Lec 14 Combine Multiple Models

Good for getting slightly better acc

- Multiple architecture
- .. initialisation
- .. fine tuning

Model Ensembling

M1 → $\hat{y}_1$
M2 → $\hat{y}_2$
→ combine → result

- Reduce bias
- Models may have seen different data

* Errors tend to be uncorrelated btwn models, ensembling can even them out

→ Can also ensemble across checkpoints

Ways to combine models:
(Doesn't need many data to train; Can be context dependant ↓)

1. Linear interpolation btwn model probs
   - interpolation coefficient can be constant or learnt
   - acts like logical OR btwn the two models
   - handles 0 prob
2. Log linear interpolation — on log probs then renormalise (softmax)
   - likewise, can be constant or learnt coefficients
   - acts like logical AND — high probs if all models high prob
   - allows negative coefficient — some model serve as negative evidence
     - e.g. MT model + α(domain LM) − β(out of domain LM)
       LM + α(nontoxic LM) − β(toxic LM)

→ At test, drop out n times then combine
→ Bagging — resample dataset and train

Efficient multi model

cost ∝ amount of model

▷ Param averaging — average params of multiple models
   - needs same archi & shapes (obviously)
   - NNs have permutation invariance, so need same init.
   → Average together checkpoints (like last 5 of them)
   → Merge fine-tuned models

Model Soups paper
▷ Uniform averaging
▷ Greedy averaging — merge model if it improves
* Averaging perf correlates with ensembling perf usually

▷ Task vectors
   - $v_{task} = \theta_{fine\ tuned} - \theta_{original}$
   - vector arithmetics to change task
▷ TIES to resolve conflict btwn multiple θ fine-tuned

→ Can use mergekit to do all these

▷ Ensemble Distillation
   - make student model match the ensemble

▷ Sparse Mixture of Experts
   - make use of 0·[matrix] = [0]
   - nvidia's cuSPARSE
   ▷ Sparsely Gated Mix of Experts Layer

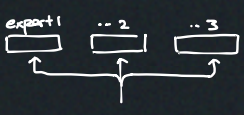

← gate these experts, zero out some of them

Pipeline Systems

E2E can be hard: - Data availability
- Interpretability

Cascading $X \rightarrow [M_1] \rightarrow \hat{Y}_1 \rightarrow [M_2] \rightarrow \hat{Y}$

Stacking $X \rightarrow [M_1] \rightarrow \hat{Y}_1 \rightarrow [M_2] \rightarrow \hat{Y}$ (with X also fed to M_2)

Iterative refinement