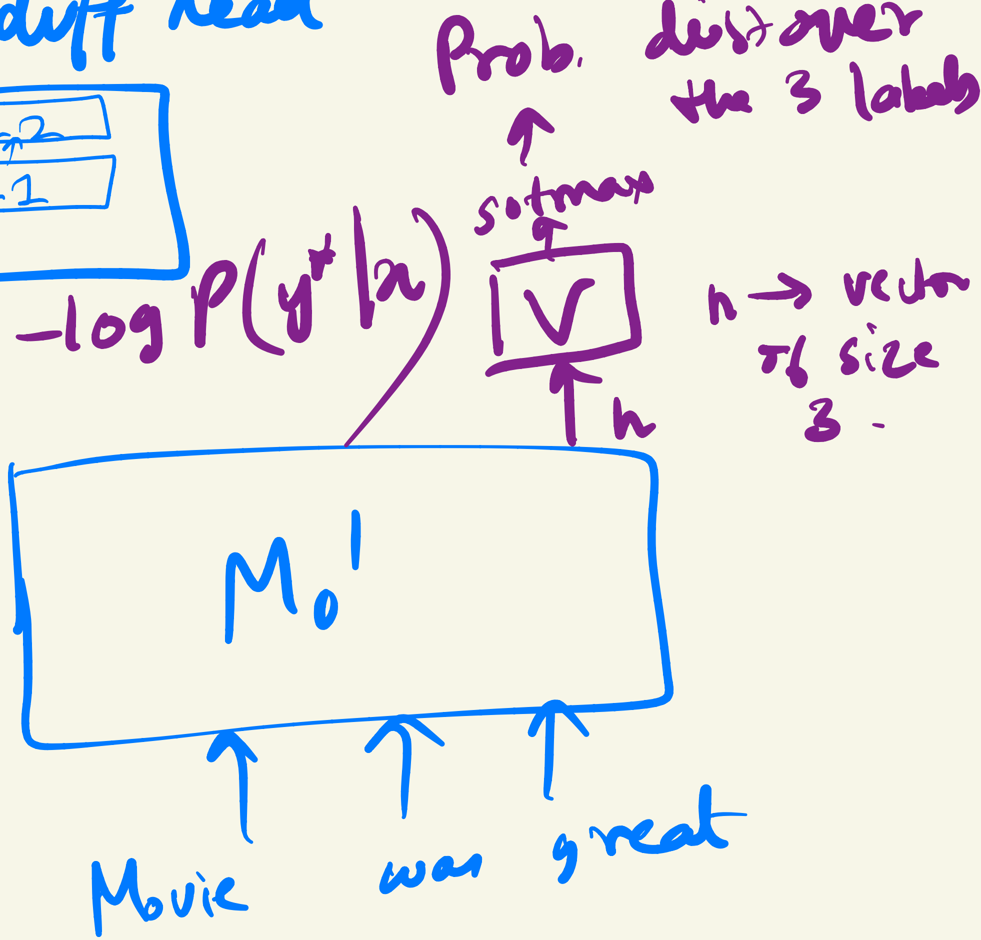

② Use the same LM head W

Prob dist over entire voca

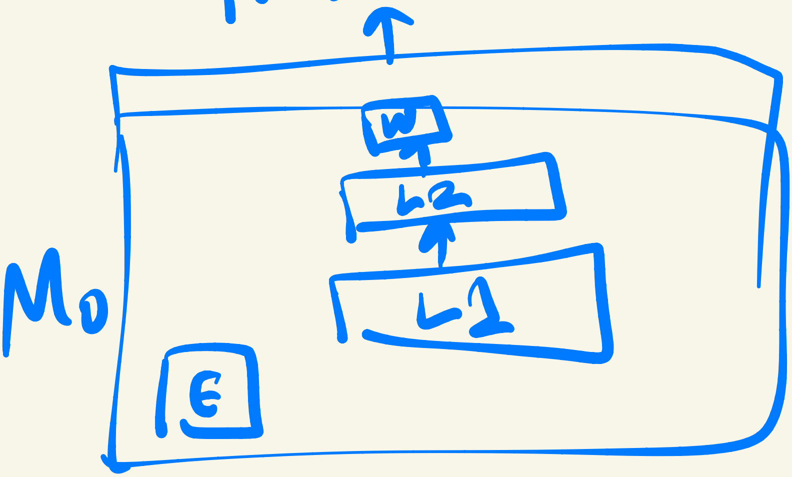

token in
9

0 → negative
1 → pos.
2 → neutr

neg | pos | neutr

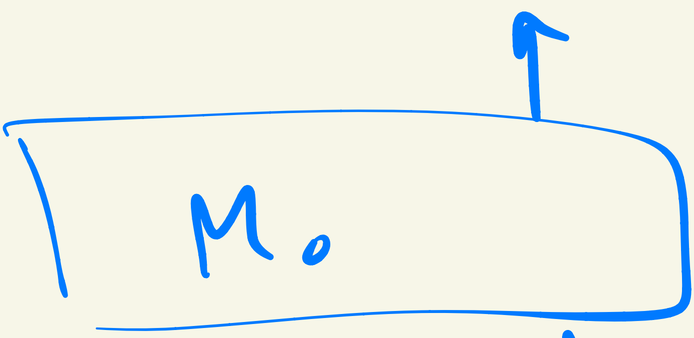

Movie was great

GOAL | Collect data (x, y^*) for inst. fine-tuning

(1) How do we get x?

(2) Given x, how do we get y^*?

---

(2) Distill from a Stronger Model (GPT-4)

$X = \{x\}$

Goal — Generate y^* for x

$$x \rightarrow \boxed{\text{GPT4}} \rightarrow y$$

This is y^*

$$(x, y^*)$$

Train your smaller model on data distilled from GPT4

DISTILLATION

① How do we get $X = \{x\}$

Use GPT-4

<10 seed instructions>

Input: Q — — —

Output:

I:

GPT —

Train RM w/ BT

$\{\underbrace{x, y_0, y_1}_{\text{Input}}, \underbrace{0|1}_{\text{Output}}\}$

0 mean $y_0 > y_1$

1 " $y_1 > y_0$

$\boxed{x, \hat{y} \longrightarrow r}$ — RM

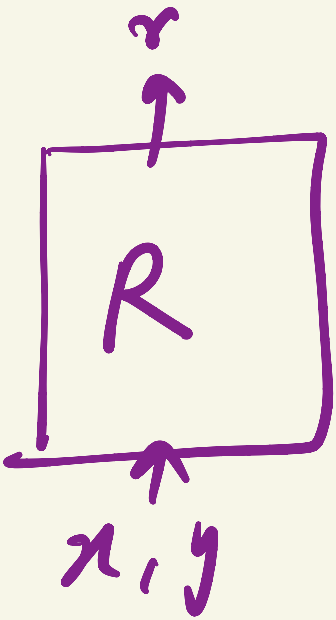

At train time

$\max \log P(y_0 > y_1)$

true label = 0

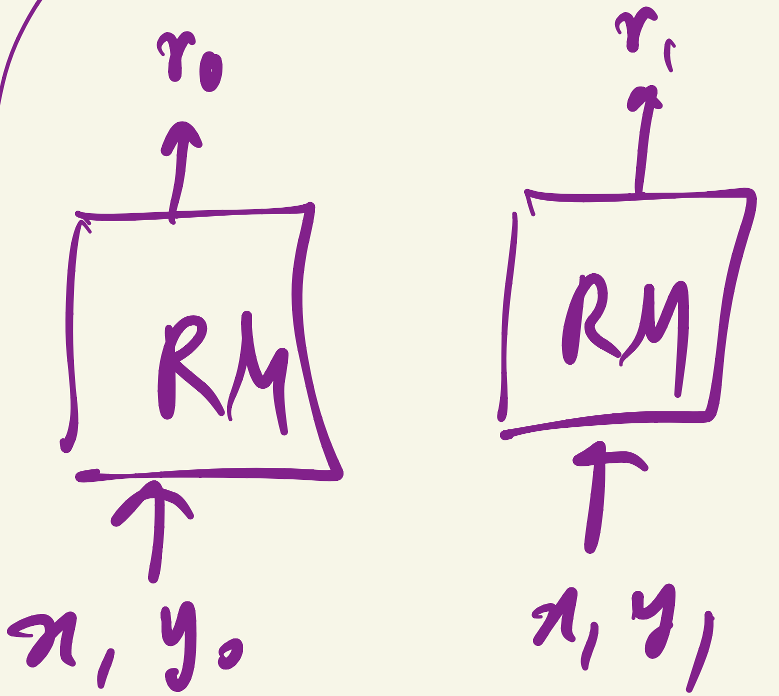

$= \log \sigma(\underline{r_0} - \underline{r_1})$

Goal RM

$(x, y) \longrightarrow r$

$y_+ > y_1$

Who won the WC? India Aust

Why do adults roll off the bed? O_1 O_2

RM_θ is a transmodel

$R_\phi(x, \underline{\text{India}}) >$

$R(x, \text{Aust})$

$$R_\phi(x_2, o_1) > R(x_2, o_2)$$

$$P(\text{India} > \text{Austr.} \mid x)$$

$$= \sigma\left(R_\phi(x, \text{India}) - R_\phi(x, \text{Aus})\right)$$