Machine Learning: Hot Topics

Hao Tang

March 27, 2024

What are representations?

What are representations?

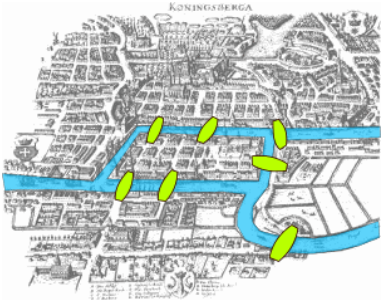

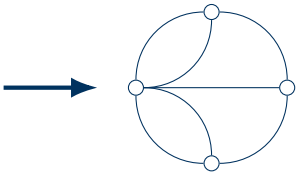

What are representations?

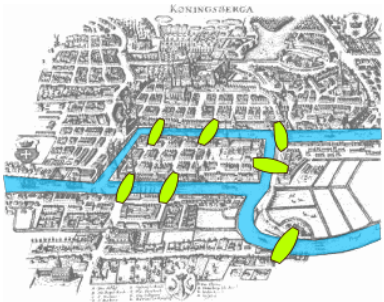

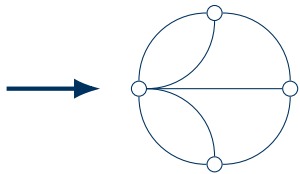

Representations are data structures for solving a particular task.

Representations inside a neural network

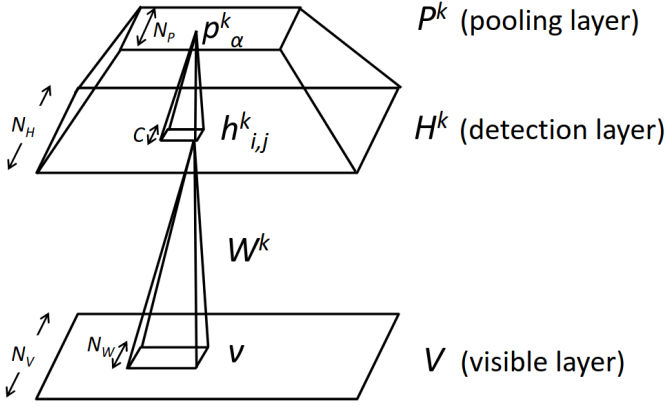

Image credit: Lee *et al.*, 2009

Representations inside a neural network

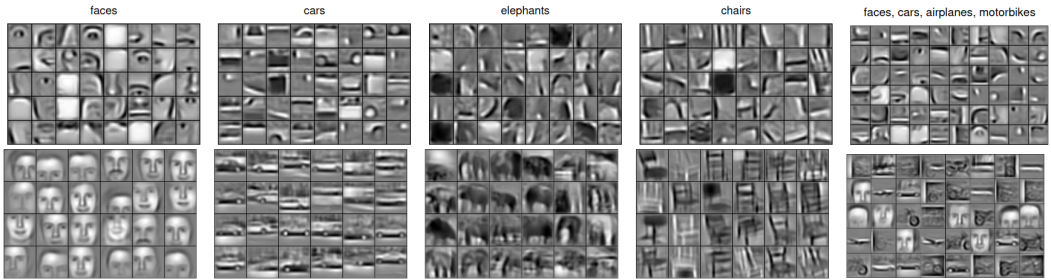

Image credit: Lee *et al.*, 2009

Supervised pre-training

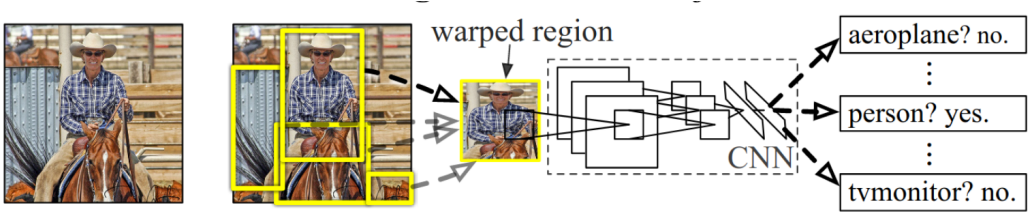

Image credit: Girshik *et al.*, 2014

Supervised pre-training

We show that it is highly effective to pre-train the network—with supervision—for a auxiliary task with abundant data (image classification) and then to fine-tune the network for the target task where data is scarce (detection). We conjecture that the "supervised pre-training/domain-specific fine-tuning" paradigm will be highly effective for a variety of data-scarce vision problems.

Supervised pre-training

labeled
data

Supervised pre-training

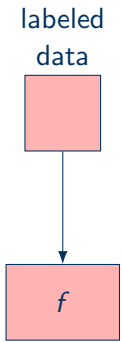

Supervised pre-training

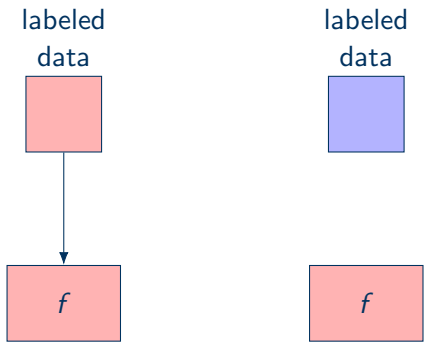

Supervised pre-training

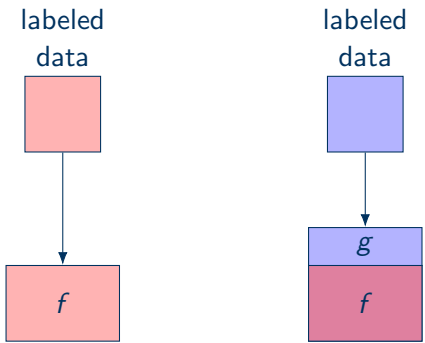

Supervised pre-training

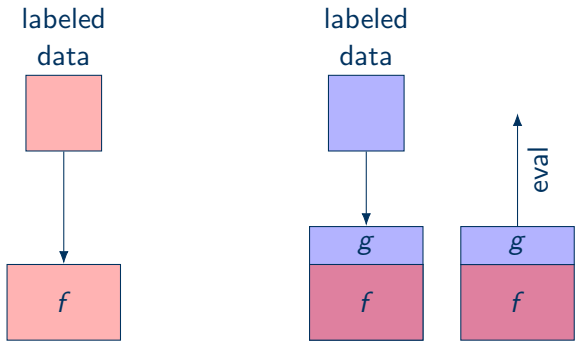

Supervised pre-training

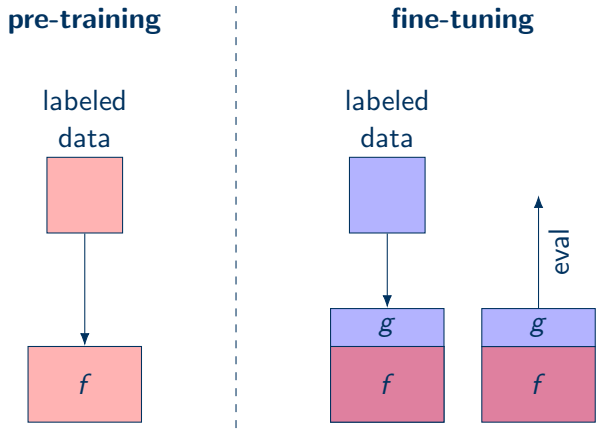

Self-supervised pre-training

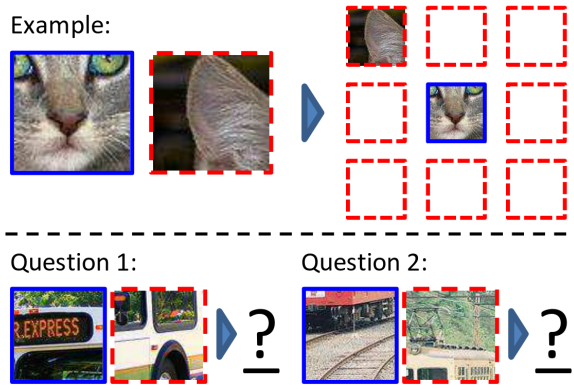

Image credit: Doersch *et al.*, 2015

Self-supervised pre-training

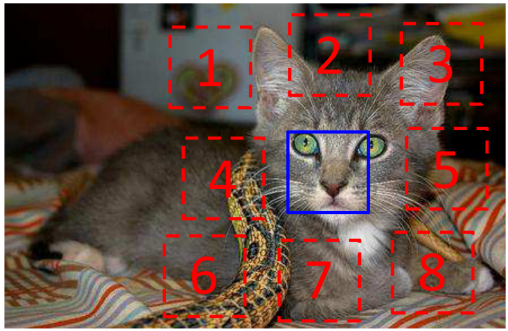

Image credit: Doersch *et al.*, 2015

Self-supervised pre-training

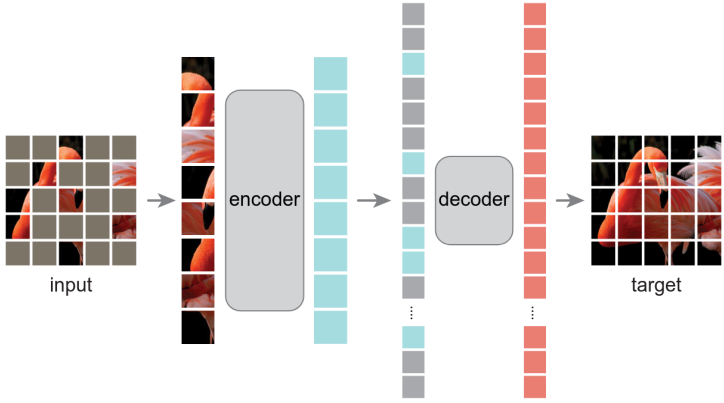

Image credit: He *et al.*, 2022

Self-supervised pre-training

Vision
(Doersch *et al.*, 2015)
BYOL (Grill *et al.*, 2020)
SimCLR (Chen *et al.*, 2020)
MAE (He *et al.*, 2022)

Text
word2vec (Mikolov *et al.*, 2013)
ELMo (Peters *et al.*, 2018)
BERT (Devlin *et al.*, 2018)

Speech
CPC (van den Oord *et al.*, 2018)
APC (Chung *et al.*, 2010)
wav2vec 2.0 (Baevski *et al.*, 2020)
HuBERT (Hsu *et al.*, 2021)

Self-supervised pre-training

unlabeled
data

f

Self-supervised pre-training

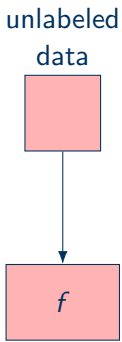

Self-supervised pre-training

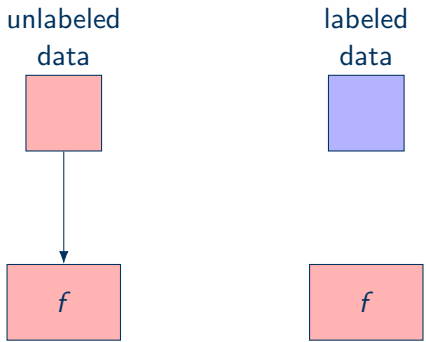

Self-supervised pre-training

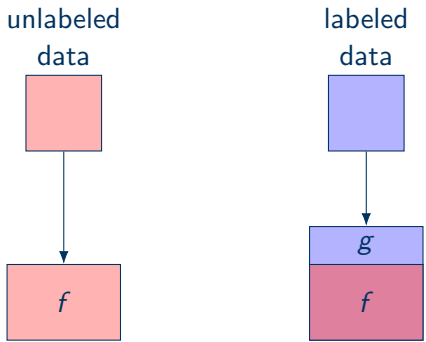

Self-supervised pre-training

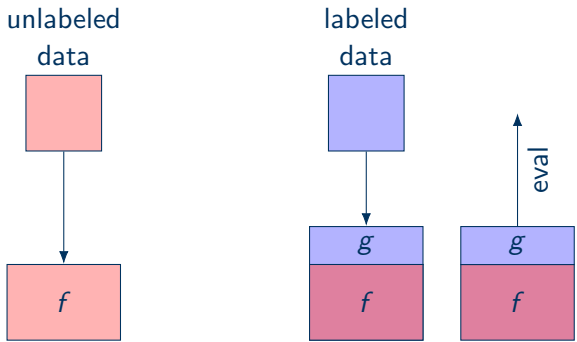

Self-supervised pre-training

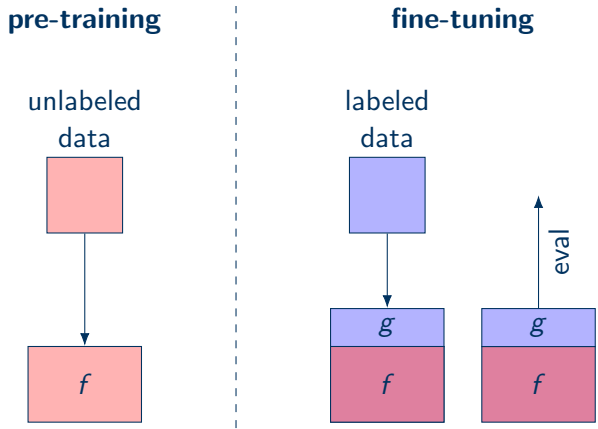

Probing

Is this a face?

Probing

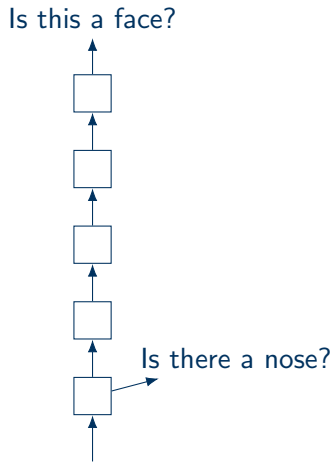

Probing

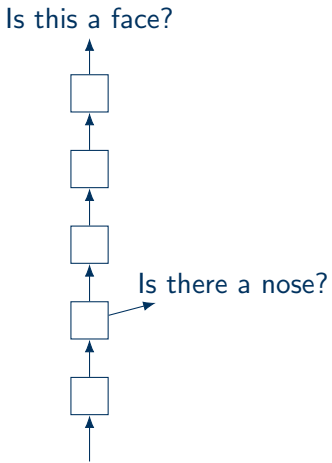

Probing

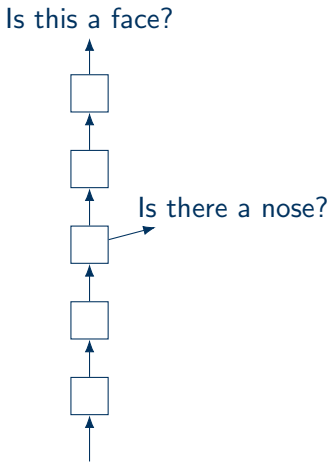

Probing

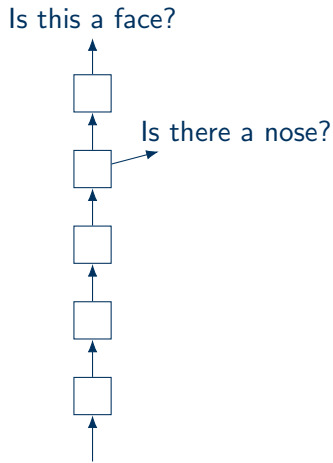

Probing

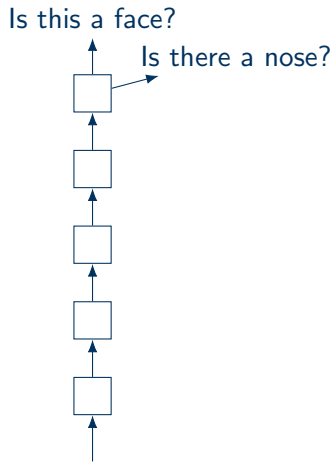

Probing

Is this a car?

Probing

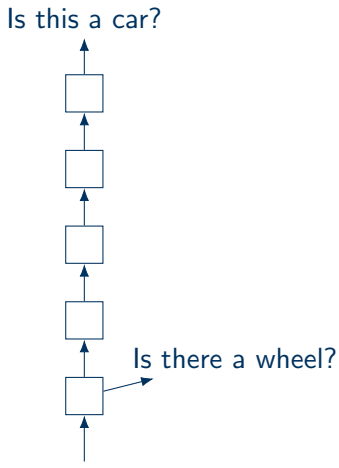

Probing

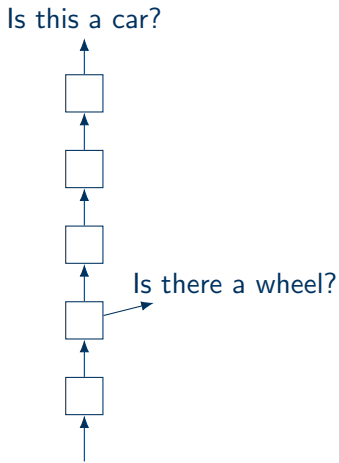

Probing

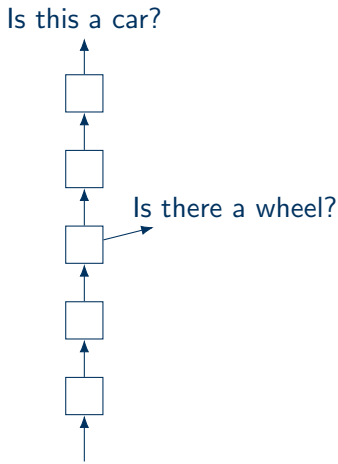

Probing

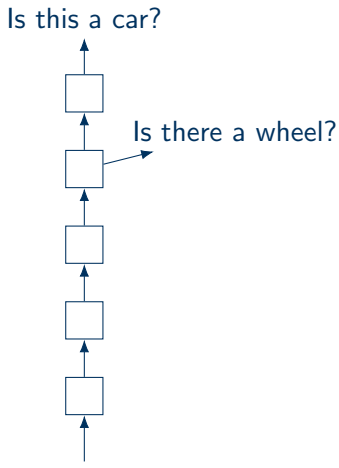

Probing

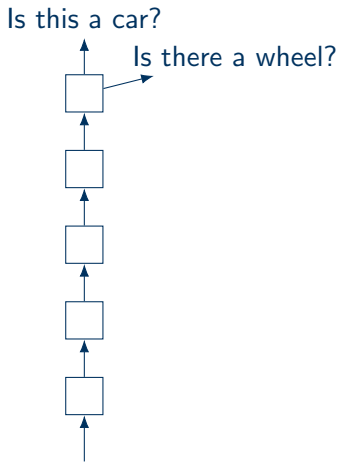

Probing

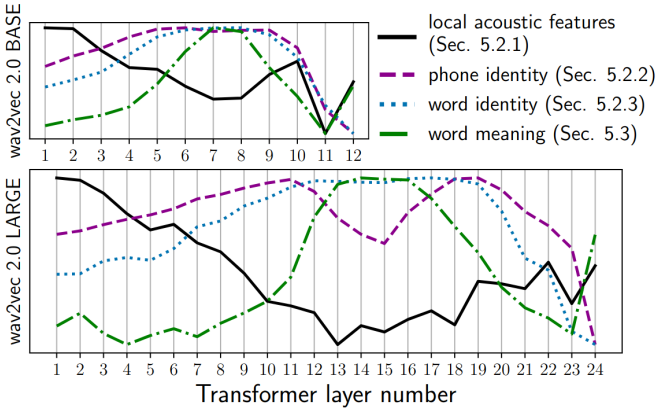

Image credit: Pasad *et al.*, 2021

Probing

25 August 2003 League of Extraordinary Gentlemen: Sean Connery is one of the all time greats and I have been a fan of his since the 1950's. I went to this movie because Sean Connery was the main actor. I had not read reviews or had any prior knowledge of the movie. The movie surprised me quite a bit. The scenery and sights were spectacular, but the plot was unreal to the point of being ridiculous. In my mind this was not one of his better movies it could be the worst. Why he chose to be in this movie is a mystery. For me, going to this movie was a waste of my time. I will continue to go to his movies and add his movies to my video collection. But I can't see wasting money to put this movie in my collection

I found this to be a charming adaptation, very lively and full of fun. With the exception of a couple of major errors, the cast is wonderful. I have to echo some of the earlier comments -- Chynna Phillips is horribly miscast as a teenager. At 27, she's just too old (and, yes, it DOES show), and lacks the singing "chops" for Broadway-style music. Vanessa Williams is a decent-enough singer and, for a non-dancer, she's adequate. However, she is NOT Latina, and her character definitely is. She's also very STRIDENT throughout, which gets tiresome. The girls of Sweet Apple's Conrad Birdie fan club really sparkle -- with special kudos to Brigitta Dau and Chiara Zanni. I also enjoyed Tyne Daly's performance, though I'm not generally a fan of her work. Finally, the dancing Shriners are a riot, especially the dorky three in the bar. The movie is suitable for the whole family, and I highly recommend it.

Image credit: Radford *et al.*, 2017

Probing

- Probing accuracy reflects the accessibility of information.
- If the probing accuracy is low, it does not necessarily mean the information is absent.
- The capacity of probing classifiers matters.
- Strong classifiers could potentially access more information than weak classifiers.
- However, we normally care less about information that is hard to access.
- The input itself is a representation in which much information is hard to access.

Self-supervised pre-training

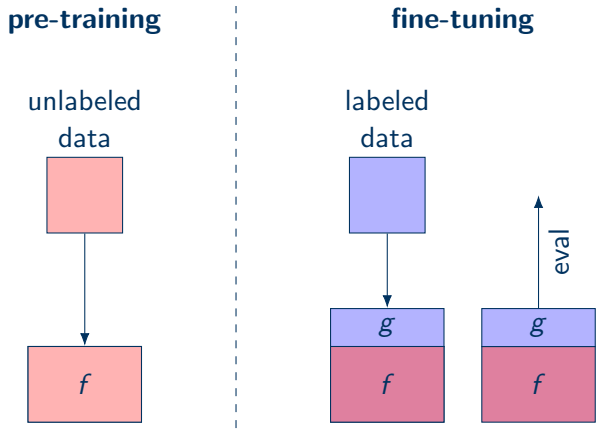

Fine-tuning

Fine-tuning

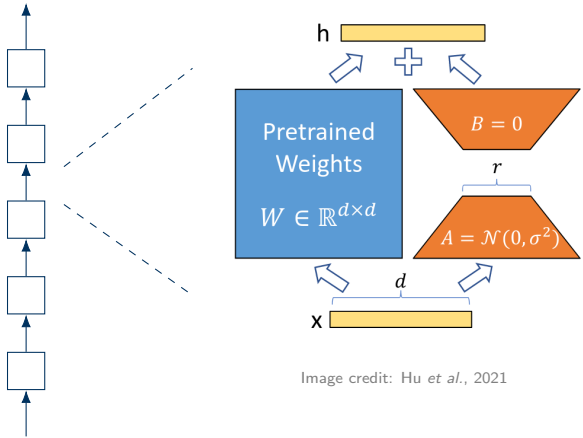

Image credit: Hu *et al.*, 2021

LoRA: Low-rank adaptation

Model&Method	# Trainable Parameters	WikiSQL Acc. (%)	MNLI-m Acc. (%)	SAMSum R1/R2/RL
GPT-3 (FT)	175,255.8M	**73.8**	89.5	52.0/28.0/44.5
GPT-3 (BitFit)	14.2M	71.3	91.0	51.3/27.4/43.5
GPT-3 (PreEmbed)	3.2M	63.1	88.6	48.3/24.2/40.5
GPT-3 (PreLayer)	20.2M	70.1	89.5	50.8/27.3/43.5
GPT-3 ($\text{Adapter}^{\text{H}}$)	7.1M	71.9	89.8	53.0/28.9/44.8
GPT-3 ($\text{Adapter}^{\text{H}}$)	40.1M	73.2	**91.5**	53.2/29.0/45.1
GPT-3 (LoRA)	4.7M	73.4	**91.7**	**53.8/29.8/45.9**
GPT-3 (LoRA)	37.7M	**74.0**	**91.6**	53.4/29.2/45.1

Image credit: Hu *et al.*, 2021

Summary

- We are not even close to understanding how neural networks solve tasks.
- There are lots of exciting open questions!
- Simple methods (learned in this class) can go really far.